

Research



Cite this article: Andrade J, Duggan J. 2023 Anchoring the mean generation time in the SEIR to mitigate biases in \mathcal{R}_0 estimates due to uncertainty in the distribution of the epidemiological delays. *R. Soc. Open Sci.* **10**: 230515.

<https://doi.org/10.1098/rsos.230515>

Received: 19 April 2023

Accepted: 13 July 2023

Subject Category:

Mathematics

Subject Areas:

computational biology/mathematical modelling/
statistics

Keywords:

SEIR, basic reproduction number, distributions of the latent and infectious period, mean generation time, Stan, inference

Author for correspondence:

Jair Andrade

e-mail: j.andradeortiz1@universityofgalway.ie

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.6753860>.

Anchoring the mean generation time in the SEIR to mitigate biases in \mathcal{R}_0 estimates due to uncertainty in the distribution of the epidemiological delays

Jair Andrade¹ and Jim Duggan²

¹Data Science Institute and School of Computer Science, and ²Insight Centre for Data Analytics and School of Computer Science, University of Galway, Galway, Republic of Ireland

JA, 0000-0002-1412-7868; JD, 0000-0002-7507-8617

The basic reproduction number, \mathcal{R}_0 , is of paramount importance in the study of infectious disease dynamics. Primarily, \mathcal{R}_0 serves as an indicator of the transmission potential of an emerging infectious disease and the effort required to control the invading pathogen. However, its estimates from compartmental models are strongly conditioned by assumptions in the model structure, such as the distributions of the latent and infectious periods (epidemiological delays). To further complicate matters, models with dissimilar delay structures produce equivalent incidence dynamics. Following a simulation study, we reveal that the nature of such equivalency stems from a linear relationship between \mathcal{R}_0 and the mean generation time, along with adjustments to other parameters in the model. Leveraging this knowledge, we propose and successfully test an alternative parametrization of the SEIR model that produces accurate \mathcal{R}_0 estimates regardless of the distribution of the epidemiological delays, at the expense of biases in other quantities deemed of lesser importance. We further explore this approach's robustness by testing various transmissibility levels, generation times and data fidelity (overdispersion). Finally, we apply the proposed approach to data from the 1918 influenza pandemic. We anticipate that this work will mitigate biases in estimating \mathcal{R}_0 .

1. Introduction

The analysis of any infectious disease's dynamics will inevitably lead to the *basic reproduction number* (\mathcal{R}_0). Initially developed in

the study of demographics [1], this quantity has been interpreted in the epidemiological context as the average number of secondary infections arising from the introduction of one infected individual into a totally susceptible population [2]. The usefulness and importance of \mathcal{R}_0 lie primarily in its *threshold phenomenon* [3]. That is, a pathogen can invade a totally susceptible population only if $\mathcal{R}_0 > 1$ [4]. Furthermore, the magnitude of \mathcal{R}_0 gauges the transmission potential of an emerging infectious disease [3] and the effort required to control the invading pathogen [5]. Thus, accurate estimation of \mathcal{R}_0 is crucial for understanding and managing infectious diseases.

Another reason for the popularity of \mathcal{R}_0 is that one can estimate it from epidemiological data [6] using a number of methods. For diseases that allow the assumption of endemic equilibrium and homogeneous mixing, one can follow Mollison's method [7] or Dietz's approach [8]. The former requires prevalence data, whereas the latter leverages readily available information such as age at infection and average life expectancy. On the other hand, if an infection leads to either immunity or death in a closed population, seroprevalence studies can inform the fraction of the population that acquired the disease during an epidemic, i.e. the *final epidemic size*. In their seminal paper, Kermack & McKendrick [9,10] formulated a relationship between the final epidemic size and \mathcal{R}_0 , from which the latter can be calculated. Unlike the previous methods, which require the epidemic to reach a steady state, \mathcal{R}_0 may be determined from the intrinsic growth rate of the infected population [3,5] using incidence data of the early stages of the epidemic, as long as the growth of new cases exhibits pure exponential behaviour. Alternatively, we can employ the entire report of daily case notifications if \mathcal{R}_0 is formulated as a function of a compartmental model's parameters [11–13]. These models can be stochastic [14], semi-deterministic [15,16] or deterministic [9,17].

These compartmental models are said to be *mechanistic* [18], namely, structures based on a scientific understanding of infectious disease dynamics [19]. The relevance of that mechanistic property lies in the role of the model. Rather than being a merely mathematical artefact to produce a desired output, the model also embeds a dynamic hypothesis of the underlying process that generates the observed data. Hence, the parameters, states and interactions that comprise a particular formulation represent their counterparts in the real world. If the model accurately captures the properties of the actual phenomenon, finding an adequate *configuration* (assign values to parameters) should yield a behaviour over time of infections that resembles the observed trajectory. The values of such parameters can be obtained from individual-level observations [4] or via statistical inference [20–22], a process also known as *trajectory matching* or *model fitting*.

Furthermore, matching simulated and observed behaviour can be regarded as a validation test on the dynamic hypothesis that links structure to behaviour [23]. Nevertheless, one should understand this validation step as a falsification test [24]. That is, if the model fails to reproduce the observed behaviour, it can certainly be rejected. On the contrary, obtaining an accurate match (or fit) does not immediately validate the dynamic hypothesis inasmuch as there may be other competing hypotheses that fit the data equally well. Indeed, this circumstance impacts the estimation of \mathcal{R}_0 from compartmental models (and the intrinsic growth rate method), where different assumptions can yield accurate fits [25]. However, estimates vary according to the specific assumptions embedded in each fitting model [3,26].

For instance, the choice of the distributions of the latent and infectious periods (epidemiological delays) in the deterministic susceptible–exposed–infectious–recovered (*SEIR*) framework plays an essential role in the inference of \mathcal{R}_0 [25]. Briefly put, misspecifying the structure of such delays leads to biases in the estimates. That is, a systematic difference between true and estimated parameters. Although there are techniques [27,28] to construct models with realistic distributions, modellers do not know exactly which distribution to incorporate in their formulation. In the view of this drawback, Wearing *et al.* [25] fitted various *SEIR* models (with different delay distributions) to a single incidence dataset to select the best structure based on a goodness-of-fit measure. Nevertheless, the results appear inconclusive. Notwithstanding that Krylova & Earn [29] assume their validity, no further research establishes the reliability of such an approach. This assessment immediately warrants the need for the work presented here: a systematic study oriented to determine whether it is possible to infer \mathcal{R}_0 accurately from *SEIR* models fitted to incidence data in light of the uncertainty in the distributions of the epidemiological delays. We describe the steps of this study in the sections below. All the analysis is performed in R. The code is freely available at <https://github.com/jandraor/delays>.

2. Data generating process

2.1. The system (latent) component

In order to undertake a systematic study, experimenters must have access to a sizeable set of observations. In this case, multiple time series of daily case notifications of a particular disease under

various conditions. Equally important, such conditions need to be known *a priori*. To meet these conditions, we leverage the mechanistic property of compartmental models and employ the *SEIR* framework as a synthetic data generator [30,31]. This framework has been widely applied to studying various infectious diseases, such as measles [29,32,33], COVID-19 [16,34,35] and influenza [12,17,22,36–38]. In this work, we restrict our attention to the simplest version of this family of models. The rationale for this decision is straightforward; conceptual models entail efficiency inasmuch as they facilitate the understanding and identification of the underlying causes of a particular result. Moreover, it is often the case that principles that stem from basic models apply to more elaborated extensions.

$$\left. \begin{aligned} \dot{S}_t &= -\frac{\beta S_t \sum_{k=1}^j I_t^k}{N} \\ \dot{E}_t^1 &= \frac{\beta S_t \sum_{k=1}^j I_t^k}{N} - i\sigma E_t^1 \\ \dot{E}_t^2 &= i\sigma E_t^1 - i\sigma E_t^2 \\ &\vdots \\ \dot{E}_t^i &= i\sigma E_t^{i-1} - i\sigma E_t^i \\ \dot{I}_t^1 &= i\sigma E_t^i - j\gamma I_t^1 \\ \dot{I}_t^2 &= j\gamma I_t^1 - j\gamma I_t^2 \\ &\vdots \\ \dot{I}_t^j &= j\gamma I_t^{j-1} - j\gamma I_t^j \\ \text{and} \quad \dot{R}_t &= j\gamma I_t^j. \end{aligned} \right\} \quad (2.1)$$

Specifically, the *SEIR* (equation (2.1)) stratifies individuals as susceptible (S_t), exposed (E_t), infectious (I_t) and recovered (R_t) and describes the transitions between states ($S_t \rightarrow E_t \rightarrow I_t \rightarrow R_t$) in terms of differential equations. Susceptible individuals acquire infection, $S_t \rightarrow E_t$, through contact with infectious individuals, where the number of contacts is independent of the population size (N). Formally, one refers to this assumption as the frequency-dependent (or mass action) transmission: $\beta S_t I_t / N$. Here, β corresponds to the effective contact rate or transmission parameter. The movement of individuals from the class E_t to class R_t is modelled using a well-known mathematical procedure [39] to achieve realistic distributions [40,41] of the time that individuals spend in states E_t and I_t , otherwise known as the latent and infectious periods, respectively. Such a procedure corresponds to the subdivision of a class into stages arranged in series. For instance, one can divide the exposed class into i stages. Newly infected individuals enter the first exposed stage, E_t^1 , pass through each in turn and become infectious upon leaving the i th stage (E_t^i). The progression between stages is assumed to occur at a constant *per capita* rate ($i\sigma$), leading to an exponential waiting time with mean $1/i\sigma$ in each stage [11]. This formulation implies that the lapse between infection and becoming infectious is described by the sum of i independent exponential random variables with equal rates, a convolution resulting in a gamma-distributed random variable [42]. Therefore, the subdivision of the exposed class into various stages is equivalent to formulating the latent period in terms of a gamma distribution with mean σ^{-1} and shape i . Similarly, one can divide the infectious class into j stages to formulate a gamma-distributed infectious period.

$$\mathfrak{R}_0 = \beta\gamma^{-1}. \quad (2.2)$$

Overall, we refer to equation (2.1) as the *SEⁱI^jR* framework. Notice that the standard *SEIR* corresponds to the *SE¹I¹R* instance. Moreover, as the parameter i increases, the distribution becomes more closely centred on its mean (tighter), to the extent that if $i \rightarrow \infty$, the variance is removed. That is, in the limit, all individuals have the same latent period. An equivalent argument applies to the infectious period. No less important, as indicated by Lloyd [11], irrespective of the values of i and j that the *SEⁱI^jR* may take, the basic reproduction number depends exclusively on the transmission rate and the mean

infectious period (equation (2.2)). Furthermore, it is noteworthy to mention that subdividing a class is a mathematical device that allows the incorporation of additional distributions in a system of differential equations, and the number of stages may not correspond to biological features of the infection process [32]. Lastly, we assume that the disease leads to permanent immunity and that the outbreak's time scale is much faster than the characteristic times for demographic processes (births and deaths), therefore their effects are not included. This last assumption implies that the population remains constant over the simulation period.

$$\text{and } \left. \begin{aligned} \dot{C}_t &= \rho i \sigma E_t^i \\ x_{t^*} &= C_{t^*+1} - C_{t^*}, \quad t^* \in \mathbb{N}_0. \end{aligned} \right\} \quad (2.3)$$

Subsequently, we define the link between the SE^iIR and incidence data (equation (2.3)). Based on the literature [15,16,22,36], we posit that incidence (\dot{C}) is proportional to the rate at which individuals become infectious ($E_t^i \rightarrow I_t^i$). Such proportional effect or reporting rate (ρ) stems from the fact that individuals experience various degrees of symptom severity [43]. In particular, individuals with low severity levels (asymptomatic and mild symptoms) may not seek health care attention, resulting in case reports that most likely miss a significant fraction of infected individuals. As opposed to the continuous nature of differential equation models, case notifications occur at discrete times. To reconcile this tension, we define the report of new cases (x_{t^*}) as the change in the total number of cases (C_t) in 1-day intervals.

Furthermore, we tailor the synthetic data generator towards influenza given that this virus causes unpredictable but recurring pandemics that can have significant global consequences [44]. As a matter of fact, there have been four influenza pandemics over the past 100 years, including the H1N1 pandemic in 1918, with an estimated 50 million deaths [45]. Adapting the SE^iIR framework to this choice involves the selection of plausible parameter values or ground truths [46]. For simplicity, we restrict the synthetic data generator to eight instances: $i = \{1, 3\} \times j = \{1, 2, 3, 4\}$. These instances share constants σ , γ , β , ρ and N , which are configured identically. In particular, we configure parameters σ and γ from the assumed values (1/2 for both) in the Cumberland case study [12,22]. Following this choice, we select a value of β that yields a basic reproduction number (2.5) within a plausible range (2–4) of pandemic influenza [47]. Regarding ρ , we choose a value (0.75) consistent with reported estimates in the literature [12,22]. The remaining constant, N , has only a scaling effect, and any particular value (10 000 in this case) does not alter the model dynamics provided that $N = S_0 + E_0 + I_0 + R_0$, where $E_0 = \sum_{k=1}^i E_0^k$ and $I_0 = \sum_{k=1}^j I_0^k$. In relation to initial conditions, we assume that a patient zero triggers the outbreak of a novel influenza pathogen. In mathematical terms, $S_0 = N - 1$ and $I_0^1 = 1$. The remaining initial conditions of the within-host profile are set to zero.

Having delimited the SE^iIR framework and configured its instances, we run simulations (figure 1) that illustrate the impact of the delay structure on the incidence dynamics. In agreement with the literature [4,25], note in figure 1 that if we fix the latent period distribution (i) and vary that for the infectious period (j), incidence reports that stem from more tightly distributed infectious periods (larger j) reach the incidence peak earlier and end more abruptly. This difference in behaviour over time occurs despite the fact that these instances share identical \mathfrak{R}_0 and equal average latent and infectious periods. On the other hand, if we fix the infectious period (compare two lines of the same colour across panels), decreasing the latent period's variance (increasing i from 1 to 3) produces the opposite effect. Namely, tighter latent period distributions (larger i) push forward the peak time and extend the outbreak's duration.

2.2. Measurement component

$$Y_{t^*} \sim Nbin(x_{t^*}, \phi). \quad (2.4)$$

Borrowing terminology from the state-space literature [18,48], one can frame the output produced by the SE^iIR framework as predictions obtained from a system or *latent* component. In practice, though, continuous and smooth predictions from ODE models differ from noisy and discrete incidence reports collected by public health surveillance. Moreover, given that a system component is merely a partial representation of a more complex reality, some elements are necessarily omitted. Consequently, it is required to equip the data generating process with a structure that accounts for the discrepancies between model prediction and actual data. We refer to this structure as the measurement component. In epidemiology, one can formulate the measurement of new infections via the *negative binomial* distribution, considering that this function does not tie the observation mean to the variance, offering

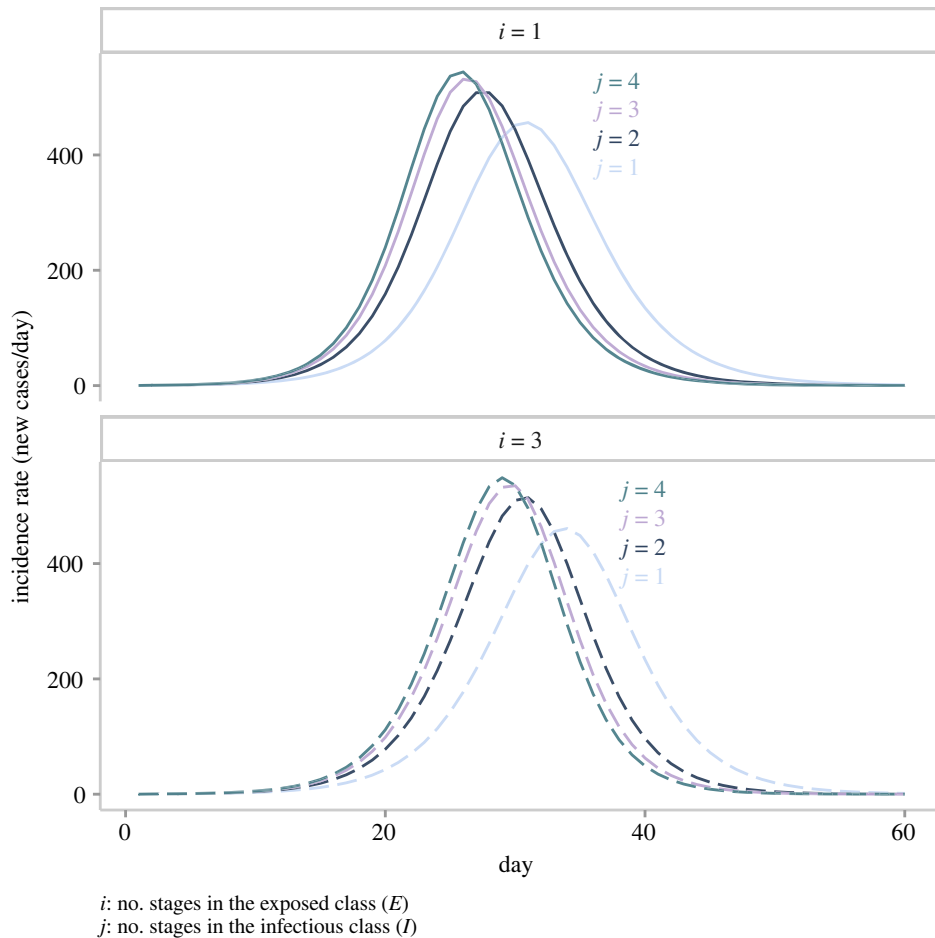


Figure 1. Incidence reports generated by various instances of the SE^iIR framework. In this plot, we present two distributions of the latent period and four distributions of the infectious period. The colour of a line corresponds to a particular value of j (infectious period distribution). Solid lines indicate that the incidence report stems from an SEIR model with an exponentially distributed latent period ($i = 1$). Dashed lines indicate that the incidence report stems from an SEIR model with a gamma-distributed latent period ($i = 3$).

the flexibility to account for overdispersion [19]. Accordingly, we define the observation of new cases (Y_t) in terms of a negative binomial distribution (equation (2.4)) specified by location (mean) and diffusion parameters. The former corresponds to the predicted incidence by the system component (x_t), whereas the latter (ϕ) modulates the concentration of measurements. Note that the inverse of the concentration parameter (ϕ^{-1}) represents overdispersion inasmuch as an increase in its magnitude leads to greater diffusion in the data.

Defining a measurement component completes the formulation of the data generating process. Consequently, we draw samples from equation (2.4) using statistical simulation (*rnbinom* in R). For each SE^iIR instance, we generate 40 *noisy* time series. We perform this process for two levels (*high* and *low*) of data fidelity, a feature measured by ϕ^{-1} . High-fidelity data ($\phi^{-1} = 0$) imply that the measurement component applies only a slight distortion on the original signal (incidence). Notice that this configuration of the negative binomial (with no overdispersion) is equivalent to the *Poisson* distribution. Conversely, a positive value (overdispersion) of ϕ^{-1} (such as $1/3$) distorts the original signal to such an extent that one cannot easily discern the underlying incidence dynamics (low-fidelity data). We generated a total of 320 incidence reports, of which figure 2 presents a sample of four representative reports. The reader can find the complete details in the electronic supplementary material, S1. To facilitate the communication of results, we introduce the notation D^j , which indicates the origin of a given set of time series. For example, D^{14} indicates that the observed incidence was obtained from the SE^1I^4R instance.

3. Inference

The synthetic incidence reports described in the previous section allow us to assess the performance of various candidate models in recovering ground truths, particularly R_0 , our quantity of interest.

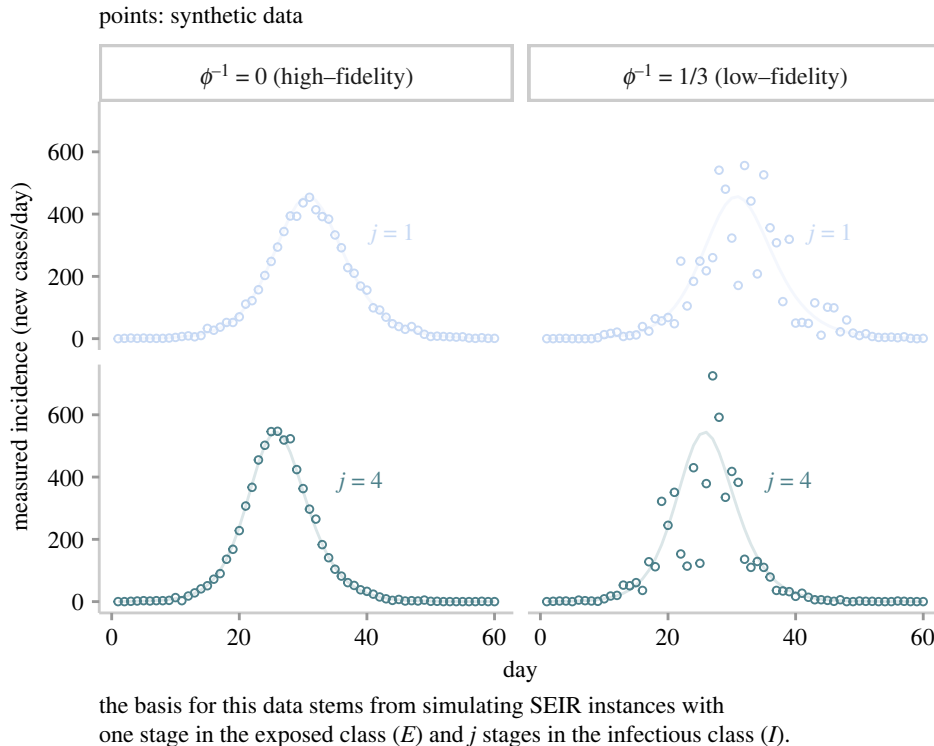


Figure 2. Sample of synthetic data. This plot shows four representative incidence reports (dots) obtained from the simulation of two SE^1I^jR instances (lines). To obtain each report, we sample from the negative binomial distribution.

Specifically, we fit model candidates to incidence data following a *Bayesian* approach [22,49]. That is, each candidate's unknown parameters are treated as random variables, which describe the knowledge (or uncertainty) about their actual values [50], expressed in terms of a probability distribution. This distribution is updated in light of new information summarized by a likelihood function. This function evaluates the compatibility between a given incidence report and multiple configurations of a model candidate [51]. Such updating process yields the target or posterior distribution, an information device whereby we derive answers for our inferential questions. We approximate the posterior distribution via sampling using Hamiltonian Monte Carlo or HMC [52], an algorithm successfully employed to perform statistical inference from epidemiological models [16,17,22,53,54]. This algorithm is provided by the statistical package *Stan* [55].

3.1. Three unknowns (traditional): β , ρ , I_0

For simplicity, we initially restrict the inference analysis to D^{1j} high-fidelity observations. To fit each incidence report, we postulate four instances, $j = \{1, 2, 3, 4\}$, from the SE^1I^jR framework, which share identical mean latent and infectious periods. We refer to the approach of fixing the means of the epidemiological delays to values obtained from the literature, regardless of their distribution, as the *traditional* parametrization. Moreover, it is assumed that the measurement component is fully known. Consequently, discrepancies between estimated and actual values are ascribed to misspecification in the infectious period distribution. To avoid confusion between the origin of data and the fitting model, we denote the latter as M^j . As a consequence, this design requires the estimation of 320 posterior distributions. Given this process's computational burden, we limit the number of random variables in each model to three: the transmission rate (β), the reporting rate (ρ) and the initial number of infected individuals in stage one (I_0^1). The remaining parameters and initial conditions are considered to be known, i.e. they are fixed to their actual values. Based on this set-up, we fit each candidate to a given dataset using HMC sampling, with *four* Markov chains and *1000* iterations (plus *1000* for warm-up) each, checking for convergence and effective sample sizes. The complete set of results can be found in the electronic supplementary material, S2, §1.

The results presented in figure 3 replicate a finding previously reported in the literature [25,56]: the existence of a subtle yet fundamental interaction between the assumed model structure and estimated R_0 .

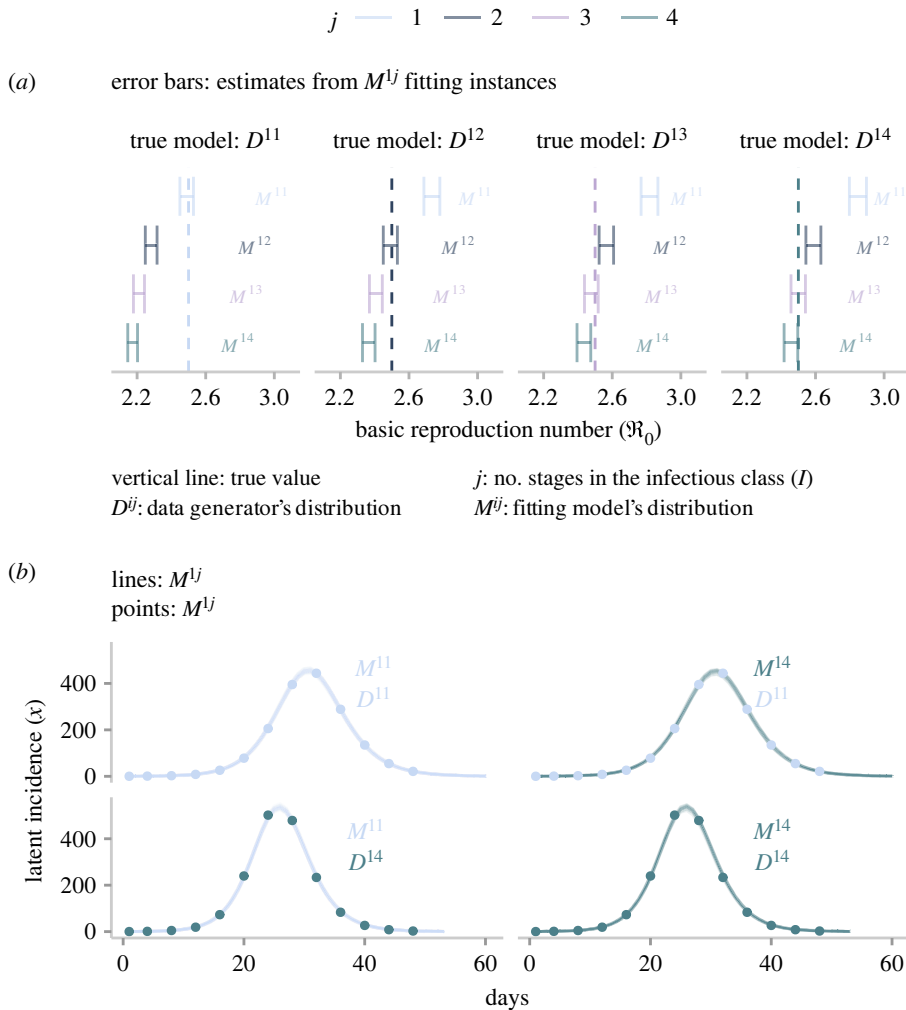


Figure 3. Inference results obtained from the three-unknown parametrization. This plot shows the results of fitting model candidates to incidence reports. (a) Comparison of estimates for the basic reproduction number obtained from fitting four candidate models to four incidence reports. Error bars correspond to 95% credible intervals, and the vertical line denotes the true value. (b) Comparison between inferred incidence (lines) obtained from two candidate models fitted to two incidence reports (dots). Twenty time series represent inferred incidence. Given the high-fidelity data, all inferred incidences are nearly identical, giving the impression of only one line in each panel.

Misspecifying the infectious period distribution with a tighter distribution (higher j) generates lower \mathcal{R}_0 estimates (figure 3a). Furthermore, regardless of the assumed distribution of the infectious period, all candidate models fit the data equally well. To emphasize the importance and implications of this observation, we compare inferred and actual latent incidences in figure 3b. Recall that fitting a candidate model to a given incidence (y_t) produces a set of samples that describes the posterior distribution. Then, we use those samples to simulate the candidate's system component, thereby generating inferred latent incidences (lines in figure 3b). Then, those lines are compared with x_t , the true latent incidence (figure 1). Notice that by definition, we do not have access to x_t in practical applications, but by virtue of this simulation study, such an impediment is overcome. The comparison reveals a symmetry shared among the candidate models. That is, any of these formulations can match the true latent incidence provided that β , ρ and I_0^1 are configured appropriately. It is important to remark that this symmetry is restricted to the latent incidence and does not extend to the dynamics of other states. For instance, candidates with different delay distributions that yield equivalent incidences will not reach the same long-term equilibrium, given the differences in their \mathcal{R}_0 .

Logically, such symmetry should render the approach of comparing fit scores impractical. A fit score, such as the maximum-likelihood estimate (MLE), measures the consistency between a dataset and the output generated by a model. Since candidates produce equivalent output, differences among MLEs will solely reflect the stochasticity (noise) of the measurement component. We empirically verify this

conjecture by selecting the candidate with the largest MLE for each incidence report (see the electronic supplementary material, S2, §1.3.2). We observe that M^{11} candidates attain the largest MLE in only 12 out of 20 times when matching D^{11} incidence reports. Even worse, M^{13} instances are always outperformed in fitting D^{13} datasets. Overall, no candidate passes the 60% mark. Similarly, the *mean absolute scaled error* (MASE), a metric specifically designed for evaluating the accuracy of time-series forecasts [57], indicates that candidates produce virtually identical scores when fitting any given incidence report. In the light of this evidence, one can safely conclude that score comparison is not a reliable approach to determining the correct distribution of epidemiological delays from incidence data. To further complicate matters, information criteria (such as Akaike (*AIC*) and Bayesian (*BIC*)) and cross-validation methods cannot assist in this task, considering that the evaluated structures produce equivalent output and share an equal number of unknown parameters.

3.2. Four unknowns: β, ρ, I_0, γ

The reason for such inherent symmetry is the *generation time*, the time between the infection of a primary case and one of its secondary cases [58]. This quantity's shape, in tandem with \mathcal{R}_0 , determines the initial dynamics of an infectious disease [5]. Interestingly, these elements also characterize long-term behaviour. Krylova & Earn [29] found that *SEIR* models that account for demographic processes with different delay distributions produce equivalent dynamics of epidemiological transitions (e.g. from annual to biennial epidemic cycles) if they share identical \mathcal{R}_0 and mean generation time (τ). An analytical expression for this last quantity can be obtained using the method described by Svensson [58]. In particular, for the *SEⁱIR* framework, τ can be expressed as a function of the average delays (σ^{-1}, γ^{-1}) and the infectious period distribution (j).

$$\tau = \sigma^{-1} + \frac{j+1}{2j} \gamma^{-1}. \quad (3.1)$$

In this analysis, we have, until now, fixed the mean generation time on each candidate model by excluding σ and γ from the inference process. Taking note of the effects on short- and long-term dynamics that produce the interaction between τ and \mathcal{R}_0 , we now promote γ to the category of *estimated parameter* in order to explore the impact of a variable mean generation time. The reason for choosing γ as the extra parameter is based on the fact that it interacts with both quantities of interest (equations (2.2) and (3.1)). This choice implies the need for estimating four parameters per model instance. To do so, we follow the approach described in the previous section. The reader can find the full set of results in the electronic supplementary material, S2, §2. Unsurprisingly, given the extra degree of freedom, all candidates fit any of the incidence data equally well. In this design, though, the match between synthetic data and fitting model's output is achieved at the expense of less precision, although greater accuracy. *Precision* refers to the width of uncertainty intervals, and *accuracy* to whether the interval captures the actual value. To illustrate this phenomenon, in figure 4, we present the results of fitting four candidates models (M^{ij}) to four incidence reports that stem from different distributions of the infectious period (D^{ij}). Here, we see that the range of \mathcal{R}_0 widened (figure 4a) compared with that presented in the previous section (figure 3a).

Undoubtedly, the primary insight from allowing γ to vary is the unravelled interaction between \mathcal{R}_0 and τ . We visualize this interaction by plugging samples of β and γ into equations (2.2) and (3.1) to obtain an approximation of the expected values of \mathcal{R}_0 and τ . When these two quantities are displayed on a scatter plot (figure 4b), a linear relationship appears, regardless of the data's origin or the fitting model's structure. The interpretation of such linear association indicates that for a given fitting model, infinite pairs of \mathcal{R}_0 and τ yield equivalent incidence dynamics. However, in virtue of their linear relationship, each value of τ corresponds to exactly one value of \mathcal{R}_0 .

3.3. Three unknowns (alternative): $\mathcal{R}_0^{-1}, \rho, I_0$

More importantly, the linear relationships shown in figure 4b reveal an intriguing insight. Notice that irrespective of the structure (M^{ij}) fitting data of any origin (D^{ij}), the true values of \mathcal{R}_0 and τ as a pair (the intersection between the dotted and dashed lines) are subsumed into any of the linear associations. This observation implies that the true \mathcal{R}_0 can correspond only to the right τ . Therefore, it could be possible to accurately estimate \mathcal{R}_0 from a model whose mean generation time is fixed to the true underlying value, but the shape of the epidemiological delays may differ from that of the data

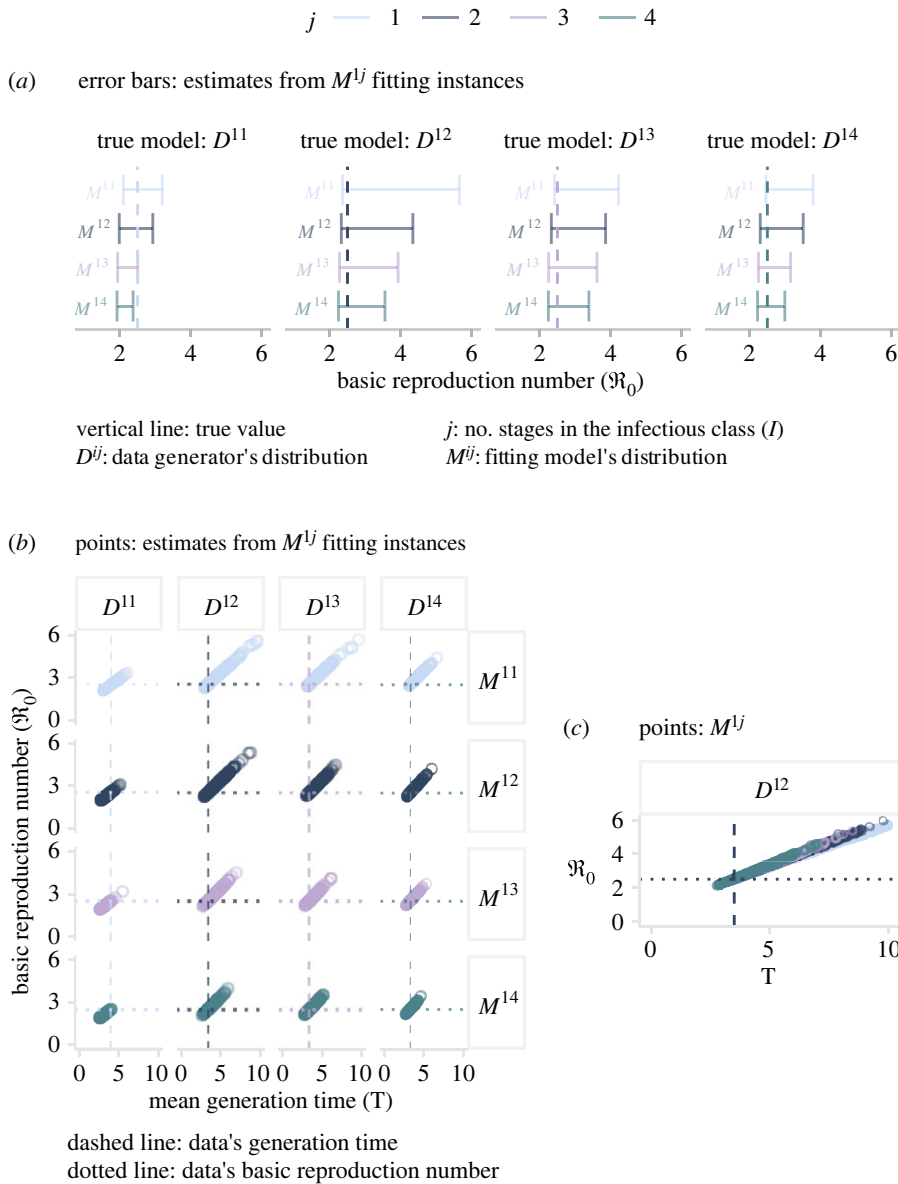


Figure 4. Inference results obtained from the four-unknown parametrization. (a) Comparison of estimates for the basic reproduction number obtained from fitting four candidate models to four incidence reports. Error bars correspond to 95% credible intervals, and the vertical line denotes the true value. (b) Linear relationship between the basic reproduction number and the mean generation time estimated from posterior distributions obtained from fitting four candidate models to four incidence reports. These distributions are represented via samples. From each sample, we compute the predicted \mathfrak{R}_0 and τ (dots). (c) This plot collapses the second column in *b* into a single panel.

generating process. To test this hypothesis, we reformulate the SE^iIR framework so that τ becomes a parameter of every model instance. Consequently, we combine equations (2.2) and (3.1) into (3.2), which expresses β as a dependent variable of four parameters: j , σ , τ and \mathfrak{R}_0 .

$$\left. \begin{aligned} \beta &= \frac{2j(\tau - \sigma^{-1})}{\mathfrak{R}_0^{-1}(j+1)} \\ \gamma &= \beta \mathfrak{R}_0^{-1}. \end{aligned} \right\} \quad (3.2)$$

and

Parameter j is based on the fitting model's structure, whereas σ and τ are fixed to the true values that produced the incidence reports. For instance, a D^{12} report stems from a structure whose σ and τ are equal to 0.5 and 3.5 (applying equation (3.1)), respectively. Therefore, an M^{14} candidate fitting this report has j , σ and τ fixed to 4, 0.5 and 3.5, respectively. An immediate consequence of this procedure is the need to constrain γ in order to maintain logical consistency. Accordingly, we define γ as a

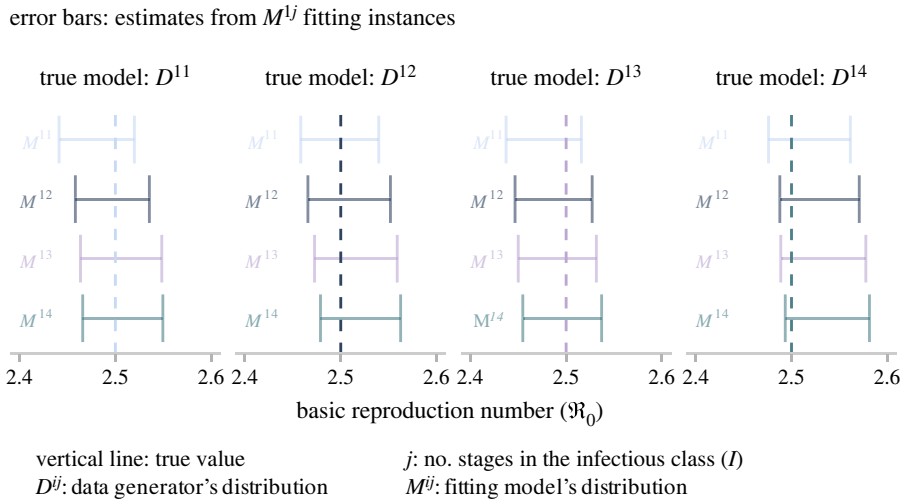


Figure 5. Inference results obtained from the three-unknown alternative parametrization. This plot compares estimates for the basic reproduction number obtained from fitting four candidate models with four incidence reports. Error bars correspond to 95% credible intervals, and the vertical line denotes the true value.

function of β and \mathcal{R}_0 (3.2). This approach is analogous to fixing γ to an arbitrary value that yields the desired τ . Such a value may not correspond to that of the data generating process. Lastly, the remaining parameter, \mathcal{R}_0 , is subject to inference. We opt to estimate its inverse for a practical reason. Taking into account the threshold phenomenon and the fact that all incidence reports exhibit outbreak-like behaviour, any estimated value of \mathcal{R}_0 must fall within the interval $(1, \infty)$. It then logically follows that its inverse (\mathcal{R}_0^{-1}) spans over the range $(0, 1)$. This transformation permits the inference algorithm to operate in a much smaller parameter space, which enhances sampling efficiency.

We subsequently incorporate the redefined components (β and γ) into the $SE^{ij}R$ framework to produce an alternative set of four candidate models with three unknowns: \mathcal{R}_0^{-1} , ρ and I_0^j . Similarly as before, we estimate the posterior distribution for each candidate fitted to an incidence report. The reader can find the complete set results in the electronic supplementary material, S2, §3. These results once more highlight the intrinsic symmetry of SEIR formulations. Specifically, provided there is an adequate configuration, any candidate structure can accurately match the observed incidence despite differences in the infectious period distribution. Nevertheless, this alternative parametrization exhibits a distinctive and crucial feature: the estimation of \mathcal{R}_0 is less sensitive to the assumed distribution of the infectious delay. To support this claim, we present in figure 5 the results of fitting the four alternative candidates to four incidence reports of dissimilar origin. Here, it can be seen that all candidates recover (via 95% credible intervals) the underlying true \mathcal{R}_0 , notwithstanding the origin of the data or the fitting model.

Recovering the underlying \mathcal{R}_0 is not exclusive to this sample of four datasets but is generalized across the 80 high-fidelity D^{lj} datasets. To summarize this insight, we borrow a concept from the *frequentist* tradition. Such a concept known as *coverage* [59] means that if one collects a large number of samples from the same process and constructs the corresponding confidence intervals, then a certain percentage of the intervals will contain or cover the true parameter. This percentage is given by the confidence level. For instance, if one fits a model to 100 datasets and estimates an equal number of confidence intervals at the 95% significance level, then 95 of those intervals will cover the true value. Admittedly, it is implicitly assumed that our 95% credible intervals (obtained from posterior distributions) are proportional to 95% confidence intervals. Indeed, estimated intervals for \mathcal{R}_0 and ρ conform to this concept (see electronic supplementary material, S2, §3.3), where minor deviations are justified by the fact that coverage is defined asymptotically (infinite measurements). However, asymptotics does not account for the large deviance observed in the estimates of I_0 . We explain this inconsistency in the section below where I_0 becomes more prominent.

To conclude this section, we report the analysis of the low-fidelity datasets (right column in figure 2). The reader can find the results in the electronic supplementary material, S3. Overall, we obtain similar insights in comparison to those derived from the high-fidelity datasets. In the absence of structural differences, it is unsurprising that the effect of larger noise in the signal (overdispersion) results in greater uncertainty in parameter estimates. This decrease in precision (wider credible intervals) can

obscure or accentuate features of the inference process. On the one hand, overdispersion masks biases in estimates. For instance, noisier measurements cause I_0 estimates from the *alternative* parametrization to conform to the expected coverage, which should not occur based on the results obtained from the high-fidelity datasets. On the other hand, overdispersion exacerbates identifiability issues. Under the *four-unknown* parametrization, some \mathfrak{R}_0 estimates reach values up to 40. This result is a reminder that choosing an adequate number of unknowns is not a trivial decision. Setting more unknowns than the data can tolerate renders models unidentifiable. In this context, unidentifiability occurs because the incidence data does not provide enough information to update the prior distribution of γ . As discussed above, many values of γ are consistent with the observed incidence, an insight that holds for both levels of data fidelity. Finally, we note that overdispersion estimates are robust to the choice of the infectious period distribution.

3.4. Misspecifying the latent period distribution

Thus far, we have conducted the inference process assuming that the latent period distribution (i) is known. Lifting this constraint would strain our computational resources, producing a fourfold increase in the pool of candidates fitting a single report (assuming $i, j \in \{1, 2, 3, 4\}$). Instead of undertaking such costly exploration, one could leverage the fact that the mean generation time depends solely on the mean latent period rather than its particular distribution (equation (3.1)). To test this idea, we compare the estimates obtained from candidate models with the *right* and *wrong* latent period distribution. We illustrate this process with the 80 D^{3j} low-fidelity ($\phi^{-1} = 1/3$) datasets. For each dataset, we fit eight candidates M^{ij} from the traditional three-unknown parametrization, where the latent period distribution can take the wrong ($i=1$) and the right ($i=3$) values, and the infectious period distribution varies as before, namely, $j \in \{1, 2, 3, 4\}$. The reader can find the complete results in the electronic supplementary material, S5.

To facilitate the presentation of the results, we first focus on candidates M^{13} and M^{33} fitting one D^{33} incidence report. Figure 6a shows that both models predict similar, although not identical, latent incidence dynamics. Further inspection reveals that the slight difference in the predicted incidence due to dissimilar latent period distributions does not lead to variation in \mathfrak{R}_0 estimates. To corroborate this assessment, we expand the analysis to the eight candidates matching the same incidence report. The right-hand side of figure 6b shows that \mathfrak{R}_0 estimates are sensitive to variation in the structure of the infectious period but are indifferent to the latent period distribution. In compliance with the literature, the more dispersed latent period ($i=1$) leads to an earlier incidence peak compared with the tighter distribution ($i=3$) in the context of identical \mathfrak{R}_0 .

Nevertheless, the mechanism that enables models with heterogeneous distributions to produce analogous incidence dynamics remains unexplained. The left-hand side of figure 6b, which displays I_0 estimates, provides the first hint. This plot shows that instances with the wrong latent period distribution (M^{1j}) systematically underestimate (via 95% credible intervals) the actual value (vertical line). To explain this phenomenon, we draw on a broader view of the posterior distribution. It is commonplace to restrict inference analyses to one parameter at a time (i.e. marginal distributions), neglecting the information provided by the full posterior distribution. To redress this shortcoming, we visualize the full distribution via pair plots. Specifically, figure 6c corresponds to the summary of the posterior distribution obtained from fitting M^{13} to one D^{33} incidence report. The upper triangular elements of this plot indicate that the three estimated parameters are strongly correlated. Especially β and I_0 , or more compellingly, \mathfrak{R}_0 and I_0 . Recall that the basic reproduction number is directly proportional to β . Therefore, although the mean generation time determines which \mathfrak{R}_0 corresponds to the observed incidence, I_0 (and ρ to a lesser extent) regulates the flexibility of \mathfrak{R}_0 to reach such a desired value. Interestingly, I_0 provides such a degree of flexibility that unrealistic adjustments in its estimates allow us to equate dissimilar model structures. Note that the only discernible difference between figure 6c,d (M^{33} fitted to D^{33}) is seen in the marginal distributions of I_0 . In fact, this phenomenon explains the *failure* of the alternative parametrization to recover the true value of I_0 .

In the view of these symmetries, it is not unreasonable to expect that candidates from the four-unknown and the alternative parametrizations, too, are indifferent to the latent period distribution once I_0 and ρ correct for any misspecification. To verify this premise, we fit the parametrizations mentioned above to the D^{3j} low-fidelity incidence reports. As anticipated, the inference results indicate that the four-unknown parametrization (electronic supplementary material, S5, §2) uncover the linear association between τ and ρ due to the unidentifiability of γ . Likewise, the alternative

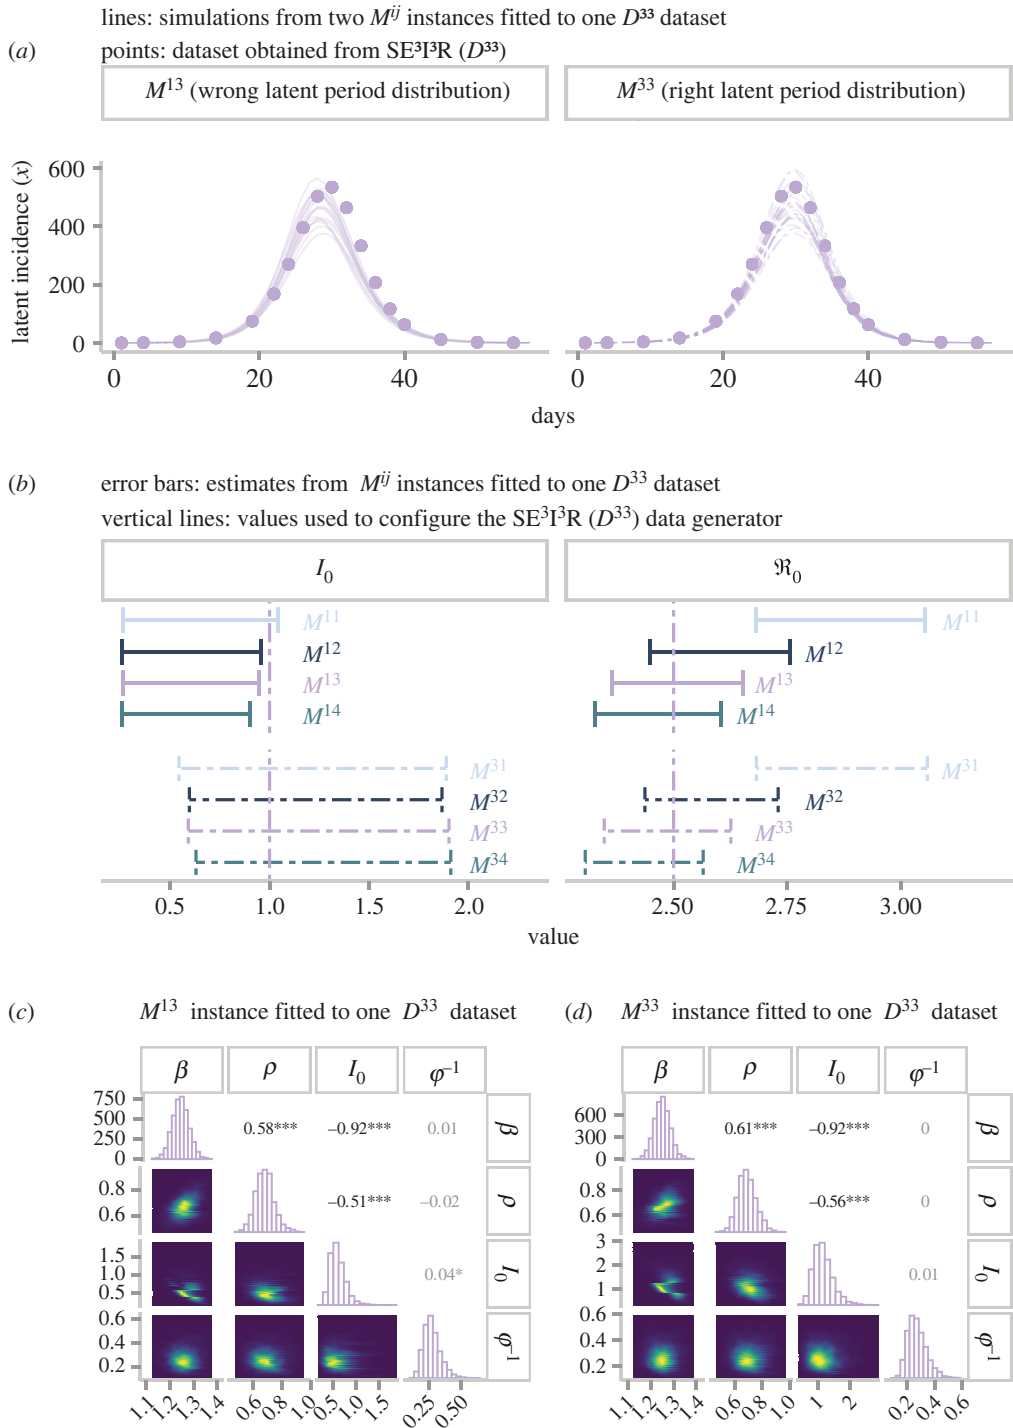


Figure 6. Comparing estimates from candidate models with the wrong and correct distributions of the latent period. These models stem from the three-unknown parametrization (a) Incidence fit from two candidate models matching a D^{33} incidence report. Candidate on the left has an exponentially distributed (wrong) latent period, whereas candidate on the right has Gamma-distributed latent period (correct). (b) Comparison of estimates of the initial number of infectious individuals at stage 1 and the basic reproduction number by fitting model. Error bars correspond to 95% credible intervals, and vertical lines denotes the true value. (c) Joint posterior distribution from an M^{13} candidate fitting a D^{33} report. The diagonal shows posterior marginal distributions. In the lower triangular part, each possible pairwise conditional distribution is displayed, whereas the upper triangular part presents the correlation among parameters.

parametrization (electronic supplementary material, S5, §3) recovers the true \mathcal{R}_0 irrespective of the formulation of the epidemiological delays. Furthermore, we replicate these results using the D^{3j} high-fidelity datasets (see electronic supplementary material, S4).

Table 1. Scenarios.

scenario	\mathcal{R}_0	τ_e	\mathcal{R}_0 recovered? (high-fidelity)	\mathcal{R}_0 recovered? (low-fidelity)
1	2.5	4	yes	yes
2	2.5	8	yes	yes
3	2.5	13	yes	yes
4	9.0	4	no	yes
5	15.0	4	no	yes

3.5. Sensitivity analysis

So far, model candidates have been amalgamated with the appropriate measurement component. In this section, we explore the implications that can arise from ignoring overdispersion. That is, equipping model candidates with a Poisson measurement component. We perform such exploration by inferring \mathcal{R}_0 from M^{1j} candidates (alternative parametrization) fitted to the D^{3j} low-fidelity datasets discussed in the previous section. As expected, the results indicate that employing the Poisson distribution (see electronic supplementary material, S5, S4) leads to overconfident (too precise) and biased (inaccurate) estimates in the context of overdispersion. We observe these features with narrow uncertainty intervals that do not cover the true value. This result implies that the wrong choice of the measurement component can offset any gains in accuracy due to the alternative parametrization.

On the other hand, the synthetic data used for the analysis presented in the previous sections stems from models configured to identical \mathcal{R}_0 and similar mean generation times (variation due to the infectious period distribution). Naturally, one wonders whether the usefulness of the alternative parametrization holds in other conditions. To answer this question, we repeat the workflow described in this paper for additional scenarios of τ and \mathcal{R}_0 . For simplicity, we restrict this sensitivity analysis to datasets derived from models with an exponentially distributed latent period (D^{1j}). Additionally, we equip the fitting candidates with the appropriate measurement component. The complete set of results is presented in the electronic supplementary material, S6. We present these results in terms of scenarios (table 1). For instance, the base case scenario, *Scenario 1*, corresponds to data generated from $SE^{ij}R$ configured to $\mathcal{R}_0 = 2.5$ and $\tau_e = 4$ (results presented in §3.3), where τ_e serves as a scenario identifier and denotes the mean generation time obtained from an exponentially distributed infectious period ($j = 1$).

For *Scenario 2*, we increase the reference mean generation time ($\tau_e = 8$), while keeping \mathcal{R}_0 at 2.5. First, we focus on the high-fidelity datasets. Overall, the greater the divergence between the fitting model's infectious period distribution and the distribution that generated the data, the greater the loss in accuracy; namely, lower coverage. To provide an example, the 95% credible intervals constructed from M^{14} candidates fitting D^{11} incidence reports only attain coverage of 30% for \mathcal{R}_0 . Closer inspection, though, reveals that such accuracy loss is more statistical than practical. To support this statement, we calculate the average relative difference between the actual and estimated \mathcal{R}_0 , finding that misspecification of the infectious period distribution leads to a maximum average relative error of 2%. In contrast, we would obtain discrepancies up to 15% if we adopted the traditional approach. Simply put, it is costlier to misspecify the mean generation time than the mean infectious period. Furthermore, such slight differences in the alternative parametrization are erased by overdispersion. That is, overdispersion masks minor misspecification in the process component. Moreover, in *Scenario 3* ($\mathcal{R}_0 = 2.5$, $\tau_e = 13$), we observe that further increasing of the mean generation time does not lead to significant drops in the coverage of \mathcal{R}_0 under both levels of data fidelity. In a nutshell, it is reasonable to suggest that the alternative parametrization is robust to various levels of the mean generation time.

Conversely, we cannot maintain the same assertion for various values of \mathcal{R}_0 . Indeed, figure 4c provided the first hint. This plot shows that the straight lines do not overlap as \mathcal{R}_0 reaches relatively high values. Consequently, in scenarios 4 ($\mathcal{R}_0 = 9$, $\tau_e = 4$) and 5 ($\mathcal{R}_0 = 17$, $\tau_e = 4$), we test the implications of larger transmissibility levels. The results indicate that as we increase the underlying \mathcal{R}_0 for generating the data, the equivalency among fitting models dissipates and misspecification in the infectious period distribution leads to biased estimates of \mathcal{R}_0 . The size of such bias is proportional to the misspecification of the infectious period and the underlying \mathcal{R}_0 . This feature is primarily seen in the estimates derived from high-fidelity datasets, where coverage levels are low, and the average

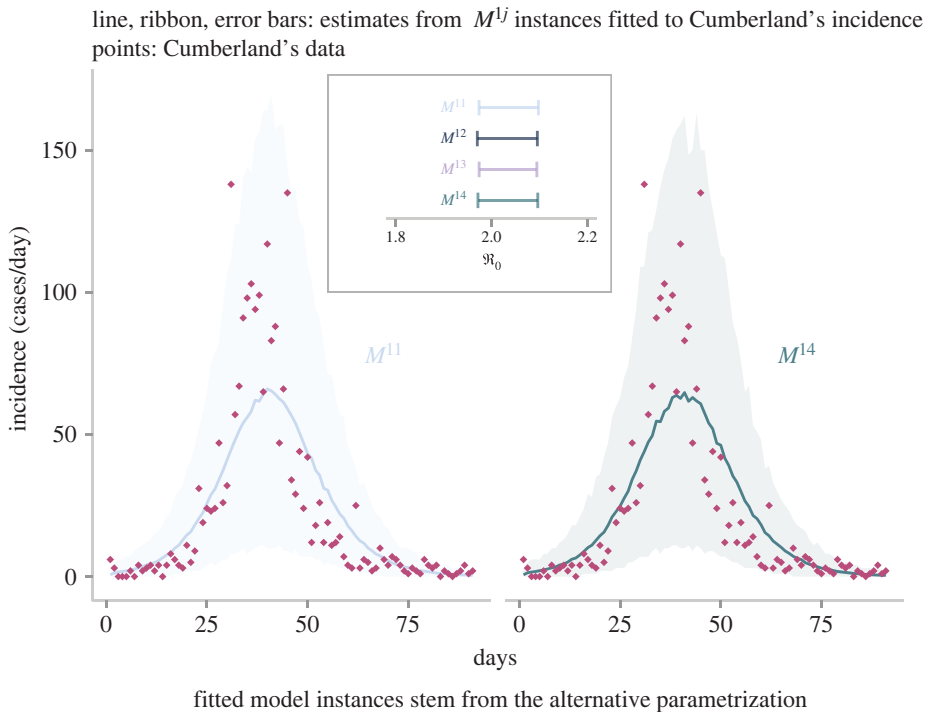


Figure 7. Application of the three-unknown alternative parametrization. This plot shows the estimates obtained from fitting four candidates (M^{1j}) to the daily number of influenza cases (rhombi) detected by the US Public Health Service in Cumberland (Maryland) during the 1918 influenza pandemic, from 22 September to 30 November 1918. Ribbons correspond to 95% credible intervals of the predicted measured incidence by two candidate models. The solid line denotes the median. The box inside the plot shows the estimates for the basic reproduction number by fitting model. Error bars correspond to 95% credible intervals.

relative error between actual and estimated values cannot be overlooked. However, when we examine the posterior distributions obtained from fitting the low-fidelity data, it is seen that, once again, overdispersion masks misspecification in the process component, as evidenced by the high coverage levels. This is not to say that overdispersion is a desired feature in the data, but rather to emphasize that its presence hinders the attainment of precise estimates. Undoubtedly, having this understanding is of practical importance, given that it allows us to discern the necessary effort in data collection and model improvement.

4. Application to influenza A

Leveraging the knowledge gained from the synthetic data, the last step in this work consists of exploiting the relationship between the basic reproduction number and the mean generation time to update the R_0 estimate of an outbreak of the 1918 influenza pandemic. The reader can find the full set of results in the electronic supplementary material, *S7*. In particular, we focus on an outbreak that occurred in the city of Cumberland (Maryland) during the autumn of 1918, for which the US Public Health Service organized special surveys [60] to determine the proportion of the population infected. Previous studies [12,22] employed the default heuristic of adopting an *SEIR* with exponentially distributed epidemiological delays whose means were configured to values reported in the literature. Moreover, in these studies, the *SEIR* was coupled with the Poisson distribution resulting in a 95% CI [2.5–2.6] for R_0 . However, adopting a more realistic measurement component, such as the negative binomial distribution, produces lower and wider estimates: 95% CI [2.2, 2.4]. Further, if we jettison the assumption of an exponentially distributed infectious period for a more realistic distribution, such as the gamma distribution, we obtain even lower estimates. For instance, a gamma-distributed infectious period with four stages (*SEI⁴R*) returns a 95% CI of [2.0, 2.2]. As noted earlier, the estimates obtained from this default heuristic or traditional approach are sensitive to the uncertainty in the *SEIR* (alternative

parametrization) to a value (2.85 days) obtained from the literature [5,61], we derive nearly identical R_0 estimates (95% CI [2.0, 2.1]) regardless of the infectious period distribution (figure 7). Notice that this estimate is similar to that obtained from the $SEIR^4R$, bolstering the fact that the actual infectious period is far from being exponentially distributed.

5. Conclusion

The misspecification of various assumptions within the $SEIR$ framework can negatively impact the estimation of R_0 . In recognition of this risk, we ran a simulation study comprising approximately 1000 synthetic datasets and 8000 model fits, whereby we identified the relative influence of some of those assumptions. Specifically, we found that fixing the mean generation time to a reliable estimate is of paramount importance. By contrast, one can be more lenient on the specification of the latent period distribution and the mean infectious period provided that other estimated parameters (I_0 and ρ) redress the misspecification. We leveraged this knowledge to formulate an alternative parametrization that is more robust to the uncertainty of the epidemiological delays. However, there is a caveat with this alternative formulation. Although it exploits a local symmetry (incidence dynamics) of the $SEIR$ framework, such symmetry does not extend to the other states of the system. Therefore, the usefulness of the alternative parametrization is confined to the estimation of R_0 , and it is not a substitute for other kinds of analyses. For instance, if, on the contrary, our variable of interest were I_0 , we would obtain unreliable estimates. Furthermore, the alternative $SEIR$ with exponentially distributed delays will be as overoptimistic as its traditional counterpart in predicting the critical vaccination proportion or the effectiveness of an imperfect HIV treatment in the context of within-host dynamics [56]. Therefore, the alternative parametrization is a mitigation strategy in the absence of complete information. Furthermore, its usefulness is abated by highly transmissible pathogens (§3.5). Nevertheless, biases in the estimates due to large R_0 are only detected with high-fidelity data. That is, data with little or no overdispersion.

Despite the significant computational effort of simulation analyses, a single study cannot offer overarching statements. Further work is required to test the validity of these insights in stricter or more elaborated contexts. For instance, we assumed the complete availability of the incidence time-series throughout this study. This assumption restricts the validity of the approach to retrospective analyses. However, other situations exist where only fewer incidence measurements are available to the modellers, such as the early phase of a pandemic response. Hence, it remains to be seen the effect of various levels of data availability on the performance of the suggested approach. Furthermore, for simplicity, we ignored age-related effects in the dynamics of the infectious disease as well as process stochasticity (demographic and environmental) and time-varying contact rates. We expect that future research builds on the findings provided by this study and addresses the aforementioned challenges to construct ever more reliable inference approaches.

Ethics. This work did not require ethical approval from a human subject or animal welfare committee.

Data accessibility. Data and relevant code for this research work are stored in GitHub: <https://github.com/jandraor/> delays and have been archived within the Zenodo repository: <https://doi.org/10.5281/zenodo.8134814> [62].

The data are provided in electronic supplementary material [63].

Authors' contributions. J.A.: conceptualization, data curation, formal analysis, investigation, visualization, writing—original draft; J.D.: resources, supervision, writing—review and editing.

Both authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. The project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 883285. The material presented and views expressed here are the responsibility of the author(s) only. The EU Commission takes no responsibility for any use made of the information set out. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

1. Dublin LI, Lotka AJ. 1925 On the true rate of natural increase. *J. Am. Stat. Assoc.* **20**, 305. (doi:10.2307/2965517)
2. Anderson RM, May RM. 1992 *Infectious diseases of humans: dynamics and control*. Oxford, UK: Oxford University Press.
3. Heffernan JM, Smith RJ, Wahl LM. 2005 Perspectives on the basic reproductive ratio. *J. R. Soc. Interface* **2**, 281–293. (doi:10.1098/rsif.2005.0042)

4. Keeling MJ, Rohani P. 2011 *Modeling infectious diseases in humans and animals*. Princeton, NJ: Princeton University Press.
5. Wallinga J, Lipsitch M. 2007 How generation intervals shape the relationship between growth rates and reproductive numbers. *Proc. R. Soc. B* **274**, 599–604. (doi:10.1098/rspb.2006.3754)
6. Kucharski A. 2020 *The rules of contagion: why things spread—and why they stop*. New York, NY: Basic Books.
7. Mollison D. 1995 The structure of epidemic models. In *Epidemic models: their structure and relation to data* (ed. D Mollison), pp. 17–33. Cambridge, UK: Cambridge University Press.
8. Dietz K. 1975 Transmission and control of arbovirus diseases. In *Epidemiology* (eds D Ludwig, KL Cooke), pp. 104–121. Philadelphia, PA: Society for Industrial and Applied Mathematics.
9. Kermack WO, McKendrick AG. 1927 A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. A* **115**, 700–721. (doi:10.1098/rspa.1927.0118)
10. Diekmann O, Heesterbeek H, Britton T. 2013 *Mathematical tools for understanding infectious disease dynamics*. Princeton, NJ: Princeton University Press.
11. Lloyd AL. 2009 Sensitivity of model-based epidemiological parameter estimation to model assumptions. In *Mathematical and statistical estimation approaches in epidemiology* (eds G Chowell, JM Hyman, LMA Bettencourt, C Castillo-Chavez), pp. 123–141. Dordrecht, The Netherlands: Springer.
12. Vynnycky E, White R. 2010 *An introduction to infectious disease modelling*. Oxford, UK: Oxford University Press.
13. Brouwer AF. 2022 Why the spectral radius? An intuition-building introduction to the basic reproduction number. *Bull. Math. Biol.* **84**, 96. (doi:10.1007/s11538-022-01057-9)
14. He D, Ionides EL, King AA. 2010 Plug-and-play inference for disease dynamics: measles in large and small populations as a case study. *J. R. Soc. Interface* **7**, 271–283. (doi:10.1098/rsif.2009.0151)
15. Dureau J, Kalogeropoulos K, Baguelin M. 2013 Capturing the time-varying drivers of an epidemic using stochastic dynamical systems. *Biostatistics* **14**, 541–555. (doi:10.1093/biostatistics/kxs052)
16. Andrade J, Duggan J. 2022 Inferring the effective reproductive number from deterministic and semi-deterministic compartmental models using incidence and mobility data. *PLoS Comput. Biol.* **18**, 1–25. (doi:10.1371/journal.pcbi.1010206)
17. Andrade J, Duggan J. 2020 An evaluation of Hamiltonian Monte Carlo performance to calibrate age-structured compartmental SEIR models to incidence data. *Epidemics* **33**, 100415. (doi:10.1016/j.epidem.2020.100415)
18. Bretó C, He D, Ionides EL, King AA. 2009 Time series analysis via mechanistic models. *Ann. Appl. Stat.* **3**, 319–348. (doi:10.1214/08-AOAS201)
19. Bretó C. 2018 Modeling and inference for infectious disease dynamics: a likelihood-based approach. *Stat. Sci.* **33**, 57–69. (doi:10.1214/17-STS636)
20. Ionides EL, Bretó C, King AA. 2006 Inference for nonlinear dynamical systems. *Proc. Natl. Acad. Sci. USA* **103**, 18 438–18 443. (doi:10.1073/pnas.0603181103)
21. Endo A, van Leeuwen E, Baguelin M. 2019 Introduction to particle Markov-chain Monte Carlo for disease dynamics modellers. *Epidemics* **29**, 100363. (doi:10.1016/j.epidem.2019.100363)
22. Andrade J, Duggan J. 2021 A Bayesian approach to calibrate system dynamics models using Hamiltonian Monte Carlo. *Syst. Dyn. Rev.* **37**, 283–309. (doi:10.1002/sdr.1693)
23. Barlas Y. 1996 Formal aspects of model validity and validation in system dynamics. *Syst. Dyn. Rev.* **12**, 183–210. (doi:10.1002/(SICI)1099-1727(199623)12:3%3C183::AID-SDR103%3E3.0.CO;2-4)
24. Oliva R. 2003 Model calibration as a testing strategy for system dynamics models. *Eur. J. Oper. Res.* **151**, 552–568. (doi:10.1016/S0377-2217(02)00622-7)
25. Wearing HJ, Rohani P, Keeling MJ. 2005 Appropriate models for the management of infectious diseases. *PLoS Med.* **2**, e174. (doi:10.1371/journal.pmed.0020174)
26. Gostic KM *et al.* 2020 Practical considerations for measuring the effective reproductive number, R_t . *PLoS Comput. Biol.* **16**, 1–21. (doi:10.1371/journal.pcbi.1008409)
27. Hurtado PJ, Kirosging AS. 2019 Generalizations of the ‘linear chain trick’: incorporating more flexible dwell time distributions into mean field ODE models. *J. Math. Biol.* **79**, 1831–1883. (doi:10.1007/s00285-019-01412-w)
28. Greenhalgh S, Rozins C. 2021 A generalized differential equation compartmental model of infectious disease transmission. *Infect. Dis. Modell.* **6**, 1073–1091. (doi:10.1016/j.idm.2021.08.007)
29. Krylova O, Earn DJD. 2013 Effects of the infectious period distribution on predicted transitions in childhood disease dynamics. *J. R. Soc. Interface* **10**, 20130098. (doi:10.1098/rsif.2013.0098)
30. Gelman A *et al.* 2020 Bayesian workflow 2020. *arXiv*. (<http://arxiv.org/abs/2011.01808>)
31. King AA, Domenech de Cellés M, Magpantay FMG, Rohani P. 2015 Avoidable errors in the modelling of outbreaks of emerging pathogens, with special reference to Ebola. *Proc. R. Soc. B* **282**, 20150347. (doi:10.1098/rspb.2015.0347)
32. Lloyd AL. 2001 Realistic distributions of infectious periods in epidemic models: changing patterns of persistence and dynamics. *Theor. Popul. Biol.* **60**, 59–71. (doi:10.1006/tpbi.2001.1525)
33. Keeling MJ, Grenfell BT. 2002 Understanding the persistence of measles: reconciling theory, simulation and observation. *Proc. R. Soc. Lond. B* **269**, 335–343. (doi:10.1098/rspb.2001.1898)
34. Davies NG, Klepac P, Liu Y, Prem K, Jit M, Eggo RM. 2020 Age-dependent effects in the transmission and control of COVID-19 epidemics. *Nat. Med.* **26**, 1205–1211. (doi:10.1038/s41591-020-0962-9)
35. Gleeson JP, Brendan Murphy T, O’Brien JD, Friel N, Bargarly N, O’Sullivan DJP. 2022 Calibrating COVID-19 susceptible-exposed-infected-removed models with time-varying effective contact rates. *Phil. Trans. R. Soc. A* **380**, 20210120. (doi:10.1098/rsta.2021.0120)
36. Chowell G, Ammon CE, Hengartner NW, Hyman JM. 2006 Transmission dynamics of the great influenza pandemic of 1918 in Geneva, Switzerland: assessing the effects of hypothetical interventions. *J. Theor. Biol.* **241**, 193–204. (doi:10.1016/j.jtbi.2005.11.026)
37. Chowell G, Nishiura H, Bettencourt LMA. 2007 Comparative estimation of the reproduction number for pandemic influenza from daily case notification data. *J. R. Soc. Interface* **4**, 155–166. (doi:10.1098/rsif.2006.0161)
38. Vynnycky E, Edmunds WJ. 2008 Analyses of the 1957 (Asian) influenza pandemic in the United Kingdom and the impact of school closures. *Epidemiol. Infect.* **136**, 166–179. (doi:10.1017/S0950268807008369)
39. Anderson D, Watson R. 1980 On the spread of a disease with gamma distributed latent and infectious periods. *Biometrika* **67**, 191–198. (doi:10.1093/biomet/67.1.191)
40. Bailey N. 1954 A statistical method of estimating the periods of incubation and infection of an infectious disease. *Nature* **174**, 139–140. (doi:10.1038/174139a0)
41. Sartwell P. 1995 The distribution of incubation periods of infectious disease. *Am. J. Epidemiol.* **141**, 386–394. (doi:10.1093/oxfordjournals.aje.a117440)
42. Blitzstein JK, Hwang J. 2019 *Introduction to probability*, 2nd edn. Boca Raton, FL: CRC Press.
43. Gamado KM, Stretfaris G, Zachary S. 2014 Modelling under-reporting in epidemics. *J. Math. Biol.* **69**, 737–765. (doi:10.1007/s00285-013-0717-z)
44. World Health Organization. 2017 *Pandemic influenza risk management: a WHO guide to inform and harmonize national and international pandemic preparedness and response*. World Health Organization.
45. Monto AS, Webster RG. 2013 Influenza pandemics: history and lessons learned. In *Textbook of influenza* (eds RG Webster, AS Monto, TJ Braciale, RA Lamb), pp. 20–34. Chichester, UK: Wiley Blackwell.
46. Talts S, Betancourt M, Simpson D, Vehtari A, Gelman A. 2018 Validating Bayesian inference algorithms with simulation-based calibration. (<http://arxiv.org/abs/1804.06788>)
47. Mills CE, Robins JM, Lipsitch M. 2004 Transmissibility of 1918 pandemic influenza. *Nature* **432**, 904–906. (doi:10.1038/nature03063)
48. Arulampalam MS, Maskell S, Gordon N, Clapp T. 2002 A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Trans. Signal Process.* **50**, 174–188. (doi:10.1109/78.978374)
49. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. 2013 *Bayesian data analysis*, 3rd edn. Boca Raton, FL: CRC Press LLC.
50. McElreath R. 2020 *Statistical rethinking: a Bayesian course with examples in R and Stan*. Boca Raton, FL: CRC Press LLC.
51. Lambert B. 2018 *A student’s guide to Bayesian statistics*. Los Angeles, CA: SAGE.

52. Betancourt M. 2017 A conceptual introduction to Hamiltonian Monte Carlo. *arXiv*. (<http://arxiv.org/abs/1701.02434>)
53. Chatzilena A, Leeuwen Evan, Ratmann O, Baguelin M, Demiris N. 2019 Contemporary statistical inference for infectious disease models using stan. *Epidemics* **29**, 100367. (doi:10.1016/j.epidem.2019.100367)
54. Grinsztajn L, Semenova E, Margossian CC, Riou J. 2021 Bayesian workflow for disease transmission modeling in Stan. *Stat. Med.* **40**, 6209–6234. (doi:10.1002/sim.9164)
55. Carpenter B *et al.* 2017 Stan: a probabilistic programming language. *J. Stat. Softw.* **76**, 1–32. (doi:10.18637/jss.v076.i01)
56. Lloyd AL. 2001 The dependence of viral parameter estimates on the assumed viral life cycle: limitations of studies of viral load data. *Proc. R. Soc. Lond. B* **268**, 847–854. (doi:10.1098/rspb.2000.1572)
57. Hyndman RJ, Koehler AB. 2006 Another look at measures of forecast accuracy. *Int. J. Forecast.* **22**, 679–688. (doi:10.1016/j.ijforecast.2006.03.001)
58. Svensson Å. 2007 A note on generation times in epidemic models. *Math. Biosci.* **208**, 300–311. (doi:10.1016/j.mbs.2006.10.010)
59. Freedman D, Pisani R, Purves R. 2007 *Statistics*. International student ed. 4th edn. New York, NY: Norton.
60. Frost WH, Sydenstricker E. 1919 Influenza in Maryland: preliminary statistics of certain localities. *Public Health Reports (1896–1970)* **34**, 491–504. (doi:10.2307/4575056)
61. Hirotsu N *et al.* 2004 Effects of antiviral drugs on viral detection in influenza patients and on the sequential infection to their family members—serial examination by rapid diagnosis (capilia) and virus culture. *Int. Congress Series* **1263**, 105–108. (doi:10.1016/j.ics.2004.02.020)
62. Andrade J, Duggan J. 2023 Anchoring the mean generation time in the SEIR to mitigate biases in \mathcal{R}_0 estimates due to uncertainty in the distribution of the epidemiological delays. *Zenodo*. (doi:10.5281/zenodo.8134814)
63. Andrade J, Duggan J. 2023 Anchoring the mean generation time in the SEIR to mitigate biases in \mathcal{R}_0 estimates due to uncertainty in the distribution of the epidemiological delays. Figshare. (doi:10.6084/m9.figshare.c.6753860)