

Estimation of HIV burden through Bayesian evidence synthesis

Daniela De Angelis*, Anne M. Presanis, Stefano Conti and A. E. Ades

Medical Research Council Biostatistics Unit and Health Protection Agency, Medical Research Council Biostatistics Unit, Health Protection Agency, University of Bristol

Abstract. Planning, implementation and evaluation of public health policies to control the human immunodeficiency virus (HIV) epidemic require regular monitoring of disease burden. This includes the proportion living with HIV, whether diagnosed or not, and the rate of new infections in the general population and in specific risk groups and regions. Estimation of these quantities is not straightforward: data informing them directly are not typically available, but a wealth of indirect information from surveillance systems and ad hoc studies can inform functions of these quantities. In this paper we show how the estimation problem can be successfully solved through a Bayesian evidence synthesis approach, relaxing the focus on ‘best available’ data to which classical methods are typically restricted. This more comprehensive and flexible use of evidence has led to the adoption of our proposed approach as the official method to estimate HIV prevalence in the United Kingdom since 2005.

AMS 2000 subject classifications: Primary Bayesian inference.

Key words and phrases: Evidence synthesis, Graphical model, HIV, Disease burden.

1. INTRODUCTION

The HIV disease is associated with serious morbidity, high costs of treatment and care, and, in developing countries, with significant mortality and a high number of potential years of life lost (Joint United Nations Programme on HIV/AIDS, 2010). Planning for care provision and for implementation and evaluation of public health policies to reduce transmission relies crucially on robust monitoring of disease burden. This burden includes the proportion (prevalence) living with HIV; the proportion of infections remaining undiagnosed; and the rate at which new

*University Forvie Site, Robinson Way, Cambridge CB2 0SR, United Kingdom
(e-mail: daniela.deangelis@mrc-bsu.cam.ac.uk)*

*We thank the HIV division at the Health Protection Agency for providing the data on which this work is based. Special thanks go to Dr. Valerie Delpech and Prof. Noel Gill for discussion on the interpretation of data and results. We are indebted also to Prof. David Spiegelhalter for many useful methodological discussions over the last few years. This work was supported by the Medical Research Council [grant number G0600675, Unit Programme Number U105260566] and the Health Protection Agency.

infections occur (incidence), in both the general population and in specific groups at high risk of infection and in different locations. To acquire robust evidence on these quantities is not easy. The assessment of HIV prevalence is complicated by the absence of symptoms for a long time after infection. Incidence is even more difficult to measure, requiring, at least, longitudinal follow-up of uninfected individuals, with all the complications of cohort studies.

Devising appropriate methods for estimation of prevalence and incidence has generated a rich literature in the last 30 years (Brookmeyer, 2010; Presanis, 2010). For HIV prevalence ‘direct’ methods have been particular popular amongst the medical community (e.g. McGarrigle et al., 2006; Lyerla et al., 2006, and references therein) for their apparent transparency. The underlying idea is that the general population, of size N , is subdivided into G non-overlapping groups at different risk of acquiring HIV. Estimates of proportions ρ_g of risk group g ($g = 1, \dots, G$) in the population are multiplied by estimates of the prevalence π_g of HIV to produce a point estimate of the number of infected individuals $N\pi_g\rho_g$ in each group and in the population $N\sum_g\rho_g\pi_g = N\sum_g(\rho_g\pi_g\delta_g + \rho_g\pi_g(1 - \delta_g))$. Here δ_g denotes the proportion of infected individuals diagnosed in group g , and $N\sum_g\rho_g\pi_g\delta_g$ and $N\sum_g\rho_g\pi_g(1 - \delta_g)$ represent the number of diagnosed and undiagnosed infections in group g , respectively. Typically, at least in developed countries with concentrated epidemics like the United Kingdom (UK), the number of diagnosed infections is known from surveillance schemes, so the problem is to estimate the number of undiagnosed infections. Provided direct data that measure size and prevalence for each group are available, these methods are very simple and, consequently, appealing. However, *direct* information on all parameters is not typically available, whereas there is a wealth of *indirect* information, from a variety of sources, which may inform functions of the parameters of interest. This additional indirect information is generally discarded as difficult to incorporate in this simplistic framework. As a result, on one hand, unverifiable assumptions and ad hoc adjustments are made to compensate for the lack of information. On the other hand, an inefficient use is made of the information that is available, with consequent imprecise and biased results due to the selective nature of the data used in the estimation. Finally, in the ‘direct’ methods there is no explicit model formulation, so it is not possible to quantify formally the uncertainty surrounding the resulting estimates or to validate them.

The statistical challenge is then to provide an inferential approach capable of combining *direct* and *indirect* information from multiple sources and appropriately accounting for any uncertainty in the data and parameters. The Bayesian paradigm naturally offers the most appropriate framework to address this challenge (see Section 5). Bayesian synthesis of evidence from different studies, perhaps even those with different designs, is not new (e.g. Eddy, Hasselblad and Shachter, 1992; Dominici et al., 1999; Ades and Sutton, 2006) and is attracting increasing attention with applications in various fields (e.g. Spiegelhalter and Best, 2003; Clark et al., 2010; Govan et al., 2010; Birrell et al., 2011).

In this paper, we describe how such an approach has been successfully adopted to estimate HIV prevalence and incidence in England and Wales (E&W) in the population aged 15-44. The remainder of the paper is organised as follows: the concept of Bayesian evidence synthesis is defined in Section 2; the model to estimate HIV prevalence is presented in Section 3; a joint model for prevalence and

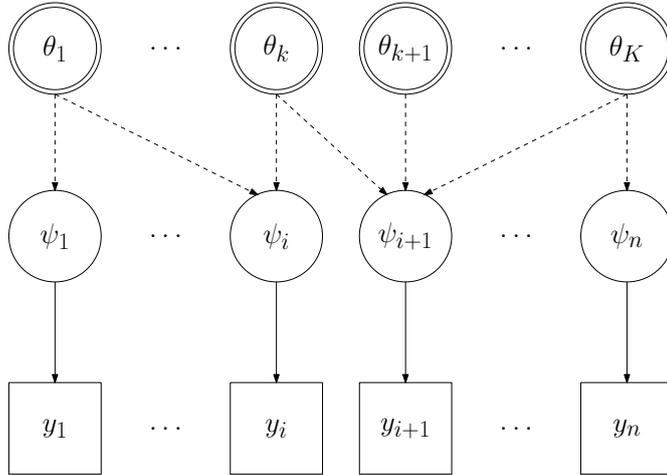


FIGURE 1. DAG representation of a generic evidence synthesis model.

incidence is described in Section 4; and Section 5 offers a concluding discussion.

2. BAYESIAN EVIDENCE SYNTHESIS

Let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_K)$ represent the parameter vector we are interested in estimating. We refer to $\boldsymbol{\theta}$ as *basic* parameters. Denote by $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$ a collection of $n \geq K$ independent data items available for the estimation of $\boldsymbol{\theta}$. Each \mathbf{y}_i provides either *direct* information on a single component θ_k of $\boldsymbol{\theta}$ or *indirect* information, i.e., on functional parameters, expressed in terms of one or more component(s) of $\boldsymbol{\theta}$. Denote by $\psi_i = \psi_i(\boldsymbol{\theta})$ a generic function of $\boldsymbol{\theta}$, which may represent the identity function, i.e., $\psi_i = \theta_k$, a function of a single parameter $\psi_i = \psi_i(\theta_k)$ or a function of multiple components of $\boldsymbol{\theta}$, $\psi_i = \psi_i(\boldsymbol{\theta})$. Indicating by $L_i(\psi_i(\boldsymbol{\theta}); \mathbf{y}_i)$ the likelihood contribution of \mathbf{y}_i to the basic parameter vector $\boldsymbol{\theta}$, from the independence of the \mathbf{y}_i , the full data likelihood is $L(\boldsymbol{\theta}; \mathbf{y}) = \prod_{i=1}^n L_i(\psi_i(\boldsymbol{\theta}); \mathbf{y}_i)$. From a Bayesian perspective, expressing the prior knowledge on $\boldsymbol{\theta}$ through a prior distribution $p(\boldsymbol{\theta})$, inference is conducted on the basis of the posterior distribution $p(\boldsymbol{\theta} | \mathbf{y}) \propto p(\boldsymbol{\theta})L(\boldsymbol{\theta}; \mathbf{y})$, which summarises all information, both *direct* and *indirect*, on $\boldsymbol{\theta}$. Such a distribution fully reflects the uncertainty about $\boldsymbol{\theta}$, including sampling variability and parameter uncertainty, which automatically percolates through to any function of the basic parameters $\boldsymbol{\theta}$. Figure 1 provides a direct acyclic graph (DAG) (Lauritzen, 1996) representation of the generic formulation above and shows schematically the dependency between data and parameters as well as the flow of information within the system. Here stochastic ‘nodes’ are represented by circles and observed ‘nodes’ by squares. The basic parameters, in double circles, are given prior (possibly hierarchical) distributions. Solid arrows represent distributional assumptions, and dashed arrows indicate functional relationships. Note the examples of functional parameters that inform multiple components of $\boldsymbol{\theta}$, such as $\psi_i = \psi_i(\theta_1, \theta_k)$. Information flows along the arrows, from the prior and from the data. The posterior distribution of each θ_k is based on its prior distribution and on direct and indirect information available on it, as well as the priors and information on other components of $\boldsymbol{\theta}$.

3. HIV PREVALENCE ESTIMATION

Extending the notation introduced in Section 1, HIV prevalence $\pi_{t,r}$ in the general population at a single point in time t in location r may be expressed as $\pi_{t,r} = \sum_g \rho_{t,g,r} \pi_{t,g,r} = \sum_g \rho_{t,g,r} \pi_{t,g,r} \delta_{t,g,r} + \sum_g \rho_{t,g,r} \pi_{t,g,r} (1 - \delta_{t,g,r})$. The aim is to estimate the basic parameters $\theta_{t,g,r} = (\rho_{t,g,r}, \pi_{t,g,r}, \delta_{t,g,r})$. Having obtained the posterior distribution of these, it is possible to obtain the posterior distribution of any function of interest, e.g., the total number of infections $N_{t,r} \sum_g \rho_{t,g,r} \pi_{t,g,r}$ or the total number of undiagnosed infections $N_{t,r} \sum_g \rho_{t,g,r} \pi_{t,g,r} (1 - \delta_{t,g,r})$, where $N_{t,r}$ is the location- and time-specific total population. There are 13 mutually exclusive risk groups defining a hierarchy of risk. Men are classified into: men who have sex with men (MSM attending sexually transmitted infection (STI) clinics; MSM not attending STI clinics; and past MSM); injecting drug users (IDU, current and past); heterosexual men born in sub-Saharan Africa (SSA); heterosexual men attending STI clinics; and heterosexual men at low risk (LR) of infection. Heterosexual women are classified in the same way as heterosexual men. Geographically, there are three locations (Inner London, Outer London, Rest of E&W), and t refers to the year 2008. In total there are $11 \times 3 + 13 \times 3 + 13 \times 3 = 111$ parameters as $\sum_g \rho_{t,g,r} = 1$ for each gender.

3.1 Data

Different types of data are available on: group sizes, HIV prevalence, prevalence of undiagnosed infections, proportion of infections diagnosed, total number of diagnosed infections and group distribution amongst diagnosed cases. Data sources are described in full and commented upon elsewhere (Goubar et al., 2008; Preanis et al., 2010, and references therein), and are only briefly reviewed here. Mid-year population estimates provide information on $N_{t,r}$ and some risk group proportions $\rho_{t,g,r}$. The remaining $\rho_{t,g,r}$ are derived from a behavioural survey. Unlinked anonymous sero-prevalence surveys amongst STI clinic attendees inform the prevalence of undiagnosed infection $\pi_{t,g,r} (1 - \delta_{t,g,r})$. The analogous surveys amongst pregnant women and IDUs inform prevalence $\pi_{t,g,r}$ and proportion diagnosed $\delta_{t,g,r}$, both directly for some groups, and indirectly through functions of $\pi_{t,g,r}$ and $\rho_{t,g,r}$. The pregnant women's survey, in particular, measures prevalence in those born in SSA and the remainder (NSSA). These NSSA are a mixture of STI clinic attendees, IDUs, and lower risk women. The observed data, therefore, provide information on a complex function of HIV prevalence in these groups and account for the probability of each group being included in the sample. An annual cross-sectional survey of diagnosed individuals collects information on functional parameters representing both the total number living with diagnosed HIV $(N_{t,r} \sum_g \rho_{t,g,r} \pi_{t,g,r} \delta_{t,g,r})$ and the distribution of risk groups amongst these individuals $(\rho_{t,g,r} \pi_{t,g,r} \delta_{t,g,r}) / (\sum_g \rho_{t,g,r} \pi_{t,g,r} \delta_{t,g,r})$ for each group g . Table 1 summarises the spread and the type of information available as well as the relationship between the available data and the basic parameters, expressed here through generic functions ψ . Note the sparseness of information on heterosexual men and the multiplicity of data on heterosexual women.

3.2 Inference

Sampling distributions Information $\mathbf{y}_{t,g,r}$ from most sources can be expressed in the form of count data $x_{t,g,r}$ with an associated denominator $n_{t,g,r}$.

TABLE 1
Relationship between available data and parameters for a generic location r .

	Risk group	N	ρ	π	δ	$\psi(\rho, \pi)$	$\psi(\pi, \delta)$	$\psi(\rho, \pi, \delta)$
Men	MSM		✓				✓	✓
	IDUs		✓	✓	✓			✓
	Born sub-Saharan Africa		✓					✓
	STI clinic attendees		✓				✓	✓
	Lower risk							
	ALL	✓						✓
Women	IDUs		✓	✓	✓	✓		✓
	Born sub-Saharan Africa		✓	✓	✓			✓
	STI clinic attendees		✓			✓	✓	✓
	Lower risk					✓		
	ALL	✓		✓				✓

These data can be assumed to naturally be realisations of a binomial random variable

$$X_{t,g,r} \sim \text{binomial}(n_{t,g,r}, \psi_{t,g,r})$$

where $\psi_{t,g,r}$ equals any of $\rho_{t,g,r}$, $\pi_{t,g,r}$ and $\delta_{t,g,r}$ if $\mathbf{y}_{t,g,r}$ provides direct information, or is a function of these basic parameters.

The observed total numbers of diagnosed men and women in each location, $x_{t,m,r}$ and $x_{t,f,r}$ respectively, are assumed to be realisations of Poisson random variables $X_{t,m,r} \sim \text{Poisson}(\mu_{t,m,r})$ and $X_{t,f,r} \sim \text{Poisson}(\mu_{t,f,r})$ where

$$\begin{aligned} \mu_{t,m,r} &= N_{t,m,r} \sum_{g_m} (1 - \nu_{t,g_m}) \delta_{t,g_m,r} \pi_{t,g_m,r} \rho_{t,g_m,r} \\ \mu_{t,f,r} &= N_{t,f,r} \sum_{g_f} (1 - \nu_{t,g_f}) \delta_{t,g_f,r} \pi_{t,g_f,r} \rho_{t,g_f,r} \end{aligned}$$

Here g_m and g_f refer to male and female groups, respectively and ν_{t,g_m}, ν_{t,g_f} are parameters representing potential bias in the reported number of diagnosed individuals due to non-attendance, under-reporting, or duplication. The region-specific numbers diagnosed in each risk group, $x_{t,g_m,r}$ and $x_{t,g_f,r}$, are simultaneously drawn from gender-specific multinomial distributions with size parameters $\mu_{t,m,r}$ and $\mu_{t,f,r}$, and probability parameters

$$\begin{aligned} \xi_{t,g_m,r} &= ((1 - \nu_{t,g_m}) \delta_{t,g_m,r} \pi_{t,g_m,r} \rho_{t,g_m,r}) / \sum_{g_m} (1 - \nu_{t,g_m}) \delta_{t,g_m,r} \pi_{t,g_m,r} \rho_{t,g_m,r} \\ \xi_{t,g_f,r} &= ((1 - \nu_{t,g_f}) \delta_{t,g_f,r} \pi_{t,g_f,r} \rho_{t,g_f,r}) / \sum_{g_f} (1 - \nu_{t,g_f}) \delta_{t,g_f,r} \pi_{t,g_f,r} \rho_{t,g_f,r} \end{aligned}$$

The full likelihood $L_t(\boldsymbol{\theta}_t; \mathbf{y}_t)$ results from the product of each of these distributions, as generically described in Section 2.

Sparseness of information One of the challenges to the ‘direct’ methods is the lack of information on some risk groups. Table 1 clearly shows that data on $\pi_{t,g,r}$ and $\delta_{t,g,r}$ for male heterosexuals are sparse. This sparsity can be addressed by sharing information between men and women. Although $\pi_{t,g,r}$ and $\delta_{t,g,r}$ are expected to vary by gender and by location, it is reasonable to assume that their male-to-female odds ratio might be similar between regions. To borrow strength across locations and risk groups, the following hierarchical structures

are then assumed for the male-to-female log odds ratios of prevalence $lor.\pi_{t,g,r}$ and proportion diagnosed $lor.\delta_{t,g,r}$:

$$\begin{aligned} \text{logit}(\pi_{t,g_m,r}) &= \text{lor}.\pi_{t,g,r} + \text{logit}(\pi_{t,g_f,r}); \text{lor}.\pi_{t,g,r} \sim \text{normal}(P_{t,g}, \sigma_{t,\pi}^2) \\ \text{logit}(\delta_{t,g_m,r}) &= \text{lor}.\delta_{t,g,r} + \text{logit}(\delta_{t,g_f,r}); \text{lor}.\delta_{t,g,r} \sim \text{normal}(D_{t,g}, \sigma_{t,\delta}^2) \end{aligned}$$

with a further hierarchy over risk groups:

$$P_{t,g} \sim \text{normal}(\Pi_t, \omega_{t,\pi}^2); D_{t,g} \sim \text{normal}(\Delta_t, \omega_{t,\delta}^2).$$

The means Π_t and Δ_t are a priori distributed as $\text{normal}(0, 100^2)$. The standard deviations $\sigma_{t,\pi}, \sigma_{t,\delta}$ and $\omega_{t,\delta}$ are given informative priors expressing the belief that only 5% of region-specific male-to-female odds ratios (of both prevalence and proportion diagnosed) will vary from the mean by more than a 1.3 factor (Section 5.7.3 of Spiegelhalter, Abrams and Myles (2004)). The odds ratios for prevalence are assumed to vary more across risk groups than across regions, hence the prior for $\omega_{t,\pi}$ is weaker: a factor of 1.6 is used.

Bias modelling and other indirect information A further challenge to the estimation problem is represented by data that indirectly inform a specific parameter of interest. The most common example occurs where the data are known to be affected by biases, as for the above total number of diagnosed infections. The parameters ν_{t,g_m} and ν_{t,g_f} are, in this case, introduced to model the known bias. In general, this is dealt with by introducing ‘bias models’ that take the generic form $\theta' = \theta + \epsilon$ on a suitable scale, where θ is the parameter of interest and θ' is the parameter directly informed by the data. The ‘bias parameter’ ϵ is a measure of the discrepancy between θ' and θ . Where information or expert opinion on the size and/or direction of the bias is available, this is expressed as an informative prior on ϵ .

Other challenges in the data sources, such as greater spatial heterogeneity than is captured by the regional structure adopted in the model, are met by more complex modelling, such as mixed effects regression on a finer regional stratification. The parameters $\psi_{t,g,r}$ in the binomial expression above may therefore have a more complex functional structure than the examples given here, see Goubar et al. (2008); Presanis et al. (2008); Presanis (2010) for more details.

Priors Diffuse uniform priors are chosen for the basic parameters $\pi_{t,g,r}$ and $\delta_{t,g,r}$. The proportions of the male and female populations in each risk group $\rho_{t,g_m,r}$ and $\rho_{t,g_f,r}$ are given Dirichlet priors such that they sum to 1. Informative normal or uniform priors are assigned to bias parameters such as ν_{t,g_m} and ν_{t,g_f} .

Results Samples from the posterior distribution are obtained using Markov chain Monte Carlo (MCMC), implemented in WinBUGS (Lunn et al., 2000). Posterior summaries are based on 8,000 samples from two chains after convergence is achieved. Figure 2 presents the estimated number of HIV infections in E&W, by diagnosis status and risk group.

4. JOINT PREVALENCE AND INCIDENCE MODEL

Application of the prevalence model over successive years using a sequence of datasets $\{\mathbf{y}_t\}, t \in 1, \dots, T$, provides the joint posterior distribution of the proportions of the population $N_{t,g,r}$ in each of three compartments: susceptible to infection $s_{t,g,r} = \rho_{t,g,r}(1 - \pi_{t,g,r})$; HIV infected but undiagnosed $u_{t,g,r} =$

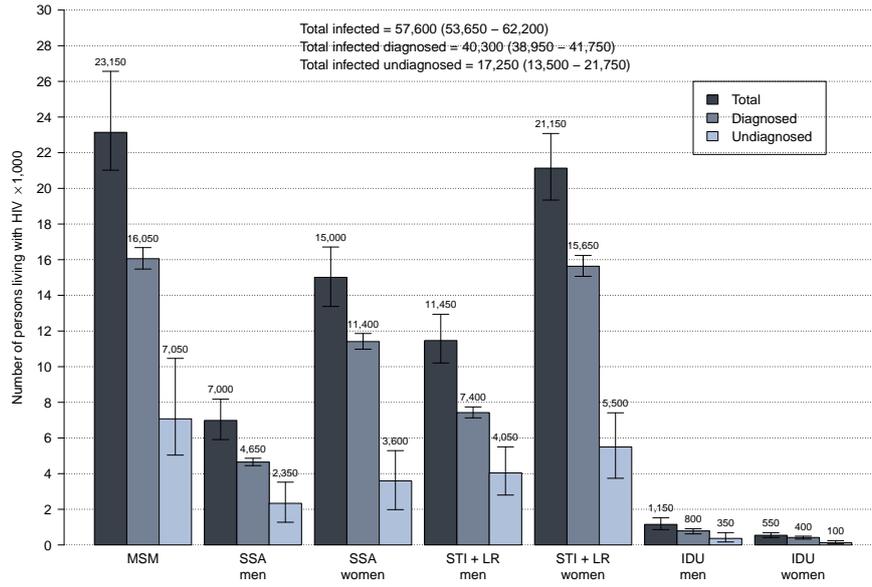


FIGURE 2. Posterior median (95% credible interval) number of HIV infections in adults aged 15-44 in E&W in 2008, by diagnosis status and risk group.

$\rho_{t,g,r}\pi_{t,g,r}(1 - \delta_{t,g,r})$; and infected and diagnosed $d_{t,g,r} = \rho_{t,g,r}\pi_{t,g,r}\delta_{t,g,r}$. These can be interpreted as estimates of the state at time t of a dynamical system describing the processes of infection and diagnosis. Such a system can be approximated by a continuous-time Markov model whose dynamics are described through a system of ordinary differential equations (ODEs). As in Presanis et al. (2011), we focus here on the MSM group and, for simplicity, drop the subscripts g and r . Let $\mathbf{c}_t = (s_t, u_t, d_t)$ and denote by $\boldsymbol{\lambda}_t$ the transition rates in the time interval $[t, t + 1)$, assumed piecewise constant over the interval. Using additional data $\{\mathbf{z}_t\}$ (Presanis et al., 2011) and prior information on demographics and risk behaviour uptake, a joint prevalence and incidence model can be formulated to allow simultaneous estimation of the prevalence parameters $\boldsymbol{\theta}_t$, the compartment proportions \mathbf{c}_t and the transition rates $\boldsymbol{\lambda}_t$ including $\lambda_t^{s,u}$, the incidence rate, i.e., the rate at which susceptible individuals enter the infected state. The DAG in Figure 3 provides a schematic representation of this joint model. The proportions \mathbf{c}_{t+1} at time $t + 1$ are defined, through the ODEs, in terms of the rates $\boldsymbol{\lambda}_t$ during the period $[t, t + 1)$ and the initial condition of the system at $t = 1$. The prevalence parameters $\boldsymbol{\theta}_t$ and the $\boldsymbol{\lambda}_t$ govern the prevalence and rate data, respectively. Note that this DAG has the same structure as that in Figure 1. Now the $\boldsymbol{\lambda}_t$ and \mathbf{c}_1 are the *basic* parameters and the $\boldsymbol{\theta}_t$ are functional parameters.

Inference is conducted as described in Section 2. The likelihood of the joint data is

$$L(\mathbf{c}_1, \boldsymbol{\lambda}; \mathbf{y}, \mathbf{z}) = \prod_{t=1}^T L_t(\mathbf{c}_1; \mathbf{y}_t) L_t(\boldsymbol{\lambda}_t; \mathbf{z}_t).$$

where $L_t(\boldsymbol{\lambda}_t; \mathbf{z}_t)$ is the likelihood contribution of the demographic and behavioural data informing transition rates. Assuming independent vague priors for $\boldsymbol{\lambda}_t$ and

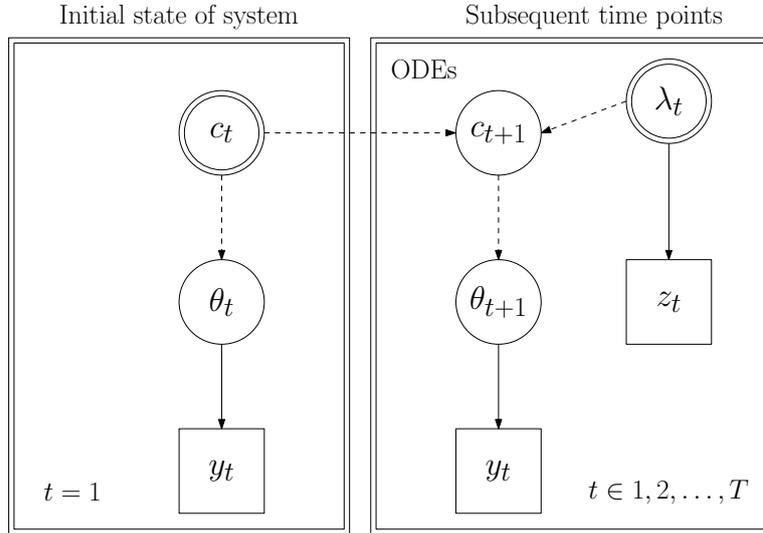


FIGURE 3. Schematic DAG of the joint prevalence and incidence model.

a Dirichlet(1, 1, 1, 1) prior for the compartment proportions at time 1, \mathbf{c}_1 , the joint posterior distribution for \mathbf{c}_1 and $\boldsymbol{\lambda}$ and therefore also of $\boldsymbol{\theta}$, is obtained through MCMC implemented in WinBUGS. Note that the likelihood contribution of the prevalence data, \mathbf{y}_t , depends on the \mathbf{c}_t , the ODEs' solutions, which are derived numerically for the current parameter values at each MCMC iteration using the Runge-Kutta algorithm in the `WBDiff` package in WinBUGS. Figure 4 shows posterior distributions resulting from the joint prevalence and incidence model.

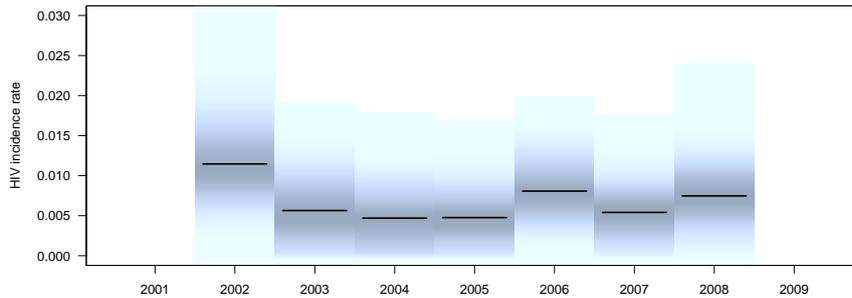


FIGURE 4. Density strip plots of posterior distributions of incidence rates in MSM (darker colour corresponds to higher density), 2002-2008.

5. DISCUSSION

From a methodological point of view, this work has responded formally to the need, perceived by epidemiologists working in the HIV arena, to ‘triangulate’ *all* information: from multiple and imperfect sources and expert opinion on the epidemiological interpretation of the data from these sources. The approach is clearly appealing: it uses data fully, minimising potential biases due to selection of information; it typically leads to more precise estimates, which are consistent

with all information; and it accounts for all sources of uncertainty, naturally reflected in the posterior distributions of parameters and quantities of interest.

5.1 Why Bayesian?

In principle, evidence synthesis does not need to be carried out in a Bayesian framework: see for instance Eddy, Hasselblad and Shachter (1992) and much of the meta-analytical work referenced in Sutton et al. (2000). Indeed, work exists to estimate HIV prevalence in a single risk group by synthesising three data sources in a classical approach, accounting for missing data (Walker et al., 2011). However, the unprecedented multiplicity of data sources, risk groups, and indirect information involved in the work described here requires a Bayesian approach, with clear benefits over classical, likelihood-based alternatives. The main advantage is the ability to (i) explicitly introduce and (ii) formally quantify expert judgements. The hierarchical model introduced in Section 3.2 to tackle data sparseness offers such an example: only through reasonably chosen informative priors on the standard deviation hyper-parameters has it been possible to overcome identifiability problems due to lack of information (see sensitivity analyses in Presanis (2010)). Secondly, a Bayesian model can be easily extended to include auxiliary ‘bias’ parameters to quantify lack of validity and relevance of data items for the estimation for any specific parameter. Expert epidemiological information on the direction and magnitude of such biases is naturally accommodated in a Bayesian set-up through carefully chosen priors (see Section 3.2 and references therein). It is not immediately obvious how a classical modelling approach would accommodate such information. Computational convenience represents a further advantage of a Bayesian approach. As the posterior distribution is estimated through simulation, it is straightforward to obtain inferences on any functional parameter of interest. The likelihood function of even a moderately sized evidence network is unlikely to be sufficiently tractable to allow comparably streamlined inference.

5.2 Impact on the real world

Since 2005, our ‘multi-parameter’ evidence synthesis has been the approach adopted to produce the official estimates of the magnitude of the HIV problem in the UK, in particular the undiagnosed component, underlying current testing recommendations (<http://www.hpa.org.uk/web/HPAweb&Page&HPAwebAutoListName/Page/1201094588821>). In 2011, estimated trends on the prevalence of undiagnosed infection and incidence in MSM informed the work of the House of Lords Select Committee on HIV/AIDS in the UK (<http://www.publications.parliament.uk/pa/ld201012/ldselect/ldaids/188/188.pdf>). Recently, dissemination of the method has also attracted the interest of international public health organisations. Funded by the World Health Organization as part of an ongoing critical review of current methods for HIV prevalence estimation in concentrated epidemics (Ghys et al., 2008), the prevalence model has been adapted to estimate HIV burden in the Netherlands for the year 2007 (Conti et al., 2011). In comparison to other direct-type methods, the evidence synthesis approach was found to be the most flexible and statistically sound (van Veen et al., 2011).

5.3 Current challenges

The model building and criticism processes in this work have led to a critical understanding of the strengths and weaknesses of the various sources of HIV infor-

mation available in the UK, often challenging common interpretation of the data. Extensive sensitivity analyses have been carried out for prior and structural assumptions, to the sampling distributions employed, as well as to the data sources included, have been carried out (Goubar et al., 2008; Presanis et al., 2008; Presanis, 2010; Presanis et al., 2011). Moreover, routine annual application of the model has led to continual model development, responding to changes in surveillance, the availability of new data sources, and ongoing model criticism in the cycle recommended by Box (1980) and O’Hagan (2003), amongst others. Some of the development required and in progress includes addressing issues of missing data, using ideas as in Walker et al. (2011), and a comprehensive model of the process of diagnosis in STI clinics based on a new surveillance system (<http://www.hpa.org.uk/Topics/InfectiousDiseases/InfectionsAZ/HIV/OverallHIVPrevalence/>).

More generally, model criticism is essential in an evidence synthesis approach. As data come from multiple sources depending on shared parameters, it becomes crucial to understand and communicate which sources (including priors) drive conclusions and whether the various items of evidence result in consistent or conflicting inference. Efforts clearly need to be focussed on the development of transparent methods for model assessment and criticism, given that evidence synthesis approaches are being increasingly employed in different areas of science.

In the same spirit, an important step towards improved communication and dissemination of Bayesian evidence synthesis would be the availability of a user-friendly computing environment, facilitating access and implementation of the methodology to non-experts. van Veen et al. (2011) also identified the lack of such a modelling interface as a restriction to the more widespread adoption of our approach.

REFERENCES

- ADES, A. E. and SUTTON, A. J. (2006). Multiparameter evidence synthesis in epidemiology and medical decision-making: current approaches. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **169** 5–35.
- BIRRELL, P. J., KETSETZIS, G., GAY, N. J., COOPER, B. S., PRESANIS, A. M., HARRIS, R. J., CHARLETT, A., ZHANG, X.-S., WHITE, P. J., PEBODY, R. G. and DE ANGELIS, D. (2011). Bayesian modeling to unmask and predict influenza A/H1N1pdm dynamics in London. *Proceedings of the National Academy of Sciences* **108** 18238–18243.
- BOX, G. E. P. (1980). Sampling and Bayes’ Inference in Scientific Modelling and Robustness. *Journal of the Royal Statistical Society. Series A (General)* **143** 383–430.
- BROOKMEYER, R. (2010). Measuring the HIV/AIDS Epidemic: Approaches and Challenges. *Epidemiologic Reviews* **32** 26–37.
- CLARK, J. S., BELL, D., CHU, C., COURBAUD, B., DIETZE, M., HERSH, M., HILLERISLAMBERS, J., IBÁÑEZ, I., LADEAU, S., MCMAHON, S., METCALF, J., MOHAN, J., MORAN, E., PANGLE, L., PEARSON, S., SALK, C., SHEN, Z., VALLE, D. and WYCKOFF, P. (2010). High-dimensional coexistence based on individual variation: a synthesis of evidence. *Ecological Monographs* **80** 569–608.
- CONTI, S., PRESANIS, A. M., VAN VEEN, M. G., XIRIDOU, M., DONOGHOE, M. C., RINDER STENGAARD, A. and DE ANGELIS, D. (2011). Modeling of the HIV infection epidemic in the Netherlands: A multi-parameter evidence synthesis approach. *The Annals of Applied Statistics* **5** 2359–2384.
- DOMINICI, F., PARMIGIANI, G., WOLPERT, R. L. and HASSELBLAD, V. (1999). Meta-Analysis of Migraine Headache Treatments: Combining Information from Heterogeneous Designs. *Journal of the American Statistical Association* **94** 16–28.
- EDDY, D. M., HASSELBLAD, V. and SHACHTER, R. (1992). *Meta-Analysis by the Confidence Profile Method*. Academic Press.

- GHYS, P. D., WALKER, N., MCFARLAND, W., MILLER, R. and GARNETT, G. P. (2008). Improved data, methods and tools for the 2007 HIV and AIDS estimates and projections. *Sexually Transmitted Infections* **84** i1–i4.
- GOUBAR, A., ADES, A. E., DE ANGELIS, D., MCGARRIGLE, C. A., MERCER, C., TOOKEY, P., FENTON, K. and GILL, O. N. (2008). Estimates of HIV prevalence and proportion diagnosed based on Bayesian multi-parameter synthesis of surveillance data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **171** 541–580.
- GOVAN, L., ADES, A. E., WEIR, C. J., WELTON, N. J. and LANGHORNE, P. (2010). Controlling ecological bias in evidence synthesis of trials reporting on collapsed and overlapping covariate categories. *Statistics in Medicine* **29** 1340–1356.
- JOINT UNITED NATIONS PROGRAMME ON HIV/AIDS, (2010). *Global report UNAIDS report on the global AIDS epidemic. 2010*. UNAIDS.
- LAURITZEN, S. L. (1996). *Graphical Models. Oxford Statistical Science Series*. Oxford University Press, USA.
- LUNN, D. J., THOMAS, A., BEST, N. and SPIEGELHALTER, D. (2000). WinBUGS - A Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing* **10** 325–337.
- LYERLA, R., GOUWS, E., GARCÍA-CALLEJA, J. M. and ZANIEWSKI, E. (2006). The 2005 Workbook: an improved tool for estimating HIV prevalence in countries with low level and concentrated epidemics. *Sexually Transmitted Infections* **82** iii41–iii44.
- MCGARRIGLE, C. A., CLIFFE, S., COPAS, A. J., MERCER, C. H., DE ANGELIS, D., FENTON, K. A., EVANS, B. G., JOHNSON, A. M. and GILL, O. N. (2006). Estimating adult HIV prevalence in the UK in 2003: the direct method of estimation. *Sexually Transmitted Infections* **82** 78–86.
- O'HAGAN, A. (2003). *HSSS model criticism (with discussion)*. In *Highly Structured Stochastic Systems*, first ed. *Oxford Statistical Science Series*. Oxford University Press, USA.
- PRESANIS, A. M. (2010). Evidence synthesis methods to estimate disease prevalence, incidence and transmission, with application to HIV. PhD thesis, MRC BSU, University of Cambridge.
- PRESANIS, A. M., DE ANGELIS, D., SPIEGELHALTER, D. J., SEAMAN, S., GOUBAR, A. and ADES, A. E. (2008). Conflicting evidence in a Bayesian synthesis of surveillance data to estimate HIV prevalence. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **171** 915–937.
- PRESANIS, A. M., GILL, O. N., CHADBORN, T. R., HILL, C., HOPE, V., LOGAN, L., RICE, B. D., DELPECH, V. C., ADES, A. E. and DE ANGELIS, D. (2010). Insights into the rise in HIV infections, 2001 to 2008: a Bayesian synthesis of prevalence evidence. *AIDS* **24** 2849–2858.
- PRESANIS, A. M., DE ANGELIS, D., GOUBAR, A., GILL, O. N. and ADES, A. E. (2011). Bayesian evidence synthesis for a transmission dynamic model for HIV among men who have sex with men. *Biostatistics* **12** 666–681.
- SPIEGELHALTER, D. J., ABRAMS, K. R. and MYLES, J. P. (2004). *Bayesian Approaches to Clinical Trials and Health-Care Evaluation. Statistics in Practice*. Wiley.
- SPIEGELHALTER, D. J. and BEST, N. G. (2003). Bayesian approaches to multiple sources of evidence and uncertainty in complex cost-effectiveness modelling. *Statistics in Medicine* **22** 3687–3709.
- SUTTON, A. J., ABRAMS, K. R., JONES, D. R., SHELDON, T. A. and SONG, F. (2000). *Methods for Meta-Analysis in Medical Research. Wiley Series in Probability and Statistics*. John Wiley & Sons Ltd., Chichester, UK.
- VAN VEEN, M. G., PRESANIS, A. M., CONTI, S., XIRIDOU, M., STENGAARD, A. R., DONOGHOE, M. C., VAN SIGHEM, A. I., VAN DER SANDE, M. A. and DE ANGELIS, D. (2011). National estimate of HIV prevalence in the Netherlands: comparison and applicability of different estimation tools. *AIDS* **25** 229–237.
- WALKER, K., SEAMAN, S. R., DE ANGELIS, D., PRESANIS, A. M., DODDS, J. P., JOHNSON, A. M., MERCEY, D., GILL, O. N. and COPAS, A. J. (2011). A synthesis of convenience survey and other data to estimate undiagnosed HIV infection among men who have sex with men in England and Wales. *International Journal of Epidemiology* **40** 1358–1366.