

## RESEARCH ARTICLE

## Addressing delayed case reporting in infectious disease forecast modeling

Lauren J. Beesley <sup>\*</sup>, Dave Osthus <sup>1</sup>, Sara Y. Del Valle <sup>1</sup>

Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America

<sup>\*</sup> [lvandervort@lanl.gov](mailto:lvandervort@lanl.gov)

## Abstract

Infectious disease forecasting is of great interest to the public health community and policy-makers, since forecasts can provide insight into disease dynamics in the near future and inform interventions. Due to delays in case reporting, however, forecasting models may often *underestimate* the current and future disease burden.

In this paper, we propose a general framework for addressing reporting delay in disease forecasting efforts with the goal of improving forecasts. We propose strategies for leveraging either historical data on case reporting or external internet-based data to estimate the amount of reporting error. We then describe several approaches for adapting general forecasting pipelines to account for under- or over-reporting of cases. We apply these methods to address reporting delay in data on dengue fever cases in Puerto Rico from 1990 to 2009 and to reports of influenza-like illness (ILI) in the United States between 2010 and 2019. Through a simulation study, we compare method performance and evaluate robustness to assumption violations. Our results show that forecasting accuracy and prediction coverage almost always increase when correction methods are implemented to address reporting delay. Some of these methods required knowledge about the reporting error or high quality external data, which may not always be available. Provided alternatives include excluding recently-reported data and performing sensitivity analysis. This work provides intuition and guidance for handling delay in disease case reporting and may serve as a useful resource to inform practical infectious disease forecasting efforts.

 OPEN ACCESS

**Citation:** Beesley LJ, Osthus D, Del Valle SY (2022) Addressing delayed case reporting in infectious disease forecast modeling. PLoS Comput Biol 18(6): e1010115. <https://doi.org/10.1371/journal.pcbi.1010115>

**Editor:** Alex Perkins, University of Notre Dame, UNITED STATES

**Received:** October 27, 2021

**Accepted:** April 18, 2022

**Published:** June 3, 2022

**Copyright:** © 2022 Beesley et al. 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 are publicly available. ILInet data are available through the DELPHI API at <https://github.com/cmu-delphi/delphi-epidata>. Puerto Rico dengue fever data are provided as part of R package NobBS (<https://cran.r-project.org/web/packages/NobBS/index.html>). Example code is provided at [https://github.com/lbeesleyBIOSTAT/Reporting\\_Delay\\_Example\\_Code](https://github.com/lbeesleyBIOSTAT/Reporting_Delay_Example_Code).

**Funding:** LJB was funded by Los Alamos National Laboratory Laboratory Directed Research and Development (LDRD, <https://www.lanl.gov/>)

## Author summary

The public health community and policymakers are interested in using models to predict *future* disease rates using information about disease rates in the *past*. However, our data about the recent past are less reliable than older data, due to a time lag between someone getting sick and their subsequent diagnosis being officially reported. In this paper, we describe strategies to correct reported disease rates from the recent past to account for disease diagnoses that haven't yet been reported. Using more accurate information about the recent past, we can do a better job predicting what will happen in the future.

[science-innovation/science-programs/ldrd/index.php](https://science-innovation/science-programs/ldrd/index.php)) Award 20210761PRD1. DO was funded by the National Institutes of Health/National Institute of General Medical Sciences (NIH/NIGMS, <http://nigms.nih.gov>) under grant R01GM130668-01. SYD was funded by the National Institutes of Health/National Institute of General Medical Sciences (NIH/NIGMS, <http://nigms.nih.gov>) under grant R01GM130668-01. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

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

This is a *PLOS Computational Biology Methods* paper.

## 1 Introduction

Accurate forecasts of future disease rates are useful for informing public health policy and response to infectious disease outbreaks [1]. Strategies for improving forecasts are of great interest to the public health and epidemic modeling communities, and the most up-to-date reports of case diagnoses can often provide useful insight into disease dynamics in the near future. For many diseases, however, surveillance and case identification take time, and there is often a delay between when people become *infected* with an infectious disease and when that infection is officially *reported* and made publicly available [2].

Often, reporting at the national level requires many intermediate reporting steps; diagnostic labs and health care systems report suspected and/or confirmed cases to local public health systems, which then report to states, which then report to the national level [2]. Delays in each one of these reporting steps can result in substantial lag in the information available for forecasting, and these delays may vary on a variety of factors, including state infrastructure and threshold for reporting suspected vs. confirmed cases. National, state, and local reporting systems produce regular reports providing the most up-to-date statistics on disease diagnoses (e.g., weekly statistics provided by the National Notifiable Diseases Surveillance System), and these provisional reports may be revised/corrected through a process called backfill as new data become available. Real-time reports of current disease burden often substantially underestimate the number of cases eventually reported (hereafter referred to as the validation case counts) for a given time period [3].

Reporting delay can also be a result of the time indexing of the forecasted endpoint. During the COVID-19 pandemic, a common problem is to predict the number of COVID-19 positive cases that will eventually be reported, relative to the date of symptom onset (e.g., [4]). Restated, we may want to predict how many newly symptomatic cases there are today based on tests that have *not yet been administered*, which naturally results in a delay between symptom onset and official case reporting. Reporting delay also occurs when predicting hospitalization rates among newly-diagnosed patients, since hospitalization due to disease takes time to occur [5]. In this way, the problem of delayed event or case reporting can arise in a variety of ways, either through official reporting processes or implicitly through our definitions of what it means to be a “case”.

When the goal is to forecast future disease burden, it is not always clear how one should use these error-prone real-time case reports. When the reporting delay is minimal, we may expect reporting to have little impact on model forecasts and use these real-time reports without correction. When reporting delay is substantial and when real-time estimates and forecasts of disease burden are relied upon to inform policy interventions, additional thought is needed to address the reporting delay.

One common strategy for handling reporting delay is to rescale real-time case reports using the expected amount of under-reporting (e.g., [6]). These scaling factors (hereafter referred to as reporting factors) can be estimated from historical data on real-time case reporting, if available, and sophisticated strategies have been developed to model reporting factors and the disease process jointly (e.g., [3, 7–9]). Other work has used external data to predict the validation cases (e.g., social media, Google Trends) and account for reporting delay using a weighted version of the real-time data and the validation data predictions [10]. Backcast imputation approaches handle the reporting error by generating past weeks’ validation data using the

most recent data reports for those weeks and accounting for uncertainty in reporting patterns [11]. These methods have been studied across multiple disciplines beyond infectious disease modeling such as in actuarial claims prediction [9] and in correction of cancer registry case reporting [12]. During the COVID-19 pandemic, a substantial literature on “nowcasting” has arisen to tackle reporting delay due to delays in testing after symptom onset and to leverage data on individual disease cases when available, where estimation is further complicated by incomplete data on date of symptom onset among confirmed cases [4, 13, 14]. Some of the studies describing how to handle reporting delay compare two or three methods head to head in narrow modeling settings, but there are *no recommendations* for how reporting delay should be handled for infectious disease forecasting in *general practice*.

In this paper, we propose a general framework for conceptualizing and addressing reporting delay in terms of a two-stage estimation problem: (1) estimation of the reporting factors as a function of time since the initial report and (2) disease modeling and forecasting, accounting for under- or over-reporting. This framework highlights and extends existing methods in the literature for handling delayed case reporting, and we demonstrate how these approaches can be implemented for infectious disease forecast modeling in general. We then propose several novel strategies for leveraging either historical data on case reporting or external social media/Google trends data to estimate the amount of reporting error. We apply these methods to address reporting delay in data on dengue fever cases in Puerto Rico from 1990 to 2009 and to reports of influenza-like illness (ILI) in the United States between 2010 and 2019. Through a simulation study, we compare how these methods perform when all assumptions are satisfied and evaluate their robustness when assumptions are violated. This work provides intuition and guidance for the comparative performance of various methods for handling delay in disease case reporting and may serve as a useful resource to inform practical infectious disease forecasting efforts.

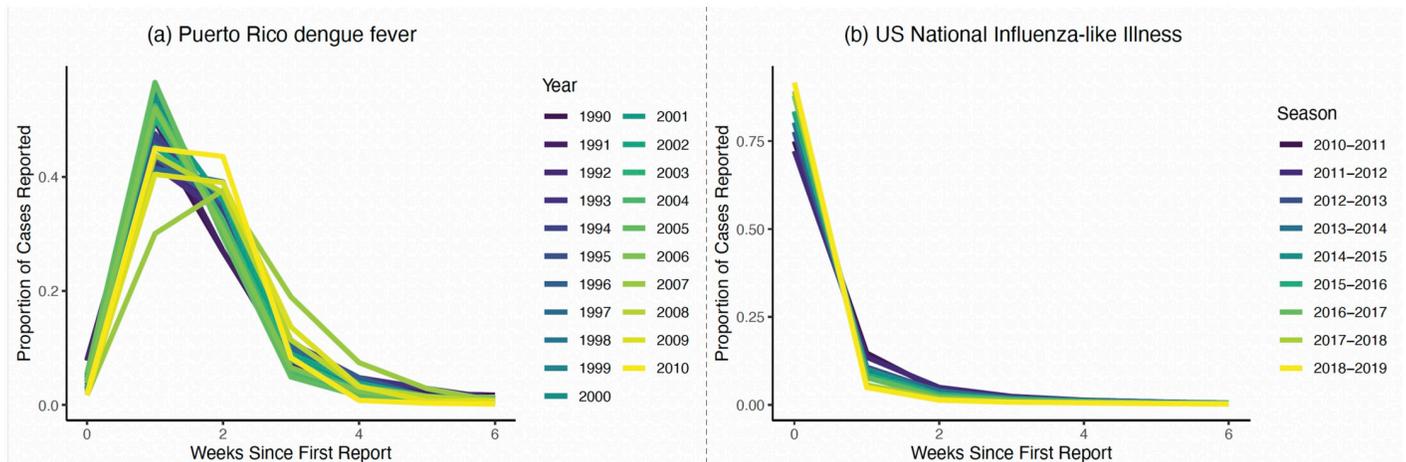
In **Section 2**, we introduce the data examples, and we describe the various conceptual approaches for handling reporting delay in **Section 3**. We implement various approaches for several data examples in **Section 4**, and we explore performance and robustness of these methods in greater detail in **Section 5**. We present a discussion in **Section 6**.

## 2 Disease surveillance data with reporting delay

We consider two motivating datasets in which delayed reporting of disease cases results in discrepancies between the number of cases reported in real-time and the number of cases eventually reported. In our first data example, we model cases of dengue fever reported between 1990 and 2009 in Puerto Rico. In our second data example, we explore influenza-like illnesses in the United States between 2010 and 2019.

### 2.1 Dengue fever surveillance data in Puerto Rico, 1990–2009

We obtain publicly-available data on dengue fever cases in Puerto Rico between 1990 and 2009 that are provided alongside McGough et al. (2020) [3] as part of the R package *NobBS*. These data provide the official diagnosis week and the week of case reporting for over 53,000 cases of laboratory-confirmed dengue fever in Puerto Rico as collected by the Puerto Rico Department of Health and the Centers for Disease Control and Prevention. Using these person-level data, we calculate the number of cases *diagnosed* in each calendar week, which we call the validation data. For each calendar week, we also calculate the number of these validation cases that were actually *reported* as part of the initial case counts from the first week and the number that were reported in each of the following 6 weeks. These weekly reporting data were constructed to mimic weekly “official” case reports. We will refer to these constructed “reports” throughout,



**Fig 1. Proportion of eventually-reported cases that were reported in each week**<sup>1</sup>. <sup>1</sup> Estimates of  $\pi_{ts}(d) - \pi_{ts}(d-1)$ , obtained using Eq 6 and stratified by season/year.

<https://doi.org/10.1371/journal.pcbi.1010115.g001>

but it should be understood that these are hypothetical weekly reports rather than officially-released governmental statistics. A subset of the resulting case counts are shown in Fig A in S1 Text. In Fig 1a, we show that only a very small proportion ( $\sim 5\%$ ) of eventually-reported cases were included in the initial case “report” (i.e., officially reported within 1 week of diagnosis). The vast majority of cases were reported in the following six weeks.

## 2.2 Seasonal influenza-like illness in the United States, 2010–2019

We also consider publicly-available data on influenza-like illness (ILI) in the United States between 2010 and 2019 compiled as part of the U.S. Outpatient Influenza-like Illness Surveillance Network (ILINet). Both validation and officially-reported real-time data (i.e., data that were available up to  $d$  weeks after the initial case report) are available at <https://github.com/cmu-delphi/delphi-epidata> using the Delphi API [15]. Validation data were defined as the case counts for each calendar week that were reported as of June 13th, 2021. In modeling of ILI cases, it is common to break up the data recorded by calendar year into flu seasons, where weeks are re-defined relative to the start of the season. For our analysis, we identified the start of each 35-week flu season as the 40th week in the calendar year as defined by the Morbidity and Mortality Weekly Report (i.e., MMWR week 40) produced by the US National Notifiable Diseases Surveillance System [16]. For example, the 2018–2019 *season* spans the end of 2018 and the beginning of 2019. In addition to the total number of ILI cases nationally, we also downloaded data stratified by state from 2010 through the 2018–2019 season. Real-time reporting data were available for the 2017–2018 and 2018–2019 seasons and also for the 2016–2017 season for some states. Florida was excluded from analysis due to the lack of real-time reporting data provided by the Delphi API. ILI cases occurring after the 2018–2019 season were not included in this analysis to avoid inclusion of COVID-19 diagnoses.

A mild amount of case reporting delay was evident for these data, where roughly 70–90% of eventually-reported cases were reported initially (Fig 1b). There was a clear seasonal trend in reporting at the *national level*, where more recent seasons had a higher proportion of validation cases reported initially. Unlike the dengue fever data, we did not obtain disease diagnosis and reporting dates for individual patients; rather, we considered the total number of reported cases for each week. Due to excluded providers or other reasons, it was possible for real-time case counts to be larger than the validation counts. Therefore, the discrepancy between the

real-time reported and validation case numbers was due to a combination of under-reporting (i.e., not-yet-reported cases) and over-reporting (cases currently reported that were eventually excluded). This phenomenon was particularly evident in the 2017–2018 season data for Vermont, where the number of reported cases often decreased as the number of weeks since the initial report increased (Figs B-C in S1 Text).

### 3 Methods

In this section, we develop the notation and describe various methods for handling reporting delay in general. Table A in S1 Text provides a summary of the notation. Suppose we consider time series data of reported cases for an infectious disease of interest, where cases are reported across multiple calendar time-points  $t$  for each of multiple disease seasons  $s$ . For example, we may consider data on weekly case reports of seasonal influenza-like illness diagnoses, reported across several seasons. Throughout this paper, we will use “week” to refer to the calendar time-point within each season and “year” to refer to the disease seasons, but these methods can be applied to more general time-scales. For a given calendar week  $t$  in a given year/season  $s$ , the first reported case counts are often inaccurate, and the final “validation” counts result from revisions at later dates. Let  $d$  denote the number of weeks since the first reported case counts for a given calendar week. The number of weeks since the week’s first report will also be called the “lag” in reporting. Define  $N_{ts}(d)$  to be the reported case counts for week  $t$  in season  $s$  reported with lag  $d$ , where  $N_{ts}(0)$  corresponds to the very first report and  $N_{ts}(\infty)$  corresponds to the final/validation case counts.

In this paper, we provide and compare strategies for forecasting future *validation* case counts using error-prone real-time reported case data,  $N_{ts}(d)$ . For example, we may want to obtain a forecast of validation cases at calendar time  $t = 5$  for season  $s$  (denoted  $N_{5s}(\infty)$ ) using the most recently-reported data for the previous 4 calendar weeks ( $N_{1s}(3)$ ,  $N_{2s}(2)$ ,  $N_{3s}(1)$ , and  $N_{4s}(0)$ ) along with case data from past seasons. The impact of and correction for reporting delay will naturally depend on how accurate the real-time case data is. Let  $\pi_{ts}(d)$  denote *expected* proportion of validation cases for calendar week  $t$  in season  $s$  that are reported by lag  $d$  weeks as follows

$$\pi_{ts}(d) = E\left(\frac{\mathbf{N}_{ts}(d)}{\mathbf{N}_{ts}(\infty)}\right), \quad (1)$$

where  $\mathbf{N}_{ts}(d)$  and  $\mathbf{N}_{ts}(\infty)$  may be viewed as random variables with data realizations  $N_{ts}(d)$  and  $N_{ts}(\infty)$ . A value of  $\pi_{ts}(d)$  less than 1 indicates under-reporting and a value greater than 1 indicates over-reporting, relative to the validation counts. As  $d \rightarrow \infty$ , we expect that  $\pi_{ts}(d) \rightarrow 1$  and  $N_{ts}(d) \rightarrow N_{ts}(\infty)$ . The *reporting factor* for week  $t$  and season  $s$  at lag  $d$  is defined as  $\frac{1}{\pi_{ts}(d)}$ .

We emphasize that this reporting factor is a function of the lag week  $d$ , and we will assume there is some threshold  $\tau$  such that  $\pi_{ts}(d) = 1$  for all  $d > \tau$ . We will not know the reporting factors in practice, and additional work will be needed to estimate the degree of reporting error and to account for reporting delay in the analysis.

In the remainder of this section, we describe a general two-stage strategy for addressing reporting delay, including (stage 1) estimating reporting factors and (stage 2) forecast modeling, accounting for under- or over-reporting. First, we will describe several conceptual strategies for how to implement forecast modeling while addressing reporting delay given estimates of the reporting factors. Then, we will provide corresponding strategies for estimating reporting factors. These methods are summarized in Fig 2. We provide an extension of these methods for forecasting proportions rather than case counts in S1 Text Section 10.

(a) Method details

Conceptual methods for handling reporting delay					
	Ignore the Problem	Exclude recent data	Rescaling	Modeling with offset	Imputation
Estimate reporting factors	No	No	Yes	Yes	Yes
Forecast model limitations	No	No	No	Yes, log link	No
Summary of method	Analyze observed cases without correction	Exclude most recently observed data from analysis	Rescale observed cases using estimated reporting factors	Add offset (function of reporting factors) to mean model	Draw imputations of validation cases from distribution given observed cases
Strategies for estimating reporting factors					
	Lag-based	Model-based	Local estimation	Proxy shrinkage	Sensitivity analysis
Uses validation data	Yes	Yes	No	Yes	N/A
Other data requirements	Historical real-time data	Historical real-time data	Recent real-time data	Historical external proxy predictors	None
Reporting assumptions	Constant in $s, t$	May vary with $s, t$	May vary with $s, t$	None	None
Summary of method	Reporting factors estimated as a function of lag only, using <i>historical real-time and validation data</i>	Reporting factors estimated using regression model for validation data given real-time data, using historical data	Reporting factors estimated as a function of lag only, using <i>recent real-time data</i>	Reporting factors estimated using weighted average of current real-time data and prediction from external data model.	When no data are available to estimate reporting factors, repeat analysis for several plausible reporting factors.

(b) Flowchart of method recommendations

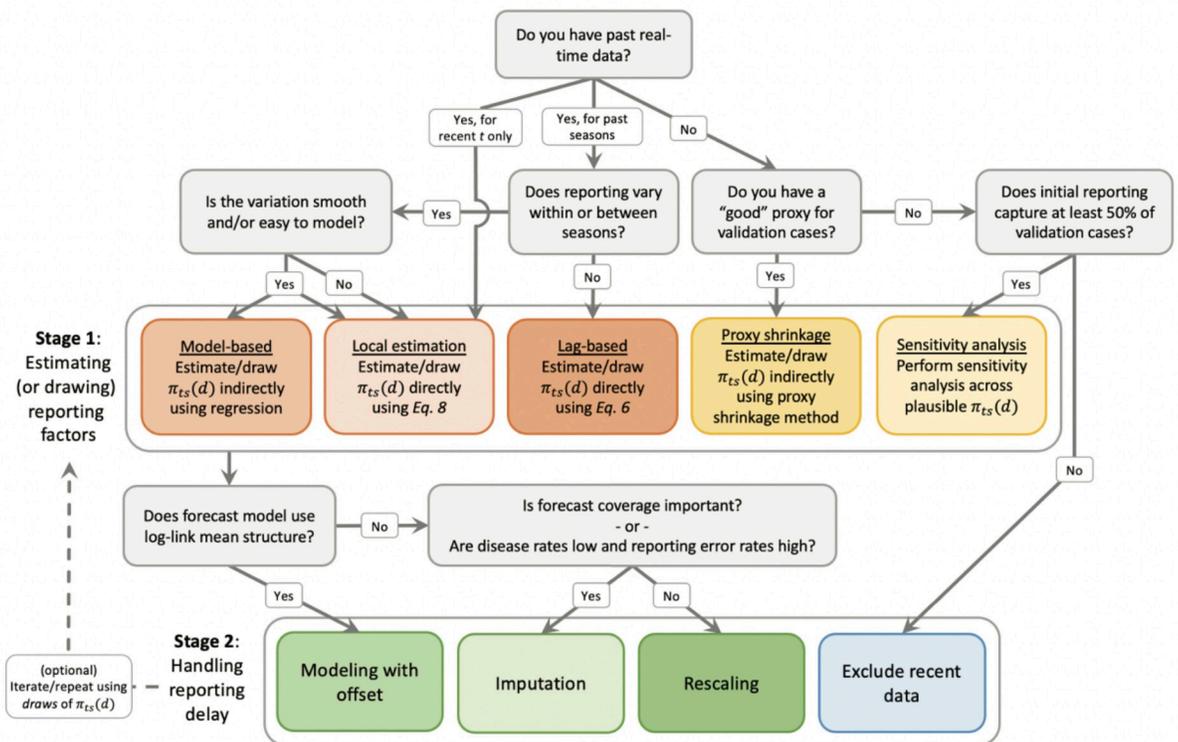


Fig 2. Summary of methods for handling reporting delay.

<https://doi.org/10.1371/journal.pcbi.1010115.g002>

### 3.1 Conceptual strategies for handling reporting delay

In this section, we provide several conceptual strategies for handling reporting delay. In implementing each of these methods, we assume that the forecasting model structure is fixed beforehand. Several of these methods require estimates of  $\pi_{ts}(d)$ . We discuss how to obtain these estimates in the next section.

**Rescaling of real-time data.** One general strategy for correcting reporting delay is to use the most recently-reported case counts,  $N_{ts}(d)$ , and assumptions about the relationship between real-time and validation data (i.e., assumptions about  $\pi_{ts}(d)$ ) to predict  $N_{ts}(\infty)$  as follows:

$$\hat{N}_{ts}(\infty) = \frac{N_{ts}(d)}{\pi_{ts}(d)}, \quad (2)$$

where  $\pi_{ts}(d)$  is replaced with an estimate [6]. Rescaling strategies are often used to correct for general case under- or over-reporting in disease modeling literature, but a key difference from simple rescaling approaches is that the correction factors are functions of the lag  $d$  and possibly  $t$ ,  $s$ , and additional covariates. Given a prediction  $\hat{N}_{ts}(\infty)$  from Eq 2, we can then fit our forecasting model using  $\hat{N}_{ts}(\infty)$  in place of the real-time data.

**Modeling real-time data with mean model offset.** The rescaling method involves pre-processing the real-time data to obtain predictions of the validation data for week  $t$  in season  $s$ . Then, the predictions are input into the forecast modeling pipeline. An alternative strategy involves specifying a forecasting model for the real-time data directly, accounting for the reporting delay in the model structure [3, 9]. Usual implementations involve specifying a model for the *incremental* cases reported for a given calendar week across lag times  $d$ , conditional on the cases reported up to lag time  $d - 1$  [9]. These incremental case models are linked together into a chain that ultimately can be used to predict future values for  $N_{ts}(\infty)$ . Letting  $n_{ts}(d) = N_{ts}(d) - N_{ts}(d - 1)$  be the *data realization* of the random incremental counts  $\mathbf{n}_{ts}(d)$  reported on lag week  $d$ , we can model the random incremental counts using a negative binomial or Poisson regression model with the following mean structure:

$$\log(E(\mathbf{n}_{ts}(d))) = \log(E(\mathbf{N}_{ts}(\infty))) + \log(y_{ts}(d)), \quad (3)$$

where  $y_{ts}(d) = \pi_{ts}(d) - \pi_{ts}(d - 1)$  corresponds to the expected proportion of eventually-reported cases that are reported on lag week  $d$ . This model has the same mean structure as a corresponding model for the validation case counts,  $\mathbf{N}_{ts}(\infty)$ , but with an offset term. A key limitation of this formulation is that  $y_{ts}(d)$  must be greater than zero; in other words, we can only apply this method when we have positive increments. In general, however, we may have over-reporting of cases in addition to under-reporting. Using logic in S1 Text Section 3, we instead propose modeling the cumulative counts directly using a negative binomial or Poisson regression model with mean structure as follows:

$$\log(E(\mathbf{N}_{ts}(d))) = \log(E(\mathbf{N}_{ts}(\infty))) + \log(\pi_{ts}(d)). \quad (4)$$

This mean structure formulation requires only that  $\pi_{ts}(d) > 0$  for all  $d > 0$ . Individual increments, however, may increase or decrease. In implementing mean model offset approaches like Eq 3, researchers often estimate  $\pi_{ts}(d)$  and the disease model parameters *jointly* and specify a prior distribution for  $y_{ts}(d)$  or  $\pi_{ts}(d)$  (e.g., [3]). To make the estimation strategy more compatible with existing out-of-the-box software, however, we propose pre-estimating  $\pi_{ts}(d)$  and then fitting the disease model treating  $\log(\hat{\pi}_{ts}(d))$  as a fixed offset in the mean model. Alternative distributional assumptions for  $\mathbf{N}_{ts}(d)$  may also be considered (e.g., lognormal as in [17]), but this offset-based method is limited to settings where mean counts are modeled using some

type of log link function. Compartmental susceptible-infectious-recovered models, for example, are not compatible with this method.

**Imputation of validation case counts.** The rescaling method involves replacing the real-time data with a predicted value of  $N_{ts}(\infty)$  prior to disease modeling. However, this approach does not capture the *uncertainty* in this prediction. Recasting this approach in a missing data framework, the rescaling method predicts  $N_{ts}(\infty)$  using the *expected value* of  $N_{ts}(\infty)$  given the observed data. It is well-known in the missing data literature, however, that so-called conditional mean imputation may result in under-coverage of model parameters [18]. Similarly, we could see under-coverage of resulting disease model forecasts.

To address this challenge, we can perform disease modeling based on *draws* or *imputations* of the validation data (e.g., [7, 11]). In order to keep the handling of reporting delay separate from the disease modeling, we propose obtaining  $M$  multiple imputations/draws of the validation data for the current season up to the current calendar time. Fixing these imputations, we can then fit the forecasting model to each of the  $M$  imputed datasets and obtain  $M$  forecasts (e.g., 1-week forecasts) and corresponding standard errors. We then combine results *across* imputed datasets using Rubin's multiple imputation combining rules [18]. Additional details about this algorithm can be found in [S1 Text Section 4](#).

Many different approaches can be used to define a distribution for imputing validation case counts given the available data. Using results from the actuarial literature as in [S1 Text Section 3](#), we propose imputing validation case counts from the following truncated normal distribution:

$$N_{ts}(\infty)|N_{ts}(d) \sim \text{TruncNormal}\left(\frac{N_{ts}(d)}{\pi_{ts}(d)}, \frac{|1 - \pi_{ts}(d)|}{\pi_{ts}(d)^2} N_{ts}(d); l, u\right), \quad (5)$$

where  $\pi_{ts}(d)$  is replaced with an estimate and where truncation limits  $l$  and  $u$  restrict imputed values to be greater than  $N_{ts}(d)$  if  $\pi_{ts}(d) < 1$  and less than  $N_{ts}(d)$  if  $\pi_{ts}(d) > 1$ . A key property of this distribution is that the variance of the imputed validation data decreases as  $\pi_{ts}(d)$  increases, meaning that imputed  $N_{ts}(\infty)$  will be closer to  $N_{ts}(d)$  when the expected amount of reporting error is small. The truncated normal distribution in [Eq 5](#) will have expectation different than  $\frac{N_{ts}(d)}{\pi_{ts}(d)}$  due to non-symmetric bounds. In [Fig E](#) in [S1 Text](#), we demonstrate that the percent error in this expectation relative to  $\frac{N_{ts}(d)}{\pi_{ts}(d)}$  is small, particularly in settings with a large amount of reporting error.

A primary limitation of the multiple imputation approach is that it requires the forecasting model to be fit multiple times. For many slow estimation methods, this may not be feasible. When the forecasting model involves Markov Chain Monte Carlo (MCMC) estimation (e.g., some Bayesian models), however, a simple alternative is to handle the missing data within the MCMC estimation algorithm, where a *single* imputed value for each missing  $N_{ts}(\infty)$  is generated using [Eq 5](#) within each MCMC iteration as in Hohle et al. (2014) [7]. Parameters are then drawn within each iteration, conditioning on the imputed validation data. The resulting posterior forecast distributions can then be used directly.

**Exclusion of most recent data reports.** In some cases, the most recently reported real-time data may be too error-prone to contribute meaningfully to forecasting given the information available to the forecaster. In this setting, we propose excluding the most recently reported data from modeling and forecasting.

### 3.2 Estimating reporting factors

For many of the methods described in [Section 3.1](#), the challenge of handling reporting delay boils down to specification of a relationship between the real-time and validation data, i.e.,

$\pi_{ts}(d)$ . Below, we describe several different strategies for estimating  $\pi_{ts}(d)$  using either real-time case reports for past seasons or by leveraging external data to directly nowcast  $N_{ts}(\infty)$ .

**Reporting factors as a function of lag only.** A common assumption in the literature is that the reporting factors are constant in  $t$  and  $s$ . Using data from past seasons/weeks in which *both* real-time and validation data are available, we estimate the inverse reporting factors as

$$\hat{\pi}_{ts}(d) = \hat{\pi}(d) = \frac{\sum_{i,j} N_{ij}(d)}{\sum_{i,j} N_{ij}(\infty)}, \quad (6)$$

with  $\hat{\pi}(d) = 1$  for all  $d$  greater than some fixed threshold,  $\tau$ . An equivalent estimator has been explored in detail elsewhere (e.g., [9, 19]). More sophisticated estimation strategies using parametric assumptions for  $\pi(d)$  can also be used, as discussed in England and Verrall (2002) [9]. An alternative strategy for estimating  $\pi(d)$  would be to average  $\frac{N_{ts}(d)}{N_{ts}(\infty)}$  across  $s$  and  $t$ .

Although not shown, simulations demonstrated that this proportion averaging approach tended to produce unstable  $\hat{\pi}(d)$  with larger variability.

**Reporting factors based on regression modeling of real-time data.** The estimator in Eq 6 assumes that the reporting factor depends on  $d$  only, but we can imagine settings where the reporting may vary across seasons ( $s$ ) and/or within seasons ( $t$ ). Additionally, we may have covariates  $X$  predictive of the difference between real-time and validation case reports (e.g., number of patients tested for the disease or disease-related Google search volumes). In this setting, we propose estimating  $\pi_{ts}(d)$  indirectly by modeling the validation case counts as a function of available data, the lag  $d$ , the season  $s$ , the week  $t$ , and/or covariates  $X$ . For example, we may fit a Poisson regression model with a log link, adjusting for lag, season, week, and  $X$  and using the log of the available counts as an offset in the mean model. See [S1 Text Section 6](#) for an example. Through this model, we can incorporate seasonality and inter-/intra-season trends into our estimation for  $\pi_{ts}(d)$ . Let function  $f$  represent the estimated mapping between the real-time data and the validation data. We can estimate the inverse reporting factors as

$$\hat{\pi}_{ts}(d) = \frac{N_{ts}(d)}{f(N_{ts}(d); X)}, \quad (7)$$

where  $\hat{\pi}_{ts}(d) = 1$  for  $d > \tau$  and where  $\hat{N}_{ts}(\infty) = f(N_{ts}(d); X)$ .

The above strategy for estimating  $\pi_{ts}(d)$  does not require reporting information for individual people; rather, it uses aggregate case counts by week  $t$  and lag  $d$ . In many modern nowcasting applications, however, line-list data with dates of disease diagnosis and reporting for individual people are available (e.g., [4]). In this setting, we can view the time between disease diagnosis and reporting as a time-to-event outcome, perhaps incorporating censoring or truncation. Classical modeling strategies for event time outcomes can be applied to estimate a delay distribution as a function of person-level predictors, which can then be aggregated to obtain estimates of  $\pi_{ts}(d)$  overall. Approaches for handling settings where baseline disease diagnosis or symptom onset dates are unknown for patients in the historical dataset have been explored elsewhere [13, 14, 20].

**Reporting factors using “local” real-time data.** The previously-described methods leverage past season data on the relationship between real-time and validation case counts to estimate reporting factors. However, these methods may perform poorly if reporting in the current season differs substantially from reporting in previous seasons in an unpredictable way (i.e., not easily predicted based on past-seasons’ reporting trends). In this setting, we would like to use real-time data from the most recent weeks to estimate reporting factors that better capture current reporting practices. The challenge, however, is that validation data are not yet available for recent weeks.

We propose using the *changes* in available case counts across  $d$  to obtain a conservative estimate of  $\pi_{ts}(d)$  as follows:

$$\hat{\pi}_{ts}(d) = \frac{\sum_{i=t-K}^{t-d-1} N_{ts}(d)}{\sum_{i=t-K}^{t-d-1} N_{ts}(t-i)}, \quad (8)$$

where the numerator contains past weeks' available data and the denominator contains the most recently-reported data for the previous  $K$  weeks. Define  $\tau$  to be the lag value after which we assume true  $\pi_{ts}(d) = 1$  for all  $d > \tau$ . For this estimator, we define  $\hat{\pi}_{ts}(d) = 1$  for  $d > \min(K, \tau)$ .  $K$  can take any positive integer value, but we recommend setting  $K = \tau$  in practice.

The estimator in Eq 8 is conservative in that it is *biased toward 1* relative to the true  $\pi_{ts}(d)$  when reporting is such that (1)  $N_{ts}(j) \leq N_{ts}(j+1)$  for  $j \geq 0$  (monotone under-reporting) or (2)  $N_{ts}(j) \geq N_{ts}(j+1)$  for  $j \geq 0$  (monotone over-reporting). This bias comes from using the most recent case counts rather than the corresponding validation values in the denominator. Despite this bias, Eq 8 may provide a more accurate estimate of  $\pi_{ts}(d)$  than the previously-described approaches when the current weeks' reporting is very different from reporting in past seasons.

We could further generalize this local estimation strategy to incorporate seasonality, covariate associations, and/or residual autocorrelation to estimate  $\pi_{ts}(d)$ . We propose viewing historical observed values of  $\pi_{ts}(d)$  and local estimated values from Eq 8 as themselves a time series that can be used to forecast  $\pi_{ts}(d)$  for each fixed value of  $d$ . For example, we can construct a time series of observed/estimated values of  $\pi_{ts}(0)$  for *prior weeks*, fit an ARMA (S1 Text Section 5) model to some function of  $\pi_{ts}(0)$  (e.g.,  $\text{logit}(\pi_{ts}(0))$ ) adjusting for any covariates or seasonality we want to account for, and obtain a one-week-ahead forecast to estimate unknown  $\pi_{ts}(0)$  for the current week. In this way, we can simultaneously combine historical information on seasonality and other trends while leveraging information on reporting from the recent past. When applied to our real data examples, this approach did not yield improved performance over Eq 8, so we did not evaluate this approach in additional detail. However, we hypothesize that there may be other infectious disease settings where this generalization could provide additional flexibility and corresponding improvements in forecast performance. Kline et al. (2021) [21] provides a related approach for modeling increments in reported COVID-19 cases as a function of reporting lag and day of week with an autoregressive structure.

**Reporting factors based on proxy shrinkage.** The preceding methods assume some prior information is available about the relationship between the real-time and validation case counts. For many diseases however, historical data on validation case counts are readily available but the real-time reports for past seasons are either unavailable or incomplete. In this setting, we propose leveraging external data sources to estimate  $\pi_{ts}(d)$ .

Suppose we have historical data on the relationship between validation case counts  $N_{ts}(\infty)$  and some external data, collectively denoted  $p_{ts}$ , that are associated with  $N_{ts}(\infty)$ . Examples of external data include Google search volumes, social media mentions (e.g., from Twitter), and Wikipedia searches [22, 23]. Using these historical data, we fit a nowcast model for validation case counts given  $p_{ts}$ . This model could be complicated and include data from many different sources, incorporate penalization, etc. Let  $g$  denote the mapping between  $p_{ts}$  and model-predicted validation case counts. These predictions may be viewed as an error-prone *proxy* for the current season validation cases.

Using this external data model proxy, we propose estimating the validation case counts as a weighted average of the proxy and the observed real-time data as follows

$$\hat{N}_{ts}(\infty) = w_d g(p_{ts}) + (1 - w_d) N_{ts}(d), \quad (9)$$

and we estimate corresponding inverse reporting factors as  $\hat{\pi}_{ts}(d) = \frac{N_{ts}(d)}{N_{ts}(\infty)}$ . For large  $d$ , we expect  $N_{ts}(d)$  and  $N_{ts}(\infty)$  to be very close. Therefore, we will define weights  $w_d$  such that  $w_d \rightarrow 0$  as  $d \rightarrow \infty$  and impose  $w_d = 0$  for all  $d > \tau$ . As the number of lag weeks increases, the shrinkage estimator will put more and more weight on the observed data. In a related method, Osthus et al. (2019) [10] proposed defining  $w_d = \left(\frac{\omega}{d+1}\right)^2$  with  $\omega = 0.75$ . The choice of weight and performance of this method will depend on the degree of reporting delay and how well the proxy model  $g$  predicts the validation case counts.

**Sensitivity analysis.** All of the above methods for estimating reporting factors require either historical real-time data or a good proxy for the validation case counts based on external data. In practice, we may not have either available. In this case, one approach for implementing the methods in **Section 3.1** is to perform multiple parallel analyses assuming different plausible reporting mechanisms in a sensitivity analysis framework. This approach can be used to evaluate the robustness of forecasts to assumptions about the delay mechanism.

**On drawing reporting factors.** Thus far, we have proposed a two-stage analysis procedure where we first estimate reporting factors  $\pi_{ts}(d)$  and then apply the methods in **Fig 2a** to obtain forecasts. However, this approach does not account for uncertainty in the estimated values for  $\pi_{ts}(d)$ . Jones et al. (2020) [24] and others have highlighted the potential for poor performance of rescaling or multiplier strategies as in **Eq 2** when a single estimate of  $\pi_{ts}(d)$  is used without accounting for uncertainty. In Bayesian forecast modeling, a common solution is to build reporting delay into the model structure and forecast assuming a prior distribution for  $\pi_{ts}(d)$  rather than fixing its value (e.g., [3]). All of the methods discussed in **Section 3.1** be similarly modified to incorporate uncertainty in estimated  $\pi_{ts}(d)$ . As shown by our simulation results, however, concerns raised about uncertainty propagation for rescaling methods can be substantially mitigated by implementing imputation strategies as in **Eq 5**.

## 4 Application to dengue fever and seasonal influenza forecasting

We considered data on dengue fever cases in Puerto Rico (1990–2009) and for seasonal influenza-like illness (ILI) cases in the United States (2010–2011 through 2018–2019 seasons). Our goal was to evaluate how the quality of disease forecasts is impacted by the handling of reporting delay. With the exception of the mean model offset method, all of the methods in **Fig 2** for handling reporting delay can be applied in conjunction with any forecasting model. For demonstration, we considered two reasonable forecasting model settings: an auto-regressive moving average (ARMA) model for log-cases (e.g., [25]) and a Bayesian Gaussian process-based model (hereafter, referred to as Inferno from Osthus (2021) [26]). Both models assumed a log link mean structure. Additional details about these forecast models are provided in **S1 Text Section 5**.

### 4.1 Obtaining weekly forecasts, accounting for reporting delay

For each week after the first two seasons, we constructed an artificial dataset consisting of the real-time and historical observations that would have been available to use for forecasting. First assuming that reporting does not vary over  $t$  and  $s$ , we estimated reporting factors  $\hat{\pi}(d)$  from **Eq 6** using the previous 2 years' data. To allow estimated reporting factors to vary by  $t$  and  $s$ , we also estimated  $\hat{\pi}_{ts}(d)$  indirectly as in **Eq 7** by fitting a Poisson regression model with a log link to the prior observed validation data, including  $d$ ,  $s$ , and a 3-degree natural spline of  $t$  as predictors and treating the available data  $\log(N_{ts}(d))$  as an offset. Additional details and diagnostics can be found in **S1 Text Section 6.1**. Finally, we estimate “local”  $\pi_{ts}(d)$  using **Eq 8** and the previous  $K$  weeks' real-time reported data (using  $K = 15$  for ILI and  $K = 6$  for dengue

fever). As an additional sensitivity analysis, we also estimate “local”  $\pi_{ts}(d)$  across a range of  $K$  values between 2 and 100 for each dataset.

Fixing these reporting factors and using either an ARMA(2,2) or Inferno forecasting model structure, we then fit the forecasting model using each of the rescaling, model offset, and imputation methods in **Section 3.1**. For ARMA, the imputation method was implemented using 10 imputed datasets and [Eq 5](#). For Inferno, imputation using [Eq 5](#) was implemented within a Bayesian MCMC estimation algorithm. We also fit the forecasting model to the observed real-time data without correction and after excluding the most recent 1–3 weeks of data. To benchmark the performance of these methods, we also fit the forecasting model using the validation data. Inferno model estimation was based on 2500 MCMC iterations after a burn-in of 2500 iterations. Results are also provided for an ensemble based on an equal weight linear combination of forecasts from the 13 methods (excluding validation data analysis).

For each forecast fit, we obtained a prediction for the validation case counts for the current week (nowcast) and for 1 and 4 weeks ahead (forecasts). We also calculated corresponding 50%, 67%, 95%, and 99% prediction intervals. For the Inferno model, ensemble forecasts and corresponding prediction intervals were obtained using the stacked draws across methods. For the ARMA model, we obtained 2500 draws of predicted log-case counts using the Gaussian prediction distribution for each of the methods. Ensemble forecasts and prediction intervals for the validation case counts were then obtained using these stacked draws. We repeated this process for each calendar week of available data. We then summarized the predictive performance using the metrics described in **Section 4.2**, aggregating across all calendar weeks and seasons for each dataset and forecast model structure separately.

## 4.2 Performance metrics

We evaluated the performance of the various methods based on properties of nowcasts (i.e., predicted validation values for the current week) and 1 and 4 week forecasts (i.e., predicted validation case counts for future weeks). Absolute prediction error for a given nowcast/forecast was defined as the absolute difference between the model estimate and the validation case counts. For assessing the accuracy of the point and interval estimates jointly, we calculated the weighted interval score from Bracher et al. (2021) [27] as follows:

$$WIS(y) = \frac{1}{K + 0.5} \left[ 0.5|y - \hat{f}_{ts}| + \sum_{j=1}^K \frac{\rho_j}{2} IS(l_j, u_j; y) \right] \quad (10)$$

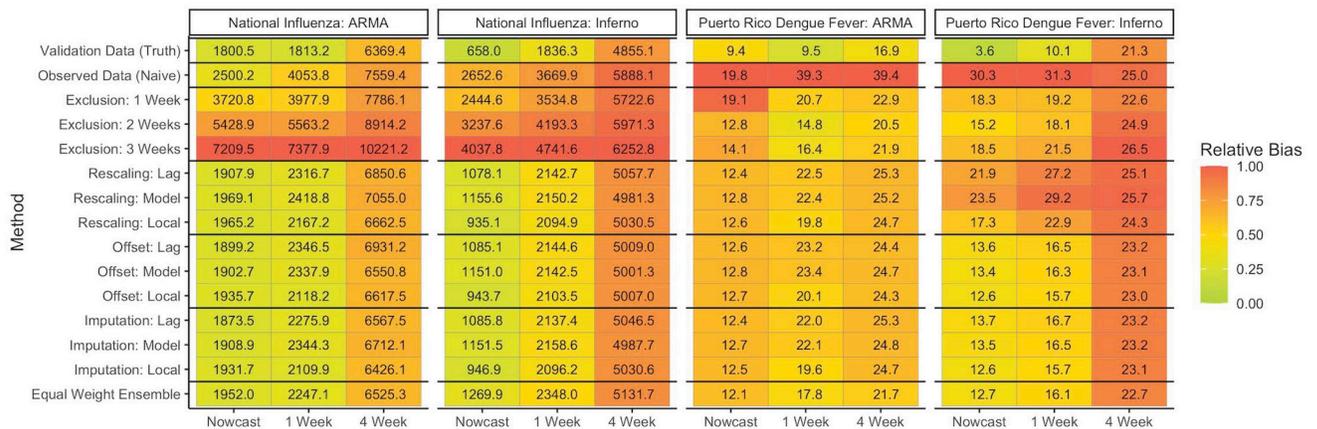
$$IS(l_j, u_j; y) = (u_j - l_j) + \frac{2}{\rho_j} (l_j - y)I(y < l_j) + \frac{2}{\rho_j} (y - u_j)I(y > u_j), \quad (11)$$

where  $K$  is the number of intervals (in our case, 11),  $\rho = (0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1, 0.05, 0.02)$ ,  $u_j$  and  $l_j$  are the upper and lower confidence/credible interval limits corresponding to level  $1 - \rho_j$ , and  $\hat{f}_{ts}$  is the median of the nowcast/forecast distribution. These values of  $\rho$  were chosen to match the quantiles used by the US COVID Forecast Hub (<https://covid19forecasthub.org/>). We evaluated the WIS using the validation values. We also calculated the empirical coverage of the 95% prediction intervals using the forecasts and validation case counts across all calendar weeks.

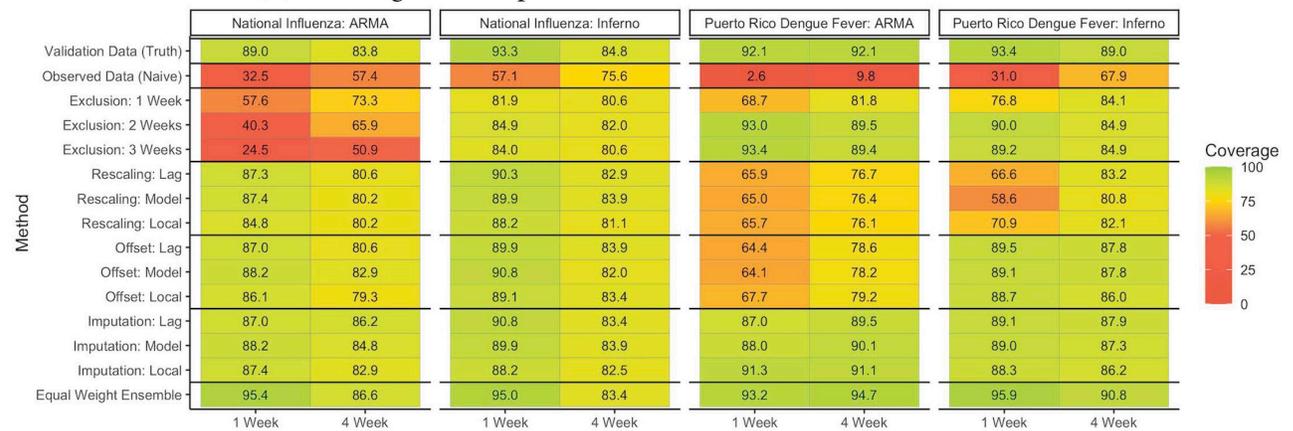
## 4.3 Results

**Prediction error and coverage.** [Fig 3](#) presents the average absolute nowcast/forecast errors for the dengue fever and US national ILI outcomes along with corresponding coverages

(a) Average absolute bias in nowcasts and 1 and 4 week forecasts



(b) Coverage of 95% prediction intervals for 1 and 4 week forecasts



**Fig 3. Performance of nowcasts and forecasts in the Puerto Rico dengue fever and national US influenza-like illness data across all weeks**<sup>1, 1</sup> Results for dengue fever are aggregated across each of 50 weeks in 18 seasons (1992–2009). Results for US national influenza are aggregated across 35 weeks in 7 seasons. The ensemble method corresponds to an equal-weight linear combination of all methods except validation data analysis. “Model” indicates that reporting factors were estimated via regression and allowed to vary by *t* and *s*. “Lag” indicates that reporting factors were estimated via Eq 6. “Local” indicates that reporting factors were estimated via Eq 8. Relative absolute biases are calculated relative to the largest value in each column.

<https://doi.org/10.1371/journal.pcbi.1010115.g003>

of 95% prediction intervals. Weighted interval scores are presented in Fig G in S1 Text. For both datasets and forecast model structures, analysis of the observed data without correction resulted in large undercoverage (2–60%) and higher prediction error relative to analysis of the validation data. Compared to uncorrected analysis of the observed data, the exclusion method (1–3 weeks) resulted in improved coverage and similar or reduced prediction error across all endpoints for the dengue fever data. In contrast, the exclusion method resulted in similar or increased prediction error in the national US ILI data.

The rescaling, offset, and imputation methods uniformly resulted in higher coverage and similar or reduced prediction error relative to uncorrected observed data analysis for the dengue fever and US national ILI data. For the dengue fever data, the imputation methods tended to give better coverage than the rescaling methods (e.g., around 65% vs nearly 90%). In general, the rescaling method can be unstable (i.e., produce high error) in the dengue fever data, particularly for the Inferno forecast model (Fig H in S1 Text). As described in S1 Text Section 7, this is due to a combination of very high under-reporting and a low disease rate, which results in

rescaled validation estimates that are very sensitive to small fluctuations in the observed case counts. Both the mean model offset method and the imputation method resulted in stabler nowcasts/forecasts in this setting.

The equal weight ensemble method attempts to produce more stable forecasts than any individual method by borrowing information across many methods. For both the Inferno dengue fever model and both ARMA and Inferno national influenza models, the ensemble estimation resulted in similar or better coverage than any of the individual methods, particularly for nowcasts. Resulting nowcast and forecast errors were generally between errors from the exclusion and rescaling/offset/imputation methods for the ILI data and similar or smaller than errors for other methods for the dengue fever data. Across both datasets, there is little evidence of improved forecast/nowcast performance for model-based vs. lag-based estimation of reporting factors. However, we observed lower forecast biases when reporting factors were estimated using *local* reporting data. We did not formally assess the performance of proxy-based reporting factor estimation for these data.

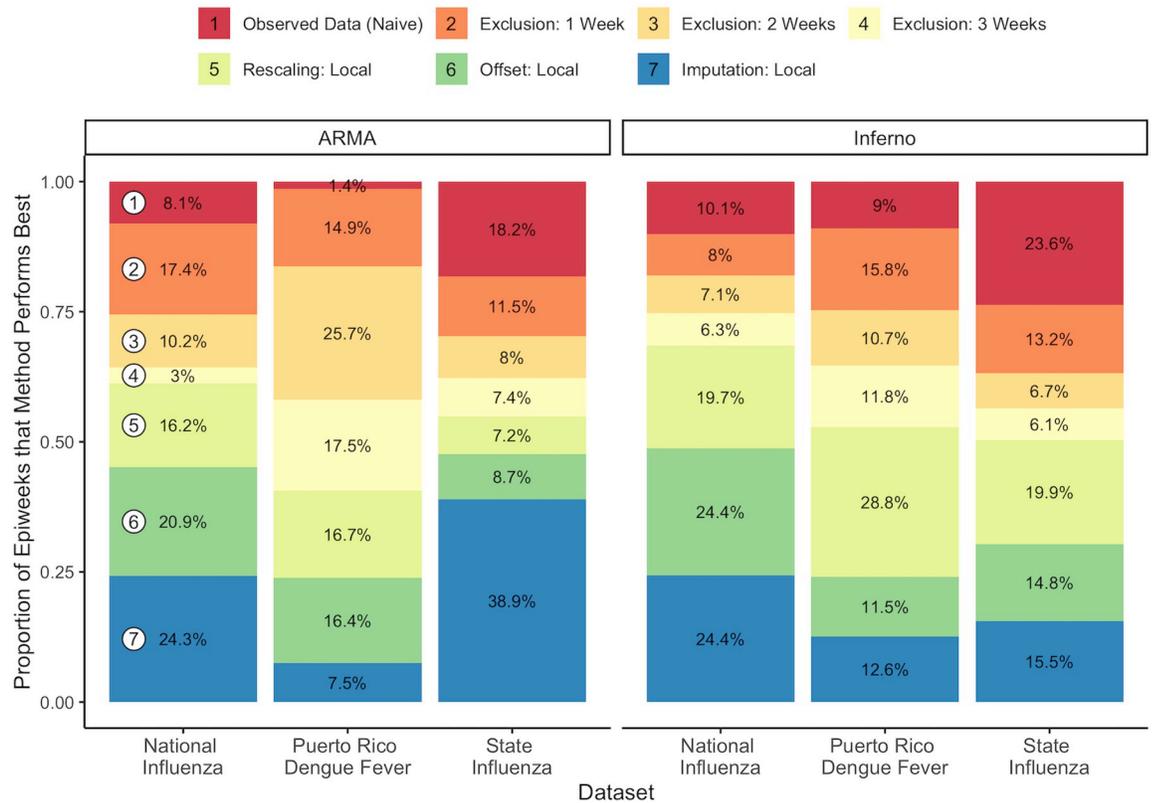
In **Figs I-J** in [S1 Text](#), we present the relative accuracy of several methods across individual weeks during follow-up. We find that the relative performance of uncorrected data analysis tended to improve for the dengue fever data when validation case counts were very small, suggesting that the correction methods may over-correct in this setting. Similarly, uncorrected analysis of the US national ILI data gave improved relative performance shortly after the season peak for seasons with sharp decreases in cases, suggesting that the correction methods may be less able to adapt to these rapid decreases in validation case counts. **Fig M** in [S1 Text](#) provides bias and weighted interval scores across methods, focusing on weeks near the season peaks and minima.

**Comparative rankings.** We then calculated the proportion of weeks across seasons in which each of 7 methods performed best in terms of weighted interval scores for 1 week forecasts. For this analysis, we focused on local estimation of  $\pi_{ts}(d)$  as it generally resulted in similar or better performance than the same methods using lag-based or model-based estimation.

If all methods performed equally, we would expect each method to be the best performer about 14% of the time (95% confidence interval of [12%,16%] for dengue fever and state-level ILI and [10%,18%] for national ILI, assuming independent trials). As shown in [Fig 4](#), however, we saw that corrected analysis tended to outperform uncorrected analysis of the observed data for dengue fever and US national ILI data. For these data sources, these results suggest that we are almost always better off performing some kind of correction strategy to handle the reporting delay. Among the correction methods, the imputation and offset methods tended to have similar or better performance than the other methods in terms of 1-week forecast interval score rankings for the national influenza-like illness datasets. In contrast, the exclusion and rescaling methods had the best performance in the dengue fever data.

Uncorrected analysis results for state-level ILI tell a different story than seen for national US ILI and dengue fever, where analysis of the observed state ILI data without correction had the best weighted interval score performance at least 17%+ of the time. As shown in **Figs K-L** in [S1 Text](#), the relative performance of the delay correction methods varied between states, but the imputation-based correction method tended to have the best performance on average. For the Inferno model, uncorrected analysis outperformed all correction methods on average (lowest forecast weighted interval score 25% of the time).

Unlike in the dengue fever and national ILI data, the reporting delay mechanism sometimes changed *substantially* between seasons for individual states (**Figs B and C**). Although not shown, this resulted in reduced forecast performance for the strategies that relied on past season reporting data to correct delay in the current season. As shown in **Fig N** in [S1 Text](#), the local  $\pi_{ts}(d)$  estimation method produced much better estimates of the inverse reporting factors



**Fig 4. Proportion of weeks in which each of 4 methods performs best in terms of 1-week forecast weighted interval scores.** <sup>1</sup> Results for Dengue fever in Puerto Rico are aggregated across all 50 weeks in 18 calendar years (1992–2009). Results for US national influenza are aggregated across 35 weeks in 7 seasons (2012–2013 through 2018–2019). Results for US state-level influenza are aggregated across 35 weeks in 49 states (excluding Florida) for the 2018–2019 season.

<https://doi.org/10.1371/journal.pcbi.1010115.g004>

for the current season. This serves as a cautionary note against using lag-based and model-based  $\pi_{ts}(d)$  estimation methods when past season information on reporting practices is poorly representative of the current season.

For all datasets, we observed better comparative performance for naive observed data analysis with the Inferno model than with the ARMA model. This was due to (1) a tendency for the observed data Inferno model to perform particularly well relative to the other methods in weeks where the validation case counts were very low (Figs I, J, and O in S1 Text) and (2) poor ARMA forecast performance across the board when analysis was based on real-time observed data (Fig P in S1 Text). These results suggest that ARMA forecast models may tend to perform poorly when applied to real-time data subject to a large amount of reporting error.

**Weeks  $K$  used for local  $\pi_{ts}(d)$  estimation.** In the analyses presented in Fig 3, we apply the local estimation strategy in Eq 8 with  $K$  set equal to  $\tau$ , the number of weeks beyond which we assume  $\pi_{ts}(d) = 1$ . However, any positive integer value for  $K$  could theoretically have been chosen. In Figs Q–S in S1 Text, we present the  $\pi_{ts}(0)$  estimates and corresponding nowcast/forecast performance obtained as a function of  $K$  for the dengue fever, national US ILI, and Vermont ILI datasets. We demonstrate that the estimate in Eq 8 looks more and more like the simple lag-based estimate using Eq 6 as  $K$  becomes large. The best nowcast and forecast performance was generally seen for  $K$  near  $\tau$ , with larger  $K$  resulting in particularly large increases in prediction errors for the Vermont data.

## 5 Performance of methods for simulated real-time data

The analyses in **Section 4** provide intuition on the performance of the various reporting delay correction methods in three real data settings. A key limitation of several of the methods explored in **Section 4** is that they require high-quality data on real-time case reports from past seasons along with corresponding assumptions about how reporting varies across and within seasons. Since reporting delay for the current season is not well-understood at the time of forecasting in practice, it is important to evaluate the impact of these assumptions and violations on method performance. To further develop our intuition for when these methods will and will not perform well, we conducted a simulation study.

### 5.1 Simulation set-up

We simulated validation disease counts to mimic observed disease rates of dengue fever in Puerto Rico between 1990 and 2009 as follows:

1. Let  $\beta_{ts}$  be the three-week moving average of the validation case counts in the actual dengue fever data, and let  $\tau_t$  be the average of  $N_{ts}(\infty)/\beta_{ts}$  across all seasons. Define  $\theta_{ts} = \beta_{ts} * \tau_t$  to be the mean disease rate in the simulated data.
2. We generated 200 simulated versions of the validation data for each  $t$  and  $s$  by drawing from a negative binomial( $r, p$ ) distribution with  $p = \frac{r}{r+\theta_{ts}}$  and  $r = 100$ . The hyperparameter  $r$  was chosen to mimic the variability in actual dengue fever validation data.

**Fig Va** in [S1 Text](#) provides a visualization of the resulting case counts across these 200 simulated datasets. For each simulated dataset, we then generated real-time data for each calendar week under different assumed reporting mechanisms. All reporting mechanisms followed the reporting profiles in **Fig Vb** in [S1 Text](#), which were parameterized by constant  $a$  between 0.05 and 1 representing the proportion of eventually reported events reported at lag 0. We considered the following simulation scenarios:

1. *Constant*: reporting was constant in  $t$  and  $s$  and corresponded to  $a = 0.05$ .
2. *Vary by week*: reporting varied by  $t$  and was constant in  $s$ .  $a$  initially increased from 0.05 to 1 during weeks 1 to 25 in each season and then decreased from 1 to 0.05 thereafter.
3. *Large improvement between seasons*: reporting substantially improved in the last season, with  $a = 0.05$  for 1990–2008 and  $a = 0.50$  for 2009.
4. *Large worsening between seasons*: reporting substantially worsened in the last season, with  $a = 0.50$  for 1990–2008 and  $a = 0.05$  for 2009.
5. *All season combinations*: Each strata of 10 simulation replicates was assigned a different value for  $a$  between 0.05 and 1. Within each strata, reporting (i.e.,  $a$ ) was constant in  $s$  and  $t$ .

For each set of simulated validation data, we also simulated 4 external proxy variables such that  $p_{ts} = 2 \log(N_{ts}(\infty) + 0.1) + e_{ts}$ , where  $e_{ts} \sim N(0, \sigma^2)$  and where  $\sigma^2$  took values in (0.01, 1, 4, 16). These error rates corresponded to correlations between transformed  $p_{ts}$  and  $N_{ts}(\infty)$  of 0.99, 0.80, 0.50, and 0.13.

In order to clarify general properties of the various methods, we also repeated the above simulation procedure with data simulated to mimic observed disease rates of national US influenza-like illness for the 2010–2011 to the 2018–2019 seasons. Detailed results for ILI are presented in [S1 Text Section 9](#), and we focus on the results for dengue fever-like data in the main paper.

## 5.2 Performance under Scenarios 1–4 in the simulated dengue fever 2009 season

In this simulation study, our goal was to compare the performance of each of the reporting delay correction methods when the reporting factor estimates were correctly specified (Scenario 1 and model-based  $\pi_{ts}(d)$  for Scenario 2) and when they were mis-specified (lag-based  $\pi_{ts}(d)$  for Scenario 2 and Scenarios 3–4). For this exploration, we focused our attention on the first 100 simulated datasets corresponding to 2009.

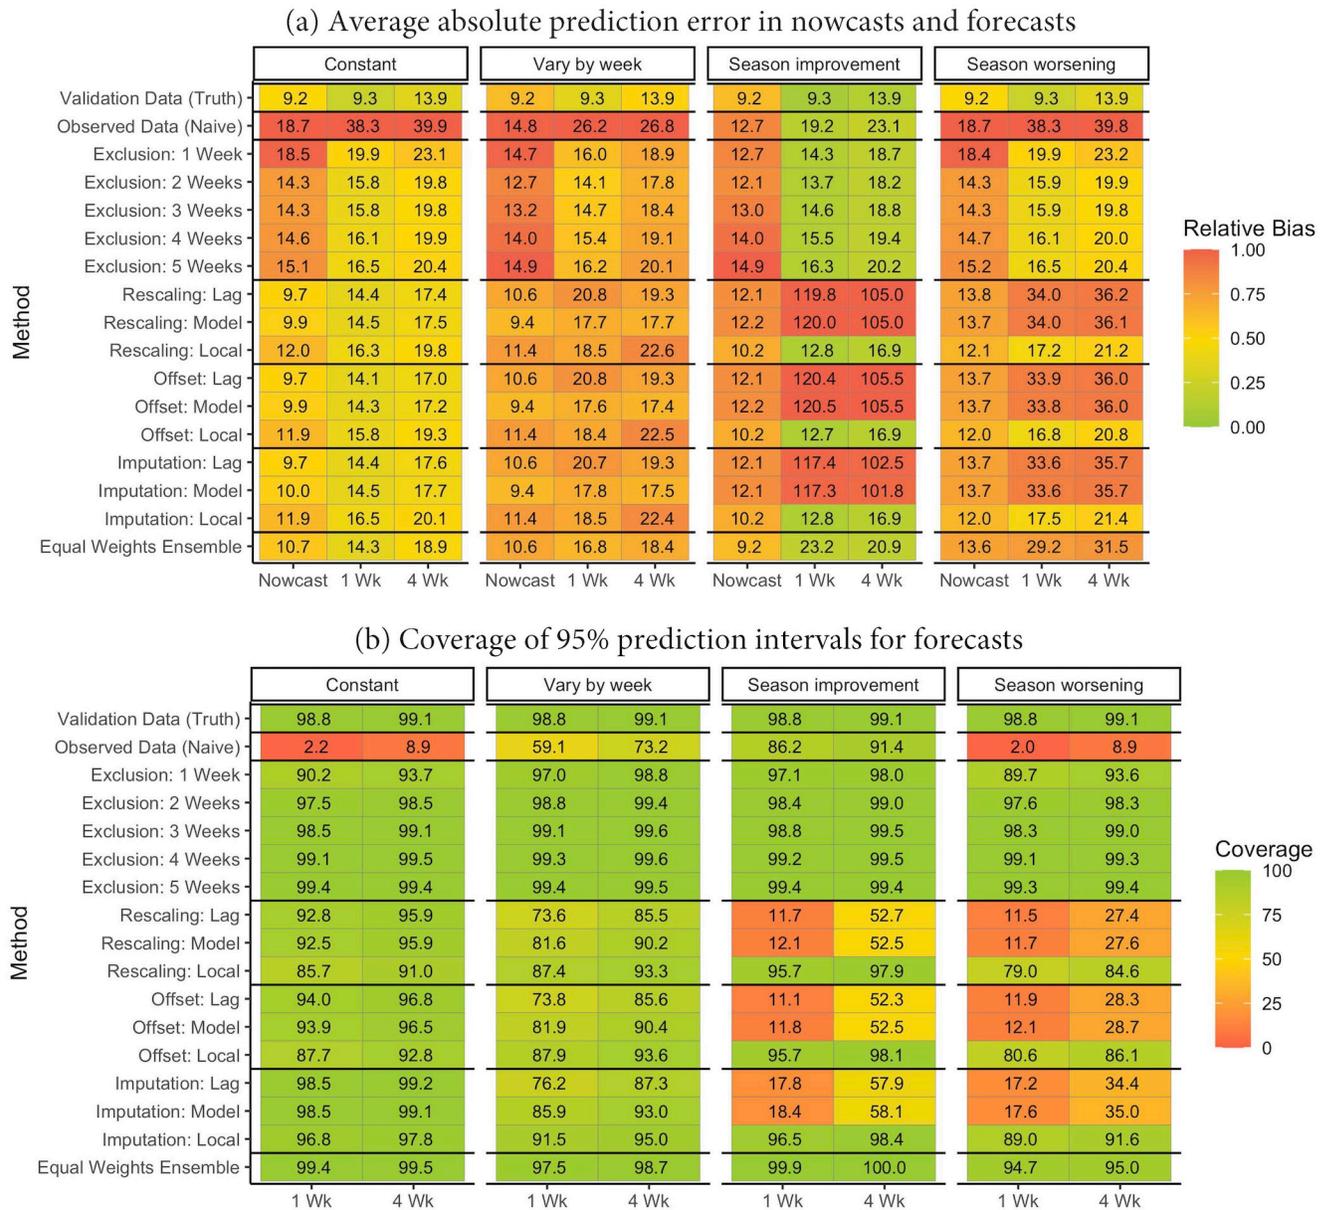
For each calendar week in 2009 and simulated datasets under Scenarios 1–4, we applied the procedure described in **Section 4.1** to construct an artificial real-time dataset that would have been available for forecasting. For each artificial dataset, we applied the methods in **Fig 2** exactly as we did in the actual dengue fever data, but this time we repeated this analysis across 100 simulated versions of the data. We obtained nowcasts, forecasts, and corresponding 95% prediction intervals, and we summarized these nowcasts/forecasts using the performance metrics in **Section 4.2**. We then aggregated these metrics across the 100 simulated datasets and 50 calendar weeks in 2009 for each of the simulation scenarios.

**Fig 5** presents the average absolute nowcast and forecast prediction errors and corresponding coverage of 95% prediction intervals obtained from applying the methods with a ARMA (2,2) forecast model. Weighted interval scores are shown in **Fig W** in **S1 Text**. We summarize these results for each of the 4 simulation scenarios below.

**Reporting factors constant (Scenario 1).** All reporting correction strategies resulted in lower prediction error and higher coverage relative to analysis of the observed data without correction (e.g., coverages < 10% for uncorrected analysis compared to coverages over 90%). We also see improvements over observed data analysis in terms of weighted interval scores for all reporting correction methods. The exclusion methods produced higher nowcast and forecast errors on average relative to the other correction methods. This indicates that the most recently reported cases can contribute useful information for nowcasting/forecasting even in the presence of extreme under-reporting as long as the reporting factors are correctly specified. Imputation methods resulted in higher nowcast and forecast coverage than the rescaling and model offset methods, although this coverage was overly conservative (greater than 95). Model-based estimation of  $\pi_{ts}(d)$  resulted in a small increase in prediction error relative to lag-based estimation. More flexible estimation of  $\pi_{ts}(d)$ , therefore, may come with some small price in terms of prediction error when reporting is truly constant across  $t$  and  $s$ . Local estimation of  $\pi_{ts}(d)$  resulted in small additional increases in prediction error and weighted interval scores.

**Reporting factors vary by week (Scenario 2).** When reporting factors truly varied within each season, correction methods using model-based estimates of  $\pi_{ts}(d)$  outperformed use of lag-based estimates, which incorrectly assumed reporting did not vary over  $s$  or  $t$ . However, the difference in performance was modest, even in this setting with extreme intra-season variation in reporting. The local  $\pi_{ts}(d)$  estimation method performed well in this setting, with higher coverages than the model-based method but also with slightly higher prediction error and weighted interval scores. As before, the exclusion method resulted in higher nowcast error than the other correction methods, but the exclusion method did provide better coverages than the rescaling, offset, and imputation methods.

**Large reporting factor change between seasons (Scenarios 3–4).** When reporting for the current season was very different than the previous seasons, reporting factor estimation using past season data (lag-based and model-based) resulted in poor estimates of the current season under-reporting. When the current season under-reporting was much better than previous seasons, use of *past-season* reporting factors resulted in over-correction and over-estimation of



**Fig 5. Performance of proposed methods for handling reporting delay in 2009 across 100 simulated datasets using ARMA models<sup>1</sup>.** <sup>1</sup> Results are aggregated across all 50 weeks in 100 replicate seasons. Each result, therefore, represents aggregates 5000 nowcasts or forecasts. When reporting factors varied across seasons,  $\pi_{2007} = \pi_{2008} = (0.01, 0.05, 0.55, 0.85, 0.95, 0.98, 1)$  and  $\pi_{2009} = (0.04, 0.54, 0.84, 0.0, 0.94, 0.97, 0.99, 1)$ . The ensemble method corresponds to an equal-weight linear combination of all methods except validation data analysis and exclusions of 4 and 5 weeks' data. "Model" indicates that reporting factors were estimated via regression and allowed to vary by  $t$  and  $s$ . "Lag" indicates that reporting factors were estimated via Eq 6. "Local" indicates that reporting factors were estimated via Eq 8. Relative absolute biases are calculated relative to the largest value in each column. Model-based  $\pi_{ts}(d)$  estimation assumes reporting factors vary across weeks but incorrectly models how reporting factors vary across weeks.

<https://doi.org/10.1371/journal.pcbi.1010115.g005>

the validation case counts, and we saw huge prediction errors and low coverage (e.g., 10–60%) resulting from very poorly specified estimates of  $\pi_{ts}(d)$  for the rescaling, offset, and imputation methods. When the current season under-reporting was much worse than previous seasons, use of past-season reporting factors resulted in under-correction and under-estimation of validation case counts, but resulting prediction errors were still less than those from analysis of

uncorrected observed data. Combined, these results indicate that we may be better off under-estimating the degree of under-reporting than severely over-estimating it.

In contrast to the model-based and lag-based methods for estimating  $\pi_{ts}(d)$ , the local estimation method performed well with much higher coverages (e.g., < 20% vs 95% for 1 week forecasts) and much lower prediction errors for both Scenarios 3 and 4. The exclusion methods also had good performance. To summarize, the best performing methods were those that did not rely on past season data for handling of under-reporting.

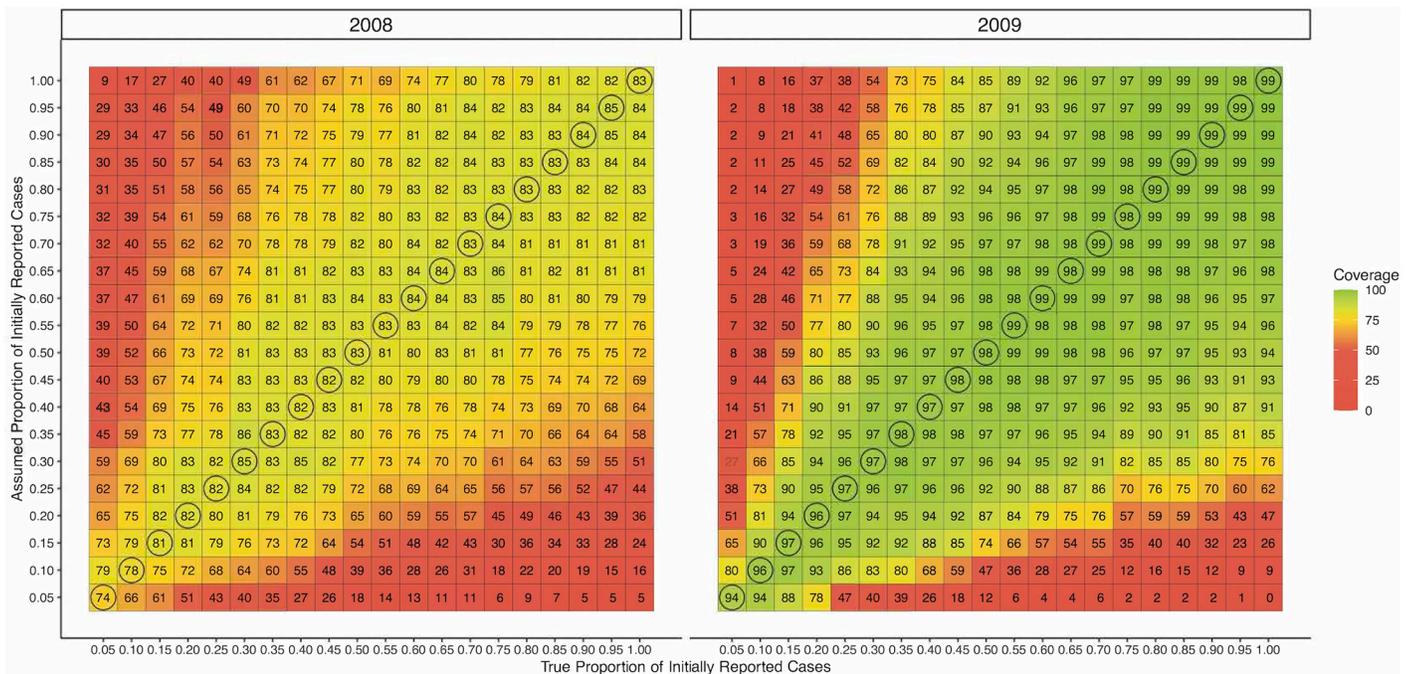
### 5.3 Performance as a function of the reporting rate (Scenario 5)

Previous simulations demonstrate that all methods can perform well in ideal settings when reporting factors are correctly specified. However, the methods relying on reporting factor estimates can run into trouble when the reporting factors are poorly specified. We now suppose that we do not have good historical data from which to estimate these reporting factors. This may be the case, for example, if we are forecasting a newly-emerging disease. We want to evaluate how other methods (e.g., the exclusion method and the rescaling method with  $\pi_{ts}(d)$  estimated using proxy shrinkage) perform for data with different amounts of under-reporting.

We considered simulation Scenario 5, where the true reporting mechanism for each season was varied between  $a = 0.05$  (strong under-reporting) and  $a = 1$  (no under-reporting). Reporting correction, estimation, and forecasting proceeded as before. Performance diagnostics were aggregated across weeks within each season and stratified by the true value of  $a$ . Each result, therefore, represents an aggregation across 50 weeks in each year and across 10 simulated datasets, resulting in 500 repeated comparisons.

**Exclusion method performance as a function of  $a$ .** In Fig X in S1 Text, we provide prediction errors and weighted interval scores for 1 week forecasts in 2006–2009 based on the uncorrected observed data, the validation data, and results from the exclusion method after excluding between 1 and 5 weeks of recent data. Results are shown as a function of true  $a$ . For the exclusion method, there is a trade-off: excluding recent data may avoid some bias due to inclusion of error-prone real-time data but will lose out on the information in the most recent data, which may also negatively impact forecast efforts. For roughly  $a > 0.50$ , the forecast error due to throwing out information outweighed the error due to including the real-time data, and prediction error for the exclusion method was higher than uncorrected analysis. For roughly  $a < 0.50$ , exclusion tended to outperform uncorrected analysis of the observed data, particularly in terms of weighted interval scores. These results may help guide decisions for whether or not to throw out the recent data as a function of the plausible range for  $a$ .

**Performance of proxy shrinkage reporting factors as a function of  $a$ .** Fig Y in S1 Text provides a similar exploration for the rescaling method with  $\pi_{ts}(d)$  estimated using proxy shrinkage based on proxies of varying quality (i.e., correlation with  $N_{ts}(\infty)$ ). The proxy shrinkage method was implemented defining  $w_d = \frac{\omega}{d+1}$  where  $\omega$  took values 0.5 or 1. When the correlation between the proxy and the validation data was very small (e.g., <0.15), proxy shrinkage served to increase prediction error relative to uncorrected observed data analysis. It also increased weighted interval scores relative to uncorrected analysis unless the degree of under-reporting was very high (e.g.,  $a < 0.3$ ). When the proxy correlation was high (e.g., 0.80), rescaling using proxy shrinkage reduced prediction error relative to uncorrected analysis and reduced weighted interval scores when the degree of under-reporting was moderate to high (e.g.,  $a < 0.7$ ). Any advantage of incorporating the proxy to address reporting delay depends on (1) the amount of reporting error and (2) the quality of the proxy. As a conservative rule-of-thumb, we suggest applying the proxy shrinkage method only if  $a < 0.5$  and the proxy is of sufficient quality (e.g., correlation > 0.50).



**Fig 6. Coverage of 95% prediction intervals for 1 week forecasts across various assumed reporting profiles in 2008 and 2009 simulated datasets<sup>1</sup>.** <sup>1</sup> Circled coverages correspond to correctly-specified reporting factors.

<https://doi.org/10.1371/journal.pcbi.1010115.g006>

**Sensitivity analysis with rescaling method.** In settings where the degree of reporting error may be comparatively modest but we still want to evaluate the impact of under-reporting on forecasts, a sensitivity analysis approach may be appealing. In performing a sensitivity analysis, we calculate performance metrics across different assumed values for the reporting factors. In Fig 6, we performed such an analysis across the simulated datasets and across different true and assumed values for *a* to demonstrate how coverage in 2008 and 2009 forecasts varied as a function of assumed *a*. In practice, this exploration would be conducted a single time on the observed data, where the true reporting mechanism would be unknown. Unsurprisingly, we obtained the highest coverages when the value of *a* was correctly specified. Coverages were fairly robust to misspecification of *a* when the true and working values of *a* were both fairly high. When the true amount of under-reporting was very strong (e.g., *a* < 0.10), however, even mildly misspecified *a* resulted in a large loss of coverage.

## 6 Discussion

Delayed reporting of infectious disease cases presents a challenge to forecasting efforts, and no general recommendations for addressing reporting delay in the forecasting setting currently exist. In this work, we synthesize many existing strategies for handling reporting delay into a single unified framework and propose a two-stage estimation procedure that can be applied to forecast modeling in general. In the first stage of the proposed estimation procedure, we leverage either (a) historical data on the accuracy of real-time reporting or (b) external data correlated with disease rates (e.g., social media or Google trends) to estimate the amount of disease under- or over-reporting. In the second stage, we describe how to implement forecast modeling, accounting for the over- or under-reporting using existing methods in the literature (sometimes, with modification).

We applied this methodology to address reporting delay in data on influenza-like illness diagnoses in the US in 2010–2019 and data on dengue fever cases in Puerto Rico between 1990 and 2009 [3]. This analysis demonstrates the potential for improving forecast performance by accounting for delayed case reporting. However, analysis of state-level influenza-like illness data highlights the need for good estimates of the current season reporting mechanisms for rescaling, mean model offset, and imputation methods; when reporting mechanisms are very poorly specified, these correction methods can sometimes do more harm than good.

The comparative performance of the various methods was further evaluated through a simulation study. In ideal settings with correctly specified reporting mechanisms, the rescaling, mean model offset, and imputation methods resulted in reduced forecast errors relative to exclusion of the most error-prone recently reported data. This indicates that, even in the setting with very high under-reporting, the observed data can contribute useful information for forecasting when the reporting mechanism is well-understood. Caution is needed when applying the rescaling method in the setting where disease rates and anticipated inverse reporting factors are low, since this method can be sensitive to small fluctuations in the observed case data in this setting. The mean model offset and imputation methods provide additional stability in this setting and tended to result in better forecast coverage than rescaling in general.

In implementing the rescaling, mean model offset, and imputation methods in this paper, we used a fixed estimate of  $\pi_{ts}(d)$ . However, this strategy does not account for our uncertainty in the estimate of  $\pi_{ts}(d)$ . One reasonable approach for addressing this issue is to view  $\pi_{ts}(d)$  as an estimated parameter with a corresponding distribution (e.g., normal with mean  $\hat{\pi}_{ts}(d)$  and variance based on estimate uncertainty) and apply the imputation correction method using separate *draws* of  $\pi_{ts}(d)$  for generating each imputation. We did not implement this approach in this paper, but we hypothesize that this approach could be useful in settings where there is a lot of unexplained variability in past reporting data.

Several strategies were proposed for estimating  $\pi_{ts}(d)$ . When past season data were available and reporting varied predictably across weeks and seasons, the lag-based and model-based estimation strategies in Eqs 6 and 7 performed well. In settings where the current season's reporting practices differed dramatically from past seasons, however, these methods failed to capture current season reporting dynamics. To address this issue, we proposed an estimator that uses the “local” recent real-time data from the past few weeks to estimate reporting factors for the current week, providing more accurate estimates.

In practice, historical data on reporting delay may be unavailable or incomplete and reporting may be very poorly understood. In this setting, the strategy of excluding the most recent data from forecast modeling produced more reliable forecast performance than methods that used fixed estimates of  $\pi_{ts}(d)$  based on real-time data. However, this exclusion strategy involves a trade-off between loss of information in recently reported data and protection against bias caused by reporting delay. Our simulations suggest that the exclusion method is best applied when the reporting error is high (e.g., when the initial case reports are less than 50% of the validation case counts).

In addition to excluding the most recent data, an alternative strategy is to leverage external data to estimate the amount of under-reporting indirectly by independently predicting the validation case counts. When these predictions are good enough, the comparison between these predicted validation case counts (here, called proxies) and the observed cases provides insight into reporting error. As with the exclusion method, this proxy approach features a trade-off between providing information for estimating reporting errors and adding more noise into the estimation. Through simulation, we evaluated how the performance of the proxy-based estimation strategy varied as a function of the amount of reporting error and the quality of the proxy (i.e., the correlation with validation case counts). We found that even poor proxies (e.g.,

correlation 0.13) contributed some information to under-reporting correction when the amount of under-reporting was very high. As the amount of under-reporting diminished, better and better proxies were needed to improve on uncorrected analysis, particularly in the setting of ILI forecasting.

Synthesizing the insights from our data analysis and simulations, we produced some general recommendations for the handling of reporting delay in Fig 2b. These recommendations provide guidance for handling delay in disease case reporting and may serve as a useful resource to inform practical infectious disease forecasting efforts. Overall, our results clearly demonstrate potential benefits of accounting for reporting delay to improve forecast performance. However, the best-performing methods require either high-quality data on historical reporting mechanisms or strongly predictive proxies of validation case counts. Historical reporting data are often not recorded or not publicly available for many commonly-studied diseases. This motivates additional thinking about disease reporting infrastructure, particularly as new initiatives are underway to expand reporting and forecasting capabilities at the US national level [28].

This work focuses on the problem of handling reporting delay and defines the “true” case rates we want to predict as the eventually-reported validation case rates. In reality, many diseases may be poorly or incompletely captured by official reporting mechanisms [29], as plainly demonstrated by the COVID-19 pandemic. Given estimates of the amount of reporting error relative to the unobserved “true” case counts, the rescaling, mean model offset, and imputation methods discussed in this paper can be applied. The challenge is then to estimate the amount of reporting error, a discussion of which is beyond the scope of this work. Alternatively, this problem can be handled in a sensitivity analysis framework, where we can compare the robustness of our forecasts across multiple plausible assumptions about the amount of reporting error relative to the “true” case counts. Additional work is needed to leverage multiple data streams to inform estimation of reporting error when the target is defined as the never-observed “true” case counts.

## Supporting information

**S1 Text. Supporting information 1: Figures and Tables. Fig A:** Comparing initially reported and validation case counts for Puerto Rico dengue fever and US national influenza-like illness. In all plots, the highest (black) curve corresponds to validation case counts. The lowest (red) curve corresponds to the cases initially reported (lag = 0). **Fig B:** Reported ILI cases in Vermont for each calendar week in the 2017–2018 and 2018–2019 flu seasons. Data for Vermont ILI cases were downloaded on June 13th, 2021. Lines correspond to the first 35 calendar weeks in the corresponding flu season. **Fig C:** Estimated proportion of validation cases initially reported by state using data from the 2018–2019 season and data from the previous two seasons. Data were downloaded on June 13th, 2021. Results from the current season may be viewed as the “truth” here, but data analysis and forecasting is conducted using estimates from the previous two seasons (except for local lag-based estimation). **Table A:** Summary of key notation. **Fig D:** Diagram of Multiple Imputation Algorithm. **Fig E:** Percent error between truncated normal expectation and  $E(\mathbf{N}_{ts}(\infty)|N_{ts}(d))$ . **Fig F:** Comparison of regression model estimates of  $\pi_{ts}(d)$  from Eq 7 based on historical real-time case reporting and the observed values of  $\pi_{ts}(d) = N_{ts}(d)/N_{ts}(\infty)$  for national US ILI and for dengue fever in Puerto Rico. Model-based estimates of  $\pi_{ts}(d)$  are obtained by fitting the regression model in Eq. h in S1 Text. to historical reporting data, excluding the most recent 16 weeks (ILI) or the most recent 6 weeks (dengue fever) from the estimation. Values along the black line indicate model-based estimates of  $\pi_{ts}(d)$  that closely align with observed values of  $\pi_{ts}(d)$ . **Fig G:** Weighted interval scores of forecasts in the Puerto

*Rico dengue fever and national US influenza-like illness data across all weeks.* Results for dengue fever are aggregated across each of 50 weeks in 18 seasons (1992–2009). Results for US national influenza are aggregated across 35 weeks in 8 seasons (2012–2018). The ensemble method corresponds to an equal-weight linear combination of all methods except validation data analysis. Relative weighted interval scores (WIS) are calculated relative to the largest value in each column. **Fig H:** *Boxplots of performance metrics across all weeks.* **Fig I:** *Relative accuracy (1/absolute prediction error, scaled across methods) of 1 week ahead forecasts across rolling 5 week window for US National ILI.* Results based on absolute prediction error for 5 week rolling window centered at plotted week. Results for 35 weeks per season are shown. The black line represents observed validation case counts for each week. **Fig J:** *Relative accuracy (1/absolute prediction error, scaled across methods) of 1 week ahead forecasts across rolling 5 week window for Puerto Rico dengue fever.* Results based on absolute prediction error for 5 week rolling window centered at plotted week. Results for 50 weeks per season are shown. The black line represents observed validation case counts for each week. **Fig K:** *Relative accuracy (1/absolute prediction error, scaled across methods) of 1 week ahead forecasts across rolling 5 week window for state-level ILI in 2018–2019 season.* Results based on absolute prediction error for 3 week rolling window centered at plotted week. Results for 35 weeks per state are shown. **Fig L:** *State-level ILI: Proportion of 50 weeks in which each of 7 methods performs best in terms of nowcasts and 1-week forecast weighted interval scores (ARMA model).* Results for US national influenza are aggregated across 35 weeks in the 2018–2019 season. Florida is excluded from this analysis due to lack of data. **Fig M:** *Forecast bias and weighted interval scores for 3 weeks just before and just after season peak and for 3 weeks surrounding the season minimum/valley.* Results based on aggregating nowcast and forecast performance for 3 weeks just before and 3 weeks just after each season's peak. Results also provided after aggregating 4 weeks before and after season minimum/valley. For ILI, the season valley corresponds to the minimum case counts within the first 35 weeks of the flu season. **Fig N:** *Estimated  $\pi_{ts}(0)$  obtained from different lag estimation methods for state-level ILI in the 2018–2019 flu season.* Data were downloaded on June 13th, 2021. The local lag method used data from the previous 15 weeks to estimate the inverse reporting factors. The standard lag method used data from the previous 2 seasons to estimate the inverse reporting factors. Observed inverse reporting factors for each week are also plotted. **Fig O:** *Relative accuracy (1/absolute prediction error, scaled across methods) of 1 week ahead forecasts for dengue fever using real-time or validation data for modeling.* Results based on absolute prediction error for 5 week rolling window centered at plotted week. Results for 50 weeks per season are shown. The black line represents observed validation case counts for each week. **Fig P:** *Dengue fever data nowcasts and forecasts from ARMA and Inferno modeling of observed data, compared to validation and initial case reports.* **Fig Q:** *Impact of K on estimated  $\pi_{ts}(0)$  using local estimation in dengue fever and national US ILI data (ARMA model)* Local estimates for  $\pi_{ts}(d)$  are obtained using Eq 8. **Fig R:** *Impact of local estimation K on nowcast/forecast errors in dengue fever and national US ILI data using rescaling method (ARMA model).* Results based on aggregation across 50 weeks (dengue fever) or 35 weeks (influenza) across all seasons. Local estimates for  $\pi_{ts}(d)$  are obtained using Eq 8. **Fig S:** *Impact of K for local  $\pi_{ts}(d)$  estimation for forecasting Vermont ILI in the 2018–2019 season (ARMA model).* Results based on aggregation across 35 weeks in the 2018–2019 season. Local estimates for  $\pi_{ts}(d)$  are obtained using Eq 8. **Fig T:** *Simulations: average absolute prediction error for nowcasts and 1 week forecasts by proportion of initially-reported cases in simulated 2009 dengue fever data.* Results aggregated across 50 weeks and 10 simulation replicates for each proportion. **Fig U:** *Data Analysis: nowcast prediction error in actual dengue fever data for 2009.* All values represent differences between nowcast and true validation values. The gray line corresponds to equality with validation data. **Fig V:** *Visualization of simulated dengue fever data.* **Fig W:** *Median forecast weighted interval scores*

in dengue fever simulated data (ARMA models, Scenarios 1–4). Results correspond to the 2009 simulated data and are aggregated across 100 simulated datasets and 50 weeks. Relative weighted interval scores (WIS) are calculated relative to the largest value in each column. **Fig X:** Forecast performance of exclusion method by proportion of initially-reported cases in dengue fever simulated data (ARMA models, Scenario 5). Results aggregated across 50 weeks and 10 simulation replicates for each season/proportion. **Fig Y:** Forecast performance of rescaling method with reporting factors based on proxy shrinkage by proxy quality and proportion of initially-reported cases in dengue fever simulated data (ARMA models, Scenario 5). Results aggregated across 50 weeks and 10 simulation replicates for each season/proportion. **Fig Z:** Average nowcast and forecast biases and coverage of 95% confidence intervals for US ILI simulated data (ARMA models, Scenarios 1–4). Results correspond to the 2018 flu season simulated data and are aggregated across 100 simulated datasets and 35 weeks. Relative biases are calculated relative to the largest value in each column. **Fig AA:** Median forecast weighted interval score for US ILI simulated data (ARMA models, Scenarios 1–4). Results correspond to the 2018 flu season simulated data and are aggregated across 100 simulated datasets and 35 weeks. Relative weighted interval scores (WIS) are calculated relative to the largest value in each column. **Fig BB:** Average absolute error for 1 week forecasts, scaled by season peak case counts. Results aggregated across 35 weeks and 10 simulation replicates for each season/proportion. **Fig CC:** Forecast performance of rescaling method with reporting factors based on proxy shrinkage by proxy quality and proportion of initially-reported cases in ILI simulated data (ARMA models, Scenario 5). Results aggregated across 35 weeks and 10 simulation replicates for each season/proportion. (PDF)

## Acknowledgments

This work is approved for distribution under LA-UR-21-30640. The findings and conclusions in this report are those of the authors and do not necessarily represent the official position of Los Alamos National Laboratory. Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is managed by Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract 89233218CNA000001.

## Author Contributions

**Data curation:** Lauren J. Beesley.

**Formal analysis:** Lauren J. Beesley.

**Methodology:** Lauren J. Beesley, Dave Osthus, Sara Y. Del Valle.

**Software:** Lauren J. Beesley.

**Supervision:** Dave Osthus, Sara Y. Del Valle.

**Visualization:** Lauren J. Beesley.

**Writing – original draft:** Lauren J. Beesley.

**Writing – review & editing:** Lauren J. Beesley, Dave Osthus, Sara Y. Del Valle.

## References

1. Lutz CS, Huynh MP, Schroeder M, Anyatonwu S, Dahlgren FS, Danyluk G, et al. Applying infectious disease forecasting to public health: A path forward using influenza forecasting examples. *BMC Public Health*. 2019; 19(1):1–12. <https://doi.org/10.1186/s12889-019-7966-8> PMID: 31823751
2. Jajosky RA, Groseclose SL. Evaluation of reporting timeliness of public health surveillance systems for infectious diseases. *BMC Public Health*. 2004; 4(29):1–9. <https://doi.org/10.1186/1471-2458-4-29> PMID: 15274746
3. McGough SF, Johansson MA, Lipsitch M, Menzies NA. Nowcasting by Bayesian smoothing: A flexible, generalizable model for real-time epidemic tracking. *PLoS Computational Biology*. 2020; 16(4):1–20. <https://doi.org/10.1371/journal.pcbi.1007735> PMID: 32251464
4. Abbott S, Hellewell J, Thompson RN, Sherratt K, Gibbs HP, Bosse NI, et al. Estimating the time-varying reproduction number of SARS-CoV-2 using national and subnational case counts. *Wellcome Open Research*. 2020; 5:112. <https://doi.org/10.12688/wellcomeopenres.16006.2>
5. Bracher J, Wolfram D, Deuschel J, Görden K, Ketterer JL, Ullrich A, et al. A pre-registered short-term forecasting study of COVID-19 in Germany and Poland during the second wave. *Nature communications*. 2021; 12(1):5173. <https://doi.org/10.1038/s41467-021-25207-0> PMID: 34453047
6. Lawless JF. Adjustments for reporting delays and the prediction of occurred but not reported events. *Canadian Journal of Statistics*. 1994; 22(1):15–31. <https://doi.org/10.2307/3315826.n1>
7. Hohle M, an der Heiden M. Bayesian Nowcasting during the STEC O104: H4 Outbreak in Germany, 2011. *Biometrics*. 2014; 70(4):993–1002. <https://doi.org/10.1111/biom.12194> PMID: 24930473
8. Bastos LS, Economou T, Gomes MFC, Villela DAM, Coelho FC, Cruz OG, et al. A modelling approach for correcting reporting delays in disease surveillance data. *Statistics in Medicine*. 2019; 38(22):4363–4377. <https://doi.org/10.1002/sim.8303> PMID: 31292995
9. England PD, Verrall RJ. Stochastic Claims Reserving in General Insurance. *British Actuarial Journal*. 2002; 8(3):443–518. <https://doi.org/10.1017/S1357321700003809>
10. Osthus D, Daughton AR, Priedhorsky R. Even a good influenza forecasting model can benefit from internet-based nowcasts, but those benefits are limited. *PLoS Computational Biology*. 2019; 15(2):1–19. <https://doi.org/10.1371/journal.pcbi.1006599>
11. Brooks LC, Farrow DC, Hyun S, Tibshirani RJ, Rosenfeld R. Nonmechanistic forecasts of seasonal influenza with iterative one-week-ahead distributions. *PLoS Computational Biology*. 2018; 14(6):1–29. <https://doi.org/10.1371/journal.pcbi.1006134> PMID: 29906286
12. Midthune DN, Fay MP, Clegg LX, Feuer EJ. Modeling reporting delays and reporting corrections in cancer registry data. *Journal of the American Statistical Association*. 2005; 100(469):61–70. <https://doi.org/10.1198/016214504000001899>
13. Günther F, Bender A, Katz K, Küchenhoff H, Höhle M. Nowcasting the COVID-19 pandemic in Bavaria. *Biometrical Journal*. 2021; 63(3):490–502. <https://doi.org/10.1002/bimj.202000112> PMID: 33258177
14. Li T, White LF. Bayesian back-calculation and nowcasting for line list data during the COVID-19 pandemic. *PLoS Computational Biology*. 2021; 17(7):1–22. <https://doi.org/10.1371/journal.pcbi.1009210> PMID: 34252078
15. Farrow DC, Brooks LC, Rumack A, Tibshirani RJ, Rosenfeld R. Delphi Epidata API; 2015. Available from: <https://github.com/cmu-delphi/delphi-epidata>.
16. NNDSS. National Notifiable Diseases Surveillance System; Available from: <https://www.cdc.gov/nndss/index.html>.
17. Kuang D, Nielsen B. Generalized log-normal chain-ladder. *Scandinavian Actuarial Journal*. 2020; 2020(6):553–576. <https://doi.org/10.1080/03461238.2019.1696885>
18. Little RJA, Rubin DB. *Statistical analysis with missing data*. 2nd ed. Hoboken, NJ: John Wiley and Sons, Inc; 2002.
19. Verrall RJ. An investigation into stochastic claims reserving models and the chain-ladder technique. *Insurance: Mathematics and Economics*. 2000; 26(1):91–99.
20. Hildebrandt AK, Bob K, Teschner D, Kemmer T, Leclaire J, Schmidt B, et al. CorCast: A Distributed Architecture for Bayesian Epidemic Nowcasting and its Application to District-Level SARS-CoV-2 Infection Numbers in Germany. *medRxiv*. 2021; p. 2021.06.02.21258209.
21. Kline D, Hyder A, Liu E, Rayo M, Malloy S, Root E. A Bayesian spatio-temporal nowcasting model for public health decision-making and surveillance. *arXiv*. 2021; p. 1–19.
22. Dugas AF, Jalalpour M, Gel Y, Levin S, Torcaso F, Igusa T, et al. Influenza Forecasting with Google Flu Trends. *PLoS ONE*. 2013; 8(2):1–7. <https://doi.org/10.1371/journal.pone.0056176> PMID: 23457520

23. Hickmann KS, Fairchild G, Priedhorsky R, Generous N, Hyman JM, Deshpande A, et al. Forecasting the 2013-2014 Influenza Season Using Wikipedia. *PLoS Computational Biology*. 2015; 11(5):1–29. <https://doi.org/10.1371/journal.pcbi.1004239> PMID: 25974758
24. Jones HE, Harris RJ, Downing BC, Pierce M, Millar T, Ades AE, et al. Estimating the prevalence of problem drug use from drug-related mortality data. *Addiction*. 2020; 115(12):2393–2404. <https://doi.org/10.1111/add.15111> PMID: 32392631
25. Kandula S, Shaman J. Near-term forecasts of influenza-like illness: An evaluation of autoregressive time series approaches. *Epidemics*. 2019; 27(January):41–51. <https://doi.org/10.1016/j.epidem.2019.01.002> PMID: 30792135
26. Osthus D. Fast and Accurate Influenza Forecasting in the United States with Inferno. *bioRxiv*. 2021; p. 1–22.
27. Bracher J, Ray EL, Gneiting T, Reich NG. Evaluating epidemic forecasts in an interval format. *PLoS Computational Biology*. 2021; 17(2):1–15. <https://doi.org/10.1371/journal.pcbi.1008618>
28. Department of Health and Human Services U. CDC Stands Up New Disease Forecasting Center; 2021. Available from: <https://www.cdc.gov/media/releases/2021/p0818-disease-forecasting-center.html>.
29. Gibbons CL, Mangen MJJ, Plass D, Havelaar AH, Brooke RJ, Kramarz P, et al. Measuring underreporting and under-ascertainment in infectious disease datasets: A comparison of methods. *BMC Public Health*. 2014; 14(1):1–17. <https://doi.org/10.1186/1471-2458-14-147> PMID: 24517715