

# Bayesian Markov switching models for the early detection of influenza epidemics

M. A. Martínez-Beneito<sup>1,2</sup>, D. Conesa<sup>2</sup>, A. López-Quílez<sup>2</sup> and A. López-Maside<sup>1</sup>

<sup>1</sup>*Àrea de Epidemiologia, Conselleria de Sanitat. Generalitat Valenciana.*

<sup>2</sup>*Departament d'Estadística i Investigació Operativa. Universitat de València.*

February 21, 2007

## Abstract

The early detection of the outbreaks of diseases is one of the most challenging objectives of epidemiological surveillance systems. In this paper, a Markov switching model is introduced to determine the epidemic and non-epidemic periods from influenza surveillance data: the process of differenced incidence rates is modelled either with a first-order autoregressive process or with a Gaussian white noise process depending on whether the system is in an epidemic or a non-epidemic phase. The transition between phases of the disease is modelled as a Markovian process. Bayesian inference is carried out on the former model to detect influenza epidemics at the very moment of their onset. Moreover, the proposal provides the probability of being in an epidemic state at any given moment. The methodology is evaluated on influenza illness data obtained from the Sanitary Sentinel Network of the Comunitat Valenciana, one of the 17 autonomous regions in Spain.

**Keywords:** Autoregressive modelling; Bayesian inference; Influenza; Hidden Markov models; Outbreak detection; Public health; Temporal surveillance.

## 1 Introduction

Influenza is an acute respiratory illness that affects the upper and/or lower parts of the respiratory tract and is caused by the influenza virus. Influenza epidemics occur virtually every year during the winter in temperate areas of the Northern and the Southern Hemispheres and result in substantial disease, death and expense. Although the extent and severity of such epidemics vary greatly, it is worth noting that approximately 10-15% of people get influenza around the world each year and that the disease is responsible for up to 50 million illnesses and up to 47,200 deaths in the United States each year, with a similar situation in Europe (Simonsen et al., 1997; Fleming et al., 1999; Monto, 1999) [20, 3, 15]. All these figures clearly turn influenza surveillance into a challenging

issue in public health practice.

In general, public health surveillance can be defined as “the ongoing, systematic collection, analysis, and interpretation of health data essential to the planning, implementation, and evaluation of public health practice, closely integrated with the timely dissemination of these data to those who need to know” (Thacker, 1994) [25]. Le Straat (2005) [11] comprehensively reviews to that date the statistical methods that have been applied for detecting or monitoring outbreaks and monitoring trends of diseases. Some of them have been implemented in the R-package `surveillance`, in particular, those by Stroup et al. (1989) [24], Farrington et al. (1996) [2] and Höhle and Riebler (2005) [10]. The list of contributions in surveillance is still growing (see, for instance, recent papers by Held et al. (2005, 2006) [8, 9]).

An important matter of concern when dealing with the surveillance of infectious diseases such as influenza is that of detecting the onset of an epidemic as soon as possible. This would imply prompt intervention which could have a great impact on the number of lives saved. Surveillance systems have also recently gained increasing importance due to the threat of emerging infections (like the outbreaks of the H5N1 bird-flu strain) and the increased potential for bioterrorist attacks.

The influenza surveillance system in Europe is run by the European Influenza Surveillance Scheme (EISS), while in the United States it is managed by the Centers for Disease Control and Prevention (CDC) Influenza Branch. In both cases weekly information on influenza activity is collected and reported from September to May. The current approach to influenza surveillance used by the EISS is based on historical limits methods. In particular, each national network depending on the EISS assesses the intensity of activity based on the historical data at its disposal and establishes numerical thresholds that define the intensity of influenza activity. In the United States, the approach to influenza surveillance is based on Serfling’s method (Serfling, 1963) [19], which monitors the normal pattern of susceptibility to death from pneumonia and influenza when there is no epidemic, with the final objective of determining a national baseline defined by a fixed threshold.

These approaches have some drawbacks: the need for a predefinition of epidemic and non-epidemic periods to model the baseline distribution, the fact that observations are treated as independent and identically distributed, and that a national baseline does not provide a useful threshold if we are interested in local influenza surveillance (Rath et al., 2003) [17]. Moreover, as Goddard et al. (2003) [6] state, the use of fixed threshold values to describe the levels of influenza activity can be misleading due to the decline in consultations for influenza in recent years.

LeStrat and Carrat (1999) [12] pioneered the use of hidden Markov Models (MacDonald and Zucchini, 1997) [13] to segment the time series of influenza indicators into epidemic and non-epidemic phases. This approach has two advantages, the first being that the method can be applied to historical data without the need to distinguish between epidemic and non-epidemic periods in the data. The second one deals with the fact that the observations are supposed to be

independent given our knowledge about the epidemic, whereas Serfling's method assumes marginal independence of the data (Rath et al., 2003) [17]. In subsequent papers, Rath et al. (2003) [17] and Madigan (2005) [14] presented further exploration of this modelling, incorporating the latter the Bayesian perspective. Current efforts on influenza surveillance include papers by Muglin et al. (2002) [16] and by Sebastiani et al. (2006) [18]. In the first one, disease counts are treated as a realization from an underlying multivariate autoregressive process where the relative risk of infection incorporates a space-time dynamic, while in the second one dynamic Bayesian networks are used to integrate four different data streams into a multivariate model for influenza surveillance. This latter paper is in line with what has recently been called syndromic surveillance, in which a large number of symptoms and data regarding nontraditional sources of information are used to detect diseases even before diagnoses can be confirmed through unmistakable signs (see, for instance, Goldenberg et al. (2002) [7]).

Our main goal in this paper is to introduce an alternative approach to influenza surveillance that avoids some of the above-mentioned disadvantages. Our proposal is to use a Markov switching model in order to determine the epidemic and non-epidemic periods from influenza surveillance data. This approach differs from those previously mentioned in the sense that we work with the series of differenced incidence rates rather than with the series of incidence rates. The advantage of this new approach lies in the fact that the differenced series is stationary and thus it allows us to take advantage of autoregressive modelling to analyze the data. In particular, depending on whether the system is in an epidemic or a non-epidemic phase, we model the differenced series either with a first-order autoregressive process or with a Gaussian white noise process. The transition between the phases of the disease is also introduced via a Markovian process. Bayesian paradigm is used in order to provide the probability of being in an epidemic phase at any given moment, which is the key to detecting influenza epidemics at their onset.

The remainder of this paper is organized as follows. Section 2 presents the data we have used in order to illustrate our findings. Section 3 summarizes our modelling of the influenza incidence rates. Section 4 describes the results obtained when applying our methodology in our particular setting, while incorporating a validation of the predictive performance of our model and finally Section 5 is devoted to concluding remarks.

## 2 Influenza surveillance data

In what follows, we introduce the data set we have used to illustrate the behaviour of our model in a real setting. In particular, we present the Valencian Sentinel Network (VSN) for influenza surveillance, a system which collects information about influenza-like illness (ILI) in the Comunitat Valenciana, one of the 17 autonomous regions in Spain. VSN has been reporting information on ILI cases since 1996. It collects data in seasons that enlarge on two consecutive years, as

the epidemic activity usually extends across both of them. Each season lasts at least 30 weeks (from the 42nd week of one year to the 19th week of the following). The Network is formed by volunteer practitioners that report weekly the number of ILI cases jointly with personalized, clinical, epidemiological and microbiological information about the cases. Although the number of participating physicians in the network varies from season to season, it has always been around 30 practitioners (covering 1% of the population of the Comunitat Valenciana).

As a result, our data consist of nine time series formed by the weekly ILI incidence rates (per 100000 inhabitants) provided by the VSN during the seasons from 1996-1997 to 2004-2005. It is worth mentioning that each weekly rate has been standardized by only considering the population from those sentinels that have reported information. The series are displayed in Figure 1. Note that they show a mixture of two dynamics: a non-epidemic dynamic in which the incidence rate does not present big changes and varies randomly around small values, and an epidemic dynamic in which the incidence rate increases and decreases sharply at irregular intervals. This behaviour is very characteristic of influenza and certain other infectious diseases (i.e. chicken pox) in which the shape of the epidemic wave depends greatly on whether we are or are not in an epidemic phase of the disease. It is also worth noting that the behaviour of the incidence rates cannot be considered as seasonal because of the low rates observed throughout the 2000-2001 season.




Figure 1: *Influenza incidence rates (per 100000 inhabitants) during the nine seasons analyzed.*

The former series are not stationary, which could lead to some difficulties in the data analysis. This suggests working with the first order differenced series (formed by the differences of rates between weeks) displayed in Figure 2. Note that the two above-mentioned dynamics can also be appreciated in these stationary series. The non-epidemic dynamic is characterized by small random changes around zero, while in the epidemic dynamic changes are greater and inter-related (positive and negative values are usually followed by positive and negative values respectively). But more importantly, the fact that this new series has a zero mean allows us to restrict our study to its variability at each moment, while in the former series we had differences not only in the

variability but also in the means, thus making its analysis more difficult. This comment is in line with that of Barón (2002) [1], who also used differenced series to distinguish between epidemic and non-epidemic phases.




Figure 2: *Increments of incidence rates between weeks during the nine seasons analyzed.*

Based on all these considerations about our data, in the following section we present our modelling for the series of differences based on a two-stage Markov switching model and the way it can be used to detect changes between the two dynamics.

### 3 The model

Basically, our modelling is based on a segmentation of the series of differences into an epidemic and a non-epidemic phase using a two-stage Markov switching model (see Frühwirth-Schnatter (2006) [4], for a broad monograph of this kind of models). Let  $Y = \{Y_{i,j}, i = 1, \dots, 29; j = 1, \dots, 9\}$  denote the difference between the rates of weeks  $i + 1$  and  $i$  in year  $j$ . The underlying idea of Markov switching models is to associate each  $Y_{i,j}$  with a random variable  $Z_{i,j}$  that determines the conditional distribution of  $Y_{i,j}$  given  $Z_{i,j}$ . In our case, each  $Z_{i,j}$  is an unobserved random variable that indicates which phase the system is in (1, epidemic; 0, non-epidemic). Moreover, the unobserved sequence of  $Z_{i,j}$  (note that we do not know which phase the system is in at any given moment) follows a two-state Markov chain of order 1 with transition probabilities:

$$P_{k,l} = P(Z_{i+1,j} = l | Z_{i,j} = k) \quad \text{where } k, l \in \{0, 1\}, \quad i \in \{1, \dots, 29\} \text{ and } j \in \{1, \dots, 9\}.$$

This kind of Markov switching models in which the variables that determine the conditional distribution are unobserved are usually known as hidden Markov models (MacDonald and Zucchini, 1997) [13].

Then, taking into account that the non-epidemic dynamic is characterized by small random changes around zero, while in the epidemic dynamic changes are greater and inter-related, we model

the conditional distribution of  $Y_{i,j}$  either as an autoregressive process of order 1 or as a Gaussian white noise process depending on whether the system is in an epidemic or in a non-epidemic phase:

$$\begin{aligned}
Y_{1,j}|(Z_{1,j} = 0) &\sim N(0, \sigma_{0,j}^2) \\
Y_{1,j}|(Z_{1,j} = 1) &\sim N(0, \sigma_{1,j}^2) \\
Y_{i,j}|(Z_{i,j} = 0) &\sim N(0, \sigma_{0,j}^2) & i = 2, \dots, 29, j = 1, \dots, 9, \\
Y_{i,j}|(Z_{i,j} = 1) &\sim N(\rho Y_{i-1,j}, \sigma_{1,j}^2) & i = 2, \dots, 29, j = 1, \dots, 9,
\end{aligned} \tag{3.1}$$

where the first subindex of the variance  $\sigma_{k,j}^2$  represents whether the system is in the epidemic phase ( $k = 1$ ) or not ( $k = 0$ ). Note that we assume a different variance for each season. This has been done in order to reflect the behaviour observed in Figure 2, in which it can be appreciated that the length of the variations is not the same in the different years. It is also worth noting that the conditional distribution of the first difference of rates cannot be modelled as an autoregressive process as there is no previous value to condition to.

Once the model is determined, the following step is to estimate its parameters. Taking into account that  $P_{k,0} + P_{k,1} = 1$  for  $k \in \{0, 1\}$ , then we only need to estimate  $\rho$ ,  $P_{0,0}$ ,  $P_{1,1}$  and  $\{\sigma_{0,j}^2, \sigma_{1,j}^2; j = 1, \dots, 9\}$ . To do so, we take advantage of the Bayesian paradigm, which requires specification of the prior distributions of each parameter involved in the model. In this case, with the aim of expressing our initial vague knowledge about them, we consider the usual noninformative prior distributions for  $\rho$ ,  $P_{0,0}$  and  $P_{1,1}$ :

$$\begin{aligned}
\rho &\sim \text{Unif}(-1, 1) \\
P_{1,1} &\sim \text{Beta}(0.5, 0.5) \\
P_{0,0} &\sim \text{Beta}(0.5, 0.5)
\end{aligned} \tag{3.2}$$

Moreover, taking into account that  $\sigma_{0,j}^2$  should be lower than  $\sigma_{1,j}^2$  as it only responds to random variations instead of the effect of the epidemic, we then express our prior knowledge about  $\{\sigma_{0,j}^2, \sigma_{1,j}^2; j = 1, \dots, 9\}$  via the following hierarchical structure:

$$\begin{aligned}
\sigma_{0,j} &\sim \text{Unif}(\theta_{\text{low}}, \theta_{\text{mid1}}) \\
\sigma_{1,j} &\sim \text{Unif}(\theta_{\text{mid2}}, \theta_{\text{sup}}) \\
\theta_{\text{low}} &\sim \text{Unif}(a, b) \\
\theta_{\text{mid1}} &\sim \text{Unif}(\theta_{\text{low}}, b) \\
\theta_{\text{mid2}} &\sim \text{Unif}(\theta_{\text{mid1}}, b) \\
\theta_{\text{sup}} &\sim \text{Unif}(\theta_{\text{mid2}}, b)
\end{aligned} \tag{3.3}$$

where  $a$  and  $b$  are hyperparameters to be fixed (based on the knowledge of the problem) in a way that they do not interfere with the results. We choose uniform distributions as priors for the

standard deviations of the random effects, as suggested for example in Gelman (2006) [5]. These priors, unlike the usual non-informative Gamma prior distributions on the precision parameters, allow us to make inference on the limits of the uniform distributions and hence learn about the appropriate range of variation for the epidemic and non-epidemic periods. Moreover, the suggested uniform priors avoid the identifiability problem (Stephens, 2000) [23] between both periods in the MCMC process, as the prior distribution for the standard deviation in the epidemic period has been restricted to be greater than in the non-epidemic one.

Expressions (3.1), (3.2) and (3.3) contain all our knowledge of the system but they do not yield to analytical estimates. Therefore, we have to resort to Markov Chain Monte Carlo (MCMC) methods. In particular, we have used WinBUGS (Spiegelhalter et al., 1999) [22] to carry out the inference. Note that, in order to avoid some traps in WinBUGS, the sampling method for log-concave variables has been changed to slice sampling (see user’s manual for more details about how to change the sampling methods for certain classes of distributions in order to avoid possible traps). We have reported the WinBUGS code of our modelling in the Appendix.

## 4 Results

Figure 3 and Table 1 show the results of our analysis (based on the formulation described in the previous section) of the weekly ILI incidence rates observed in the Comunitat Valenciana. The results have all been obtained by using 3 independent chains of simulations with 30000 iterations, discarding the first 15000 and selecting 1 of each 45 (3000 simulated values in total). Values for  $a$  and  $b$  in (3.3) were set to be 10 and 300. The selection of this latter value was based on the fact that the biggest difference between weekly rates was 212, while we selected  $a = 10$  in order to avoid the posterior distribution of  $\sigma_{0,j}^2$  converging to 0.

In particular, at the top of Figure 3, we present the posterior probability of being in an epidemic phase for each week in the analysis. These values correspond to the posterior mean of the state variable  $Z_{i,j}$ . Note that these probabilities can be very useful for decision makers (in our case, the Valencian Regional Health Authorities), as they allow them to quantify the possibility of being in either phase each week. In particular, values exceeding 0.5 indicate that, in that week, we are observing a higher probability of being in an epidemic phase than of being in a non-epidemic one, and so an alarm could be triggered if it is considered necessary. This information can be appreciated in more detail in the graph of the influenza incidence rates that we present at the bottom of Figure 3, in which we have plotted with black spots those weeks with a posterior probability of being in an epidemic phase higher than 0.5.

Table 1 shows the posterior mean and the 95% posterior credible interval of the parameters involving the model obtained when using WinBUGS. As can be seen, jumps from an epidemic to a non-epidemic phase are very robust: there is no more than one epidemic period per season although

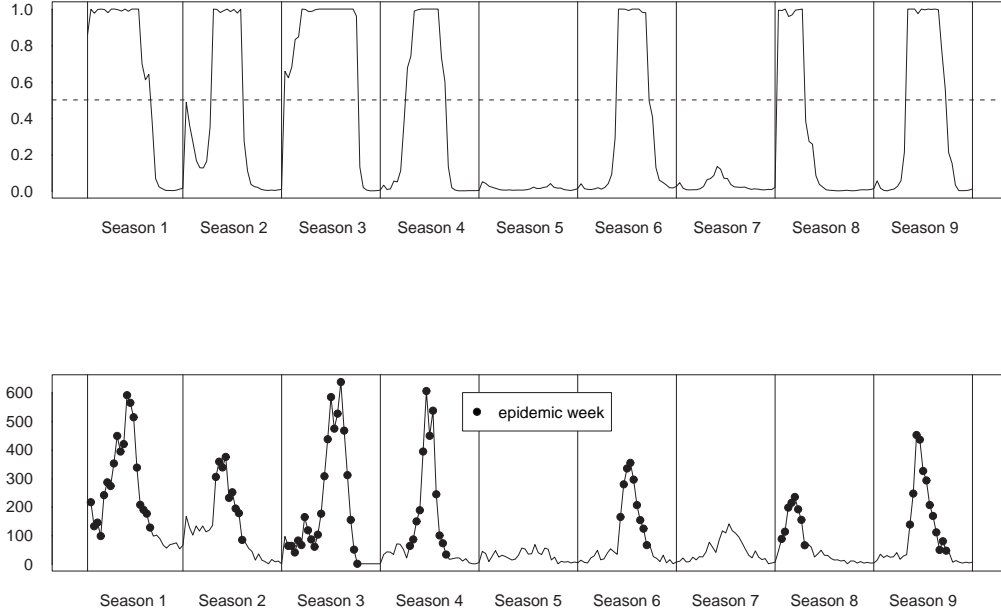


Figure 3: *Posterior probability of being in the epidemic phase (at the top) and a representation of the influenza incidence rates per 100000 inhabitants in which the spots indicate those weeks where the posterior probability of being in an epidemic phase exceeds 0.5 (at the bottom).*

the model allows for any number of changes. This is a consequence of the Markovian behaviour defined in the model. Indeed, as stated in Table 1, the model estimates that it is very likely to be in the epidemic phase if we were in it during the previous week (mean posterior probability equals 0.906) and it is also very likely to be in a non-epidemic state if we were in it in during the previous week (0.964 being its mean posterior probability). Of particular interest is the strictly positive value of the correlation  $\rho$ , which seems to justify the autoregressive process used for the epidemic period. It is also important to note that the posterior distributions of  $\theta_{mid1}$  and  $\theta_{mid2}$  seem to justify the selection of disjoint intervals for representing the variances in the epidemic and non-epidemic phases.

One of the main aims of the present study is to help managers to raise an alarm at those precise moments in which there is a high probability of being in an epidemic phase. Clearly, at each moment, available information only comes from those weeks preceding the particular moment we are analyzing, but there is no known information about what it is going to happen in the following weeks. In order to validate the predictive performance of our methodology, we have carried out the following procedure. For each week of the study we have predicted the probability



Parameter	Posterior Mean	95% Credible Interval
$\theta_{\text{low}}$	12.716	[10.130, 16.410]
$\theta_{\text{mid1}}$	19.196	[15.270, 24.952]
$\theta_{\text{mid2}}$	49.093	[20.614, 87.295]
$\theta_{\text{sup}}$	134.711	[92.883, 212.582]
$\rho$	0.344	[0.140, 0.560]
$P_{0,0}$	0.964	[0.925, 0.989]
$P_{0,1}$	0.036	[0.011, 0.075]
$P_{1,0}$	0.094	[0.040, 0.164]
$P_{1,1}$	0.906	[0.836, 0.960]

Table 1: Posterior mean and the 95% credible interval of the parameters involving the model.

of being in an epidemic phase by only taking into account information from the previous weeks of the season in which the analyzed week is included, and that from the other eight seasons in the study.

For each week, Figure 4 shows the posterior probability of being in an epidemic phase when using the complete information (9 seasons) and when only using information from the previous weeks in the season of the week analyzed and information from the remaining 8 seasons. As can be seen, although probabilities are not as close to 0 (indication of non-epidemic) or 1 (epidemic) as they are when we use the complete information, they are close enough to those values in order to properly indicate the phase the system is in. Moreover, if we use the horizontal dotted line representing the 0.5 probability to classify weeks into both dynamics, the results indicate that 93% of the weeks (234 out of 252) were classified in the same way when using all the information and when only using information from the previous weeks jointly with information from the remaining seasons.

It is worth noting that this on-line methodology has already been incorporated into the weekly bulletin published by the Valencian Sentinel Network. In particular, the probability of being in the epidemic phase is included in their report, which is distributed to all the Primary Health Care Centers in the Comunitat Valenciana.

Finally, with the aim of checking that all the components of our model expressed through (3.1), (3.2) and (3.3) improve model performance, we have carried out a comparison between the model in Section 3 and three other proposals in which some parts of our modelling have been removed. In particular, in the first alternative, we model the system using the same standard deviation for all seasons, while in the second one we do not incorporate the auto-regressive part of the epidemic phase into the model. Finally, in the last competitive proposal, we use a non-Markovian model.

In order to compare competing models, we use the deviance information criterion (DIC) introduced by Spiegelhalter et al. (2002) [21]. Table 2 summarizes the obtained results. As can be




Figure 4: *Posterior probability of being in an epidemic phase when using all the information (solid line) and when only using information from the previous weeks jointly with information from the remaining 8 seasons (dashed line).*

appreciated, our proposal outperforms (by providing a better fit, which can be seen by a lower DIC) the other three models. In fact, the Markovian hypothesis emerges as the most important feature of our modelling in terms of DIC, as its elimination provokes the highest increase in this criterion. The auto-regressive structure of the epidemic phase also shows a marked effect on the modelling. Finally, the variability on the variances between seasons seems to be the least important factor in our modelling, although the DIC criterion prefers to include such variability in the model.

MODEL	DIC
Model in Section 3	2542.3
Same standard deviations for all seasons	2545.7
Avoiding the auto-regressive part of the epidemic phase	2551.6
Using a non-Markovian model	2628

Table 2: Comparison of the DIC value obtained when using our model without its principal components.

## 5 Concluding remarks

Our interest in this paper has been to describe a methodology to detect influenza epidemics at the very moment of their onset, implemented in a local setting. An important novelty with respect to previous studies in epidemiological surveillance is that our proposal tackles this question from the perspective of the changes in the differences of the rates. This makes our proposal complementary to existing ones which are based on the analysis of the value of the weekly rates instead of their

variability. As far as we know, Barón (2002) [1] was the first approximation to influenza surveillance in which the interest was to analyze changes in the detrended mortality rates. Nevertheless, his modelling did not include an auto-regressive term and it was not based on a Markovian structure, which can be very helpful to model influenza surveillance as we have just shown in the previous section. As in other models based on Markov switching models (LeStraat and Carrat, 1999; Rath et al., 2003; Madigan, 2005) [12, 17, 14], another characteristic of this kind of proposal is that of providing the probability of being in an epidemic phase at any given moment. But, as happens with any modelling based on past data, it is so important to have a large amount of data from previous seasons in order to have correct results.

We now comment on possible extensions to this study. A first possibility could be to explore whether the probability of being in an epidemic phase could depend not only on the rate of the previous week but also on the particular moment in the season (maybe at its early stages or at its final ones). Another potential extension of our model would include a multivariate or a spatial component that could help us to explore any geographical disaggregation of the rates (in this case, we would have to consider as many series as practitioners). Moreover, as Goldenberg et al. (2002) [7] state “the next generation of biosurveillance systems will incorporate information from multiple sources, including public-health and nontraditional data”. This makes it important to develop and refine methodologies that could incorporate these characteristics, which would help in identifying outbreaks and thus facilitate timely actions to decrease unnecessary morbidity and mortality.

Finally, we would like to stress that the coherence of the results obtained using our method on the data from the Valencian Sentinel Network, jointly with its good on-line behaviour, seem to make feasible its more general application for monitoring epidemiologic surveillance data based on Sentinel Networks.

## Appendix

```

model{
  for (j in 1:nyear)
  {
    dif.rates[1, j] ~ dnorm(0,tau[1, j])
    tau[1, j] <- pow(lambda[comp[1, j], j],-2)
  }
  for (j in 1:nyear)
  {
    for (i in 2:nweek[j]) {
      dif.rates[i, j] ~ dnorm(mu[i, j],tau[i, j])
      tau[i, j] <- pow(lambda[comp[i, j], j],-2)
    }
  }
}

```

```

        mu[i, j] <- ro*dif.rates[i-1, j]*equals(comp[i, j],2)
    }
}
ro ~ dunif(-1,1)
for (j in 1:nanyo)
{
    comp[1, j] ~ dcat(P0[])
    lambda[1, j] ~ dunif(linf,lmed1)
    lambda[2, j] ~ dunif(lmed2,lsup)
}
linf ~ dunif(a,b)
lmed1 ~ dunif(linf,b)
lmed2 ~ dunif(lmed1,b)
lsup ~ dunif(lmed2,b)
for (j in 1:nyear)
{
    for (i in 2:nweek[j])
    {
        comp[i, j] ~ dcat(P.mat[comp[i-1, j], ])
    }
}
P0[1]<-0.5
P0[2]<-0.5
P.mat[1,1] ~ dbeta(0.5,0.5)
P.mat[2,2] ~ dbeta(0.5,0.5)
P.mat[1,2]<- 1-P.mat[1,1]
P.mat[2,1]<- 1-P.mat[2,2]
}

```

## Acknowledgements

The authors would like to thank the generous work of the practitioners of the Valencian Sentinel Network. Financial support from the Conselleria de Sanitat of the Generalitat Valenciana (the Valencian Regional Health Authority) is gratefully acknowledged. Authors would also like to acknowledge the financial support from the Ministerio de Educación y Ciencia (the Spanish Ministry of Education and Science) via the research grant MTM2004-03290 (jointly financed with European Regional Development Fund) and from the Generalitat Valenciana via research grants

## References

- [1] M. I. Barón. Bayes and asymptotically pointwise optimal stopping rules for the detection of influenza epidemics. In C. Gatsonis, R. E. Kass, A. Carriquiry, A. Gelman, D. Higdon, D. K. Pauler, and I. Verdinelli, editors, *Case Studies in Bayesian Statistics*, volume VI, pages 153–163. Springer, 2002.
- [2] C. P. Farrington, N. J. Andrews, A. D. Beale, and M. A. Catchpole. A statistical algorithm for the early detection of outbreaks of infectious disease. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 159(3):547–563, 1996.
- [3] D. M. Fleming, M. Zambon, A. I. Bartelds, and J. C. de Jong. The duration and magnitude of influenza epidemics: a study of surveillance data from sentinel general practices in England, Wales and the Netherlands. *European Journal of Epidemiology*, 15:467–473, 1999.
- [4] S. Frühwirth-Schnatter. *Finite mixture and Markov switching models*. Springer, 2006.
- [5] A. Gelman. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, 1:515–533, 2006.
- [6] N. L. Goddard, J. Kyncl, and J. M. Watson. Appropriateness of thresholds currently used to describe influenza activity in England. *Communicable disease and public health*, 6(3):238–245, 2003.
- [7] A. Goldenberg, G. Shmueli, R. A. Caruana, and S. E. Fienberg. Early statistical detection of anthrax outbreaks by tracking over-the-counter medication sales. *Proceedings of the National Academy of Sciences*, 99:5237–5249, 2002.
- [8] L. Held, M. Höhle, and M. Hofmann. A statistical framework for the analysis of multivariate infectious disease surveillance data. *Statistical Modelling*, 5:187–199, 2005.
- [9] L. Held, M. Hofmann, and M. Höhle. A two-component model for counts of infectious diseases. *Biostatistics*, 7(3):422–437, 2006.
- [10] M. Höhle and A. Riebler. The R-package ‘surveillance’. Technical Report SFB386 discussion paper 422, Department of Statistics, University of Munich, 2005.
- [11] Y. LeStrat. Overview of temporal surveillance. In A. B. Lawson and K. Kleinman, editors, *Spatial and syndromic surveillance for public health*, pages 13–29. John Wiley and Sons, Ltd., 2005.

- [12] Y. LeStrat and F. Carrat. Monitoring epidemiological surveillance data using hidden Markov models. *Statistics in Medicine*, 18:3463–3478, 1999.
- [13] I. L. MacDonald and W. Zucchini. *Hidden Markov and other models for discrete-valued time series*. Chapman and Hall, 1997.
- [14] D. Madigan. Bayesian data mining for health surveillance. In A. B. Lawson and K. Kleinman, editors, *Spatial and syndromic surveillance for public health*, pages 203–221. John Wiley and Sons, Ltd., 2005.
- [15] A. S. Monto. Individual and community impact of influenza. *Pharmacoeconomics*, 16:1–6, 1999.
- [16] A. S. Mugglin, N. Cressie, and I. Gemmell. Hierarchical statistical modelling of influenza epidemic dynamics in space and time. *Statistics in Medicine*, 21:2703–2721, 2002.
- [17] T. M. Rath, M. Carreras, and P. Sebastini. Automated detection of influenza epidemics with hidden Markov models. In M. R. Berthold, H. J. Lenz, E. Bradley, R. Kruse, C. Borgelt, and F. Pfennig, editors, *Advances in Intelligent data analysis V*, pages 521–531. Springer-Verlag, Berlin, 2003.
- [18] P. Sebastiani, K. D. Mandl, P. Szolovits, I. S. Kohane, and M. F. Ramoni. A Bayesian dynamic model for influenza surveillance. *Statistics in Medicine*, 25:1803–1816, 2006.
- [19] R. E. Serfling. Methods for current statistical analysis of excess pneumonia-influenza deaths. *Public Health Reports*, 78:494–506, 1963.
- [20] L. Simonsen, M. J. Clarke, G. D. Williamson, D. F. Stroup, N. H. Arden, and L. B. Schonberger. The impact of influenza epidemics on mortality: introducing a severity index. *American Journal of Public Health*, 87:1944–1950, 1997.
- [21] D. J. Spiegelhalter, N. G. Best, B. Carlin, and A. Van der Linde. Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B*, 64:583–640, 2002.
- [22] D. J. Spiegelhalter, A. Thomas, and N. G. Best. Winbugs version 1.2 user manual. *MRC Biostatistics Unit*, 1999.
- [23] M. Stephens. Dealing with label switching in mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62:795–809, 2000.
- [24] D. F. Stroup, G. D. Williamson, J. L. Herndon, and J. M. Karon. Detection of aberrations in the occurrence of notifiable diseases surveillance data. *Statistics in Medicine*, 8:323–329, 1989.

- [25] S. B. Thacker. Historical development. In S. M. Teutsch and R. E. Churchill, editors, *Principles and Practice of public health surveillance*, pages 3–17. Oxford University Press, Inc., 1994.