

## RESEARCH ARTICLE

## Improved estimation of time-varying reproduction numbers at low case incidence and between epidemic waves

Kris V. Parag \*

MRC Centre for Global Infectious Disease Analysis, Imperial College London, London, United Kingdom

\* [k.parag@imperial.ac.uk](mailto:k.parag@imperial.ac.uk) OPEN ACCESS

**Citation:** Parag KV (2021) Improved estimation of time-varying reproduction numbers at low case incidence and between epidemic waves. *PLoS Comput Biol* 17(9): e1009347. <https://doi.org/10.1371/journal.pcbi.1009347>

**Editor:** Sergei L. Kosakovsky Pond, Temple University, UNITED STATES

**Received:** October 5, 2020

**Accepted:** August 13, 2021

**Published:** September 7, 2021

**Copyright:** © 2021 Kris V. Parag. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** All data and code for EpiFilter available at: <https://github.com/kpzoo/EpiFilter>.

**Funding:** This work is jointly funded to KVP under grant reference MR/R015600/1 by the UK Medical Research Council (MRC) and the UK Department for International Development (DFID) under the MRC/DFID Concordat agreement and is also part of the EDCTP2 programme supported by the European Union. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

## Abstract

We construct a recursive Bayesian smoother, termed EpiFilter, for estimating the effective reproduction number,  $R$ , from the incidence of an infectious disease in real time and retrospectively. Our approach borrows from Kalman filtering theory, is quick and easy to compute, generalisable, deterministic and unlike many current methods, requires no change-point or window size assumptions. We model  $R$  as a flexible, hidden Markov state process and exactly solve forward-backward algorithms, to derive  $R$  estimates that incorporate all available incidence information. This unifies and extends two popular methods, EpiEstim, which considers past incidence, and the Wallinga-Teunis method, which looks forward in time. We find that this combination of maximising information and minimising assumptions significantly reduces the bias and variance of  $R$  estimates. Moreover, these properties make EpiFilter more statistically robust in periods of low incidence, where several existing methods can become destabilised. As a result, EpiFilter offers improved inference of time-varying transmission patterns that are advantageous for assessing the risk of upcoming waves of infection or the influence of interventions, in real time and at various spatial scales.

## Author summary

Inferring changes in the transmissibility of an infectious disease is crucial for understanding and controlling epidemic spread. The effective reproduction number,  $R$ , is widely used to assess transmissibility.  $R$  measures the average number of secondary cases caused by a primary case and has provided insight into many diseases including COVID-19. An upsurge in  $R$  can forewarn of upcoming infections, while suppression of  $R$  can indicate if public health interventions are working. Reliable estimates of temporal changes in  $R$  can contribute important evidence to policymaking. Popular  $R$ -inference methods, while powerful, can struggle when cases are few because data are noisy. This can limit detection of crucial variations in transmissibility that may occur, for example, when infections are waning or when analysing transmissibility over fine geographic scales. In this paper we improve the general reliability of  $R$ -estimates and specifically increase robustness when cases are few. By adapting principles from control engineering, we formulate EpiFilter, a novel method for inferring  $R$  in real time and

**Competing interests:** The authors have declared that no competing interests exist.

retrospectively. EpiFilter can potentially double the information extracted from epidemic time-series (when compared to popular approaches), significantly filtering the noise within data, to minimise both bias and uncertainty of  $R$ -estimates and enhance the detection of salient change-points in transmissibility.

This is a *PLOS Computational Biology Methods* paper.

## Introduction

During an unfolding epidemic, one of the most commonly available and useful types of surveillance data is the daily (or weekly) number of newly reported cases. This time-series of case counts, also known as the incidence curve, not only measures the epidemic size and burden, but also provides information about trends or changes in its transmissibility [1, 2]. These trends are captured by the time-varying effective or instantaneous reproduction number, denoted  $R_s$  at time  $s$ , which defines how the number of secondary cases generated per primary case varies across the outbreak [3]. Broadly, when  $R_s > 1$  we can expect and prepare for growing incidence, whereas sustained  $R_s < 1$  signifies that the epidemic is waning and likely to enter a more controlled phase [4].

Inferring changes in  $R_s$  given an observed incidence curve is therefore crucial, both to understanding transmissibility and to forecasting upcoming case loads, especially for an ongoing epidemic, where it can help inform policymaking and intervention choices or predict healthcare demands [1, 5]. Real-time and retrospective  $R_s$  estimates have been used to characterise rates and patterns of spread in various diseases such as malaria [6] and Ebola virus disease [7]. Such estimates have proven valuable throughout the COVID-19 pandemic, providing updating synopses of global transmission [8] and evidencing the impact of past control actions (e.g., lockdowns and social distancing) [9] or the likelihood of a resurgence in infections when those controls are relaxed [10].

Most studies that infer  $R_s$  or related quantities either apply the Wallinga-Teunis (WT) method [2] or the Cori *et al* method, known as EpiEstim [11]. Both methods take complementary viewpoints on how incidence data inform on transmissibility and hence have diverse use-cases. The WT method reconstructs the average number of new cases caused by infectious individuals at  $s$  and so requires incidence data beyond time  $s$  for its estimate. It computes the case or cohort reproduction number,  $R_s^c$ , which is a function of  $R_{s+j}$  for future times  $j \geq 0$ , and is suited for retrospective analyses [12]. Alternatively, EpiEstim infers how past infections propagate to form the incidence observed at  $s$ , only requiring data prior to time  $s$ . EpiEstim directly computes instantaneous reproduction numbers,  $R_s$ , and is preferred for real-time investigations [3].

While both methods provide useful and important estimators of transmission, they are not perfect. Two main limitations exist. First, each suffers from data censoring or edge-effects [3]. Because the WT method is forward-looking i.e., depends on data later than  $s$ , its estimates are right censored when  $s$  is close to the last observed time point [12]. In contrast, EpiEstim looks backward in time and suffers edge-effects when  $s$  is near the first observed time point [11]. Estimates in the vicinity of the start and end of the incidence time-series are therefore unreliable under EpiEstim and the WT method, respectively. Techniques have been proposed to limit this unreliability [5, 13], but the problem is intrinsic, and inevitable near the actual start and end times of an epidemic, where there is necessarily no or sparse data.

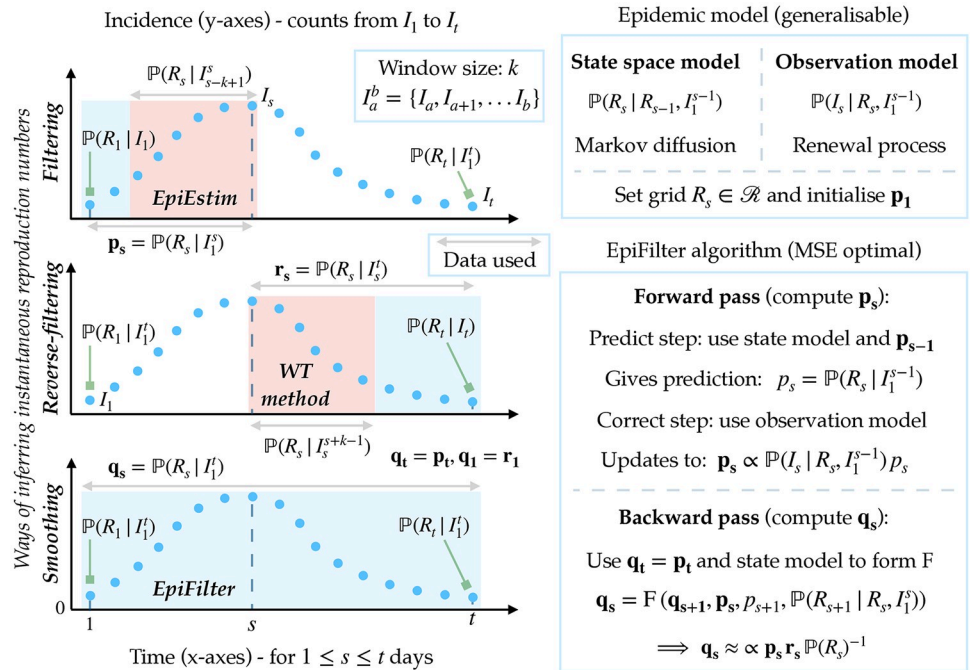
This leads into the second key limitation: inference in periods of small incidence. This presents a fundamental challenge for any  $R_s$  estimation approach and effectively creates additional edge-effects. When few or no case counts are available to constrain inference, methods are largely driven by their inherent prior distributions and assumptions. This can result in misleading or unreliable estimates and mask subcritical (i.e.,  $R_s < 1$ ) transmission patterns [11, 15]. Small incidence may occur during sustained periods of interventions, when pathogens invade new susceptible regions and also naturally arise by division when analyses are performed at smaller spatial scales (e.g., regional or community levels). Understanding how to best mediate the trade-off between prior assumptions and data when incidence is small is of both statistical and epidemiological significance.

Following a period of low incidence, two important outcomes are possible: either the epidemic continues to exhibit small or zero case counts until it goes extinct, or a resurgence in infections, also termed a second wave, occurs. Inferring, in real time, which of these conditions is likely presents a key challenge for infectious disease epidemiology given the information bottleneck at low incidence [16]. Better inference of transmission under these conditions is currently considered central to designing data-informed COVID-19 intervention exit or relaxation strategies [17]. With many countries facing multiple resurgent waves in this pandemic, estimating fluctuations in transmission during suspected epidemic troughs could be essential to achieving sustained control [10].

Here we present and develop a novel method, termed EpiFilter, for reliably estimating  $R_s$  in real time, which ameliorates the above limitations. We take an engineering inspired approach and construct an exact, recursive and deterministic (i.e., EpiFilter produces the same output for fixed input data and requires no Monte Carlo steps) inference algorithm that is quick and easy to compute both across an unfolding outbreak and in retrospect. Our method solves what is called the smoothing problem in control engineering [14]. This means we compute instantaneous reproduction number,  $R_s$ , estimates that formally integrate both forward and backward looking information. This unifies the WT method and EpiEstim, and largely nullifies their edge-effect issues.

Further, EpiFilter only makes a minimal Markov assumption for  $R_s$ , which allows it to avoid the strong prior window size and change-point assumptions that existing methods may apply to infer shifts in transmission [9, 11]. Using simulated and empirical data, we demonstrate that EpiFilter accurately tracks changes in  $R_s$  and provides reliable one-step-ahead incidence predictions. Moreover, we show that EpiFilter is appreciably more robust and statistically efficient than even optimised versions of EpiEstim [13]. We find that EpiFilter does not easily destabilise when performing real-time inference in periods of low incidence, such as in lulls between epidemic waves, and that generally (i.e., for outbreaks of any size) it minimises the mean squared error (MSE) of  $R_s$  estimates, while maintaining good coverage and predictive performance.

We illustrate the practical utility of EpiFilter on the COVID-19 incidence curve of New Zealand, which exhibits a second wave that was seeded during a prolonged low incidence period. We find stark improvements in the transmission patterns EpiFilter uncovers. Our method, which is outlined in Fig 1, provides a straightforward yet formally optimal (in mean squared error sense [18]) solution to real-time and retrospective instantaneous reproduction number estimation. Because it couples minimal prior assumptions with maximum information extraction, it more gracefully handles periods with scarce data. Hopefully our approach will serve as a useful inference tool for investigating the transmissibility risks of COVID-19 and other epidemics. Matlab and R implementations of EpiFilter are available at <https://github.com/kpzoo/EpiFilter> and its mechanics are explored and validated in the S1 Appendix.



**Fig 1. EpiFilter algorithm and relationship to other methods.** In the left panels we consider three ways of inferring the instantaneous or effective reproduction number at time  $s$ ,  $R_s$ , from the incidence curve,  $I_t^s$  (blue dots). The filtering solution produces the posterior distribution  $\mathbf{p}_s$  from all data prior to time  $s$ . EpiEstim approximates this solution by using the subset of data in a window of size  $k$  into the past. Reverse-filtering considers the complementary part of the incidence curve, leading to  $\mathbf{r}_s$ , which utilises data beyond  $s$ . The WT method, with future window  $k$ , approximates this type of solution. Smoothing uses all information from  $I_t^s$  to generate  $\mathbf{q}_s$ , which is precisely computed by EpiFilter. Blue windows show the portions of  $I_t^s$  that inform on  $R_s$  for each of  $\mathbf{p}_s$ ,  $\mathbf{r}_s$  and  $\mathbf{q}_s$  while red windows highlight the subsets used by EpiEstim and the WT method. Double arrows indicate data used for constructing various posterior distributions, while square arrows pinpoint instances of those distributions at the edges of  $I_t^s$ . In the right panels we summarise the construction of EpiFilter. We outline the main assumptions (the model box) and computations (the algorithm box) necessary for realising EpiFilter, which allow us to obtain the most informative and minimum mean squared error (MSE) smoothing posterior distribution  $\mathbf{q}_s$ . See the main text for the specific equations employed in our implementation [14].

<https://doi.org/10.1371/journal.pcbi.1009347.g001>

## Methods

### Renewal models and inference problems

We consider an infectious disease epidemic observed over some time period  $1 \leq s \leq t$  in a homogeneous and well-mixed population. While epidemics actually spread on dynamic networks involving stratified contact structures, homogeneous models can provide useful real-time insight into key transmission patterns and are more easily fit and verified with routine surveillance data such as incidence curves [19]. If the incidence or number of newly infected cases at time  $s$  is  $I_s$  then the set  $I_t^s = \{I_1, I_2, \dots, I_t\}$  is the incidence curve of the epidemic. We assume that incidence is available on a daily scale so that  $I_t^s$  is a vector of  $t$  daily counts but weeks or months could be used instead. A common problem in infectious disease is the inference of the transmissibility of the epidemic given this curve. The renewal model [1, 20] presents a general and popular framework for investigating this problem and its estimates, in some instances, can even approximate those from detailed network models [21].

The renewal model posits that epidemic transmissibility, summarised by the effective or instantaneous reproduction number,  $R_s$ , generates the observed incidence as in Eq (1). We

assume incidence counts are Poisson distributed with  $\text{Pois}$  indicating Poisson noise. Here  $\stackrel{d}{=}$  signifies equality in distributions and  $|$  means ‘conditioned on’.

$$\mathbb{P}(I_s | R_s, I_1^{s-1}) \stackrel{d}{=} \text{Pois}(\Lambda_s R_s) \quad (1)$$

While Eq (1) does not directly model how susceptible individuals become infected, these effects are encoded in the reproduction number  $R_s$ , which measures the secondary cases generated per effective primary case at  $s$  [3]. The quantity  $\Lambda_s := \sum_{u=1}^{s-1} I_{s-u} w_u$ , known as the total infectiousness, counts how many effective past cases are still infectious at  $s$  i.e., it describes the number of circulating cases that can actively transmit.

The generation time distribution of the epidemic,  $w_1^{s-1} = \{w_1, w_2, \dots, w_{s-1}\}$ , controls how past incidence influences  $\Lambda_s$ , with  $w_u$  as the probability that a primary case takes between  $u-1$  and  $u$  days to generate a secondary case [1]. We make the standard assumption that  $w_1^\infty$  is well approximated by the serial interval distribution of the epidemic of interest, which is known [11]. The serial interval is the time between symptom onset of a primary and its secondary case. While onset times are more practically measurable than actual infection times they do not include asymptomatic or subclinical cases. This approximation can limit inference (for example, serial intervals often have larger variances than generation times), but methods are being developed to improve its quality [22].

We focus on inferring the complete set of  $R_s$  values, denoted  $R_1^t = \{R_1, R_2, \dots, R_t\}$ , given the incidence curve  $I_1^t$  with  $t$  as the last recorded time. Estimating  $R_1^t$  is important because changes in the values of instantaneous reproduction numbers often signify key transitions in epidemic transmissibility, which might be due to the imposition or relaxation of interventions. Instantaneous reproduction numbers are also the basis of other transmissibility metrics, such as cohort reproduction numbers or growth rates [3]. We define three main inference problems, based on how information in  $I_1^t$  is recruited to construct every  $R_s$  estimate. We represent these problems in terms of the posterior distribution their solutions induce over possible  $R_s$  values. Estimates are functions of these posteriors.

The first is called filtering, where we sequentially compute the filtering posterior  $\mathbf{p}_s = \mathbb{P}(R_s | I_1^s)$  for every  $s \leq t$  [23]. Filtering only uses incidence data up to time  $s$ ,  $I_1^s$ , for inferring  $R_s$ . Solving this problem is fundamental to real-time inference [24]. Filtering solutions are commonly employed for inferring instantaneous reproduction numbers. The second problem, which we call reverse-filtering, is the complement of the first. The reverse-filtering posterior  $\mathbf{r}_s = \mathbb{P}(R_s | I_1^t)$ , is important for retrospective or backward-looking estimates and infers  $R_s$  from incidence beyond  $s$  i.e.,  $I_1^t$  [25]. Practical  $R_s$  calculations do not involve reverse-filtering. Instead, the information used by  $\mathbf{r}_s$  is implicit to deriving cohort reproduction numbers,  $R_s^c = \sum_{j=0}^{t-s} R_{s+j} w_{j+1}$  [12]. Later we show that estimates from EpiEstim and the WT method are related to the  $\mathbf{p}_s$  and  $\mathbf{r}_s$  distributions, respectively.

The last problem, which is termed smoothing, is our main interest. It asks the question: how can we construct an  $R_s$ , at every  $s \leq t$ , that integrates all available incidence information from  $I_1^t$ . To solve this problem we must formulate the smoothing posterior distribution  $\mathbf{q}_s = \mathbb{P}(R_s | I_1^t)$  [14]. Functions of this posterior would then yield maximally informed instantaneous reproduction number estimates (and cohort reproduction number estimates by extension). Note that  $\mathbf{p}_s$ ,  $\mathbf{r}_s$  and  $\mathbf{q}_s$  depend on the choices of state space model, which describes the dynamics of  $R_s$  across time, and observation model, which explains how changes in  $R_s$  lead to trends in observed incidence data [18]. These models encode our assumptions about the epidemic of interest and determine how estimates trade off assumptions against data. We next

explore how posterior distribution selection determines performance, especially when data are scarce, and examine how EpiEstim and the WT methods fit within this framework.

### Inference methods and low incidence

We mostly detail EpiEstim as instantaneous reproduction number,  $R_s$ , estimates are the main focus of this work. EpiEstim assumes that the estimate of  $R_s$  at time  $s$  depends on a rolling past window of data,  $\tau(s) = \{s, s - 1, \dots, s - k + 1\}$ , of size  $k$  [11]. Consequently, at time  $s$ , only size  $k$  subsets of the incidence,  $I_{s-k+1}^s$ , and total infectiousness,  $\Lambda_{s-k+1}^s = \{\Lambda_s, \Lambda_{s-1}, \dots, \Lambda_{s-k+1}\}$ , are considered informative about  $R_s$ . Their sums over this window are  $i_{\tau(s)} = \sum_{u \in \tau(s)} I_u$  and  $\lambda_{\tau(s)} = \sum_{u \in \tau(s)} \Lambda_u$ . When  $k = 1$ , EpiEstim only uses the most recent case count,  $I_s$ , (and  $\Lambda_s$ ) to estimate  $R_s$ . While this maximises flexibility, it usually results in over-fitting and so larger windows are often employed to trade off estimate variance with bias [11, 26].

EpiEstim therefore effectively solves a filtering problem, as discussed in the previous section. The filtering posterior distribution produced by EpiEstim is restricted to the informative window  $\tau(s)$  and denoted  $\mathbf{p}_{\tau(s)} = \mathbb{P}(R_s | I_{s-k+1}^s)$ . This is parametrised by the shape-scale gamma posterior distribution in Eq (2).

$$\mathbf{p}_{\tau(s)} \stackrel{d}{=} \text{Gam}(a + i_{\tau(s)}, (c + \lambda_{\tau(s)})^{-1}) \tag{2}$$

The posterior  $\mathbf{p}_{\tau(s)}$  results from combining a gamma prior distribution on  $R_s$ ,  $\mathbb{P}(R_s) \stackrel{d}{=} \text{Gam}(a, c^{-1})$ , with a Poisson observation likelihood for the incidence (see Eq (1)), as detailed in [11, 13]. We never explicitly include  $\Lambda_s$  terms in our notation (e.g., we could write  $\mathbf{p}_{\tau(s)}$  as  $\mathbb{P}(R_s | I_{s-k+1}^s, \Lambda_{s-k+1}^s)$ ) since they appear naturally when using Eq (1) and the key difference among our inference problems relate to  $I_s$  terms.

The posterior mean estimate from Eq (2) is constructed as  $\tilde{R}_{\tau(s)} = \int \mathbf{p}_{\tau(s)} R_s dR_s = \mathbb{E}[R_s | I_{s-k+1}^s], (a + i_{\tau(s)})(c + \lambda_{\tau(s)})^{-1}$ . The variance around this estimate is  $(a + i_{\tau(s)})(c + \lambda_{\tau(s)})^{-2}$ . The observation model is given by Eq (1) but the state space model of EpiEstim is not explicit. However, if  $R_{\tau(s)}$  is the assumed average reproduction number in  $\tau(s)$  (which is used to estimate  $R_s$ ) then  $\lambda_{\tau(s)} R_{\tau(s)} = \sum_{u \in \tau(s)} \Lambda_u R_u$  [13]. Thus, EpiEstim somewhat incorporates a linear moving average state space model, and assumes that the filtering distribution  $\mathbf{p}_s \approx \mathbf{p}_{\tau(s)}$  by deeming data outside  $\tau(s)$ ,  $I_1^{s-k}$ , as effectively uninformative [12]. Since  $\mathbf{p}_{\tau(s)}$  can be computed sequentially across an ongoing epidemic, EpiEstim provides real-time inference.

The WT method takes a complementary approach to EpiEstim, computing transmissibility over a forward-looking window  $\gamma(s) = \{s, s + 1, \dots, s + k - 1\}$  [2]. Often  $k = t - s + 1$  i.e., the window extends to the last observed incidence. The WT method uses the observation model of Eq (1) and has an implicit moving average state model  $R_{\gamma(s)} = \sum_{u \in \gamma(s)} R_u w_{u-s+1}$ , which leads to its cohort reproduction number estimates [3]. As this method effectively uses future information [12], it implicitly involves approximating the reverse-filtering distribution (see previous section) as  $\mathbf{r}_s \approx \mathbf{r}_{\gamma(s)} = \mathbb{P}(R_s | I_s^{s+k-1})$ . We illustrate the information windows employed by the WT method and EpiEstim, as well as the complete filtering and reverse-filtering windows in Fig 1. The goodness of the windowed distributions  $\mathbf{p}_{\tau(s)}$  and  $\mathbf{r}_{\gamma(s)}$  as approximations to the general posterior distributions  $\mathbf{p}_s$  and  $\mathbf{r}_s$  will depend on  $k$  and the appropriateness of the state model underlying each method [13].

While EpiEstim and the WT methods are powerful tools for inferring transmissibility in real-time and in retrospect, they have two main and related limitations, which necessarily reduce the reliability of their outputs [12]. First, their performance degrades as  $s$  gets close to 1 for EpiEstim and  $t$  for the WT method [11]. These edge or censoring effects correspond, at the extreme, to

$\mathbf{p}_1 = \mathbf{p}_{\tau(1)} = \mathbb{P}(R_1 | I_1)$  and  $\mathbf{r}_t = \mathbf{r}_{\gamma(t)} = \mathbb{P}(R_t | I_t)$ , which are weakly informed posterior distributions. As a result, at the beginning of the incidence curve EpiEstim can be unreliable (and even unidentifiable). The WT method suffers similarly at the end of that curve [1] (see Fig 1).

The second limitation occurs in phases of the epidemic where incidence is low for a prolonged period of time [17, 26]. In these periods data are sparse and the quality of estimates depend on how well the method of choice mediates between the little available information and its inherent assumptions [15]. We illustrate this with EpiEstim, using the mean and variance of the posterior estimate from Eq (2),  $\tilde{R}_{\tau(s)}$ , which are defined above. If incidence is small over the window  $\tau(s)$  then the sum of incidence,  $i_{\tau(s)}$ , and the total infectiousness  $\lambda_{\tau(s)}$  shrink, meaning that the prior hyperparameters,  $a$  and  $c$ , strongly influence the resulting estimate mean and variance. This contrasts the data-rich scenario when an epidemic is large, where  $i_{\tau(s)}$  and  $\lambda_{\tau(s)}$  overpower  $a$  and  $c$ .

Further, should a sequence of  $n \geq k$  zero-incidence days occur, then  $i_{\tau(s)} = 0$  and  $\lambda_{\tau(s)} \rightarrow 0$  as  $n$  increases, with a rate controlled by the serial interval of the epidemic. Here the shape parameter of  $\mathbf{p}_{\tau(s)}$  is exactly that of the prior distribution  $\mathbb{P}(R_s)$ . As  $\lambda_{\tau(s)}$  decays,  $\mathbf{p}_{\tau(s)} \rightarrow \mathbb{P}(R_s)$  (see Eq (2)),  $R_s$  becomes statistically unidentifiable from the window of data and inference is completely prior driven [26, 27]. While lack of data is a fundamental limitation, the point at which we lose inferential power is not fixed, and depends on the window size,  $k$ . Studies that formally optimised  $k$  for estimate reliability, found that small  $k$  is needed to infer sharp changes in transmissibility (e.g., due to lockdowns) [13], indicating that these issues can be acute. Analogous effects occur in the WT method if there are few incident cases across its forward-looking window  $\gamma(s)$  [12].

These prior-driven scenarios are realistic for epidemics in waning or tail phases, and can precede either elimination (i.e., epidemic extinction) or resurgence [28]. While some estimate degradation is guaranteed for any  $R_s$  inference method when faced with either edge-effects or low incidence, robustness can still be improved. Edge-effects can be largely overcome by constructing the smoothed posterior distribution for estimating the instantaneous reproduction number  $R_s$ , denoted  $\mathbf{q}_s$ . Solving the smoothing problem melds the advantages of the opposite looking windows of EpiEstim and the WT method, removing the vulnerability near the ends of the incidence curve  $I_t^*$ . This follows as  $\mathbf{q}_1 = \mathbf{r}_1$  and  $\mathbf{q}_t = \mathbf{p}_t$  (see Fig 1). Further, by maximising the information used for inferring every  $R_s$  and by minimising our state model assumptions, we can ameliorate the impact of low incidence. We next develop a method, termed EpiFilter, to realise these improvements.

### Bayesian (forward) recursive filtering

We reformulate the inference problem of estimating instantaneous reproduction numbers  $R_s$  from past incidence  $I_1^*$  as an optimal Markov state filtering problem. Filtering describes a general class of engineering problems aimed at optimally, usually in a MSE sense, inferring some hidden state in real time from noisy observations [14, 18]. Given some functions  $f_s$  and  $g_s$ , which describe the state ( $R_s$  in our case) space dynamics and the process of generating noisy observations (the  $I_s$  here), the filtering problem tries to construct the posterior distribution  $\mathbf{p}_s$  (see previous section) [23], which EpiEstim approximates. The conditional mean estimate  $\tilde{R}_s = \mathbb{E}[R_s | I_1^*]$  leads to the minimum MSE of  $\mathbb{E}[(R_s - \tilde{R}_s)^2]$  [23], which depends on all the past information.

The famed Kalman filter [24] was the genesis of these methods. Here we focus on Bayesian recursive filters for models with noisy count observations. These generalise the Kalman filter [23] and have been successfully applied to similar problems in phylodynamics and computational biology [29–31]. We reconsider our renewal model inference problem within this

engineering state-observation framework, as described in Eq (3) [14].

$$R_s = f_s(R_{s-1}, \epsilon_{s-1}), \quad I_s = g_s(R_s, v_s) \quad (3)$$

Here  $R_s$  is the hidden Markov state that we wish to infer. It dynamically depends on the previous state  $R_{s-1}$  and a noise term  $\epsilon_{s-1}$  via  $f_s$ . Observation  $I_s$  is then elicited due to  $R_s$  and a noise term  $v_s$ , according to  $g_s$  [32].

We develop our filter under two very mild assumptions. First, we define some closed space,  $\mathcal{R}$ , over which  $R_s$  is valid. For a given resolution  $m$ , extrema  $R_{\min}$  and  $R_{\max}$ , and grid size  $\delta = m^{-1}(R_{\max} - R_{\min})$  then  $\mathcal{R} := \{R_{\min}, R_{\min} + \delta, \dots, R_{\max}\}$ . This means the instantaneous reproduction number  $R_s$  must take a discrete value in  $\mathcal{R}$ , the  $i^{\text{th}}$  element of which is denoted  $\mathcal{R}[i]$ . We formalise this notion in Eq (4a).

$$\sum_{i=1}^m \mathbb{P}(R_s = \mathcal{R}[i]) = 1, \quad 1 \leq s \leq t \quad (4a)$$

$$R_s = R_{s-1} + (\eta\sqrt{R_{s-1}})\epsilon_{s-1} \quad (4b)$$

This is not restrictive since we can compute our filter for large  $m$  if needed and usually we are only interested in  $R_s$  on a coarse scale (e.g., policymakers may only want to know if  $R_s \leq 1$  or not). Other approaches, which depend on MCMC or related sampling methods (e.g., [9] and some implementations of EpiEstim), all implicitly assume some discretisation [29]. In Fig A of the S1 Appendix we show that convergence often occurs at small  $m$ .

Second, we propose a linear model for  $f_s$ , as defined in Eq (4b). There  $\epsilon_{s-1}$  is a standard white noise term i.e.,  $\mathbb{P}(\epsilon_{s-1}) \stackrel{d}{=} \text{Norm}(0, 1)$  with Norm signifying a normal distribution and  $\eta$  as a free parameter. We assume that a noisy linear projection of states over consecutive time-points provides a good approximation of the state trajectory. Not only is this assumption standard in engineering [23] and epidemiology [33] but it is also more flexible than the state model inherent to EpiEstim and the WT method. We scale the noise of this projection by a fraction,  $\eta < 1$ , of the magnitude of  $\sqrt{R_{s-1}}$ . This parameter controls the correlation among successive instantaneous reproduction numbers (and hence the state noise) but is set to ensure that  $R_s$  is a-priori non-negative.

Our observation model,  $g_s$ , is implicit and leads to the probability law in Eq (1). As a result, both our observations and state models are discrete (see Fig 1 for summaries). Because the state model governing  $R_s$  in Eq (4b) is stochastic, Eq (1) actually describes an over-dispersed (doubly stochastic) Poisson incidence curve. Consequently,  $\eta$ , allows us to partially model some heterogeneity in transmissibility, and may increase robustness to violations of the well-mixed assumption inherent to renewal models [21]. We can optimise our choice of  $\eta$  value by minimising the square error of the one-step-ahead predictions that result from our observation model [13].

We now define the Bayesian recursive filtering procedure, which is a main contribution of this work, and can be solved exactly, in real-time and with minimal computational effort. We adapt general recursive filtering equations from [14, 18, 25, 32], which are valid for various types of observation and state models, to our renewal model inference problem. The proof of the equations we employ can be found in these works. While we solve discrete, univariate problems (our state model is one dimensional), extensions to continuous-time, multivariate problems also exist [23, 30]. These recursive equations can also be approximately solved using particle filters [14, 32].

Recursive filtering involves two steps: prediction and correction. The first, given in Eq (5a), constructs a sequential prior predictive distribution,  $p_s = \mathbb{P}(R_s | I_1^{s-1})$ . This is informed by past



incidence data  $I_1^{s-1}$  and the last state  $R_{s-1}$ . The second step then corrects or updates this prior prediction into a posterior filtering distribution,  $\mathbf{p}_s$ , which constrains  $p_s$  using the latest observation,  $I_s$ , according to Eq (5b).

$$p_s = \int \mathbb{P}(R_s | R_{s-1}, I_1^{s-1}) \mathbf{p}_{s-1} dR_{s-1} \tag{5a}$$

$$\mathbf{p}_s \propto \mathbb{P}(I_s | R_s, I_1^{s-1}) p_s \tag{5b}$$

Here  $\mathbb{P}(R_s | R_{s-1}, I_1^{s-1}) \stackrel{d}{=} \text{Norm}(R_{s-1}, \eta^2 R_{s-1})$  is the state model from Eq (4b),  $\mathbb{P}(I_s | R_s, I_1^{s-1})$  is the observation model from Eq (1) and the constant of proportionality for Eq (5b) is simply a normalising factor.

Solving Eq (5) iteratively and simultaneously over the grid of  $\mathcal{R}$  leads to our novel real-time estimate of the time-varying effective reproduction number. We initialise this process with a uniform prior distribution over  $\mathcal{R}$  for  $\mathbf{p}_1$  and note that  $p_s$  and  $\mathbf{p}_s$  are  $m$  element vectors that sum to 1, with  $i^{\text{th}}$  term corresponding to when  $R_s = \mathcal{R}[i]$ . Eq (5) forms the first half of EpiFilter, is flexible and can be adapted to many related problems [14]. A key difference between EpiFilter and the EpiEstim-type methods [11, 13] is that the latter approximate the distributions  $\mathbb{P}(R_s | I_1^s)$  and  $\mathbb{P}(R_s | I_1^{s-1})$  with  $\mathbb{P}(R_s | I_{s-k+1}^s)$  and  $\mathbb{P}(R_s)$ , respectively. Estimators based on these approximations can be suboptimal, especially when data (i.e., cases) are scarce.

### Bayesian (backward) recursive smoothing

While Eq (5) provides a complete real-time solution to the filtering problem, it is necessarily limited at the starting edge of the incidence curve, where past data are sparse or unavailable. Further, because it does not update past estimates as new data accumulate, it cannot provide optimal retrospective estimates. Here we develop the second half of EpiFilter, which involves solving the optimal smoothing problem i.e., computing the smoothing posterior distribution  $\mathbf{q}_s = \mathbb{P}(R_s | I_1^t)$ , which provides maximally informed estimates of the instantaneous reproduction number  $R_s$ , given the complete incidence curve  $I_1^t$ . To our knowledge, smoothing has not yet been explicitly considered in infectious disease epidemiology (either exactly or approximately).

We specialise the general methodology from [14, 25] to obtain the recursive smoother of Eq (6a). This equation uses the filtering distribution,  $\mathbf{p}_s = \mathbb{P}(R_s | I_1^s)$  and the predictive distributions  $p_{s+1} = \mathbb{P}(R_{s+1} | I_1^s)$ , which we obtain from Eq (5). Our state model means that  $\mathbb{P}(R_{s+1} | R_s, I_1^s) \stackrel{d}{=} \text{Norm}(R_s, \eta^2 R_s)$ .

$$\mathbf{q}_s = \mathbf{p}_s \int \mathbb{P}(R_{s+1} | R_s, I_1^s) \mathbf{q}_{s+1} p_{s+1}^{-1} dR_{s+1} \tag{6a}$$

$$\mathbf{q}_s \propto \mathbf{r}_s \mathbf{p}_s \mathbb{P}(R_s)^{-1}, \text{ if } \mathbf{r}_s \approx \mathbb{P}(R_s | I_{s+1}^t) \tag{6b}$$

We realise Eq (6a) exactly by taking a forward-backward algorithmic approach (this is the backward pass whereas Eq (5) is the forward one). We solve this equation by noting that  $\mathbf{q}_t = \mathbf{p}_t$  and iterating backwards in time to obtain the first smoothing distribution  $\mathbf{q}_1$ . The integrals become sums over our grid  $\mathcal{R}$  and distributions are  $m$  element vectors. Eq (6a) sequentially updates our earlier filtering solutions to include future data, forms the second half of EpiFilter and can also be approximately solved using particle smoothers [14].

This approach neatly links the filtering and smoothing distributions  $\mathbf{p}_s$  and  $\mathbf{q}_s$ . If we assume that the reverse-filtering distribution  $\mathbf{r}_s$  is reasonably approximated by  $\mathbb{P}(R_s | I_{s+1}^t)$  then we can

also apply the two-filter smoothing solution of [25] to get Eq (6b). If either future ( $I_{s+1}^t$ ) or past ( $I_1^{s-1}$ ) incidence is uninformative then either  $\mathbf{r}_s$  or  $\mathbf{p}_s$  will reduce to the prior  $\mathbb{P}(R_s)$ , leading to  $\mathbf{q}_s \propto \mathbf{p}_s$  or  $\mathbf{q}_s \propto \mathbf{r}_s$ , respectively. The end and beginning of the epidemic provide important examples of each of these scenarios. Consequently, Eq (6b) shows how smoothing connects EpiEstim and the WT methods, and explains why EpiFilter, which can be used for both real-time and retrospective inference, better overcomes edge-effects and periods of low data.

Further, the smoothed posterior  $\mathbf{q}_s$  yields the conditional mean estimate  $\hat{R}_s = \mathbb{E}[R_s | I_1^t]$ , which is known to significantly improve on the MSE of the filtered equivalent  $\tilde{R}_s$  (see previous section) [23]. While filtering provides the minimum MSE estimator of every instantaneous reproduction number  $R_s$  given past knowledge, smoothing provides the minimum given all knowledge. This relationship is formal, with filtered and smoothed MSE values mapping to the amount of mutual information that  $I_1^t$  provides about  $R_1^t$  [34]. Extracting the maximum information from the incidence curve  $I_1^t$  should engender estimates that are more robust and statistically efficient in periods of low incidence. We summarise the EpiFilter algorithm in Fig 1.

While our main interest is on optimised and rigorous real-time and retrospective estimates of transmissibility, which are completely defined by the smoothing distribution  $\mathbf{q}_s$ , we may also want to predict future incidence, for informing epidemic preparedness plans and for validating past  $R_s$  estimates [13, 35]. We compute the filtered one-step-ahead posterior predictive distribution as in Eq (7) (integrals are over  $\mathcal{R}$ ) [14].

$$\mathbb{P}(I_{s+1} | I_1^s) = \int \mathbb{P}(I_{s+1} | R_s, I_1^s) \mathbf{p}_s \, dR_s \quad (7)$$

We assume, as in [36], that  $\mathbb{P}(I_{s+1} | R_s, I_1^s) \stackrel{d}{=} \text{Pois}(\Lambda_{s+1} R_s)$ . Replacing  $\mathbf{p}_s$  with  $\mathbf{q}_s$  yields the smoothed equivalent of Eq (7). We will use Eq (7) to compare EpiFilter against APEestim, which is the prediction-optimised version of EpiEstim, developed in [13]. Since predictions depend strongly on  $\mathbf{p}_s$  or  $\mathbf{q}_s$ , optimising these distributions can be important, for example, when forecasting second waves of infection.

Last, we comment on how we represent uncertainty in our estimates and predictions. Often we will provide 95% equal tailed Bayesian credible intervals. These are computed directly from the 2.5<sup>th</sup> and 97.5<sup>th</sup> quantiles of the relevant posterior distribution and used to assess performance statistics such as coverage of true values.

## Results

### Improved estimation at low incidence

The reliable estimation of time-varying effective reproduction numbers,  $R_s$ , at low incidence,  $I_s$ , is a key challenge limiting our understanding of transmission [17]. Periods with small counts of new cases contain little information and so present necessary statistical difficulties [26]. Uncovering the transmissibility patterns in these periods can evidence intervention relaxation strategies, indicate potential elimination and provide insight into mechanisms of spread [28, 50]. Here we compare EpiFilter, which allows exact inference over a state grid  $\mathcal{R}$ , with EpiEstim and APEestim at these data-poor settings. We use EpiEstim with weekly and monthly windows. Weekly windows are the default recommendation in [11]. We include monthly ones as long windows can improve robustness at low incidence [28]. APEestim, developed in [13], optimises the window choice of EpiEstim to minimise one-step-ahead prediction errors. The assumptions and choices inherent in estimation methods become important and visible when data are scarce, and can bias inference or support spurious predictions [15].

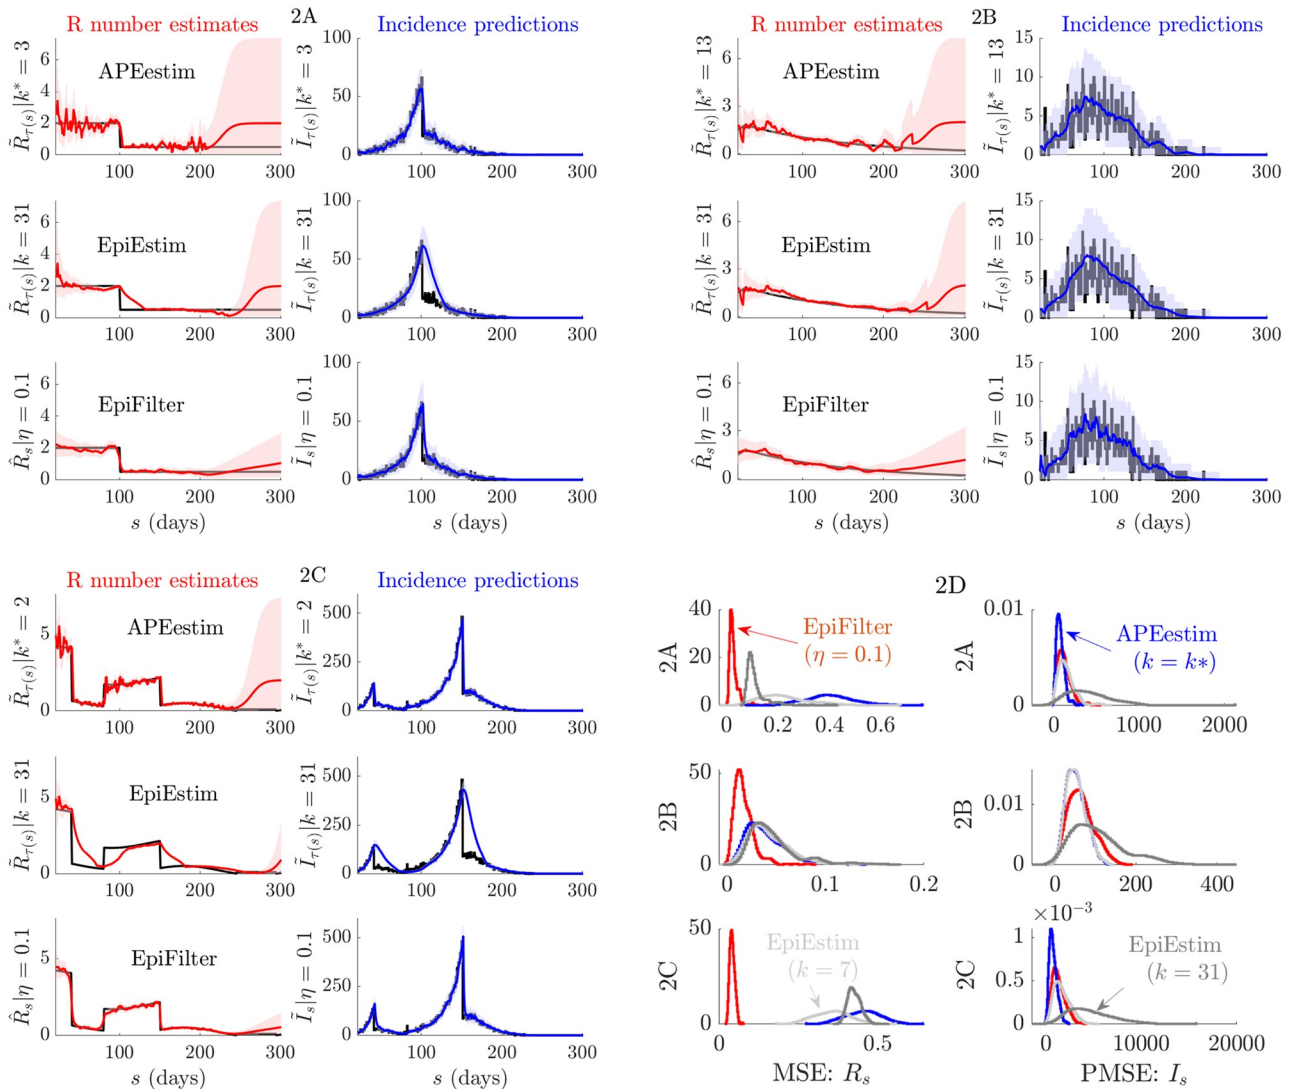
We generate epidemics using the Poisson noise model of Eq (1) under the serial interval distribution of Ebola virus disease described in [37]. We examine three diverse scenarios in Fig 2. These describe (A) rapidly controlled epidemics ( $R_s$  step changes from 2 to 0.5 at  $s = 100$ ), (B) small outbreaks with exponentially rising and falling  $R_s$  (change-point at  $s = 30$  and rates of 0.02 and -0.008 per time unit) and (C) medium outbreaks that are initially controlled ( $R_s$  changes from 4 to 0.6 at  $s = 40$ ) then resurge ( $R_s$  rebounds to 2 at  $s = 80$ ) into large epidemics, before finally being suppressed ( $R_s = 0.2$  from  $s = 150$ ). These scenarios are similar to ones investigated in [12, 28] and describe various epidemic dynamics that culminate in elimination. We simulate 200 incidence curves for each scenario and apply APEestim (optimal window  $k^*$ ), EpiEstim ( $k = 7$  and 31) and EpiFilter (state noise  $\eta = 0.1$ ) to estimate  $R_s$  and sequentially predict  $I_s$  (given data up to  $s - 1$ ) for each curve.

For the first two methods we compute mean one-step-ahead incidence predictions ( $\tilde{I}_{\tau(s)}$ ) as in [13] and instantaneous reproduction number estimates ( $\tilde{R}_{\tau(s)}$ ) from Eq (2) with  $\tau(s)$  delimiting the window times used. We obtain smoothed EpiFilter estimates ( $\hat{R}_s$ ) and filtered predictions ( $\tilde{I}_s$ ) from Eqs (6) and (7). We could use smoothed predictions ( $\hat{I}_s$ ), which implicitly include data beyond  $s - 1$ , to assess fit but as we want to test model adequacy  $\tilde{I}_s$  is more appropriate [38]. Our main focus is on real-time performance so we do not investigate the WT method. See [11, 12] for comparisons of the WT method and EpiEstim. Fig 2A–2C provide representative single runs (we only show  $k = 31$  for EpiEstim as  $k = 7$  is often similar to APEestim) while Fig 2D presents reproduction number MSE and one-step-ahead predicted incidence MSE (PMSE) distributions from the 200 runs.

In the Methods we showed that since EpiEstim-type methods group data over some window into the past, they revert to their prior distribution as the total cases in this window becomes small. During low incidence periods, e.g., when the epidemic is waning, using long windows makes sense [28]. However, this reduces predictive accuracy as fluctuations in  $R_s$  are underfit. This is the use-case for APEestim, which optimises for prediction. The consequences of these trade-offs are made clear in Fig 2A–2C where APEestim provides the best one-step-ahead  $I_s$  predictions but quickly reverts to its prior (seen as wide credible intervals) when cases are few. Longer window EpiEstim improves on this destabilisation but cannot track salient fluctuations in  $R_s$ . Fig 2D verifies these trends. APEestim has the smallest PMSE but often worse MSE than the larger  $k$  EpiEstim choices.

Interestingly, EpiFilter is able to jointly optimise both instantaneous reproduction number tracking and incidence predictions (computed via Eq (7)). In Fig 2A–2C we see that EpiFilter maintains stable and accurate estimates of  $R_s$  and only slowly reverts to its prior (which has the same support as that of EpiEstim and APEestim). However, unlike long window methods it does not sacrifice prediction fidelity. The improvement in MSE shown in Fig 2D is stark (the numerical reduction in MSE when compared to the next best method is on average at least 2-fold and often 10-fold). The PMSE, while larger than that of APEestim (which is optimised for predictions) is still good. These points are reinforced in Fig 3, which expands on the statistics of scenarios A-C.

There the MSE and PMSE computed over the 200 replicate simulations are plotted with time. We observe a clear and significant reduction in MSE at most time points (both at low and large incidence) for EpiFilter, with similar PMSE performance, as compared to EpiEstim and APEestim. This confirms the widespread benefits of smoothing solutions. Moreover, the coverage of both the true  $R_s$  and  $I_s$  of EpiFilter (i.e., the probability that  $R_s$  or  $I_s$  is contained within estimated or predicted 95% equal tailed credible intervals) is more consistent than all other approaches. We present this in Fig C of the S1 Appendix. Thus, EpiFilter combines the



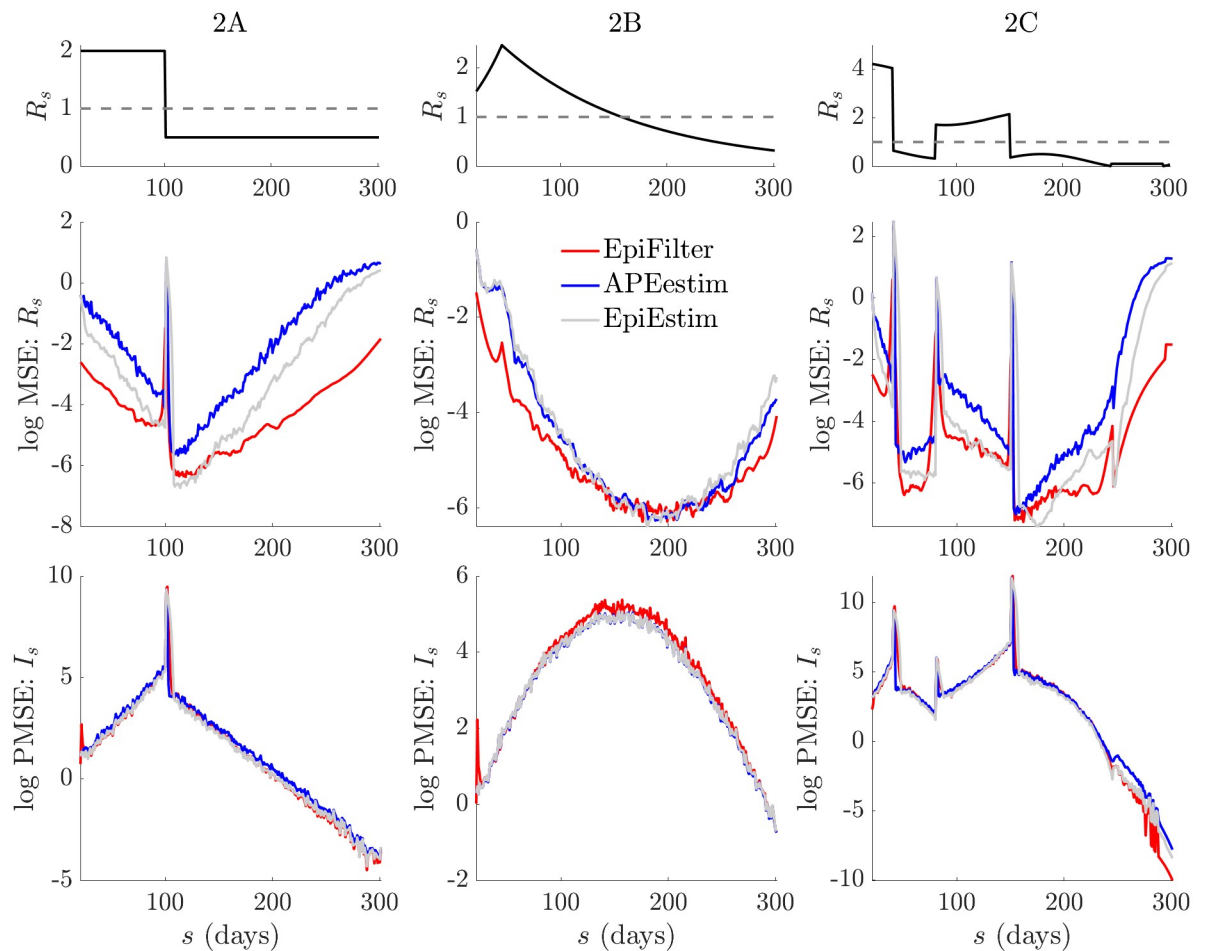
**Fig 2. Small or waning epidemics.** We compare reproduction number estimates ( $\hat{R}_{\tau(s)}$  or  $\hat{R}_s$ ) and one-step-ahead incidence predictions ( $\hat{I}_{\tau(s)}$  or  $\hat{I}_s$ ) from APEestim with optimal window  $k^*$ , EpiEstim with window  $k$  and EpiFilter with state noise  $\eta$ . We simulate 200 epidemics with low daily case numbers or long tails (long sequences of zero cases) using the standard renewal model (Eq (1)) for three scenarios, representative examples of which are given in A-C. The true  $R_s$  and  $I_s$  are in black. All mean estimates or predictions are in red and blue with 95% credible intervals. APEestim and EpiEstim use a Gam(1, 2) prior distribution and EpiFilter a grid with  $m = 2000$ ,  $R_{\min} = 0.01$  and  $R_{\max} = 10$ . In D we provide statistics of the MSE of these estimates (relative to  $R_s$ ) and the PMSE of these predictions (relative to  $I_s$ ) for all 200 runs. We find that EpiFilter is more robust to small incidence (better uncertainty), whereas the other approaches can quickly decay to their prior distribution. It achieves significantly smaller MSE (2–10 fold reductions) and comparable PMSE to APEestim (which is optimised for prediction).

<https://doi.org/10.1371/journal.pcbi.1009347.g002>

advantages of APEestim and long-window EpiEstim, strongly minimises estimation error and, further, is able to reliably detect transmission change-points automatically.

### Improved estimation between epidemic waves

Maintaining robust instantaneous reproduction number,  $R_s$ , estimation when incidence,  $I_s$ , becomes small is not just statistically important. Two possible outcomes may follow periods of small  $I_s$ : either the epidemic goes extinct (elimination occurs, as in the previous section), or an



**Fig 3. Temporal statistics of small or waning epidemics.** We expand on the results from Fig 2D by decomposing the MSE and one-step-ahead PMSE statistics across the 200 simulated trajectories for every scenario in Fig 2. We do not consider the  $k = 31$  EpiEstim example given its poor performance. We present the  $k = 7$  case, which is the generally recommended EpiEstim setting. We observe that EpiFilter significantly improves on MSE throughout the epidemic trajectory (and not only in periods of low incidence) while maintaining comparable prediction accuracies. Coverage statistics for these scenarios, which are given in Fig C of the S1 Appendix, confirm that EpiFilter also consistently contains the true  $R_s$  and  $I_s$  values within its credible intervals.

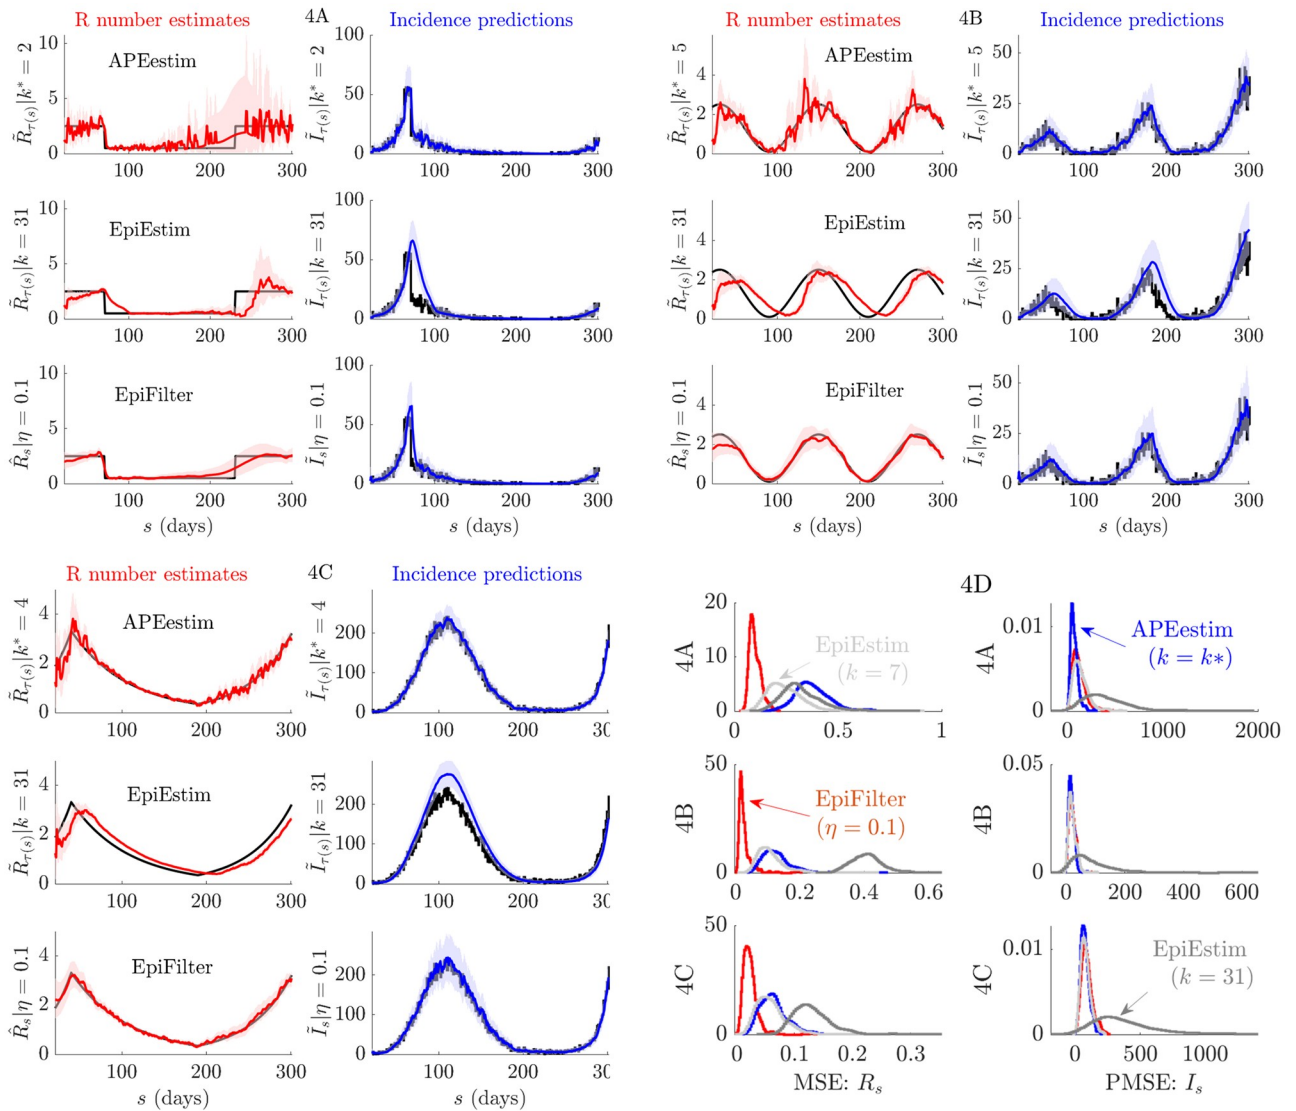
<https://doi.org/10.1371/journal.pcbi.1009347.g003>

additional wave of infection surfaces (resurgence) e.g., due to imports or unmonitored local transmission. Predicting which outcome is likely, in real-time, is of global concern as countries aim to relax interventions during the ongoing COVID-19 pandemic, while also minimising the risks of further cases [17, 39]. As changes in instantaneous reproduction numbers signal variations in transmission and hence incidence, reliably identifying and inferring  $R_s$  trends in the trough preceding potential new peaks can be crucial for preparedness, providing evidence for timely and effective epidemic interventions [10].

Reliable estimation of  $R_s$  between epidemic waves depends on the prior assumptions of the inference method used and on how that method relies on those assumptions when data are scarce [40, 41]. Here we examine this dependence and investigate cases where resurgence follows a low-incidence period. As in the above section, we compare EpiFilter ( $\eta = 0.1$ ) with APEestim (optimal window  $k^*$ ) and EpiEstim (weekly,  $k = 7$ , and monthly,  $k = 31$ , windows) over 200 simulated epidemics under the serial interval of Ebola virus. We explore scenarios

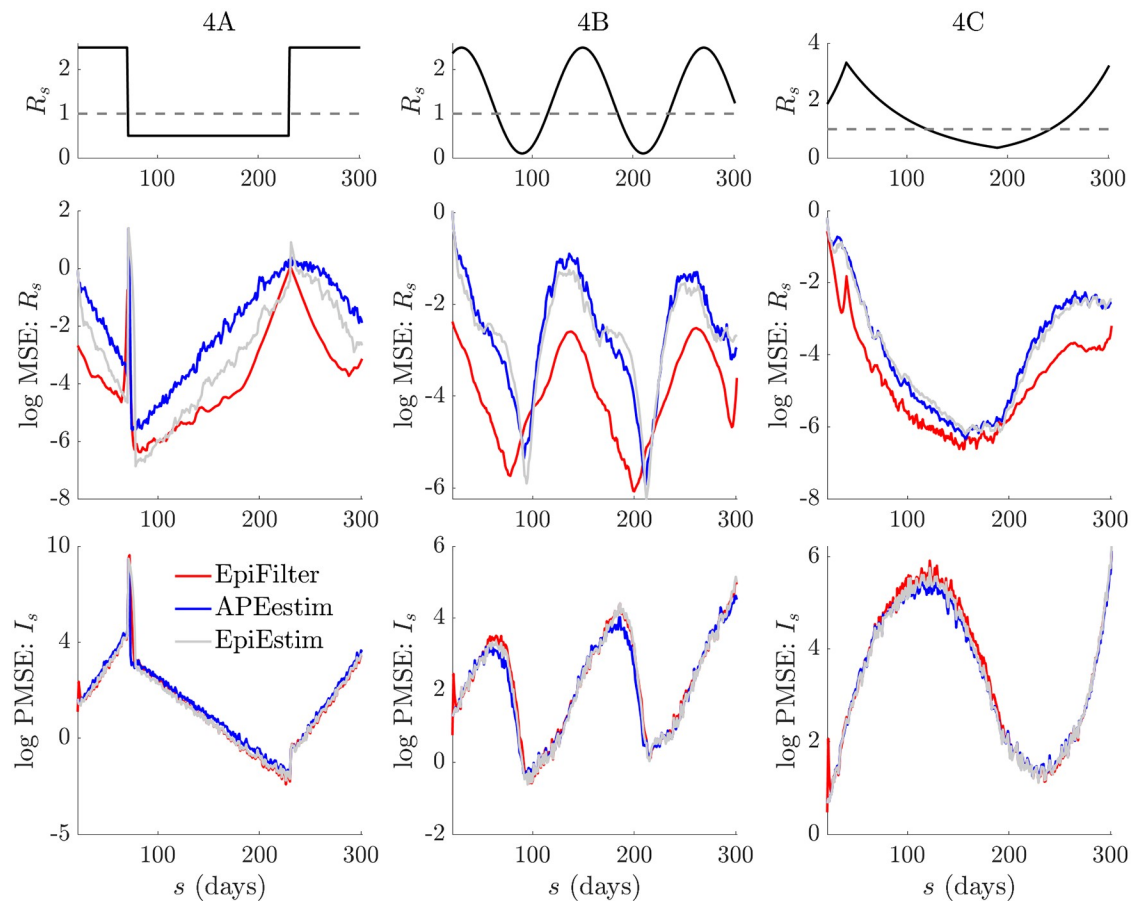
depicting (A) epidemics that are initially controlled ( $R_s$  falls from 2.5 to 0.5 at  $s = 70$ ) but which resurge just as quickly ( $R_s$  returns to 2.5 from  $s = 230$ ), (B) periodic or seasonal transmission ( $R_s$  is sinusoidal with magnitude  $1.3 \pm 1.2$  and period of 120 time units) and (C) outbreaks with exponentially rising and then falling transmissibility (change-points at  $s = 40$  and 190 and exponent rates 0.03, -0.015 and 0.02).

These examples are similar to some in [13] and describe diverse epidemics with multiple peaks and troughs. We provide representative runs of each scenario in A-C and collect MSE (for  $R_s$ ) and PMSE (relative to  $I_s$ ) distributions over the 200 runs of every scenario in Fig 4D.



**Fig 4. Epidemics with multiple waves.** We compare reproduction number estimates ( $\hat{R}_{\tau(s)}$  or  $\hat{R}_s$ ) and one-step-ahead incidence predictions ( $\hat{I}_{\tau(s)}$  or  $\hat{I}_s$ ) from APEestim with optimal window  $k^*$ , EpiEstim with window  $k$  and EpiFilter with state noise  $\eta$ . We simulate 200 epidemics with multiple waves of infection using the standard renewal model (Eq (1)) for three scenarios, representative examples of which are given in A-C. The true  $R_s$  and  $I_s$  are in black. All mean estimates or predictions are in red and blue with 95% equal tailed credible intervals. APEestim and EpiEstim use a Gam(1, 2) prior distribution and EpiFilter a grid with  $m = 2000$ ,  $R_{\min} = 0.01$  and  $R_{\max} = 10$ . In D we provide statistics of the MSE of these estimates (relative to  $R_s$ ) and the PMSE of these predictions (relative to  $I_s$ ) for all 200 runs. We find EpiFilter is best able to negotiate troughs between epidemic peaks and hence infer resurging infectious dynamics, achieving significantly smaller MSE (2–10 fold reductions) and comparable PMSE to APEestim (which is optimised for prediction).

<https://doi.org/10.1371/journal.pcbi.1009347.g004>



**Fig 5. Temporal statistics of epidemics with multiple waves.** We expand on the results from Fig 4D by decomposing the MSE and one-step-ahead PMSE statistics across the 200 simulated trajectories for every scenario in Fig 4. We do not consider the  $k = 31$  EpiEstim example given its poor performance. We present the  $k = 7$  case, which is the generally recommended EpiEstim setting. We observe that EpiFilter significantly improves on MSE throughout the resurgent epidemic trajectory (whether incidence is small or large) while maintaining comparable prediction accuracies. Coverage statistics for these scenarios, which are given in Fig D of the S1 Appendix, confirm that EpiFilter consistently contains the true  $R_s$  and  $I_s$  values within its credible intervals.

<https://doi.org/10.1371/journal.pcbi.1009347.g005>

We also provide these statistics across time in Fig 5. We observe a similar pattern in performance among the methods as in our previous analyses of Figs 2 and 3. APEestim is best able to predict upcoming incidence and achieves the lowest PMSE as expected. While the monthly window EpiEstim is less useful in these cases (since prior reversion does not occur as often, though could be an issue for more extended troughs), the weekly window version loses predictive fidelity for minor MSE improvements.


EpiFilter once again combines the advantages of the other approaches. For every scenario in Fig 4 it provides accurate tracking of changes in  $R_s$  with stable credible intervals and a MSE that is at least  $\frac{1}{2}$  and sometimes even  $\frac{1}{10}$  that of the next best method. The improvement in MSE is clear and maintained at almost every time point of each scenario, regardless of whether incidence is low or high. Concurrently, the incidence PMSE of EpiFilter rivals that of APEestim and it attains the most consistent coverage of the true  $R_s$  and  $I_s$  values, as shown in Fig D of the S1 Appendix. EpiFilter is therefore a proficient tool for inferring resurgence. We also find that the  $\eta = 0.1$  parameter value seems to be an all-purpose heuristic, meaning that usage of EpiFilter can be simpler than EpiEstim and other window or change-point based methods. The

improvements of EpiFilter likely result from its minimal assumptions (see Eq (4)) and its increased information extraction. We next test our method on empirical data.

### COVID-19 in New Zealand and H1N1 influenza in the USA

The previous sections confirmed EpiFilter as a powerful inference and prediction tool, especially in data-poor conditions, using simulated epidemics. We now confront our method with empirical data from the 1918 H1N1 influenza pandemic in Baltimore (USA) [42] and the ongoing COVID-19 pandemic in New Zealand (up to 17 August 2020) [43]. The H1N1 dataset has been well-studied and so we first use this to benchmark EpiFilter. We clean this dataset by applying a 5-day moving average filter as recommended in [42]. Previous work [11] analysed this dataset with EpiEstim and found that sensible instantaneous reproduction number,  $R_s$ , estimates are obtained when a weekly window ( $k = 7$ ) is applied. However, more recent work, using APEestim [13], showed that while  $k = 7$  provides stable estimates for this epidemic, it is a poor predictor of the incidence data. Instead, an optimised window of 2 days ( $k^* = 2$ ) yields good predictions but the resulting  $R_s$  estimates are noisy.

We reproduce the instantaneous reproduction number estimates ( $\tilde{R}_{\tau(s)}$ ) and incidence predictions ( $\tilde{I}_{\tau(s)}$ ) from both studies in Fig 6 and compare them against EpiFilter with  $\eta = 0.1$  ( $\hat{R}_s$  and  $\tilde{I}_s$ ). Top and middle rows of Fig 6 illustrate the aforementioned trade-off between



**Fig 6. H1N1 influenza transmission in Baltimore (1918).** We compare APEestim (top), EpiEstim with recommended weekly window (middle) (both with Gam(1, 2) prior distribution) and EpiFilter (with  $m = 2000$ ,  $\eta = 0.1$ ,  $R_{\min} = 0.01$  and  $R_{\max} = 10$ ) on the H1N1 influenza dataset from [42]. We use a 5-day moving average filter, as in [42], to ameliorate known sampling biases. Estimates of reproduction numbers,  $R_s$ , and corresponding 95% equal tailed credible intervals are in red. One-step-ahead predictions of incidence,  $I_s$ , (with 95% credible intervals) are in blue with the actual incidence in black. We find that EpiFilter combines the benefits of APEestim and EpiEstim, achieving both good estimates and predictions.

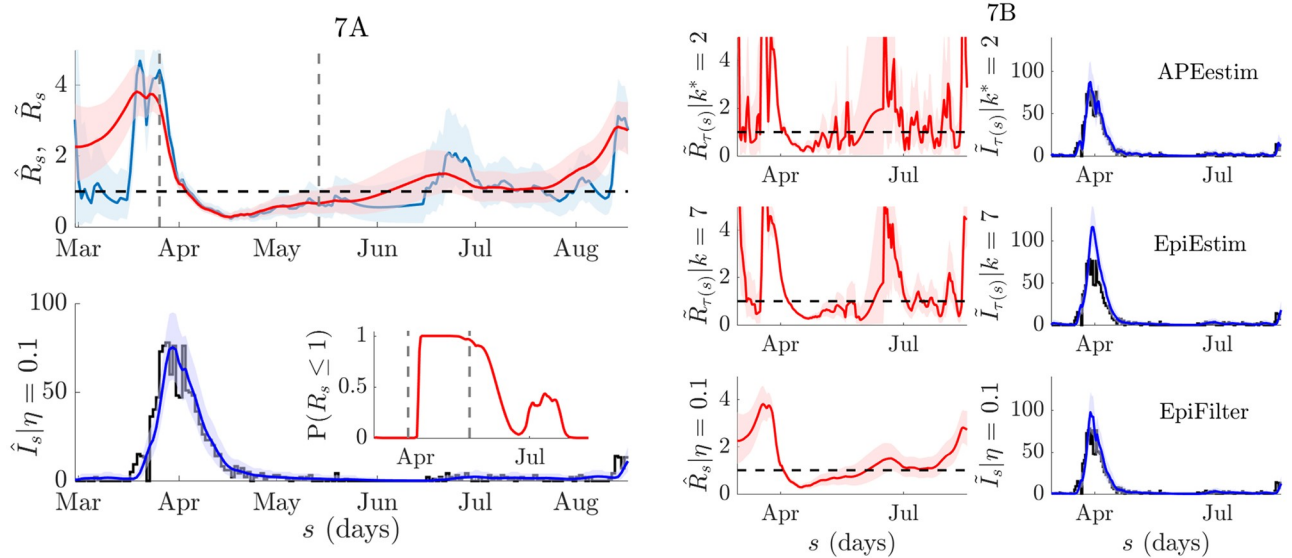
<https://doi.org/10.1371/journal.pcbi.1009347.g006>



estimate stability and prediction accuracy. The bottom row confirms the power of EpiFilter. Our  $R_s$  estimates are of comparable stability to those of EpiEstim at  $k = 7$ , yet our prediction fidelity matches that of APEestim. Our improved inference again benefits from using more information (i.e., the backward pass in Fig 1) and making less restrictive prior assumptions. We see the latter from the  $R_s$  credible intervals over  $40 \leq s \leq 60$ . There EpiEstim seems overconfident, and this results in a rigid overestimation of incidence. However, EpiFilter mediates its estimate uncertainty to a level similar to APEestim.

We explore COVID-19 transmission patterns in New Zealand using incidence data up to 17 August 2020 from [43]. New Zealand presents an insightful case study because officials combined swift lockdowns with intensive testing to achieve and sustain very low incidence levels that eventually led to local elimination of COVID-19 [44]. However, an upsurge in cases in early August inspired concerns about a second wave (which led to new interventions and is why we do not consider data beyond 17 August). Here we investigate the time-varying transmission in New Zealand to see if this uptick suggests that the epidemic was resurfacing in mid-August. We believe smoothing can confer important inferential advantages in exactly these types of low incidence scenarios.

We make the common assumptions that case under-reporting is constant [11], which seems reasonable given the intensive surveillance employed by New Zealand [45]. We ignore reporting delays, which are known to be small [46] and use the COVID-19 serial interval distribution from [47]. We do not explicitly distinguish imported from local cases in our analysis. The latter could bias our study [28, 39] but our focus is on demonstrating differences between filtering and smoothing on instantaneous reproduction number,  $R_s$ , trends and not on providing detailed  $R_s$  estimates during this period. We plot the results of our exploration in Fig 7.



**Fig 7. COVID-19 transmission in New Zealand.** We compute smoothed and filtered reproduction number estimates,  $\hat{R}_s$  (red) and  $\tilde{R}_s$  (blue) respectively, from the COVID-19 incidence curve for New Zealand (available at [43]) in the left panels. We use EpiFilter with  $m = 2000$ ,  $\eta = 0.1$ ,  $R_{\min} = 0.01$  and  $R_{\max} = 10$  with a uniform prior distribution over the grid  $\mathcal{R}$ . The top of 7A shows conditional mean estimates and 95% credible intervals for  $\hat{R}_s$  (red) and  $\tilde{R}_s$  (blue). Vertical lines indicate the start and end of lockdown, a major intervention that was employed to halt transmission. The additional ‘future’ information used in smoothing has a notable effect. The bottom of 7A provides smoothed one-step-ahead predictions  $\hat{I}_s$  (blue, with 95% credible intervals) of the actual reported cases  $I_s$  (black). The inset gives the estimated probability of  $R_s \leq 1$ . We observe a clear trend of subcritical transmission that eventually seeds a second wave by August. In 7B we compare EpiFilter with EpiEstim (using weekly windows) and APEestim (both with Gam(1, 2) priors) with all left subfigures presenting  $R_s$  estimates and right ones providing filtered  $I_s$  predictions. We observe that both APEestim and EpiEstim lead to largely unusable estimates that mask transmission trends, in sharp contrast to EpiFilter.

<https://doi.org/10.1371/journal.pcbi.1009347.g007>

We apply EpiFilter and obtain filtered ( $\tilde{R}_s$ , blue) and smoothed ( $\hat{R}_s$ , red) conditional mean estimates together with their 95% confidence intervals. These are in the left panels of Fig 7 and computed from Eqs (5) and (6) respectively. The times of lockdown and release are included for reference. Interestingly, we see a notable difference in the quality of inference between  $\tilde{R}_s$  and  $\hat{R}_s$ . The former, as expected, is unreliable at the beginning of the incidence curve and features wider uncertainty and noisier trends. The smoothed  $\hat{R}_s$ , by using both forward and backward-looking data largely overcomes these issues and clarifies transmission dynamics. In the right panels we compare EpiFilter with EpiEstim (weekly windows) and APEestim. The difference is striking. Neither APEestim nor EpiEstim recovers a clear trend and both are appreciably worse than even the filtered estimate  $\tilde{R}_s$ .

Our smoothed analysis suggests that  $R_s$  has resurged and supports the re-implementation of measures around 14 August. We recover suppression of the initial wave in April, likely associated with the implementation of key interventions, including lockdown [46]. Following this, we infer a prolonged period of subcritical transmission, where  $\mathbb{P}(R_s \leq 1) \approx 1$ . Low and then zero incidence over this period strikingly destabilises both APEestim and EpiEstim and conceals the rise in  $R_s$  that occurs after July. EpiFilter signals this upsurge, which continually grows until August, where a second wave becomes likely. We also provide one-step-ahead predictions (which are from the smoothed  $\hat{R}_s$  hence the notation  $\hat{I}_s$ ) and their equal tailed 95% credible intervals against the reported incidence from [43]. These verify that  $\hat{R}_s$  reasonably describes the data (we also present  $\tilde{I}_s$  when comparing methods). We find that  $\mathbb{P}(R_s \leq 1) \approx 0$  around July, further supporting this resurgence hypothesis.

While the analyses of the H1N1 and COVID-19 data above illustrate the advantages of EpiFilter, these estimates can be further improved. We used case data and not actual infection times (which relates to the approximation of the generation time by the serial interval distribution) and did not account for introductions, case ascertainment fractions and reporting delays. For more practical analyses it may be necessary to first compensate for these biases to obtain the best possible incidence curve. Methods from [6, 12, 48] can be applied if relevant incubation period, reporting distributions and contact tracing data are available to diagnose and correct for these issues. The resulting pre-processed incidence curve can then be input to EpiFilter to obtain more realistic  $R_s$  estimates.

## Discussion

Estimating time-varying trends in the instantaneous or effective reproduction number,  $R_s$ , reliably and in real-time is an important and popular problem in infectious disease epidemiology [5]. As the COVID-19 pandemic has unfolded, the interest in solving this problem has only elevated with  $R_s$  playing a central role both in aiding situational awareness [8] and informing policymaking [49]. Initially, interest was on understanding how changes in  $R_s$  may correlate with interventions such as lockdowns and social distancing [9, 10]. However, as countries have entered waning phases of the pandemic and vaccine deployment has begun, focus has shifted to charactering how existing interventions can be relaxed with minimum risk [50]. The literature on intervention exit strategies is, however, still in development, and several challenges remain to modelling transmission.

One key challenge lies in understanding and inferring transmissibility during periods when the incidence of new cases is small [17]. Such periods may occur under sustained control measures and necessarily contain limited data, which make inferences difficult. Moreover, it is in these lulls that information on transmission may be crucial, helping to determine if the removal of interventions will lead to resurgence or if elimination is realistic by maintaining

controls [28, 39]. While reproduction numbers are not the only analytic for assessing these outcomes, they do provide an important real-time diagnostic since upticks in  $R_s$  generally precede elevations in case loads. Unfortunately, many current approaches to estimating  $R_s$  become underpowered, unstable or prior-constrained in these data-limited conditions [11, 26, 50]. These problems are only magnified when finer-scale analyses (where cases are fewer by division) are of interest (e.g., regional versus national level estimation).

In this paper we re-examined existing methodology for inferring instantaneous reproduction numbers,  $R_s$ , from an engineering perspective. We observed that two of the most useful and popular inference approaches, EpiEstim [11] and the WT method (this computes cohort reproduction numbers, which are functions of  $R_s$ ) [2], only capitalise on a portion of the data available, deeming either upcoming or past incidence to be informative (see Fig 1) [12]. This informative portion is directly controlled by prior assumptions on the speed of possible  $R_s$  changes, which are often characterised by a window of size  $k$ . Other methods also apply similarly strong change-point or state assumptions on  $R_s$ , explicitly linking its variations with specific dates or events, for example [9, 26]. When data are scarce these assumptions can unduly control or skew inference.

In control engineering a common problem, known as filtering, involves optimally (in a MSE sense) estimating hidden Markov states, in real-time, from noisy and uncertain observations [18]. A related problem termed smoothing provides accompanying and optimal retrospective inferences [14]. By reinterpreting  $R_s$  as a Markov state Eq (4) observed through a noisy renewal process (Eq (1)) and defining  $R_s$  on a predetermined grid  $\mathcal{R}$ , we were able to construct exact filtering (Eq (5)) and smoothing (Eq (6)) solutions. This led to EpiFilter, which is our central contribution. Generally, filtering and smoothing can be involved and require sophisticated sequential Monte Carlo techniques [32]. However, because we make only minimal assumptions about  $R_s$ , modelling it as a simple diffusion, we were able to solve these problems exactly and without complex sampling algorithms [29].

Our solutions are computationally simple, often executing in a few minutes (see Fig A of the S1 Appendix), and deterministic i.e., precisely reproducible given the same data and settings. Our method replaces strong change-point or window size assumptions with one free parameter,  $\eta$ , which allows us to model some heterogeneity in transmission and sets the correlation among successive  $R_s$  values. We find that  $\eta = 0.1$  serves as a general heuristic, providing good estimates and automatically detecting change-points and salient  $R_s$  dynamics over diverse scenarios (see Figs 2 and 4, and Fig B in S1 Appendix). This heuristic is statistically justified by its good one-step-ahead predictive performance and consistent coverage of true simulated values (see Fig C-D of the S1 Appendix) [13].

Importantly, EpiFilter is able to look both forward and backward through the incidence data, and so maximise the information extracted at every time point [34]. This property means it combines advantages from both EpiEstim and the WT method (see Fig 1) and largely ameliorates their edge-effect issues [12]. These benefits, which also hold at large incidence, make EpiFilter a useful and robust tool for both real-time and retrospective  $R_s$  inference. We confirmed the advantages of EpiFilter by comparing it to EpiEstim and APEestim (a prediction optimised analogue to EpiEstim) on many simulated examples with periods of low incidence and epidemic resurgences (Figs 2 and 4). Interestingly, we found EpiFilter was able to achieve significant 2–10 fold reductions in the MSE of  $R_s$  estimates without compromising predictive power or coverage of the true  $R_s$  and  $I_s$  values.

EpiFilter was especially better at negotiating periods of low incidence, offering a graceful degradation to its prior distribution or assumptions without sacrificing predictive accuracy. When incidence is low, it can be beneficial to use longer windows with EpiEstim [28]. This keeps  $R_s$  estimates reasonably stable but often leads to poor predictions [13]. APEestim, which

optimises window size for prediction fidelity, showed that in many of the simulated scenarios short windows are necessary for describing transmission patterns. Consequently, we have a trade-off between estimate robustness and prediction accuracy. We found that EpiFilter overcomes this trade-off, concurrently achieving good estimates and predictions. In doing so, it revealed or clarified subcritical transmission trends and unmasked important signals of resurgence from noisy data in those periods.

We verified the practical utility and performance of EpiFilter on empirical data from the H1N1 pandemic of 1918 (see Fig 6) and COVID-19 in New Zealand (see Fig 7). In the first, which is a standard dataset that has been used to test previous  $R_t$  methods, we found that EpiFilter integrated the benefits of EpiEstim and APEestim to achieve simultaneously good estimates and predictions. A key use-case for EpiFilter is in estimating resurgence during low incidence. The COVID-19 epidemic in New Zealand featured precisely those dynamics [46]. While EpiEstim and APEestim were destabilised and unable to extract clear transmission trends, EpiFilter inferred subcritical  $R_t$  values and warned of upcoming cases by signalling an uptick in  $R_t$  just before a second wave became apparent. Recent, more involved COVID-19 analyses [39], have confirmed EpiFilter as a useful outbreak analytics tool.

Balancing the assumptions inherent to a model against the data it is applied on, to produce reliable inference is a non-trivial problem that is still under active investigation in several fields [15, 40, 41]. EpiFilter, by maximising the information extracted from available incidence data and minimising its state space model assumptions, appears to strike this balance as an estimator of instantaneous or effective reproduction numbers. Consequently, it performs strongly on a wide range of problems, including those involving sparse data, where other methods might struggle. Given its demonstrated advantages, straightforward formulation and theoretical underpinning, we hope that EpiFilter will prove a useful diagnostic tool for robustly signalling waves of infection across multiple scales, improving the uncertainty quantification of transmissibility estimates and assessing dynamical transmission patterns, both in real time and retrospectively. EpiFilter is freely available at <https://github.com/kpzoo/EpiFilter>.

## Supporting information

**S1 Appendix. Validating the mechanics and performance of EpiFilter.** This includes analyses of computational complexity, further examples showing performance and calculation of coverage statistics (additionally supporting Figs 2–5 of the main text). Results are embedded within Figs A–D of the S1 Appendix. **Fig A: EpiFilter mechanics.** We examine the computational complexity (as a function of grid size  $m$  on the examples in Figs 6 and 7) and accuracy (in terms of  $m$  and state noise  $\eta$ ) of EpiFilter. **Fig B: Validation against APEestim.** We compare estimates from EpiFilter with optimised ones from APEestim [13] over various general epidemic scenarios (which may have large or small incidence). **Fig C: Coverage statistics of waning epidemics** We provide the probability that the true values of instantaneous reproduction numbers and incidence from scenarios in Fig 3 are within EpiFilter's 95% credible intervals. **Fig D: Coverage statistics of resurging epidemics.** We provide the probability that the true values of instantaneous reproduction numbers and incidence from scenarios in Fig 5 are within EpiFilter's 95% credible intervals.  
(PDF)

## Acknowledgments

Thanks to Christl A Donnelly for thoughtful and helpful comments.

## Author Contributions

**Conceptualization:** Kris V. Parag.

**Formal analysis:** Kris V. Parag.

**Funding acquisition:** Kris V. Parag.

**Investigation:** Kris V. Parag.

**Methodology:** Kris V. Parag.

**Software:** Kris V. Parag.

**Validation:** Kris V. Parag.

**Visualization:** Kris V. Parag.

**Writing – original draft:** Kris V. Parag.

**Writing – review & editing:** Kris V. Parag.

## References

1. Fraser C. Estimating Individual and Household Reproduction Numbers in an Emerging Epidemic. *PLOS One*. 2007; 8:e758. <https://doi.org/10.1371/journal.pone.0000758>
2. Wallinga J, Teunis P. Different Epidemic Curves for Severe Acute Respiratory Syndrome Reveal Similar Impacts of Control Measures. *Am J Epidemiol*. 2004; 160(6):509–16. <https://doi.org/10.1093/aje/kwh255> PMID: 15353409
3. Nishiura H, Chowell G. The Effective Reproduction Number as a Prelude to Statistical Estimation of Time-Dependent Epidemic Trends. In: *Mathematical and statistical estimation approaches in epidemiology*. Springer; 2009. p. 103–21.
4. Anderson R, May R. *Infectious diseases of humans: dynamics and control*. Oxford University Press; 1991.
5. Cauchemez S, Boelle P, Thomas G, Valleron A. Estimating in Real Time the Efficacy of Measures to Control Emerging Communicable Diseases. *Am J Epidemiol*. 2006; 164(6):591–7. <https://doi.org/10.1093/aje/kwj274> PMID: 16887892
6. Churcher T, Cohen J, Novotny J, Ntshalintshali N, Kunene S, Cauchemez S. Measuring the path toward malaria elimination. *Science*. 2014; 344(6189):1230–32. <https://doi.org/10.1126/science.1251449> PMID: 24926005
7. WHO Ebola Response Team. Ebola Virus Disease in West Africa—The First 9 Months of the Epidemic and Forward Projections. *N Engl J Med*. 2014; 371(16):1481–95. <https://doi.org/10.1056/NEJMoa1411100>
8. Bhatia S, Parag K, Wardle J, Imai N, Van Elsland S, Lassmann B et al. Global predictions of short- to medium-term COVID-19 transmission trends: a retrospective assessment. *medRxiv* 2021;2021.07.19.21260746.
9. Flaxman S, Mishra S, Gandy A, Unwin H, Coupland H, Mellan T, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*. 2020; 584:257–261. <https://doi.org/10.1038/s41586-020-2405-7> PMID: 32512579
10. Hao X, Cheng S, Wu D, Wu T, Lin X, Wang C. Reconstruction of the full transmission dynamics of COVID-19 in Wuhan. *Nature*. 2020; 584:420–4. <https://doi.org/10.1038/s41586-020-2554-8> PMID: 32674112
11. Cori A, Ferguson N, Fraser C, Cauchemez S. A New Framework and Software to Estimate Time-Varying Reproduction Numbers During Epidemics. *Am J Epidemiol*. 2013; 178(9):1505–12. <https://doi.org/10.1093/aje/kwt133> PMID: 24043437
12. Gostic K, McGough L, Baskerville E, Abbott S, Joshi K, Tedijanto C, et al. Practical considerations for measuring the effective reproductive number, Rt. *PLOS Comput Biol*. 2020; 16(12):e1008409. <https://doi.org/10.1371/journal.pcbi.1008409> PMID: 33301457
13. Parag K, Donnelly C. Using information theory to optimise epidemic models for real-time prediction and estimation. *PLOS Comput Biol*. 2020; 16(7):e1007990. <https://doi.org/10.1371/journal.pcbi.1007990> PMID: 32609732
14. Sarrka S. *Bayesian Filtering and Smoothing*. Cambridge, UK: Cambridge University Press; 2013.

15. Parag K, Pybus O, Wu C. Are skyline plot-based demographic estimates overly dependent on smoothing prior assumptions? *Syst Biol*. 2021;syab037. <https://doi.org/10.1093/sysbio/syab037> PMID: 33989428
16. Britton T, House T, Lloyd A, Mollison D, Riley S, Trapman P. Five challenges for stochastic epidemic models involving global transmission. *Epidemics*. 2015; 10:54–7. <https://doi.org/10.1016/j.epidem.2014.05.002> PMID: 25843384
17. Thompson R, Hollingsworth D, Isham V, Arribas-Bel D, Ashby B, Britton T, et al. Key questions for modelling COVID-19 exit strategies. *Proc R Soc B*. 2020; 287(1932):20201405. <https://doi.org/10.1098/rspb.2020.1405> PMID: 32781946
18. Astrom K, Murray R. *Feedback Systems: An Introduction for Scientists and Engineers*. New Jersey: Princeton University Press; 2008.
19. Riley S, Eames K, Isham V, Mollison D, Trapman P. Five challenges for spatial epidemic models. *Epidemics*. 2015; 10(68–71). <https://doi.org/10.1016/j.epidem.2014.07.001> PMID: 25843387
20. Wallinga J, Lipsitch M. How generation intervals shape the relationship between growth rates and reproductive numbers. *Proc R Soc B*. 2007; 274:599–604. <https://doi.org/10.1098/rspb.2006.3754> PMID: 17476782
21. Liu Q, Ajelli M, Aleta A, Merler S, Morenno Y, Vespignani A. Measurability of the epidemic reproduction number in data-driven contact networks. *PNAS*. 2018; 115(50):12680–85. <https://doi.org/10.1073/pnas.1811115115> PMID: 30463945
22. Britton T, Tomba G. Estimation in emerging epidemics: biases and remedies. *J R Soc Interface*. 2019; 16(150):20180670. <https://doi.org/10.1098/rsif.2018.0670> PMID: 30958162
23. Snyder D, Miller M. *Random Point Processes in Time and Space*. 2nd ed. Springer-Verlag; 1991.
24. Kalman R. A new approach to linear filtering and prediction problems. *Trans ASME J Basic Eng*. 1960; 82:35–45. <https://doi.org/10.1115/1.3662552>
25. Anderson B, Rhodes I. *Smoothing Algorithms for Nonlinear Finite-Dimensional Systems*. *Stochastics*. 1983; 9:139–65. <https://doi.org/10.1080/17442508308833251>
26. Parag K, Donnelly C. Adaptive Estimation for Epidemic Renewal and Phylogenetic Skyline Models. *Syst Biol*. 2020; 69(6):1163–79. <https://doi.org/10.1093/sysbio/syaa035> PMID: 32333789
27. Rothenberg T. Identification in Parametric Models. *Econometrica*. 1971; 39(3):577–91. <https://doi.org/10.2307/1913267>
28. Parag K, Donnelly C, Jha R, Thompson R. An exact method for quantifying the reliability of end-of-epidemic declarations in real time. *PLOS Comput Biol*. 2020; 16(11):e1008478. <https://doi.org/10.1371/journal.pcbi.1008478> PMID: 33253158
29. Parag K, Pybus O. Exact Bayesian Inference for Phylogenetic Birth-death Models. *Bioinformatics*. 2018; 34(21):3638–45. <https://doi.org/10.1093/bioinformatics/bty337> PMID: 29718104
30. Parag K, Vinnicombe G. Point Process Analysis of Noise in Early Invertebrate Vision. *PLOS Comput Biol*. 2017; 13(10):e1005687. <https://doi.org/10.1371/journal.pcbi.1005687> PMID: 29077703
31. Zechner C, Koepl H. Uncoupled Analysis of Stochastic Reaction Networks in Fluctuating Environments. *PLOS Comput Biol*. 2014; 10(12):e1003942. <https://doi.org/10.1371/journal.pcbi.1003942> PMID: 25473849
32. Chen Z. Bayesian Filtering: From Kalman Filters to Particle Filters, and Beyond. *Statistics*. 2003; 182(1):1–69.
33. Allen L. A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. *Infect Dis Model*. 2017; 2:128–142. <https://doi.org/10.1016/j.idm.2017.03.001> PMID: 29928733
34. Guo D, Shamaï D, Verdu S. Mutual Information and Minimum Mean-Square Error in Gaussian Channels. *IEEE Trans Info Theo*. 2005; 51(4):1261–82. <https://doi.org/10.1109/TIT.2005.844072>
35. Funk S, Camacho A, Kucharski A, Lowel R, Eggo R, Edmunds W. Assessing the performance of real-time epidemic forecasts: A case study of Ebola in the Western Area region of Sierra Leone, 2014–15. *PLOS Comput Biol*. 2019; 15(2):e1006785. <https://doi.org/10.1371/journal.pcbi.1006785> PMID: 30742608
36. Nouvellet P, Cori A, Garske T, Blake I, Dorigattia I, Hinsley W, et al. A simple approach to measure transmissibility and forecast incidence. *Epidemics*. 2018; 22:29–35. <https://doi.org/10.1016/j.epidem.2017.02.012> PMID: 28351674
37. Van Kerkhove M, Bento A, Mills H, Ferguson N, Donnelly C. A review of epidemiological parameters from Ebola outbreaks to inform early public health decision-making. *Sci Data*. 2015; 2:150019. <https://doi.org/10.1038/sdata.2015.19> PMID: 26029377
38. Wagenmakers E, Grunwald P, Steyvers M. Accumulative prediction error and the selection of time series models. *J Math Psychol*. 2006; 50:149–166. <https://doi.org/10.1016/j.jmp.2006.01.004>

39. Parag K, Cowling B, Donnelly C. Deciphering early-warning signals of the elimination and resurgence potential of SARS-CoV-2 from limited data at multiple scales. medRxiv. 2021;2020.11.23.20236968.
40. Volz E, Didelot X. Modeling the Growth and Decline of Pathogen Effective Population Size Provides Insight into Epidemic Dynamics and Drivers of Antimicrobial Resistance. *Syst Biol*. 2018; 67(4):719–28. <https://doi.org/10.1093/sysbio/syy007> PMID: 29432602
41. Faulkner J, Magee A, Shapiro B, Minin V. Horseshoe-based Bayesian Nonparametric Estimation of Effective Population Size Trajectories. *Biometrics*. 2020; 76(3):677–90. <https://doi.org/10.1111/biom.13273> PMID: 32277713
42. Frost W, Sydenstricker E. Influenza in Maryland: preliminary statistics of certain localities. *Public Health Rep*. 1919; 34:491–504. <https://doi.org/10.2307/4575056>
43. WHO. WHO Coronavirus Disease (COVID-19) Dashboard; 2020. Available from: <https://covid19.who.int/>.
44. Cousins S. New Zealand eliminates COVID-19. *Lancet*. 2020; 395:1474. [https://doi.org/10.1016/S0140-6736\(20\)31097-7](https://doi.org/10.1016/S0140-6736(20)31097-7) PMID: 32386582
45. Ritchie H, Ortiz-Ospina E, Beltekian D, Mathieu E, Hasell J, Macdonald B, et al. Coronavirus Pandemic (COVID-19); 2020. Available from: <https://ourworldindata.org/coronavirus>.
46. Jefferies S, French N, Gilkison C, Graham G, Hope V, Marshall J, et al. COVID-19 in New Zealand and the impact of the national response: a descriptive epidemiological study. *Lancet Public Health*. 2020; 5(11):e612–23. [https://doi.org/10.1016/S2468-2667\(20\)30225-5](https://doi.org/10.1016/S2468-2667(20)30225-5) PMID: 33065023
47. Ferguson N, Laydon D, Nedjati-Gilani G, Imai N, Ainslie K, Baguelin M, et al. Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Imperial College London; 2020.
48. Ganyani T, Kremer C, Dongxuan C, Torneri A, Faes C, Wallinga J, et al. Estimating the generation interval for coronavirus disease (COVID-19) based on symptom onset data, March 2020. *Euro Surveill*. 2020; 25(17). <https://doi.org/10.2807/1560-7917.ES.2020.25.17.2000257>
49. Anderson R, Donnelly C, Hollingsworth D, Keeling M, Vegvari C, Baggaley R et al. Reproduction number (R) and growth rate (r) of the COVID-19 epidemic in the UK: methods of estimation, data sources, causes of heterogeneity, and use as a guide in policy formulation. The Royal Society; 2020.
50. Li Y, Campbell H, Kulkarni D, Harpur A, Nundy M, Wang X, et al. The temporal association of introducing and lifting non-pharmaceutical interventions with the time-varying reproduction number (R) of SARS-CoV-2: a modelling study across 131 countries. *Lancet Infect Dis*. 2020; 21(2):193–202. [https://doi.org/10.1016/S1473-3099\(20\)30785-4](https://doi.org/10.1016/S1473-3099(20)30785-4) PMID: 33729915