**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2017

Probabilistic forecasting in infectious disease epidemiology: the 13th Armitage lecture

Held, Leonhard ; Meyer, Sebastian ; Bracher, Johannes

Abstract: Routine surveillance of notifiable infectious diseases gives rise to daily or weekly counts of reported cases stratified by region and age group. From a public health perspective, forecasts of infectious disease spread are of central importance. We argue that such forecasts need to properly incorporate the attached uncertainty, so they should be probabilistic in nature. However, forecasts also need to take into account temporal dependencies inherent to communicable diseases, spatial dynamics through human travel and social contact patterns between age groups. We describe a multivariate time series model for weekly surveillance counts on norovirus gastroenteritis from the 12 city districts of Berlin, in six age groups, from week 2011/27 to week 2015/26. The following year (2015/27 to 2016/26) is used to assess the quality of the predictions. Probabilistic forecasts of the total number of cases can be derived through Monte Carlo simulation, but first and second moments are also available analytically. Final size forecasts as well as multivariate forecasts of the total number of cases by age group, by district and by week are compared across different models of varying complexity. This leads to a more general discussion of issues regarding modelling, prediction and evaluation of public health surveillance data.

DOI: <https://doi.org/10.1002/sim.7363>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-146024>

Journal Article

Originally published at:

Held, Leonhard; Meyer, Sebastian; Bracher, Johannes (2017). Probabilistic forecasting in infectious disease epidemiology: the 13th Armitage lecture. *Statistics in Medicine*, 36(22):3443-3460.

DOI: <https://doi.org/10.1002/sim.7363>

Probabilistic forecasting in infectious disease epidemiology: the 13th Armitage lecture

Leonhard Held,^{a,*†} Sebastian Meyer^{a,b} and Johannes Bracher^a

Routine surveillance of notifiable infectious diseases gives rise to daily or weekly counts of reported cases stratified by region and age group. From a public health perspective, forecasts of infectious disease spread are of central importance. We argue that such forecasts need to properly incorporate the attached uncertainty, so they should be probabilistic in nature. However, forecasts also need to take into account temporal dependencies inherent to communicable diseases, spatial dynamics through human travel and social contact patterns between age groups. We describe a multivariate time series model for weekly surveillance counts on norovirus gastroenteritis from the 12 city districts of Berlin, in six age groups, from week 2011/27 to week 2015/26. The following year (2015/27 to 2016/26) is used to assess the quality of the predictions. Probabilistic forecasts of the total number of cases can be derived through Monte Carlo simulation, but first and second moments are also available analytically. Final size forecasts as well as multivariate forecasts of the total number of cases by age group, by district and by week are compared across different models of varying complexity. This leads to a more general discussion of issues regarding modelling, prediction and evaluation of public health surveillance data. Copyright © 2017 John Wiley & Sons, Ltd.

Keywords: age-structured contact matrix; endemic–epidemic modelling; multivariate probabilistic forecasting; proper scoring rules; spatio-temporal surveillance data

1. Introduction

Forecasting is one of the key goals of infectious disease epidemiology, and mathematical and statistical models play a central role in this task. For example, Keeling and Rohani [1] (section 1.5) write that ‘models have two distinct roles, prediction and understanding. Prediction is the most obvious use of models’. However, most of traditional infectious disease epidemiology is concerned with ‘understanding’, and less so with forecasting.

Indeed, the World Health Organization concluded in 2014 that ‘forecasting disease outbreaks is still in its infancy, however, unlike weather forecasting, where substantial progress has been made in recent years’. [2]. As a consequence, the World Health Organization has recently organized an informal consultation with more than 130 global experts entitled ‘Anticipating Emerging Infectious Disease Epidemics’ in order ‘to define the elements within which epidemics of the future will occur’ [3]. Likewise, the United States Federal Government has recently sponsored two prediction competitions on dengue fever in Puerto Rico and influenza-like illness in the USA [4].

There have been various attempts to predict future infectious disease trends with mathematical or statistical models. For example, there has been success in reproducing the spread of the 2002/2003 SARS epidemic based on global flight travel data [5]. Another commonly considered problem is the prediction of influenza epidemics [6–11]. However, these approaches concentrate on data available as univariate time series. This is often of limited value because infectious diseases may spread differently in subgroups

^aEpidemiology, Biostatistics and Prevention Institute, University of Zurich, 8001 Zurich, Switzerland

^bInstitute of Medical Informatics, Biometry, and Epidemiology, Friedrich-Alexander-Universität Erlangen-Nürnberg, 91054 Erlangen, Germany

*Correspondence to: Leonhard Held, Epidemiology, Biostatistics and Prevention Institute, University of Zurich, Hirschengraben 84, 8001 Zurich, Switzerland.

†E-mail: leonhard.held@uzh.ch

of the population considered, for example, in different age groups. Strong spatio-temporal dynamics are also very common. A multivariate view is therefore required to predict incidence in different regions and age groups, and this is the scenario we are considering in this article. Of course, other stratification levels such as gender may also be relevant in other applications.

The key requirements for stratified forecasting of infectious disease incidence are threefold: First, suitable multivariate models have to be developed for stratified time series of surveillance counts. This includes spatio-temporal models [12–15] and models that reflect contact patterns between age groups [16, 17]. Second, focus should be on probabilistic forecasts, rather than on deterministic point predictions [18]. Probabilistic forecasting is the standard in many other areas of science, including economics and weather forecasting. The inherent difficulty in forecasting epidemics [19] makes the use of probabilistic forecasts in infectious disease epidemiology even more necessary to properly reflect forecasting uncertainty. The natural way to obtain probabilistic predictions is through the use of statistical models. An important aspect of the focus on predictions is the emphasis on the primacy of observables and the notion of a model as a (probabilistic) prediction device for such observables. Our models may therefore not perfectly represent (individual-based) disease transmission, but may still be useful for prediction of suitably aggregated surveillance data.

Third, the quality of the forecasts has to be assessed with appropriate predictive model criteria. As in weather forecasting, we use proper scoring rules to assess the quality of probabilistic (count) predictions. Propriety ensures that both calibration and sharpness are addressed. Calibration is defined as the statistical consistency of the probabilistic forecasts and the observations. Sharpness refers to the concentration of the predictive distributions. Thus, the goal is to ‘maximize sharpness subject to calibration’ [20]. Calibration of univariate forecasts can be visually assessed with probability integral transform (PIT) histograms for count data [21]. Calibration tests for count data can also be employed [22–24]. Related methods have been proposed to assess the quality of multivariate probabilistic forecasts [25–27].

The major contribution of this paper is that we address these issues jointly, whereas previous methods have not combined the three key requirements for stratified probabilistic forecasting of infectious disease incidence. In our application, we focus on norovirus gastroenteritis (described in Section 2), an endemic infectious disease where the amount of historical data makes it possible to predict future disease spread. Although we apply our methods for predictive validation in infectious disease epidemiology exclusively in the `hh4` [28] modelling framework (described in Section 3) using the `R` package `surveillance` [29], the proposed methodology for assessing probabilistic forecasts (described in Section 4) can also be applied to other modelling approaches (including Bayesian methods). Section 5 presents results for the norovirus surveillance data, and Section 6 concludes with some discussion.

2. Norovirus gastroenteritis surveillance data

Norovirus gastroenteritis is characterized by a ‘sudden onset of vomiting, diarrhoea and abdominal cramps lasting 2–3 days’ [30]. Its generation time is similar to seasonal influenza (3–4 days). The disease is highly contagious, transmitted not only directly from person to person but also indirectly via contaminated surfaces or food. No vaccination is available. In the northern hemisphere, norovirus incidence consistently peaks during wintertime and is possibly associated with climatic factors [31]. The seasonal pattern can vary across age groups, where outbreaks in childcare facilities and schools have been found to precede those in nursing homes, hospitals and private households [32].

Weekly (lab-confirmed) counts of norovirus infections have been downloaded from <https://survstat.rki.de>. Our training data – identical to the one described and analysed in Meyer and Held (2017) [17] – are stratified into 6 commonly used age groups (0–4, 5–14, 15–24, 25–44, 45–64 and 65+ years) and 12 city districts (of Berlin) and cover the period from week 2011/27 to week 2015/26. The following year (2015/27 to 2016/26) has been used to assess model predictions (test data). Figure 1 shows the age-specific time series.

3. Endemic–epidemic models for infectious disease counts

We start with a modelling framework for spatio-temporal infectious disease counts in Section 3.1. This is extended in Section 3.2 to include social contact patterns between age groups.

Figure 1. Reported incidence of norovirus infections in Berlin, Germany, stratified by age group. The incidence on the left axis is per 100 000 inhabitants and obeys the same $\sqrt{\cdot}$ -scale in all panels. The corresponding counts can be read off from the right axis. The population fractions of the age groups are 4.6%, 7.8%, 10.4%, 30.4%, 27.9% and 18.9%, respectively. The yearly Christmas break in calendar weeks 52 and 1 is highlighted. The last season (to the right of the dashed vertical line) is used as test data.

3.1. Space-time model

In a series of papers [13, 28, 33, 34], a modelling framework for multivariate surveillance count time series of infectious diseases has been proposed. The formulation is built upon an additive decomposition of disease incidence into an endemic and an epidemic component. The endemic component may represent seasonal and climatic variation, heterogeneity in population numbers and other socio-demographic characteristics. The epidemic component describes the force of previously infected individuals and thus spatio-temporal interaction [13].

Let Y_{rt} denote disease counts in region r and week t . Given the counts $Y_{r,t-1}$, $r = 1, \dots, R$, in the previous week, Y_{rt} is assumed to follow a negative binomial distribution, that is, $Y_{rt} \sim \text{NBin}(\mu_{rt}, \psi)$, with mean

$$\mu_{rt} = v_{rt} e_r + \phi_{rt} \sum_{r'} [w_{r'r}] Y_{r',t-1} \quad (1)$$

and overdispersion parameter ψ , so $\text{Var}(Y_{rt}) = \mu_{rt}(1 + \psi \mu_{rt})$. Here, v_{rt} is an unknown endemic log-linear predictor, and e_r is the fraction of the total population that lives in region r . The epidemic log-linear predictor ϕ_{rt} describes the force of infection from time $t-1$ to time t where $[w_{r'r}]$ denote normalized weights for r' to r transmission from $Y_{r',t-1}$ to Y_{rt} , that is, $\sum_r [w_{r'r}] = 1$ [13].

Spatio-temporal modelling has always been an important feature of infectious disease epidemiology. For example, Keeling and Rohani [1] dedicate a whole section to spatial dispersal of infections. An important discovery was that short-time travel behaviour follows approximately a power law with respect to distance [35]. Specifically, the relative frequency $\text{Pr}(d)$ of the distance d traversed by $\sim 500\,000$ dollar bills within 4 days in the USA has been shown to follow a power law: $\text{Pr}(d) \propto d^{-1.59}$. An interesting feature of the power law is that it is scale-free, that is, the power parameter (here -1.59) does not depend on the unit in which the distances d are measured. A power law for areal data has also been proposed [13, 17] and is based on the adjacency order $o_{r'r}$ between regions r' and r as distance measure:

$$w_{r'r} = (o_{r'r} + 1)^{-\rho}. \quad (2)$$

Here, adjacent regions r and r' have order $o_{r'r} = 1$, regions where we need to traverse one other region are of order 2 and higher orders are defined accordingly. The power parameter ρ is treated as unknown and estimated from the data.

In our application, we model v_{rt} and ϕ_{rt} in (1) as

$$\log v_{rt} = \alpha_r^{(\nu)} + \beta x_t + \gamma^{(\nu)} \sin(\omega t) + \delta^{(\nu)} \cos(\omega t) \quad (3)$$

$$\log \phi_{rt} = \alpha_r^{(\phi)} + \gamma^{(\phi)} \sin(\omega t) + \delta^{(\phi)} \cos(\omega t) \quad (4)$$

with region-specific effects $\alpha_r^{(\nu)}$ and $\alpha_r^{(\phi)}$ in both components, a Christmas break indicator x_t (to account for reduced reporting and school closure in calendar weeks 52 and 1) in the endemic component and sinusoidal log-rates with frequency $\omega = 2\pi/52$ in both components [34]. An alternative model replaces $\alpha_r^{(\phi)}$ by $\alpha^{(\phi)} + \tau \log(e_r)$, so it includes a parametric ‘gravity model’ [12, 36] as described in the next section in more detail.

A multivariate branching process formulation [28] is useful to derive the (time-dependent) epidemic proportion λ_t as a function of the model parameters. We omit details here but will report the average epidemic proportion $\bar{\lambda}$ during the training period, which describes the average proportion of the incidence that can be explained with the epidemic component.

3.2. Age-structured spatio-temporal model

Consider now counts Y_{grt} stratified by age group g , region r and week t . Again, we use a negative binomial likelihood with mean

$$\mu_{grt} = v_{grt} e_{gr} + \phi_{grt} \sum_{g', r'} [c_{g'g} w_{r'r}] Y_{g', r', t-1}, \quad (5)$$

now with age group-specific overdispersion parameter ψ_g . The unknown log-linear predictors v_{grt} and ϕ_{grt} , and the known population fractions e_{gr} also become age group specific. There are now two sets of weights: $c_{g'g}$ quantifies transmission from age group g' to age group g , while $w_{r'r}$ remain spatial weights for r' to r transmission. The product $c_{g'g} w_{r'r}$ is normalized such that $\sum_{g, r} [c_{g'g} w_{r'r}] = 1$.

We now specify the model as

$$\log v_{grt} = \alpha_g^{(\nu)} + \alpha_r^{(\nu)} + \beta x_t + \gamma_g^{(\nu)} \sin(\omega t) + \delta_g^{(\nu)} \cos(\omega t) \quad (6)$$

$$\log \phi_{grt} = \alpha_g^{(\phi)} + \alpha_r^{(\phi)} + \tau \log(e_{gr}) + \gamma_g^{(\phi)} \sin(\omega t) + \delta_g^{(\phi)} \cos(\omega t) \quad (7)$$

$$c_{g'g} = (\mathbf{C}^\kappa)_{g'g} \quad (8)$$

$$w_{r'r} = (o_{r'r} + 1)^{-\rho} \quad (9)$$

with age group and region-specific effects α_g and α_r in both components (ν and ϕ). Again, we include an indicator x_t for the Christmas break and let seasonal terms enter both components, where we now assume age group-specific seasonality with coefficients $\gamma_g^{(\nu)}$ and $\delta_g^{(\nu)}$ in the endemic component. Note that the formulation (7) for ϕ_{grt} extends the one used in Meyer and Held (2017) [17], because we here also allow for seasonal variation in the epidemic component [34]. The gravity model feature $\phi_{grt} \propto e_{gr}^\tau$ causes the amount of disease transmission to scale with the population size e_{gr} of the ‘importing’ age group g in region r – with unknown coefficient τ .

To specify the weights $c_{g'g}$, we estimate an age-structured contact matrix $\mathbf{C}$ from physical or conversational contacts recorded as part of the POLYMOD survey in Germany, whose 1341 participants were broadly representative of the whole population in terms of geographical spread, age and sex [37]. Note that social contacts are reciprocal at the population level, that is, there is an equal number of contacts between age group g and g' as between age group g' and g , which has been taken into account appropriately [38]. We finally specify a power transformation for the contact matrix [17],

$$c_{g'g} = (\mathbf{C}^\kappa)_{g'g} \text{ where } \mathbf{C}^\kappa := \mathbf{E} \mathbf{\Lambda}^\kappa \mathbf{E}^{-1},$$

where $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues and $\mathbf{E}$ is the corresponding matrix of eigenvectors of $\mathbf{C}$. The exponent κ adjusts the amount of transmission between the age groups and is to be estimated as part of the epidemic model component.

To investigate the relevance of the various model features, we consider four different models of increasing complexity: model 1 is based on equation (6) with $\phi_{grt} \equiv 0$ in (5), so it only includes the endemic component. Model 2 extends this with an autoregression on the number of cases $Y_{g,r,t-1}$ in the same age group and region, that is, (5) reduces to $\mu_{grt} = v_{grt} e_{gr} + \phi_{grt} Y_{g,r,t-1}$ with ϕ_{grt} as in (7), but without the gravity model component $\tau \log(e_{gr})$. Model 3 allows for spatial dispersal according to (9) and includes the gravity model in (7), but does not include contact data, so $c_{g'g} = 1$ for $g' = g$, otherwise $c_{g'g} = 0$. Finally, model 4 is based on the full formulation (5)–(9), so it allows for dispersal across space and across age groups.

3.3. Reference models for long-term forecasts

Our target quantity to assess the long-term forecasts is the number of reported cases during the test period, suitably aggregated to stratification levels of interest. At each level, we define simpler reference models that are applied directly to the respective aggregated data. The simplest cases are aggregation over age group and region to weekly totals (the so-called epidemic curve) and, aggregating in addition over time, the final size, that is, the total number of reported cases during the whole year. We obtain reference forecasts for these two quantities from model 5, a univariate negative binomial time series model for the counts $Y_t = \sum_{g,r} Y_{grt}$:

$$\begin{aligned} Y_t &\sim \text{NBin}(\mu_t, \psi) \\ \mu_t &= v_t + \phi_t Y_{t-1} \\ \log(v_t) &= \alpha^{(v)} + \beta x_t + \gamma^{(v)} \sin(\omega t) + \delta^{(v)} \cos(\omega t) \\ \log(\phi_t) &= \alpha^{(\phi)} + \gamma^{(\phi)} \sin(\omega t) + \delta^{(\phi)} \cos(\omega t). \end{aligned}$$

Likewise, aggregating the data only over regions yields a model for $Y_{gt} \sim \text{NBin}(\mu_{gt}, \psi_g)$ to additionally obtain final size predictions in the different age groups (model 6):

$$\begin{aligned} \mu_{gt} &= v_{gt} + \phi_{gt} \sum_{g'} [c_{g'g}] Y_{g',t-1} \\ \log(v_{gt}) &= \alpha_g^{(v)} + \beta x_t + \gamma_g^{(v)} \sin(\omega t) + \delta_g^{(v)} \cos(\omega t) \\ \log(\phi_{gt}) &= \alpha_g^{(\phi)} + \gamma^{(\phi)} \sin(\omega t) + \delta^{(\phi)} \cos(\omega t) \\ c_{g'g} &= (\mathbf{C}^K)_{g'g}. \end{aligned}$$

Two models for regional counts $Y_{rt} \sim \text{NBin}(\mu_{rt}, \psi)$ (aggregated over age group) have been applied. Both are based on the decomposition (1) with power law (2) and endemic model (3). Model 7 uses the epidemic model (4) with region-specific intercept $\alpha_r^{(\phi)}$. Model 8 is more parsimonious and replaces $\alpha_r^{(\phi)}$ by $\alpha^{(\phi)} + \tau \log(e_r)$, so it includes a gravity model component instead. Note that the more general formulation $\alpha_r^{(\phi)} + \tau \log(e_r)$ is not identifiable.

For comparison, we adopt the idea of naïve reference forecasts [20]. At the different levels of aggregation, we define these as (potentially multivariate) four-point distributions that assign a probability of 1/4 to the four realizations of the quantity of interest in each year of the training data.

4. Predictive model assessment

We validate the models based on probabilistic one-step-ahead and long-term predictions of the last year (2015/27 to 2016/26) given the data from the previous 4 years. The one-step-ahead predictions all follow negative binomial distributions, so the predictive probability mass function is known in closed form. Note that the underlying parameters change from prediction to prediction, as we always refit the model to all available data prior to the counts to be predicted. A simpler approach would always use the parameter estimates based on the original training data.

The model formulation implies that one-step-ahead predictions in different age groups and different regions are conditionally independent. Conditional on the first 4 years, long-term predictions, sometimes called path forecasts, are generated for the last year through Monte Carlo simulation by sequentially simulating from one-step-ahead predictions. Note that our time series model is multivariate (across regions and age groups), and thus, the path forecasts are also multivariate. Alternatively, the first two moments of multivariate path forecasts can be computed analytically in the model class considered (see Appendix A for details).

Predictions are computed for all combinations of 6 age groups $\times$ 12 districts $\times$ 52 weeks and have been suitably aggregated if required. Similarly, the first two moments of aggregated counts can be calculated analytically from the moments of the original path forecasts.

4.1. One-step-ahead forecasts

Scoring rules (also called scoring functions) are the key measures for the evaluation of probabilistic forecasts. For count data, a probabilistic forecast is a predictive probability distribution F with probability mass function $f(y)$ on the set of the non-negative integers. Scoring rules assign a numerical score based on F and on the true value y_{obs} , which has later materialized. They are typically negatively oriented, that is, the smaller, the better. They are called proper, if they do not provide any incentive to the forecaster to digress from her true belief and strictly proper if any such digress results in a penalty, that is, the forecaster is encouraged to quote her true belief rather than any other predictive distribution. Note that in the literature, inappropriate scoring methods are still often used, for example, correlation coefficients between point predictions and observations [9, 11]. It is well known in the medical literature that high correlation does not necessarily imply good agreement [39], and therefore, a very poor forecasting method may have high correlation. Furthermore, usage of metrics incorporating the whole probabilistic forecast (rather than using only point predictions) is still rare [7].

For ease of presentation, we drop the indices g , r and t in this section and let $Y \sim F$, which we compare with the actual observation y_{obs} . The logarithmic score ($\log S$) [40] is strictly proper and defined as

$$\log S(F, y_{\text{obs}}) = -\log f(y_{\text{obs}}),$$

that is, minus the log predictive probability mass function at the observed value. A normal approximation gives the (univariate) Dawid-Sebastiani score [41]. A strictly proper alternative is the ranked probability score (RPS) [20], which can be written for count data as

$$\text{RPS}(F, y_{\text{obs}}) = \sum_{k=0}^{\infty} \{ \Pr(Y \leq k) - \mathbf{1}(y_{\text{obs}} \leq k) \}^2,$$

the sum of the Brier scores for binary predictions at all possible thresholds $k \in \{0, 1, \dots\}$ [21]. An equivalent definition is

$$\text{RPS}(F, y_{\text{obs}}) = \mathbb{E}|Y - y_{\text{obs}}| - \frac{1}{2} \mathbb{E}|Y - Y'|; \quad (10)$$

here, Y and Y' are independent realizations from $f(y)$ [20].

For Poisson and negative binomial one-step-ahead predictive distributions, both scores ($\log S$ and RPS) can be computed directly based on the corresponding probability mass function, as implemented in the R package *surveillance* [29]. In order to develop a calibration test for one-step-ahead count forecasts [24], we compute the average score $\overline{\text{RPS}}$ across all regions, age groups and one-step-ahead predictions and compare it to its expected value $\mathbb{E}_0(\overline{\text{RPS}})$ under the null hypothesis of ‘forecast validity’ [23], sometimes also called ‘perfect calibration’ (H_0). The difference $\overline{\text{RPS}} - \mathbb{E}_0(\overline{\text{RPS}})$ is standardized using the variance $\text{Var}_0(\overline{\text{RPS}})$. Note that our model implies that the one-step-ahead predictions in different age groups and different districts are (conditionally) independent, so under H_0 independence holds for all components of $\overline{\text{RPS}}$ for fixed t . Furthermore, proofs of independence are available for the sequence of scores across time [22, 23], so in the computation of $\text{Var}_0(\overline{\text{RPS}})$, we can assume independence of all components of $\overline{\text{RPS}}$. Finally, a z -statistic is defined, which follows asymptotically a standard normal distribution if the null hypothesis H_0 is true:

$$z = \frac{\overline{\text{RPS}} - \mathbb{E}_0(\overline{\text{RPS}})}{\text{Var}_0(\overline{\text{RPS}})^{1/2}} \underset{H_0}{\overset{a}{\sim}} N(0, 1).$$

The sign of z indicates if the observations are overdispersed/underdispersed relative to the predictions (+/− sign of z) [24, 42]. A (two-sided) P -value can be computed to quantify the evidence for miscalibration.

The PIT $F(y_{\text{obs}})$ is often used to visually assess calibration; here, $F(y)$ denotes the predictive CDF. Under H_0 , the PITs of a sequence of absolutely continuous one-step-ahead forecasts are independent and

uniformly distributed [43]. Here, we use a modification of the PIT histogram for count data [21]. We can also apply the Diebold–Mariano test for equal predictive performance [44].

For negative binomial predictions, sharpness can be quantified with the estimated overdispersion parameter ψ , which may be age-dependent. Easier to interpret is the variance-to-mean ratio $1 + \psi \mu$, but this requires an estimate of the predictive mean μ , where we simply average (in each age group for age-dependent overdispersion parameters) the predictive means across weeks and districts.

4.2. Multivariate path forecasts

One possibility for a proper scoring rule is the multivariate Dawid-Sebastiani score (DSS) [41]

$$\text{DSS}(F, \mathbf{y}_{\text{obs}}) = \log|\Sigma| + (\mathbf{y}_{\text{obs}} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{y}_{\text{obs}} - \boldsymbol{\mu}), \quad (11)$$

which depends only on the mean vector $\boldsymbol{\mu}$ and the covariance matrix Σ of the predictive distribution. The first term in (11) involves the determinant $|\Sigma|$ of the covariance matrix Σ ; here, Σ is a $d \times d$ matrix. Transformed to

$$\text{DS} = |\Sigma|^{1/(2d)}, \quad (12)$$

this is known as the *determinant sharpness* (DS) and recommended as a multivariate measure of sharpness [25], with smaller values corresponding to sharper predictions. In higher dimensions, it is useful to report the log determinant sharpness $\log \text{DS} = \log|\Sigma|/(2d)$ to avoid unnecessarily large numbers. For presentational reasons, we report the Dawid-Sebastiani score (11) in a scaled version and divide it by $2d$. Finally, we compute an approximate P -value for forecast validity based on the DSS [27]. We note that this method assumes approximate normality of the observed data, which may be questionable if the observed counts are very low.

However, evaluation of (11) and (12) based on Monte Carlo estimates of the first two moments is not recommended, because the determinant $|\Sigma|$ is known to be very sensitive to Monte Carlo sampling error [26]. Fortunately, we can evaluate this score exactly because we have analytical results for recursive computation of $\boldsymbol{\mu}$ and Σ in our modelling framework (see Appendix A for details).

In order to take the full predictive distribution into account, the multivariate log-score $\log S(F, \mathbf{y}_{\text{obs}}) = -\log f(\mathbf{y}_{\text{obs}})$ is an obvious choice. However, the log-score is very difficult to compute based on samples from the predictive distribution F , as its approximation via the empirical distribution is not *consistent* [45]. Specifically, if there is no sample exactly equal to the observed vector $\mathbf{y}_{\text{obs}}$, the empirical estimate of $f(\mathbf{y}_{\text{obs}})$ will be zero and the log-score will be infinite. Instead, we use a generalization of the ranked probability score (10), the so-called energy score (ES) [20]

$$\text{ES}(F, \mathbf{y}_{\text{obs}}) = \mathbb{E}\|\mathbf{Y} - \mathbf{y}_{\text{obs}}\| - \frac{1}{2}\mathbb{E}\|\mathbf{Y} - \mathbf{Y}'\| \quad (13)$$

for multivariate forecasts $\mathbf{Y}$; here, $\|\cdot\|$ denotes the Euclidean norm. The energy score reduces to the Euclidean error $\|\boldsymbol{\mu} - \mathbf{y}_{\text{obs}}\|$ for a deterministic point mass forecasts at $\boldsymbol{\mu}$ and can be reported in the same unit as the observations. It has been noted [26, 46] that the ES discriminates well between forecasts with different means or variances, but less so for forecasts with different correlation structures.

Monte Carlo estimation of (13) is possible if samples from the predictive distribution are available [25]. These samples can also be used to derive the distribution of ES under the null hypothesis of forecast validity: Suppose we have n_{sim} samples $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(n_{\text{sim}})}$ from the predictive distribution $f(\mathbf{y})$. We can then generate n_{sim} samples $\text{ES}_1, \dots, \text{ES}_{n_{\text{sim}}}$ from the distribution of $\text{ES}(F, \mathbf{y}_{\text{obs}})$ under H_0 by computing the Monte Carlo approximation of (13) based on the realization $\mathbf{y}_{\text{obs}} = \mathbf{y}^{(i)}$, $i = 1, \dots, n_{\text{sim}}$ and approximate the distribution of $\mathbf{Y}$ using the remaining $n_{\text{sim}} - 1$ samples from $f(\mathbf{y})$. Computation of a (one-sided) Monte Carlo P -value is then possible based on the energy score ES_{obs} of the actual observation $\mathbf{y}_{\text{obs}}$:

$$p = \frac{1}{1 + n_{\text{sim}}} \left(1 + \sum_{i=1}^{n_{\text{sim}}} \mathbf{1}(\text{ES}_i \geq \text{ES}_{\text{obs}}) \right).$$

5. Results

5.1. Model fitting results

Selected parameter estimates together with AIC values are given in Table I for models 1–4. AIC improves for each of the features added from model 1 to model 4, where the largest impact is due to the account for autoregression ($1 \rightarrow 2$). Likewise, the average epidemic proportion $\bar{\lambda}$ increases with increasing model complexity, so in model 4, 70% of the incidence can be explained by the epidemic component. There is strong evidence that the epidemic part scales with the population size of the ‘importing’ districts and age groups with the coefficient τ estimated as 0.62 (95% CI: 0.27 to 0.97) in model 3 and 0.82 (95% CI: 0.49 to 1.14) in model 4. The power parameter ρ of the spatial power law (9) is estimated to be 2.21 (95% CI: 1.92 to 2.55) in model 3 and 2.30 (95% CI: 2.01 to 2.63) in model 4. Thus, the inclusion of social mixing between age groups in model 4 results in a slightly larger coefficient, reducing the range of spatial dispersal. Note that the results of model 4 differ slightly from the estimates reported in Meyer and Held [17] because of the additional seasonal variation in the epidemic component here. This adds two parameters but improves AIC by 18.3 points. The estimate of the power coefficient κ for the contact matrix C is 0.41 (95% CI: 0.29 to 0.60). For comparison, in model 6, the estimated exponent is nearly identical ($\hat{\kappa} = 0.40$) but has more uncertainty (95% CI: 0.23 to 0.67).

5.2. Assessing one-step-ahead forecasts

Figure 2 gives the variance-to-mean ratios of the one-step-ahead forecasts in the different age groups and models. Generally, the variance-to-mean ratio decreases with increasing model complexity. The predictions become sharper in age groups 65+ and 45–64 when we move from model 1 to 2, 3 and finally to model 4. In age group 00–04, we also see sharper predictions in models 2–4 compared with model 1.

The mean RPS with z -value and P -value is shown in Table II. The PIT histograms [21] based on all $6 \times 12 \times 52 = 3744$ one-step-ahead forecasts are shown in Figure 3. The mean RPS values are in the expected order with smaller (i.e. better) values for models of increasing complexity and the full model 4 as the best. For all models except for model 1 we see some evidence for underdispersed predictions (Figure 3) with more PITs close to one than we would expect under forecast validity. This indicates potential problems of the negative binomial distribution in the more complex model formulations.

Table I. Comparison of models 1–4 in terms of AIC with selected parameter estimates and associated 95% Wald confidence intervals.

Model	dim	Δ AIC	$\hat{\tau}$	$\hat{\rho}$	$\hat{\kappa}$	$\bar{\lambda}$
1	36	0.0	—	—	—	0.00
2	55	−514.4	—	—	—	0.48
3	57	−631.1	0.62 (0.27–0.97)	2.21 (1.92–2.55)	—	0.68
4	58	−677.7	0.82 (0.49–1.14)	2.30 (2.01–2.63)	0.41 (0.29–0.60)	0.70

Figure 2. Estimated age group-specific variance-to-mean ratios of the one-step-ahead predictive distributions.

Table II. Comparison of RPS and calibration test results, overall and in the first and last age group.

Model	Overall			Age group 00–04			Age group 65+		
	RPS	z	P -value	RPS	z	P -value	RPS	z	P -value
1	0.454	0.65	0.52	0.535	2.84	0.005	1.112	−1.83	0.067
2	0.436	1.35	0.18	0.516	2.39	0.017	1.019	−0.17	0.87
3	0.434	1.80	0.072	0.513	2.18	0.029	1.009	0.82	0.41
4	0.433	2.28	0.022	0.512	2.26	0.024	1.008	0.79	0.43

Figure 3. Probability integral transform histograms of the 3744 one-step-ahead forecasts for models 1–4.

The left and middle panels of Figure 4 give PIT histograms of the one-step-ahead forecasts for models 1–4 and 6 for age group 00–04 and 65+, respectively. The PIT histograms for the other age groups are all very close to uniformity. The PIT histograms for age group 00–04 are very similar for models 1–4 with a pronounced peak at values close to 1. This indicates a tendency to underpredict relatively large observed values. A similar pattern can be seen for model 6, where the PIT histogram is based on a smaller number of observations. The PIT histogram for age group 65+ and model 1 indicates a bias of the estimates with a tendency to predict larger values than observed. This results in overprediction of the aggregated counts shown in the right panel of Figure 4. In contrast, the PIT histogram of model 6 is inconspicuous. The other three models compensate the bump of the PIT histogram around 0.2 with another peak at values of PIT close to 1. This seems to lead to greater uncertainty of the final size predictions as shown in the right panel.

Note that the calibration tests based on the mean RPS of the one-step-ahead forecasts (Table II) give evidence for miscalibrated forecasts in age group 00–04 for all the models, but no such evidence for age group 65+. In age group 00–04, there is evidence for underdispersed predictions with values of the z -statistic larger than 2 for all four models and an even more pronounced peak of the PIT histogram close to one than for the overall set of forecasts (compare the left panel in Figure 4). However, we see some evidence for overdispersed predictions in age group 65+ for model 1 with a negative value of the z -statistic, which fits the corresponding pattern of the PIT histogram in Figure 4.

Finally, P -values from the Diebold–Mariano test comparing models 1–4 are shown in Table III. We see very strong evidence ($P < 0.0001$) for differences in predictive performance between model 1 and the other three models but no evidence for differences in predictive performance between models 2, 3 and 4.

Figure 4. Left and middle panel: probability integral transform (PIT) histograms of one-step-ahead forecasts in age groups 00–04 and 65+, respectively. The PIT histograms for model 6 are based on 52 observations each, whereas the other histograms are based on 12×52 observations. Right panel: final size predictions in the different age groups based on 1000 simulations from models 1–4 and 6. The red dots represent the observed counts.

5.3. Properties of long-term forecasts

Table IV gives values of the log determinant sharpness for the different models and the different stratification levels. The first column gives values for total final size, where we observe for models 1 to 4 a tendency for larger values of logDS, that is, more dispersed predictions, with increasing complexity. However, this pattern is less clear for the models 5–8 that work on aggregated data. An interesting pattern appears for the predictions by age group, region, or week: With only one exception, the models that anal-

Table III. *P*-values from the Diebold–Mariano test for the pairwise comparison of RPS across models 1–4.

	Model 1	Model 2	Model 3	Model 4
Model 1	—	< 0.0001	< 0.0001	< 0.0001
Model 2	—	—	0.48	0.23
Model 3	—	—	—	0.12

Table IV. Log determinant sharpness (logDS).

Model	Total		By age group		By region		By week	
	logDS	Rank	logDS	Rank	logDS	Rank	logDS	Rank
1	4.63	1	3.12	1	3.30	1	2.30	4
2	5.32	3	3.32	2	3.78	6	2.29	2
3	6.06	7	3.59	4	3.62	2	2.29	3
4	6.10	8	3.51	3	3.63	3	2.27	1
5	5.86	5					2.57	8
6	5.92	6	3.64	5			2.41	7
7	5.68	4			3.66	4	2.32	6
8	5.21	2			3.69	5	2.30	5

Figure 5. Correlations implied by the different models: between age-groups (top left), regions (top right) and consecutive weeks (bottom). Models where correlations are 0 by construction are omitted from the two top figures.

Figure 6. Final size predictions based on 1000 simulations with RPS and two-sided P -value. The white dots represent the predictive means.

use the data in the finest resolution (1–4) give sharper predictions than the models that analyse aggregate data (5–8). In particular, the most sophisticated model with both a spatial power law and social contact data (model 4) gives the sharpest prediction for the epidemic curve. This can be explained by the fact that the predictions from this model have the largest first-order autocorrelations among all models from calendar week 48 onwards (and also quite large before) (Figure 5 bottom). Note also that the logDS values by region roughly reflect the order of the correlations between regions (Figure 5 top right) with a tendency for larger values of logDS for smaller correlations. However, logDS is also a function of the marginal variances, so the agreement of logDS and the correlations is not perfect. Similarly, the correlations of the predictions between age groups, shown in the top left panel of Figure 5 for models 4 and 6, do not lead to smaller values of logDS by age group because of larger variances compared to models 1 and 2, as can be seen in the right panel of Figure 4.

5.4. Assessing long-term forecasts

5.4.1. Final size predictions. Final size predictions are shown in Figure 6. All models tend to overpredict the total number of cases with varying uncertainty. In the worst cases (models 1, 2 and 8), the observed number of cases is not or only poorly supported by the predictions. The best model (in terms of RPS) is model 4, the most complex formulation, followed by model 6. This underlines the importance of using an age-structured modelling approach supported by social contact data.

5.4.2. Final size predictions by age group. Figure 4 (right panel) displays the probabilistic forecasts for the total number of cases in the different age groups for models 1–4 and 6. While all models predict the number of cases reasonably well in the lower five age groups, there are remarkable differences for the last age group 65+. Model 1 overpredicts the number of cases in age group 65+ considerably, while models 2–4 give similar point predictions, but with increasing uncertainty. As a consequence, the observed number of cases in age group 65+ is well supported by the predictive distributions of models 3, 4 and 6, less so for model 2, and not at all for model 1. Note that predictions are correlated across age groups only for the models that incorporate social contact data (models 4 and 6), as can be seen in the top left panel of Figure 5.

5.4.3. Long-term predictions of the epidemic curve. Figure 7 compares the long-term predictions of the epidemic curve from the different models with the observed time series. The point predictions do not differ much between the models, but there is variation in the associated uncertainty with wider prediction bands for the more complex models 3, 4 and 5.

5.4.4. Scores and P -values for forecast validity. Table V gives both the energy and the Dawid–Sebastiani score by different stratification levels. Likewise, Table VI gives the corresponding P -values for forecasts validity. The two scores agree quite well with model 4 giving the best (or second best) predictions in total, by age group and by region. Somewhat surprisingly, model 4 gives quite poor predictions by week for the DSS (rank 7), but not for the ES (rank 2). This can be explained by the fact that the ES is known to be insensitive to misspecifications in the correlation structure of multivariate predictions. We have already

Figure 7. Long-term predictions of the epidemic curve in the last year (2015/27 to 2016/26). The dots represent the observed counts.

noted that the autocorrelations of the model 4 predictions are quite large. However, the observed time series (Figure 7) has an oscillating pattern around weeks 20 and 30. This discrepancy between observed and predicted correlation seems to be detected by the DSS, but not by the ES. This can also be seen from Table VI, where the test based on DSS gives strong evidence for miscalibration of the forecasts by week for all models except for models 5 and 6. Interestingly, these two models have the largest values of logDS (Table IV) but are still best in terms of DSS (Table V). Larger uncertainty of the predictions (Figure 7)

Table V. Energy score and scaled Dawid–Sebastiani score for the different models and aggregation levels.

Model	Total				By age group				By region				By week			
	ES	Rank	DSS	Rank	ES	Rank	DSS	Rank	ES	Rank	DSS	Rank	ES	Rank	DSS	Rank
1	452.5	8	17.10	8	409.7	5	5.67	5	161.8	4	4.85	6	136.3	8	3.32	8
2	418.3	7	8.67	7	377.3	4	4.13	4	162.7	5	4.25	5	133.5	7	3.12	4
3	322.5	3	6.83	3	293.0	3	3.93	2	158.8	3	3.98	1	126.2	6	3.14	6
4	242.7	1	6.58	1	227.0	1	3.88	1	132.8	1	3.99	2	109.3	2	3.21	7
5	369.3	5	7.03	4									118.3	3	3.05	1
6	269.3	2	6.68	2	277.3	2	4.07	3					108.0	1	3.07	2
7	388.5	6	7.39	5					180.2	6	4.10	3	122.4	4	3.13	5
8	364.5	4	8.41	6					136.1	2	4.22	4	123.5	5	3.11	3

Table VI. *P*-values for forecast validity.

Model	Total		By age group		By region		By week	
	Direct	DSS	ES	DSS	ES	DSS	ES	DSS
1	0.002	< 0.0001	< 0.0001	< 0.0001	0.0006	0.0002	0.0003	< 0.0001
2	0.002	0.01	0.021	0.14	0.17	0.52	0.034	0.002
3	0.084	0.21	0.19	0.66	0.16	0.72	0.15	0.001
4	0.24	0.33	0.27	0.62	0.29	0.72	0.25	0.0001
5	0.07	0.13					0.35	0.57
6	0.15	0.22	0.14	0.52			0.34	0.064
7	0.034	0.064			0.097	0.57	0.09	0.003
8	0.01	0.011			0.093	0.39	0.008	0.003

combined with relatively small autocorrelations (see bottom panel of Figure 5) seems to make model 5 the best in terms of DSS. Note that the corresponding test based on the ES identifies only model 1 and 8 as miscalibrated (and model 2 to a lesser extent). For the other stratification levels, the *P*-values based on the DSS tend to be larger than the corresponding ones based on the ES. This suggests that the calibration test based on the ES is less sensitive to misspecifications of the predictions than the test based on DSS.

Finally, we report the performance of the naïve reference forecasts from Section 3.3. The Dawid–Sebastiani scores could not be evaluated, because the covariance matrices could not be estimated reliably from only four observations. In terms of the energy score, however, these forecasts performed surprisingly well for certain quantities: They rank second both for the yearly total burden (ES = 261.4) and the burden by age group (ES = 255.4). This is because the seasons 2012/13 and 2015/16 are remarkably similar in terms of the incidence per age group. For the yearly burden per region (ES = 177.7) and the weekly incidences (ES = 137.1), however, the naïve forecasts perform poorly.

6. Discussion

We have described a multivariate time series model for infectious disease surveillance counts by borrowing strength from different regions and different age groups. This model framework provides probabilistic forecasts of infectious disease spread. Specifically, means and covariance matrices are available analytically as well as Monte Carlo samples from the full predictive distribution. Predictive model assessment helps to identify poor predictive models. The application has shown the importance of predictive model assessment in different strata. We have emphasized the importance to use appropriate methodology for predictive model assessment, including proper scoring rules and tests for forecast validity.

Our study has shown that complex modelling on the original fine resolution generally leads to better predictions of future disease incidence, even of aggregated quantities. As a consequence, the most complex model 4 was nearly always the best in predictive performance. The only exception are predictions by week, where model 4 turned out to predict poorly in terms of DSS. We were able to explain this feature by the discrepancy of strong predicted correlations on the one hand, but oscillating observed number of cases on the other hand. This model deficiency was not detected by the energy score, in accordance with similar observations made in the literature [26, 46]. Consideration of models for aggregated data showed the importance to integrate contact pattern data in the analysis of age-stratified surveillance data

(model 6), as this model was consistently among the top three models in terms of ES and DSS. The spatial dimension turned out to be less important, which may be different in applications with more districts.

Our forecasts were always based on one particular model. A possible extension of our approach would be to use a Bayesian model average framework to combine forecasts from different models into one averaged forecast. The model weights may even differ based on differing forecasts and could be based on AIC or BIC computed from the training data [47], if the models considered act on the same data resolution. Such model averages are known to have better forecast properties and have been successfully applied in weather forecasting [48]. For example, they might help to improve the PIT histograms shown in Figures 3 and 4.

Appendix A

We derive the first two moments of multivariate path forecasts in the hhh4 modelling framework. For this purpose, it is irrelevant whether strata are defined by age groups, regions or both, so we only use one index $r = 1, \dots, R$ to denote strata. The weeks are indexed such that $t = 0$ corresponds to the week on which we condition our path forecasts. Because the following derivations do not rely on model-specific choices for offsets or weights, we generalize (1) as

$$Y_{rt} \sim \text{NBin}(\mu_{rt}, \psi_r)$$

$$\mu_{rt} = v_{rt} + \sum_{r'} \varphi_{r'rt} Y_{r',t-1},$$

where $\varphi_{r'rt}$ is a function of $\phi_{r'}$ and $w_{r'r}$.

We can derive several recursive relationships for the conditional moments of $\mathbf{Y} = (\mathbf{Y}_1^\top, \mathbf{Y}_2^\top, \dots)^\top$, where $\mathbf{Y}_t = (Y_{1t}, Y_{2t}, \dots, Y_{Rt})^\top$. For the means, we obtain

$$\mathbb{E}(Y_{rt} | \mathbf{Y}_0) = \mathbb{E}(\mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1}) | \mathbf{Y}_0) = v_{rt} + \sum_{r'} \varphi_{r'rt} \mathbb{E}(Y_{r',t-1} | \mathbf{Y}_0).$$

Similarly, we can derive the second moments

$$\begin{aligned} \mathbb{E}(Y_{rt}^2 | \mathbf{Y}_0) &= \mathbb{E}(\mathbb{E}(Y_{rt}^2 | \mathbf{Y}_{t-1}) | \mathbf{Y}_0) \\ &= \mathbb{E}(\text{Var}(Y_{rt} | \mathbf{Y}_{t-1}) + \mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1})^2 | \mathbf{Y}_0) \\ &= \mathbb{E}\left(\left(1 + \psi_r \cdot \mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1})\right) \cdot \mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1}) + \mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1})^2 \mid \mathbf{Y}_0\right) \\ &= \mathbb{E}\left(\mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1}) + (1 + \psi_r) \cdot \mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1})^2 \mid \mathbf{Y}_0\right) \\ &= \mathbb{E}\left(v_{rt} + \sum_{r'} \varphi_{r'rt} Y_{r',t-1} + (1 + \psi_r) \cdot \left(v_{rt} + \sum_{r'} \varphi_{r'rt} Y_{r',t-1}\right)^2 \mid \mathbf{Y}_0\right) \\ &= v_{rt} + \sum_{r'} \varphi_{r'rt} \mathbb{E}(Y_{r',t-1} | \mathbf{Y}_0) \\ &\quad + (1 + \psi_r) \cdot \left(v_{rt}^2 + 2 v_{rt} \cdot \sum_{r'} \varphi_{r'rt} \mathbb{E}(Y_{r',t-1} | \mathbf{Y}_0) + \sum_{r', \tilde{r}} \varphi_{r'rt} \varphi_{\tilde{r}rt} \mathbb{E}(Y_{r',t-1} Y_{\tilde{r},t-1} | \mathbf{Y}_0)\right). \end{aligned}$$

A somewhat simpler result can be obtained for the product terms ($r \neq s$)

$$\begin{aligned} \mathbb{E}(Y_{rt} Y_{st} | \mathbf{Y}_0) &= \mathbb{E}(\mathbb{E}(Y_{rt} Y_{st} | \mathbf{Y}_{t-1}) | \mathbf{Y}_0) \\ &= \mathbb{E}(\mathbb{E}(Y_{rt} | \mathbf{Y}_{t-1}) \mathbb{E}(Y_{st} | \mathbf{Y}_{t-1}) | \mathbf{Y}_0) \\ &= \mathbb{E}\left(\left(v_{rt} + \sum_{r'} \varphi_{r'rt} Y_{r',t-1}\right) \cdot \left(v_{st} + \sum_{s'} \varphi_{s'st} Y_{s',t-1}\right) \mid \mathbf{Y}_0\right) \\ &= v_{rt} v_{st} + v_{rt} \cdot \sum_{s'} \varphi_{s'st} \mathbb{E}(Y_{s',t-1} | \mathbf{Y}_0) + v_{st} \cdot \sum_{r'} \varphi_{r'rt} \mathbb{E}(Y_{r',t-1} | \mathbf{Y}_0) \\ &\quad + \sum_{r', s'} \varphi_{r'rt} \varphi_{s'st} \mathbb{E}(Y_{r',t-1} Y_{s',t-1} | \mathbf{Y}_0). \end{aligned}$$

These componentwise formulations are somewhat involved, but a nice matrix formulation of the recursion exists. Define $\tilde{\mathbf{Y}}_t = (1, Y_{1t}, \dots, Y_{Rt})^\top$, $\tilde{\boldsymbol{\psi}} = (1, \psi_1, \dots, \psi_R)^\top$ and

$$\tilde{\boldsymbol{\varphi}}_t = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ \nu_{1t} & \varphi_{11t} & \varphi_{21t} & \cdots & \varphi_{R1t} \\ \nu_{2t} & \varphi_{12t} & \varphi_{22t} & \cdots & \varphi_{R2t} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \nu_{Rt} & \varphi_{1Rt} & \varphi_{2Rt} & \cdots & \varphi_{RRt} \end{pmatrix}$$

such that

$$\mathbb{E}(\tilde{\mathbf{Y}}_t | \mathbf{Y}_0) = \tilde{\boldsymbol{\varphi}}_t \mathbb{E}(\tilde{\mathbf{Y}}_{t-1} | \mathbf{Y}_0) = \tilde{\boldsymbol{\varphi}}_t \tilde{\boldsymbol{\varphi}}_{t-1} \cdots \tilde{\boldsymbol{\varphi}}_1 \tilde{\mathbf{Y}}_0.$$

With

$$\mathbf{M}_{tt} = \mathbb{E}(\tilde{\mathbf{Y}}_t \tilde{\mathbf{Y}}_t^\top | \mathbf{Y}_0) = \begin{pmatrix} 1 & \mathbb{E}(Y_{1t} | \mathbf{Y}_0) & \mathbb{E}(Y_{2t} | \mathbf{Y}_0) & \cdots \\ \mathbb{E}(Y_{1t} | \mathbf{Y}_0) & \mathbb{E}(Y_{1t}^2 | \mathbf{Y}_0) & \mathbb{E}(Y_{1t} Y_{2t} | \mathbf{Y}_0) & \cdots \\ \mathbb{E}(Y_{2t} | \mathbf{Y}_0) & \mathbb{E}(Y_{2t} Y_{1t} | \mathbf{Y}_0) & \mathbb{E}(Y_{2t}^2 | \mathbf{Y}_0) & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix},$$

the aforementioned moment equations then allow for a simple recursive procedure consisting of one matrix transformation and one modification of the diagonal:

$$\begin{aligned} \mathbf{M}_{00} &= \tilde{\mathbf{Y}}_0 \tilde{\mathbf{Y}}_0^\top, \\ \tilde{\mathbf{M}}_{tt} &= \tilde{\boldsymbol{\varphi}}_t \mathbf{M}_{t-1,t-1} \tilde{\boldsymbol{\varphi}}_t^\top, \text{ and} \\ (\mathbf{M}_{tt})_{rs} &= \begin{cases} 1 & \text{if } r = s = 1, \\ (1 + \tilde{\psi}_r) \cdot (\tilde{\mathbf{M}}_{tt})_{rr} + (\tilde{\mathbf{M}}_{tt})_{1r} & \text{if } r = s > 1, \\ (\tilde{\mathbf{M}}_{tt})_{rs} & \text{otherwise.} \end{cases} \end{aligned}$$

This is still limited to pairs of counts from the same week t . The following allows calculations across weeks:

$$\begin{aligned} \mathbf{M}_{t,t+k} &= \mathbb{E}(\tilde{\mathbf{Y}}_t \tilde{\mathbf{Y}}_{t+k}^\top | \mathbf{Y}_0) = \mathbb{E}(\mathbb{E}(\tilde{\mathbf{Y}}_t \tilde{\mathbf{Y}}_{t+k}^\top | \mathbf{Y}_t) | \mathbf{Y}_0) \\ &= \mathbb{E}(\tilde{\mathbf{Y}}_t \mathbb{E}(\tilde{\mathbf{Y}}_{t+k} | \mathbf{Y}_t)^\top | \mathbf{Y}_0) \\ &= \mathbb{E}(\tilde{\mathbf{Y}}_t \mathbb{E}(\tilde{\mathbf{Y}}_{t+k-1} | \mathbf{Y}_t)^\top \tilde{\boldsymbol{\varphi}}_{t+k}^\top | \mathbf{Y}_0) \\ &= \mathbb{E}(\tilde{\mathbf{Y}}_t \mathbb{E}(\tilde{\mathbf{Y}}_{t+k-1} | \mathbf{Y}_t)^\top | \mathbf{Y}_0) \tilde{\boldsymbol{\varphi}}_{t+k}^\top \\ &= \mathbb{E}(\tilde{\mathbf{Y}}_t \tilde{\mathbf{Y}}_{t+k-1}^\top | \mathbf{Y}_0) \tilde{\boldsymbol{\varphi}}_{t+k}^\top \\ &= \mathbf{M}_{t,t+k-1} \tilde{\boldsymbol{\varphi}}_{t+k}^\top \\ \Leftrightarrow \mathbf{M}_{t+k,t} &= \mathbb{E}(\tilde{\mathbf{Y}}_{t+k} \tilde{\mathbf{Y}}_t^\top | \mathbf{Y}_0) = \tilde{\boldsymbol{\varphi}}_{t+k} \mathbf{M}_{t+k-1,t}. \end{aligned}$$

To compute $\mathbb{E}(\mathbf{Y}\mathbf{Y}^\top)$ from the $\mathbf{M}_{tu}$'s, we only need to omit the first row and column of each $\mathbf{M}_{tu}$. The entire recursive construction of this matrix can be visualized as follows:

$$\mathbb{E}(\mathbf{Y}\mathbf{Y}^\top) = \begin{pmatrix} \mathbb{E}(\mathbf{Y}_1 \mathbf{Y}_1^\top) & \rightarrow & \mathbb{E}(\mathbf{Y}_1 \mathbf{Y}_2^\top) & \rightarrow & \mathbb{E}(\mathbf{Y}_1 \mathbf{Y}_3^\top) & \rightarrow & \cdots \\ \downarrow & \searrow & & & & & \\ \mathbb{E}(\mathbf{Y}_2 \mathbf{Y}_1^\top) & & \mathbb{E}(\mathbf{Y}_2 \mathbf{Y}_2^\top) & \rightarrow & \mathbb{E}(\mathbf{Y}_2 \mathbf{Y}_3^\top) & \rightarrow & \cdots \\ \downarrow & & \downarrow & \searrow & & & \\ \mathbb{E}(\mathbf{Y}_3 \mathbf{Y}_1^\top) & & \mathbb{E}(\mathbf{Y}_3 \mathbf{Y}_2^\top) & & \mathbb{E}(\mathbf{Y}_3 \mathbf{Y}_3^\top) & \rightarrow & \cdots \\ \downarrow & & \downarrow & & \downarrow & \searrow & \\ \vdots & & \vdots & & \vdots & & \ddots \end{pmatrix}.$$

The covariance matrix is then calculated as $\text{Cov}(\mathbf{Y}) = \mathbb{E}(\mathbf{Y}\mathbf{Y}^\top) - \mathbb{E}(\mathbf{Y})\mathbb{E}(\mathbf{Y})^\top$.

To aggregate across strata (e.g. to break our results down to stratification only by region, only by age group, etc.), we can use $\text{Cov}(\mathbf{B}\mathbf{Y}) = \mathbf{B} \text{Cov}(\mathbf{Y}) \mathbf{B}^\top$ where the matrix $\mathbf{B}$ is suitably defined.

Acknowledgements

This paper summarizes and extends the 13th Armitage lecture, presented by the first author at the Medical Research Council Biostatistics Unit in Cambridge, UK, on 12 November 2015. A videotaped version can be found at <http://www.statisticsviews.com/details/webinar/8526701/Armitage-Lecture-2015.html>. The support and hospitality of the Biostatistics Unit are gratefully acknowledged. We also acknowledge helpful comments by three anonymous reviewers and by Phil Dawid, in particular, for making us aware of his work on calibration tests [22, 23]. The authors' research was supported by the Swiss National Science Foundation (project no. 137919).

References

1. Keeling MJ, Rohani P. *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press: Princeton and Oxford, 2008.
2. World Health Organization. Anticipating epidemics. *Weekly Epidemiological Record* 2014; **89**(22):244. <http://www.who.int/wer> [Accessed on 2 June 2017].
3. World Health Organization (ed.). *Anticipating Emerging Infectious Disease Epidemics: Meeting Report of WHO Informal Consultation*. WHO Press: Geneva, Switzerland, 2016. http://www.who.int/csr/disease/anticipating_epidemics/meeting-report-2015/en/ [Accessed on 2 June 2017].
4. Centers for Disease Control and Prevention. Flu activity forecasting website launched, 2016. <https://www.cdc.gov/flu/news/flu-forecast-website-launched.htm> [Accessed on 2 June 2017].
5. Hufnagel L, Brockmann D, Geisel T. Forecast and control of epidemics in a globalized world. *Proceedings of the National Academy of Sciences of the United States of America* 2004; **101**:15 124–15 129.
6. Birrell PJ, Ketsetzis G, Gay NJ, Cooper BS, Presanis AM, Harris RJ, Charlett A, Zhang XS, White PJ, Pebody RG, De Angelis D. Bayesian modeling to unmask and predict influenza A/H1N1pdm dynamics in London. *Proceedings of the National Academy of Sciences of the United States of America* 2011; **108**(45):18 238–18 243.
7. Chretien JP, George D, Shaman J, Chitale RA, McKenzie FE. Influenza forecasting in human populations: a scoping review. *PLOS ONE* 2014; **9**(4):e94 130. <https://doi.org/10.1371/journal.pone.0094130>.
8. Nsoesie EO, Marathe M, Brownstein JS. Forecasting peaks of seasonal influenza epidemics. In *PLOS Currents Outbreaks*, 1st ed., 2013, <https://doi.org/10.1371/currents.outbreaks.bb1e879a23137022ea79a8c508b030bc>.
9. Santillana M, Nguyen AT, Dredze M, Paul MJ, Nsoesie EO, Brownstein JS. Combining search, social media, and traditional data sources to improve influenza surveillance. *PLOS Computational Biology* 2015; **11**(10):e1004 513. <https://doi.org/10.1371/journal.pcbi.1004513>.
10. Dukic V, Lopes HF, Polson NG. Tracking epidemics with Google Flu Trends data and a state-space SEIR model. *Journal of the American Statistical Association* 2012; **107**(500):1410–1426.
11. Yang S, Santillana M, Kou SC. Accurate estimation of influenza epidemics using Google search data via ARGO. *Proceedings of the National Academy of Sciences of the United States of America* 2015; **112**(47):14 473–14 478.
12. Xia Y, Bjørnstad ON, Grenfell BT. Measles metapopulation dynamics: a gravity model for epidemiological coupling and dynamics. *The American Naturalist* 2004; **164**(2):267–281.
13. Meyer S, Held L. Power-law models for infectious disease spread. *The Annals of Applied Statistics* 2014; **8**(3):1612–1639.
14. Riley S, Eames K, Isham V, Mollison D, Trapman P. Five challenges for spatial epidemic models. *Epidemics* 2015; **10**: 68–71.
15. Höhle M. Infectious disease modelling. In *Handbook of Spatial Epidemiology*, Lawson AB, Banerjee S, Haining RP, Ugarte MD (eds), chap. 26, Chapman & Hall/CRC Handbooks of Modern Statistical Methods. Chapman and Hall/CRC: Boca Raton, 2016; 477–500.
16. Baguelin M, Flasche S, Camacho A, Demiris N, Miller E, Edmunds WJ. Assessing optimal target populations for influenza vaccination programmes: an evidence synthesis and modelling study. *PLOS Medicine* 2013; **10**(10):e1001 527. <https://doi.org/10.1371/journal.pmed.1001527>.
17. Meyer S, Held L. Incorporating social contact data in spatio-temporal models for infectious disease spread. *Biostatistics* 2017; **18**(2):338–351.
18. Gneiting T, Katzfuss M. Probabilistic forecasting. *Annual Review of Statistics and Its Application* 2014; **1**(1):125–151.
19. Moran KR, Fairchild G, Generous N, Hickmann K, Osthus D, Priedhorsky R, Hyman J, Del Valle SY. Epidemic forecasting is messier than weather forecasting: the role of human behavior and internet data streams in epidemic forecast. *Journal of Infectious Diseases* 2016; **214**(suppl 4):S404–S408.
20. Gneiting T, Raftery AE. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 2007; **102**(477):359–378.
21. Czado C, Gneiting T, Held L. Predictive model assessment for count data. *Biometrics* 2009; **65**(4):1254–1261.
22. Seillier-Moiseiwitsch F, Sweeting TJ, Dawid AP. Prequential tests of model fit. *Scandinavian Journal of Statistics* 1992; **19**(1):45–60.
23. Seillier-Moiseiwitsch F, Dawid AP. On testing the validity of sequential probability forecasts. *Journal of the American Statistical Association* 1993; **88**(421):355–359.
24. Wei W, Held L. Calibration tests for count data. *Test* 2014; **23**(4):787–805.
25. Gneiting T, Stanberry LI, Gneiting EP, Held L, Johnson NA. Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *Test* 2008; **17**(2):211–235.
26. Scheuerer M, Hamill TM. Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Monthly Weather Review* 2015; **143**(4):1321–1334.

27. Wei W, Balabdaoui F, Held L. Calibration tests for multivariate Gaussian forecasts. *Journal of Multivariate Analysis* 2017; **154**:216–233.
28. Held L, Höhle M, Hofmann M. A statistical framework for the analysis of multivariate infectious disease surveillance counts. *Statistical Modelling* 2005; **5**(3):187–199.
29. Meyer S, Held L, Höhle M. Spatio-temporal analysis of epidemic phenomena using the R package surveillance. *Journal of Statistical Software* 2017; **77**(11):1–55.
30. Pringle K, Lopman B, Vega E, Vinje J, Parashar UD, Hall AJ. Noroviruses: epidemiology, immunity and prospects for prevention. *Future Microbiology* 2015; **10**(1):53–67.
31. Ahmed SM, Lopman BA, Levy K. A systematic review and meta-analysis of the global seasonality of norovirus. *PLOS ONE* 2013; **8**(10):e75922. <https://doi.org/10.1371/journal.pone.0075922>.
32. Bernard H, Höhle M, Niendorf S, Altmann D, Stark K. Epidemiology of norovirus gastroenteritis in Germany 2001–2009: eight seasons of routine surveillance. *Epidemiology & Infection* 2014; **142**(1):63–74.
33. Paul M, Held L, Toschke A. Multivariate modelling of infectious disease surveillance data. *Statistics in Medicine* 2008; **27**(29):6250–6267.
34. Held L, Paul M. Modeling seasonality in space-time infectious disease surveillance data. *Biometrical Journal* 2012; **54**(6):824–843.
35. Brockmann D, Hufnagel L, Geisel T. The scaling laws of human travel. *Nature* 2006; **439**(7075):462–465.
36. Cliff AD, Ord JK. Model building and the analysis of spatial pattern in human geography. *Journal of the Royal Statistical Society Series B (Methodological)* 1975; **37**(3):297–348.
37. Mossong J, Hens N, Jit M, Beutels P, Auranen K, Mikolajczyk R, Massari M, Salmaso S, Tomba GS, Wallinga J, Heijne J, Sadkowska-Todys M, Rosinska M, Edmunds WJ. Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLOS Medicine* 2008; **5**(3):e74. <https://doi.org/10.1371/journal.pmed.0050074>.
38. Wallinga J, Teunis P, Kretzschmar M. Using data on social contacts to estimate age-specific transmission parameters for respiratory-spread infectious agents. *American Journal of Epidemiology* 2006; **164**(10):936–944.
39. Bland JM, Altman DG. Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet* 1986; **327**(8476):307–310.
40. Good I. J. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* 1952; **14**(1):107–114.
41. Dawid A P, Sebastiani P. Coherent dispersion criteria for optimal experimental design. *Annals of Statistics* 1999; **27**(1): 65–81.
42. Riebler A, Held L. Projecting the future burden of cancer: Bayesian age-period-cohort analysis with integrated nested Laplace approximations. *Biometrical Journal* 2017; **59**(3):531–549.
43. Diebold FX, Gunther TA, Tay AS. Evaluating density forecasts with applications to financial risk management. *International Economic Review* 1998; **39**(4):863–883.
44. Diebold FX, Mariano RS. Comparing predictive accuracy. *Journal of Business & Economic Statistics* 1995; **13**(3):253–263.
45. Krüger F, Lerch S, Thorarinsdottir TL, Gneiting T. Probabilistic forecasting and comparative model assessment based on Markov chain Monte Carlo output 2017. <https://arxiv.org/abs/1608.06802>, arXiv:1608.06802 [stat.ME] [Accessed on 2 June 2017].
46. Hemri S, Lisniak D, Klein B. Multivariate postprocessing techniques for probabilistic hydrological forecasting. *Water Resources Research* 2015; **51**(9):7436–7451.
47. Claeskens G, Hjort NL. *Model Selection and Model Averaging*. Cambridge University Press: Cambridge, 2008.
48. Slughter M, Gneiting T, Raftery AE. Probabilistic wind speed forecasting using ensembles and Bayesian model averaging. *Journal of the American Statistical Association* 2010; **105**(489):25–35.