

Bayesian modeling to unmask and predict influenza A/H1N1pdm dynamics in London

Paul J. Birrell^{a,1}, Georgios Ketssetzis^b, Nigel J. Gay^c, Ben S. Cooper^d, Anne M. Presanis^a, Ross J. Harris^b, André Charlett^b, Xu-Sheng Zhang^b, Peter J. White^{b,e}, Richard G. Pebody^b, and Daniela De Angelis^{a,b}

^aMedical Research Council Biostatistics Unit, University Forvie Site, Robinson Way, Cambridge CB2 0SR, United Kingdom; ^bHealth Protection Services, Health Protection Agency, 61 Colindale Avenue, London NW9 5HT, United Kingdom; ^cFu Consulting, Hungerford, United Kingdom; ^dMahidol-Oxford Tropical Medicine Research Unit (MORU), Faculty of Tropical Medicine, Mahidol University, 420/6 Rajvithi Road, Bangkok 10400, Thailand; and ^eMedical Research Council Centre for Outbreak Analysis and Modelling, Department for Infectious Disease Epidemiology, Imperial College Faculty of Medicine, Norfolk Place, London W2 1PG, United Kingdom

Edited by Aravinda Chakravarti, Johns Hopkins University School of Medicine, Baltimore, MD, and accepted by the Editorial Board September 26, 2011 (received for review February 28, 2011)

The tracking and projection of emerging epidemics is hindered by the disconnect between apparent epidemic dynamics, discernible from noisy and incomplete surveillance data, and the underlying, imperfectly observed, system. Behavior changes compound this, altering both true dynamics and reporting patterns, particularly for diseases with nonspecific symptoms, such as influenza. We disentangle these effects to unravel the hidden dynamics of the 2009 influenza A/H1N1pdm pandemic in London, where surveillance suggests an unusual dominant peak in the summer. We embed an age-structured model into a Bayesian synthesis of multiple evidence sources to reveal substantial changes in contact patterns and health-seeking behavior throughout the epidemic, uncovering two similar infection waves, despite large differences in the reported levels of disease. We show how this approach, which allows for real-time learning about model parameters as the epidemic progresses, is also able to provide a sequence of nested projections that are capable of accurately reflecting the epidemic evolution.

Bayesian statistics | real-time modeling | general practice consultation data | infectious disease | seroepidemiology

An emerging epidemic engenders an increased demand upon health services. Resolving the extent to which this is due to high levels of disease transmission as opposed to a heightened public sensitivity is essential for determining the appropriate public health response.

This was especially crucial when estimating the course of the 2009 influenza A/H1N1pdm outbreak in England, where, unusually, the pandemic resulted in a summer peak in rates of consultation at general practices (GPs) for influenza-like illness (ILI). This is clearly demonstrated by data from the return service of the Royal College of General Practitioners (RCGP) in Fig. 1*A* where weekly GP consultation rates per 100,000 population over the 2009 pandemic are compared with rates from the three previous years. Also shown is the proportion of swabbed individuals whose swabs tested positive for the presence of any flu virus (*SI Data*). Note that the GP consultation rate for 2009 is much higher than the usual seasonal rate, whereas the corresponding positivity is comparable to that observed in the preceding winters. This suggests that a substantial proportion of the peak in consultations was not directly attributable to A/H1N1pdm. Conversely, serological studies (3) have shown a marked increase in the prevalence of influenza antibodies among the population. Therefore, the degree to which the increased demand upon GPs is due to high levels of disease transmission as opposed to heightened public sensitivity remains unclear (4). Fig. 1*B* and *C* show GP consultation rates by region and age group: consultations in Greater London and the West Midlands exhibit rapid early exponential growth, but the peak in London is much higher; rates appear to decrease markedly with age. Importantly, a first peak occurs immediately prior to the summer school holiday and the launch of the National Pandemic Flu Service (NPFSS) phone line, ntro-

duced to relieve the pressure on GPs and expedite antiviral distribution (see *SI Data*); a second, much smaller peak, is observed in the autumn. This evidence, supported by the work of ref. 5, promotes the further hypothesis of a fluctuating propensity for individuals with symptoms of ILI to seek medical attention, perhaps induced by media coverage and changing governmental advice, as well as the social distancing effects of school holidays.

Traditionally, transmission modeling is used to investigate epidemic development. In the area of infectious respiratory diseases, many approaches have been proposed (6–13), including those that account for the effects of behavioral changes upon transmission (14), explicitly model the impact of school closure (15), and incorporate a temporally varying case-detection rate (16).

Here we model these aspects simultaneously while additionally accounting for the time-varying noise in the data due to consultations for non-A/H1N1pdm ILI. This is achieved by developing a model for integrating noisy GP consultation data, virological positivity data, virologically confirmed case data, and information from serological (seroprevalence) surveys (see *SI Data*). Each dataset is available and used at daily intervals. An age-structured transmission model is embedded within a Bayesian framework, allowing incorporation of any a priori information about model parameters from previous influenza strains via probability distributions. These prior distributions are then updated by available data to provide posterior statements about parameters of interest and their uncertainty, presented here in the form of 95% credible intervals (CrIs).

Results

Fig. 2 is a schematic representation of the model used to describe the data-generating process. Three different components are knitted together: an age-structured transmission-governing component, a disease component, and a third component describing the mechanisms through which infected individuals report their symptoms to the health-care system. In the transmission component, susceptible individuals (*S*) become exposed (*E*) through an effective contact with infectious (*I*) individuals and become infective themselves after a short latent period, to be then removed (*R*) from the pool of infectious individuals after a further period. Transmission is governed both by a time- and age-varying force of infection $\lambda(t,a)$, depending on the transmissibility of

Author contributions: N.J.G., B.S.C., and D.D.A. designed research; P.J.B., G.K., and X.-S.Z. performed research; P.J.B., R.J.H., A.C., P.J.W., and R.G.P. analyzed data; and P.J.B., A.M.P., and D.D.A. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. A.C. is a guest editor invited by the Editorial Board. Freely available online through the PNAS open access option.

¹To whom correspondence may be addressed. E-mail: paul.birrell@mrc-bsu.cam.ac.uk or daniela.deangelis@mrc-bsu.cam.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1103002108/-DCSupplemental.

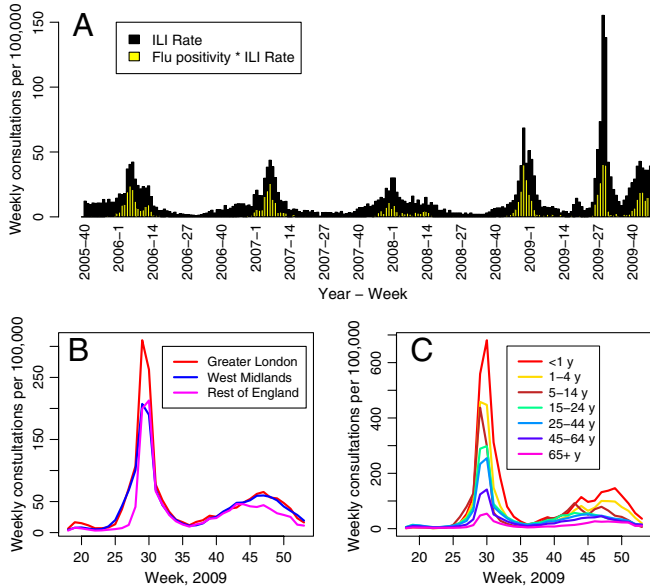


Fig. 1. (A) Time series of GP consultation rate for ILI and virological positivity within the RCGP surveillance scheme in England and Wales (1), by week, from end-2005 to end-2009. The positivity is given by the proportion of each bar shaded yellow. (B) Observed weekly consultation rates per 100,000 within each group (see ref. 2 and *SI Data*, section 1.1) by region (Greater London, West Midlands, and Rest of England) and (C) by age group, over the weeks 18–53, 2009.

the virus and the mixing patterns in the population and by the transition rates among the S , E , I , and R states (see *Material and Methods*). Only a proportion, θ , of the newly exposed individuals develop febrile symptoms, from which further proportions $p_{GP}(t,a)$ and p_{CC} consult their GP or have their illness virologically confirmed. Note that $p_{GP}(t,a)$ is calendar time and age-specific to accommodate potential fluctuations in consultation behavior. There are no direct data on transmission. However, serological surveys (see *SI Data*), carried out before and during the epidemic, provide data on indicators, $Z(t,a)$ (see *Material and Methods*), informing the level of susceptibility within the population at the epidemic onset and over time. We make direct observation of $X_{CC}(t,a)$, the symptomatic cases that are virologically confirmed, though this is limited to the early stages of the epidemic (see *SI Data*). Only *indirect* information is available on the number of symptomatic cases consulting GPs, $X_{GP}(t,a)$, in the form of routine surveillance (see *SI Data*, section 1.1) counts of GP consultations for *all* ILI. This includes a background

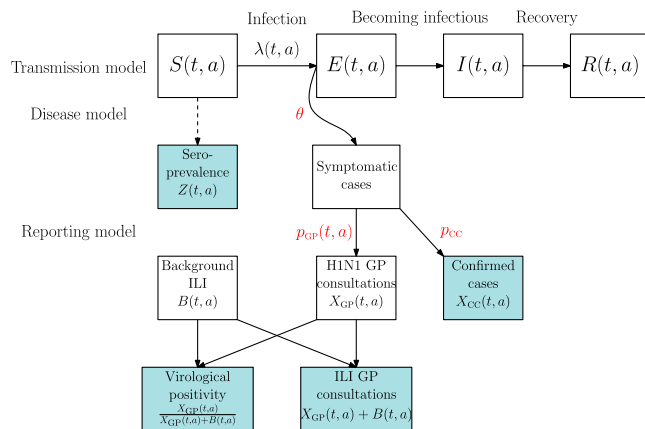


Fig. 2. Model schematic diagram representing the data-generating process. The shaded boxes represent the quantities upon which we make observation.

component, $B(t,a)$, of non-A/H1N1pdm ILI. To identify these two components we use the total number of ILI consultations $Y(t,a) = X_{GP}(t,a) + B(t,a)$ and information on the virological positivity $X_{GP}(t,a)/(X_{GP}(t,a) + B(t,a))$ (see *Material and Methods*). By combining direct, indirect, and prior information, we produce posterior distributions for the process governing parameters (see *Material and Methods*) and other quantities of interest.

Reconstructing the Epidemic. Fig. 3A shows the posterior median and pointwise 95% CrI for the total number of weekly incident infections of A/H1N1pdm in Greater London, using 245 d of epidemic data covering May 1 to December 31 (i.e., from week 18 to week 53) of 2009. Additionally, Fig. 3A also shows the estimated age-specific incidences. Much like the GP consultation data, the epidemic occurs in two waves: a summer first wave (May to end-August) and an autumn second wave (September to December). The first wave rises sharply to a peak of 109,000 (81,000–146,000) new infections in the week immediately prior to the school holidays. The second wave has a smaller peak with posterior probability 0.885. Conversely, as can be seen from Table 1, which reports estimates of the infection attack rate (i.e., the cumulative incidence expressed as a proportion of the total population), there is slightly larger cumulative incidence in the second wave, a phenomenon not at all evident from the GP consultation data (Fig. 1B and C). The discrepancy between the GP consultation data and the estimated infection pattern is clarified in Fig. 3B, which compares the cumulative consultations with the estimated cumulative infections, both calculated as proportions of their corresponding total. In this plot, a steep gradient identifies points in time during the pandemic when a relatively high density of GP consultations (or infections) occur. From the steep gradient in the GP consultations curve over weeks 27 and 28, it can be seen that the consultations were highly localized around this time. The separation between the two lines and the smaller

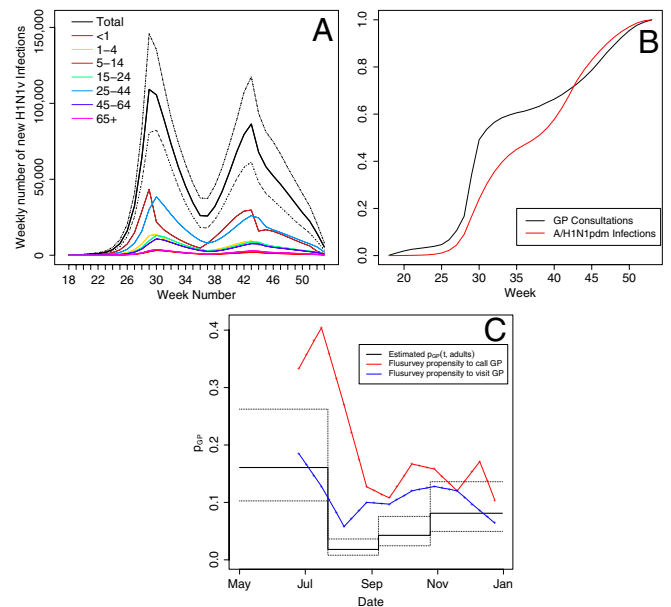


Fig. 3. (A) Estimated weekly infections for Greater London, spanning weeks 18 to 52 of 2009, as reconstructed from the model. The black line is the total incidence of infections over all ages and the dotted lines represent a 95% CrI. (B) The cumulative incidence of GP consultations (black line) versus the estimated (posterior median) cumulative incidence of infection (red line), both expressed as a fraction of the respective total incidence. (C) The estimated propensity for incident adult symptomatic cases to consult with a GP over time (black line). The red and blue lines present the propensity for ILI cases to call and visit (respectively) their GP from an Internet-based survey, FluSurvey (*SI Materials and Methods*).

breakpoint, in late October, the propensity returns to a value close to 10%. Similar results are obtained for this parameter in children. These estimates are similar to values expected during seasonal influenza epidemics (17, 18), but lower than estimates from the Internet-based Flusurvey (see Fig. 3C and *SI Data*), possibly reflecting biases in the population captured by the survey.

Predicting the Epidemic. The above results are related to an epidemic that is now over. A crucial question is whether the model can be used as a tool for inferences and predictions while an epidemic is ongoing. To assess this, further analyses were conducted based on 83, 143, and 192 d of epidemic surveillance data. The 83-d analysis contains no serological data except those used to inform the baseline prevalence of antibodies, whereas the 143- and 192-d analyses incorporate serological data collected during the epidemic (see ref. 3 and *SI Data*).

Fig. 4A illustrates how the predictions evolve as data accumulate. Fig. 4B shows how the estimated posterior densities for the parameters R_0 , m_1 , θ , and $p_{GP}(1,1)$, the consultation propensity in children in the first 83 d (see *SI Materials and Methods*), evolve over time, starting with their prior distributions. From the 143-d analysis onward, credible intervals for the future number of infections appear to enclose the estimated numbers in the subsequent analysis. However, this is not so in moving from the 83- to the 143-d analysis. This is due to the lack of serological data in the 83-d analysis. Given the large degree of dispersion, the GP consultation data are too weakly informative to overcome the informative priors placed, partly for the sake of identifiability, upon parameters such as θ and $p_{GP}(\cdot, \cdot)$ (see *SI Materials and Methods*). The densities of Fig. 4B show that, in the earliest analysis with no serological data, the posterior distributions for these parameters are near identical to the priors, centered on values far larger than the posteriors obtained from the subsequent analyses. The inclusion of the serological data in the 143-d analysis provides a clear indication of the level of cumulative incidence, which is higher than the 83-d results might suggest. With stronger information on the incidence, the data become sufficiently informative to overcome the prior distributions for θ and p_{GP} , and this is reflected in a shifting of the posterior distributions to values that remain consistent across the 143-, 192-, and 245-d analyses.

Discussion

Our approach allows reconstruction and projection of the trajectory of an epidemic by disentangling epidemic and behavioral dynamics. Combining data from different sources is crucial, as each plays an important role: virological data partition consultations between A/H1N1pdm and other ILL, GP consultations determine the temporal trend, and the serological data give the scale of the epidemic.

Our estimates of attack rates are lower than has been obtained elsewhere (19), importantly providing improved understanding of how a third wave of A/H1N1pdm infection occurred in late 2010. Although our estimate of R_0 is consistent with that obtained by others (13, 16), the estimate that only a third of infections are symptomatic is much lower than corresponding estimates from Mexico and Hong Kong [0.86 and 0.64, (13, 20)] but comparable to estimates from New Zealand [0.45 (21)] and France [0.20 (22)] and in broad agreement with the systematic review of ref. 23. Our estimate of 0.47 (= $1/2.13$) for the relative risk of transmission in infectious contacts involving at least one adult is in direct agreement with a previous estimate of 0.485 (95% CrI 0.302–0.625) (13). By combining our estimated θ and attack rates, we obtain a number of symptomatic cases, which is a fourfold increase on the central and a twofold increase on the upper bound of the official estimates for the two waves (24). Previous work (10) uses these central estimates as data, multiplying them by a factor of 10 in order to achieve a good model fit. This factor can be interpreted as a product of two components: one that accounts for

the asymptomatic infections ($1/\theta$) and one that accounts for underascertainment in the symptomatic case number estimates. Here, these two components multiply to give a factor of approximately 12.

In its transmission component, the model is similar to that used elsewhere (16), where data on laboratory-confirmed cases, modeled using a time-varying reporting rate, have been considered. Here, in addition to some lab-confirmed cases, other sources of information are used, notably noisy GP consultation data. This enables us to advance earlier efforts to model the epidemic in England. Using a hybrid of estimation approaches, ref. 10 treats estimated weekly incident cases of A/H1N1pdm (24) as data, with no propagation of the error inherent in the estimation process. We employ a more rigorous, statistical approach, which utilizes a richer array of raw data. For the same pandemic in Singapore, ref. 25 implements an algorithm for online updating of estimates arising from an S, E, I, R transmission model fitted to GP consultation data alone. This approach, which makes no stratification by age, suffers similarly from a lack of nesting in the early stages and masks the interage group transmission dynamics. In the same vein, ref. 18 develops a methodology for real-time inference, but this offers little opportunity to make any learning about many key model parameters. Our model also allows for the quantification of the impact of school holidays, the level of non-A/H1N1pdm consultation, and obtains estimates for the propensity for patients affected by the A/H1N1pdm virus to seek consultation.

Assumptions. One key strength of any model used to predict or assess epidemic impact is a robustness to (often unavoidable) modeling assumptions. In *SI Sensitivity Analyses* we investigate the impact of dropping, or changing, a number of the assumptions made in producing the epidemic estimates obtained here. Our results are generally robust to reasonable deviations from modeling assumptions, yet also suggest avenues for further investigation. Specifically, we examine three modeling components: the performance of the virological testing procedure, the assumed contact patterns, and the assumed functional form of the propensity to consult, $p_{GP}(t,a)$.

Test Sensitivity. Thus far, it has been assumed that the virological testing procedure has a sensitivity of 1; i.e., there are no false negatives. If we relax this assumption to reasonable values for the test sensitivity, say 0.8 and 0.9, the key results change very little. As expected, there is a small increase in the estimated symptomatic attack rates, as a lower sensitivity allows for more A/H1N1pdm cases among the GP consultations, but there is negligible impact on the total infection attack rates. Of the values investigated, a test sensitivity of 0.8 had larger values for the likelihood in its posterior distribution.

Mixing Matrices. The mixing matrices used to describe rates of contact between the different age groups are based upon United Kingdom (UK) data from the POLYMOD (Improving Public Health Policy in Europe through Modeling and Economic Evaluation of Interventions for the Control of Infectious Diseases) study (see ref. 26). Results are robust to small changes in both parameterization of the mixing matrices (see *Materials and Methods*) and the rates of contact themselves. More interestingly, a preliminary attempt to estimate the entire mixing matrices (*SI Sensitivity Analyses*), using the POLYMOD data only as prior information, indicates that contacts among the 5–14 age group are particularly important, while also suggesting that contacts between this age group and the 15–24 age group may be more influential than previously thought. Adopting these estimates results in a slight fall in the estimated R_0 , while increasing the total infection attack rate. This is due to a shift in the infection profile from small children to the bigger pool of susceptibles in the 15–24 age group. Developing a more rigorous approach to

estimation of contact patterns, and in particular the choices of the informative priors, constitutes a promising avenue for future work.

Propensity to Consult. A piecewise linear parameterization for $p_{GP}(t,a)$ in the post-NPFS era gave results not materially different from those featured here. Attempts to adopt this piecewise linear parameterization over the entire epidemic period resulted in lack of identifiability and undue influence of prior distributions (see *SI Sensitivity Analyses*).

Further Applicability. Our modeling approach has focused on the reconstruction of the epidemic in a globally prominent metropolitan region. However, the general methodology presented here is highly applicable to any influenza epidemic within any similar health-care network. Although it is unreasonable to carry out this modeling exercise on England as a whole, due to the unrepresentativeness of the pooled virological and serological data, we have repeated the analyses in three other disjoint regions, which, together with London cover the whole of England (see *SI Further Applicability*). London and West Midlands experience very similar epidemics, whereas the remainder of the country is split into two regions, North and South, neither of which had a substantial first wave of infection. There is mostly nonsignificant variation in the estimated parameters across the four regions, with the exception of R_0 , which suggests a possible link between the reproductive number of the epidemic and the population density within the affected region (London and West Midlands are England's two most densely populated areas).

Serological Data. As epidemic surveillance data accumulate over time, the model is capable of producing sequential epidemic estimates that converge (in the sense that successive credible intervals are nested) and could be used for real-time modeling and prediction. However, we have shown surveillance data alone are insufficient unless or until the epidemic is very far progressed. In general, to generate reliable projections early in the epidemic, the timely availability of relevant data on cumulative incidence and/or the proportion of infections reported in surveillance (a role originally envisaged for the Flusurvey during England's 2009 A/H1N1pdm outbreak; see *SI Data*) is required. Serological data, in particular, are shown in this paper to be vital to ensure convergence of sequentially obtained estimates, given our choice of priors. Analysis conducted in the absence of serological data (see *SI Serological Studies*) shows that with surveillance data alone, a realistic epidemic reconstruction is still impossible at 192 d, i.e., after the peak of the second wave. As a result, any online inference is rendered highly infeasible. This highlights the critical importance of the timely availability of serological information in an emergent epidemic, when information on key parameters may be lacking and/or priors may be misspecified, as here. This is clearly a challenge, given current limitations on test developments, facilities, and recruitment of appropriately representative populations (3, 19), but one that is very important to meet.

Materials and Methods

Our approach integrates data from a number of sources, combining information from GP surveillance networks with epidemic specific data. The *SI Data* section provides an in-depth description of the available data on GP consultations for ILI, virological positivity, virological confirmed cases, and serological surveys.

An Integrated Model. The proposed model in Fig. 2 comprises a transmission model and a disease and reporting model. The model dynamics are deterministic and discrete. We model from the period May 1 to December 31, 2009, and for the following age groups: <1 y, 1–4, 5–14, 15–24, 25–44, 45–64, and 65+ y. The epidemic is initiated with a small number of infectious individuals and a pool of susceptible individuals. At each subsequent time point the transmission model generates a number of newly infected individuals, which

enter the disease and reporting model, while the pool of susceptible individuals diminishes. The disease and reporting models then govern the proportion of these incident infections that appear in the GP consultation and confirmed case datasets and the delay inherent in doing so.

In the age-structured S, E, I, R model, transmission is dictated by a time- and age-varying force of infection $\lambda(t,a)$ and transition rates σ and γ , which describe the rates of transition between states $E \rightarrow I$ and $I \rightarrow R$, respectively. These rates are functions of the mean latent period, d_L , and the mean infectious period, d_I , the expected times spent in states E and I , respectively. The force of infection depends on two key quantities: the basic reproduction number of the virus, R_0 , and the relative rates of contact between the different age groups, introduced through the time-varying matrix, $\mathbf{M}(t)$. Details of how these quantities combine to give the incident number of infections can be found in *SI Materials and Methods, Eq. 7*. A proportion, θ , of the exposed individuals become clinical cases, with further fractions $p_{GP}(t,a)$ and p_{CC} of these symptomatic individuals consulting their GP or being virologically confirmed, respectively. Typically, there will also be a time lag from infection to either of these events. This delay is assumed to be distributed as a gamma random variable and arises from three independent processes: the incubation time until symptom onset, the delay in reporting the GP consultation or having illness virologically ascertained, and the subsequent reporting delay. These component delays are assumed to have known mean and variances, which are summed to give the mean and variance of the distribution governing the overall time from infection to each event. This is discussed in more detail in *SI Materials and Methods*. The size of the initially susceptible population within an age group, $S(0,a)$, is informed by baseline serological data from 2008 (3). Subsequently, for serological data taken at time t , the expected seropositivity is given by $1 - (S(t,a)/N_a)$, where N_a is the size of the population in age group a .

Modeling Challenges. Consultation Behavior. The propensity of individuals to consult with their GP given symptomatic ILI varied significantly over the course of the study period. Initially, this propensity was high, as seen from the marked increase in the consultation rates during the first wave, with only a modest increase in the accompanying virological positivity. However, government advice that patients were to consult through the NPFS, rather than their GP, drastically reduced this propensity. The model has to be sufficiently flexible to account for this, as well as permitting some temporal variation in the levels of adherence to the governmental guidelines over time. This impacts upon our model in two ways: (i) through the propensity to consult with a GP conditional upon symptomatic infection with A/H1N1pdm, $p_{GP}(t,a)$; this is modeled as a piecewise function over time, with differing rates for children and adults, the details of which can be found in *SI Materials and Methods*; (ii) through the "background" consultation of non-A/H1N1pdm patients with ILI symptoms. The background component of the consultation, $B(t,a)$, is parameterized as a piecewise constant function over time, with varying rates for each age group, thus allowing for temporal fluctuation in the behavior and prevalence of individuals with non-A/H1N1pdm ILI. See *SI Materials and Methods* for further details of the model and the estimation of these background rates of consultation, using informative priors derived on the basis of pandemic data heralding from other regions of England.

Mixing Rates and School Holidays. Estimated contact rates based on UK weekday data from within the POLYMOD study (26) formed the basis of the contact matrices $M(t)$ used in our analysis. In school term times, these POLYMOD matrices were modified through the introduction of a scaling factor, m_1 , applied to all matrix elements representing a contact rate involving adults. This confers an interpretation upon m_1 of a relative infectivity of adult infectious contacts in comparison to those solely involving children. Effects of school holidays upon disease transmission were accounted for by introducing further factors, m_2 and m_3 , which describe the proportionate reduction in rates of contact among 1–4 and 5–14 y-olds, respectively, during the summer holiday. During the shorter half-term holidays, additional multipliers, m_4 and m_5 , were applied to the same contact rates to permit differing effects of social distancing brought about by the two types of holiday. See *SI Materials and Methods* for further details.

Inference. Parameters. Inference is carried out within the Bayesian framework, based upon the posterior distributions of parameters and derived quantities of interest, obtained through the combination of the prior distributions and the likelihood function. We estimate posterior distributions for parameters R_0 , d_L , m_i ($i = 1, \dots, 5$), and the size of the initial spark of infection. Conversely, the mean latent period d_L is assumed to be known. Preliminary attempts to estimate both d_L and d_I highlighted that only their sum, not the individual components, is easily identified and these findings have been

formalized elsewhere (27). For the disease and reporting models, we estimate θ , p_{CC} , and the parameters describing $p_{GP}(t,a)$. Furthermore, we estimate the nuisance parameters used to model $B(t,a)$.

Likelihood. If we denote the collection of all model parameters by the vector φ , and

- i. w_{ta} is a realization of $W(t,a)$, the virological positivity at time t in age group a , based on a sample of size m_{ta}^w .
- ii. x_{ta}^{CC} is a realization of $X_{CC}(t,a)$, the number of lab-confirmed cases at time t in age group a .
- iii. y_{ta} is a realization of $Y(t,a)$, the number of GP consultations at time t in age group a .
- iv. z_{ta} is a realization of $Z(t,a)$, the seropositivity at time t in age group a , based on a sample of size m_{ta}^z .

Then, treating the above as independent data, the likelihood is given by

$$L(\varphi) = \prod_{t=1}^{n_t} \prod_{a=1}^{n_a} \{L(w_{ta}|\varphi) \times L(x_{ta}^{CC}|\varphi) \times L(y_{ta}|m_{ta}^y, \varphi) \times L(z_{ta}|m_{ta}^z, \varphi)\},$$

where n_t and n_a are the number of time points (245) and age groups (7), respectively. The third term in the product gives the likelihood of the GP consultation data. This is modeled through a negative binomial distribution to account for the overdispersion in the count data. This overdispersion is in part due to the within-week pattern of consultation characterized by very few consultations on weekends or bank holidays and a higher rate of reported consultations on Mondays, gradually declining through the week until a small increase on Fridays. The negative binomial distribution is parameterized in terms of the mean number of consultations (as found in *SI Materials and Methods*, Eq. 8) and a piecewise constant dispersion parameter, with one breakpoint at the time of NPFS launch. Otherwise, the confirmed cases, x_{CC} , are modeled as Poisson count data, and the positivity

and serological data are both treated as realizations of binomial random variables with known denominators.

Priors. A list of the model parameters comprising φ can be found in *SI Materials and Methods*, section 2.3.2. Where possible, parameters have been included as stochastic quantities; i.e., we have placed a prior upon them, so that we can learn about them through the data and so that the modeling procedure incorporates as much a priori knowledge/uncertainty as possible. Some parameters, due to reasons of identifiability, are held to fixed values. Fixed values and the majority of prior distributions are taken from the literature (see *SI Materials and Methods* for details). Such information is deemed to be unknown or unavailable for the parameters of the mixing matrices, m_i , and the overdispersion parameters, and so we place priors that are reasonably uninformative upon them.

Implementation. Posterior distributions for the unknown parameters are evaluated through Markov chain Monte Carlo methods, using a random walk Metropolis algorithm (28, 29). The algorithm was implemented using a bespoke C++ code specifically generated for this class of models. Two separate chains, each consisting of 450,000 iterations, were run in parallel, with the results presented based on a thinned subsample of the final 250,000 iterations from the two chains.

ACKNOWLEDGMENTS. The authors thank the Health Protection Agency Pandemic Influenza team for the timely availability of data; Professor E. Miller for providing serological data; the RCGP Research and Surveillance Centre; the University of Nottingham, Egton Medical Information Systems (EMIS), and EMIS practices contributing to the QSurveillance database. P.J.B., A.M.P., and D.D.A. were funded by the UK Medical Research Council (Grant G0600675). D.D.A. was funded also by the UK Health Protection Agency, as were B.S.C., R.J.H., A.C., X.S.Z., P.J.W., and R.G.P. G.K. was supported by the Commission of the European Community under the Sixth Framework Program Specific Targeted Research Project, SARS (Severe Acute Respiratory Syndrome) Control "Effective and Acceptable Strategies for the Control of SARS and new emerging infections in China and Europe" (Contract SP22-CT-2004-003824). B.S.C. acknowledges support by the Oak Foundation. P.J.W. thanks the Medical Research Council Centre for funding.

1. Fleming DM (1999) Weekly Returns Service of the Royal College of General Practitioners. *Commun Dis Public Health* 2:96–100.
2. Harcourt SE, et al. (2011) Use of a large general practice syndromic surveillance system to monitor the progress of the influenza A(H1N1) pandemic 2009 in the UK. *Epidemiol Infect*, 10.1017/S095026881100046X.
3. Miller E, et al. (2010) Incidence of 2009 pandemic influenza A H1N1 infection in England: a cross-sectional serological study. *Lancet* 375:1100–1108.
4. Fleming DM (November 16, 2009) Influenza surveillance, the swine-flu pandemic, and the importance of virology. *Clin Evid*.
5. Rubin GJ, Potts HW, Michie S (2010) The impact of communications about swine flu (influenza A H1N1v) on public responses to the outbreak: Results from 36 national telephone surveys in the UK. *Health Technol Assess (Rocky)* 14:183–266.
6. Ferguson NM, et al. (2006) Strategies for mitigating an influenza pandemic. *Nature* 442:448–452.
7. Cooper BS, Pitman RJ, Edmunds WJ, Gay NJ (2006) Delaying the international speed of pandemic influenza. *PLoS Med* 3:e212.
8. Chowell G, Ammon CE, Hengartner NW, Hyman JM (2006) Transmission dynamics of the great influenza pandemic of 1918 in Geneva, Switzerland: Assessing the effects of hypothetical interventions. *J Theor Biol* 241:193–204.
9. Riley S, et al. (2003) Transmission dynamics of the etiological agent of SARS in Hong Kong: Impact of public health interventions. *Science* 300:1961–1966.
10. Baguelin M, Van Hoek AJ, Flasche S, White PJ, Edmunds WJ (2010) Vaccination against pandemic influenza A/H1N1v in England: A real-time economic evaluation. *Vaccine* 28:2370–2384.
11. Bettencourt LMA, Ribeiro RM (2008) Real time Bayesian estimation of the epidemic potential of emerging infectious diseases. *PLoS One* 3:e2185.
12. Cauchemez S, Carrat F, Viboud C, Valleron AJ, Boëlle PY (2004) A Bayesian MCMC approach to study transmission of influenza: application to household longitudinal data. *Stat Med* 23:3469–3487.
13. Fraser C, et al. (2009) Pandemic potential of a strain of influenza A(H1N1): Early findings. *Science* 324:1557–1561.
14. Bootsma MCJ, Ferguson NM (2007) The effect of public health measures on the 1918 influenza pandemic in U.S. cities. *Proc Natl Acad Sci USA* 104:7588–7593.
15. Cauchemez S, Valleron A.-J., Boëlle P.-Y., Flahault A, Ferguson NM (2008) Estimating the impact of school closure on influenza transmission from Sentinel data. *Nature* 452:750–754.
16. Wu JT, et al. (2010) School closure and mitigation of pandemic (H1N1) 2009, Hong Kong. *Emerg Infect Dis* 16:538–541.
17. Fleming DM, Zambon M, Bartelds AI (2000) Population estimates of persons presenting to general practitioners with influenza-like illness, 1987–96: A study of the demography of influenza-like illness in sentinel practice networks in England and Wales, and in The Netherlands. *Epidemiol Infect* 124:245–253.
18. Hall IM, Gani R, Hughes HE, Leach S (2007) Real-time epidemic forecasting for pandemic influenza. *Epidemiol Infect* 135:372–385.
19. Hardelid P, et al. (2010) Assessment of baseline age-specific antibody prevalence and incidence of infection to novel influenza A(H1N1) 2009. *Health Technol Assess (Rocky)* 14:115–192.
20. Cowling BJ, et al. (2010) Comparative epidemiology of pandemic and seasonal influenza A in households. *N Engl J Med* 362:2175–2184.
21. Bandanarayake D, Bissielo A, Huang S, Wood T (2010) Seroprevalence of the 2009 influenza A(H1N1) pandemic in New Zealand. (New Zealand Ministry of Health, Wellington) Technical report.
22. Flahault A, de Lamballerie X, Hanslik T, Salez N (2009) Symptomatic infections less frequent with H1N1pdm than with seasonal strains. *Plos Curr*, RRN1140, 10.1371/currents.RRN1140.
23. Carrat F, et al. (2008) Time lines of infection and disease in human influenza: A review of volunteer challenge studies. *Am J Epidemiol* 167:775–785.
24. Evans B, et al. (2011) Has estimation of numbers of cases of pandemic influenza H1N1 in England in 2009 provided a useful measure of the occurrence of disease? *Influenza Other Respi Viruses*, 5 10.1111/j.1750-2659.2011.00259.x.
25. Ong JBS, et al. (2010) Real-time epidemic monitoring and forecasting of H1N1-2009 using influenza-like illness from general practice and family doctor clinics in Singapore. *PLoS One* 5:e10036.
26. Mossong J, et al. (2008) Social contacts and mixing patterns relevant to the spread of infectious disease. *PLoS Med* 5:e74.
27. Wallinga J, Lipsitch M (2007) How generation intervals shape the relationship between growth rates and reproductive numbers. *Proc R Soc Lond B Biol Sci* 274:599–604.
28. Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E (1953) Equations of state calculations by fast computing machines. *J Chem Phys* 21:1087–1091.
29. Hastings WK (1970) Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57:97–109.