Journal of the Royal Statistical Society
SERIES A Statistics in Society

J. R. Statist. Soc. A (2014)

Estimating prevalence of injecting drug users and associated heroin-related death rates in England by using regional data and incorporating prior information

Ruth King,

University of St Andrews, UK

Sheila M. Bird,

Medical Research Council Biostatistics Unit, Cambridge, and University of Strathclyde, Glasgow, UK

Antony M. Overstall,

University of St Andrews, UK

Gordon Hay

Liverpool John Moores University, UK

and Sharon J. Hutchinson

University of Strathclyde, Glasgow, and Health Protection Scotland, Glasgow, UK

[Received March 2012. Revised December 2012]

Summary. Injecting drug users (IDUs) have a direct social and economic effect yet can typically be regarded as a hidden population within a community. We estimate the size of the IDU population across the nine different Government Office regions of England in 2005–2006 by using capture–recapture methods with age (ranging from 15 to 64 years) and gender as covariate information. We consider a Bayesian model averaging approach using log-linear models, where we can include explicit prior information within the analysis in relation to the total IDU population (elicited from the number of drug-related deaths and injectors' drug-related death rates). Estimation at the regional level allows for regional heterogeneity with these regional estimates aggregated to obtain a posterior mean estimate for the number of England's IDUs of 195840 with 95% credible interval (181700, 210480). There is significant variation in the estimated regional prevalence of current IDUs per million of population aged 15–64 years, and in injecting drug-related death rates across the gender × age cross-classifications. The propensity of an IDU to be seen by at least one source also exhibits strong regional variability with London having the lowest propensity of being observed (posterior mean probability 0.21) and the South West the highest propensity (posterior mean 0.46).

Keywords: Drug-related deaths; Injecting drug users; Log-linear models; Model averaging; Population size; Prior information

Address for correspondence: Ruth King, Centre for Research into Ecological and Environmental Modelling, School of Mathematics and Statistics, University of St Andrews, North Haugh, St Andrews, KY16 9SS, UK.
E-mail: ruth@mcs.st-and.ac.uk

 0964–1998/14/177000

1. Introduction

We focus on estimating the prevalence in 2005–2006 of current injecting drug users (IDUs) mainly of opiates in England, and at the Government Office region level when cross-classified across gender and age (15–34 and 35–64 years). England's population of injectors rose epidemically in the (late) 1980s (de Angelis *et al.*, 2004). Opiate substitution therapy was introduced to reduce injection-related harms and to promote off-injecting. Quality assurance in methadone prescribing was achieved between 2000 and 2004 (see Strang *et al.* (2010)). A major public health reason to engage injectors in methadone substitution therapy is to reduce their risks of bloodborne virus transmission and drug-related death (DRD). Methadone clients may continue to inject but, typically, their number of injections of illicit heroin reduces considerably (Hutchinson *et al.*, 2000). Estimating the number of current IDUs at the regional and national levels (cross-classified by gender and age) permits the estimation of the injecting DRD rate by taking the ratio of the corresponding number of deaths attributed to injectors (namely heroin-related deaths (HRDs)) with the estimated population size of IDUs.

Sudden deaths (which include DRDs) in individuals within the UK are almost always subject to a post-mortem examination to determine the cause of death. Toxicology tests are conducted to identify illicit drugs within the system. Because there is no standardized protocol for conducting or reporting the toxicology of DRDs, there may be some heterogeneity in recording such deaths. Official statistics do not document whether the deceased person had a history of injection drug use, let alone whether she or he was a current injector, and so we cannot know which opiate-related DRDs occurred in current injectors. (Not all opiate-related DRDs occur in injectors, but the majority do). It is also possible for a death to be recorded as a DRD even though the drug(s) made no significant contribution to the death, but these cases are likely to be very few for opiate-related deaths. Thus, as a reasonable approximation, we shall count, or attribute, all HRDs (but no methadone-only DRDs) as having occurred in current IDUs. We use the term injecting DRD rate to denote HRDs per 100 current injectors.

To obtain estimates of the number of IDUs we use capture–recapture methods. For closed populations, capture–recapture methods have a long history in both ecological (Otis *et al.*, 1978) and epidemiological (Fienberg, 1972) applications. For an overview of the use of capture–recapture methods in epidemiology see, for example, Hook and Regal (1995) and Chao *et al.* (2001) with recommendations on the use of such methods presented by Hook and Regal (1999, 2000). Within epidemiology, capture–recapture studies involve collating data across a series of different data sources. Each source records all individuals in the target population who were observed by that source. Individuals are uniquely identifiable, which allows the construction of a contingency table wherein each cell entry corresponds to the number of individuals who were observed by each distinct combination of sources. However, there is an unobservable cell corresponding to the number of individuals who belong to the target population but were not observed by any source. Thus, failing to estimate this cell entry can significantly underestimate the true target population size, particularly with difficult-to-reach populations. To estimate the unobservable cell, a model is fitted to the observed data. Capture–recapture studies have been used in a variety of situations including for the estimation of hidden populations (Mastro *et al.*, 1994; Frischer *et al.*, 1993; Beynon *et al.*, 2001; King, Bird, Hay and Hutchinson, 2009) and disease prevalence (Hook *et al.*, 1980; Madigan and York, 1997; Chao *et al.*, 2001). We consider the commonly used log-linear models and apply a Bayesian approach that permits the use of a model-averaged estimate of the target population size, thereby accounting for both parameter and model uncertainty (Madigan and York, 1997; King and Brooks, 2001a).

Additional covariate information is often collected corresponding to individual characteristics, such as gender, location, age and marital status. Individuals with different characteristics may have different propensities to be observed by different combinations of sources (King, Bird, Hay and Hutchinson, 2009). Covariates can be introduced as additional factors within the analysis to account for covariate heterogeneity (Tilling and Sterne, 1999; Tilling *et al.*, 2001). For the nine Government Office regions of England, we adopt a similar approach to King *et al.* (2005) by considering two demographic characteristics, each with two levels: gender and age group (15–34 years and 35–64 years), by which DRDs are also cross-classified. The break between the different age groups is chosen since 35 years and above is one of the preferred age groups for reporting injecting prevalence estimates at the European Monitoring Centre for Drugs and Drug Addiction. This age group also corresponds to the aging of young initiates into England's injector epidemic from the mid- to late 1980s to be in the 35–64 years age group in 2005 and beyond. In addition, we note that there is interest in the HRD rates per 100 current IDUs and the 35 years and above age group can represent 15 or more years of injecting. We do not include the region itself as a discrete covariate but analyse the regional data independently of each other. This permits a direct comparison of important interactions identified for each region. Of particular interest is not only the estimates of IDUs within and across regions, but also injectors' HRD rates. We use expert prior information on the injecting DRD rate, combined with information on the regional number of HRDs, to elicit an informative prior on the total number of injectors. The HRDs are themselves provided across the different covariate levels, permitting the estimation of injecting DRD rates for the different joint covariate levels.

In Section 2 we describe the capture–recapture data and introduce the notation that we use before describing the Bayesian approach that we implement to analyse the data in Section 3. Section 4 presents the results, with particular focus on the number of IDUs and associated injecting DRD rates. We conclude with a discussion in Section 5.

2. Regional data

Data that were used within the capture–recapture analyses were collected nationally across England in the financial year 2005–2006. These data can be disaggregated to the drug action team (DAT) area level. In 2005–2006, there were 149 DATs in England. For each DAT area, the same four sources were used to identify IDUs uniquely from which we can construct a 2^4 contingency table with a single unknown cell. The four sources of data were

(a) probation,
(b) drug intervention programme (DIP) prison assessments,
(c) drug treatment and
(d) DIP community assessments.

To cross-classify individuals between the different sources the following set of common identifiers was used for each contributing source of data:

(i) forename initial,
(ii) surname initial,
(iii) gender and
(iv) date of birth.

For each source, only individuals with all four identifiers known were included. For two recorded cases, if all four identifers were the same (Hay *et al.*, 2009), we assumed that they related to the same person either between different sources or duplicated within a single source. Geographical

information, such as address, postcode sector or district, was used to allocate area of residence. For further information on the sources of data see, for example, Singleton *et al.* (2006) and Skodbo *et al.* (2007) with particular reference to DIPs.

DIPs are a crime reduction initiative which works across different criminal justice bodies (such as police, prison and probation) and drug treatment services. Assessments which record an individual's current drug injecting status are carried out at various points in their journey through the criminal justice system and into treatment, e.g. via drug testing while in police custody. The registers are comprehensive in their recording of clients either because they relate to formal justice processes (probation or DIP) or because they are needed for reimbursement (such as treatment numbers). However, for an individual client to be identified as an IDU it does depend on this characterization being both disclosed and recorded. In England, there has been major investment in DIPs, both in prisons and in the community, with the aim of engaging in assessment and drug treatment those who are involved in the criminal justice system who test positively for opiates or cocaine (for further details see Skodbo *et al.* (2007)). Regions where connections across services are made successfully would be revealed by the same clients tending to feature on more than one source of data and perhaps by lower injecting DRD rates if current injectors are successfully engaged in opiate substitution therapy, of which methadone accounted for 83% in England in 2005 (Strang *et al.*, 2007).

Notationally, we label the four sources S_1, S_2, S_3 and S_4, using the same order as above. We label each cell in the 2^4 contingency table by $\mathbf{k} \in \{0, 1\}^4$, which represents the combination of sources that an individual is observed by. For example, cell $\mathbf{k} = \{0, 1, 0, 0\}$ corresponds to being observed by only source S_2 (DIP prison assessments). This approach permits the individual identifiers of gender and age group to be considered as covariate data. We requested that the observed individuals be cross-classified by gender and age group (15–34 and 35–64 years) which allowed us to receive four 2^4 regional contingency tables, which can be written as a 2^6 contingency table with each cell corresponding to the number of IDUs who are observed by each combination of four sources for each gender $\times$ age group classification. These contingency table data were originally calculated at the DAT area level. However, there is a trade-off between the geographical scale that is used and the amount of information that is contained in the corresponding area-specific data. Regional estimates (and variability) are themselves of interest, yet to retain a reasonable level of information within the contingency tables for fitting models and obtaining estimates we requested for the DAT contingency table data to be aggregated to the nine Government Office region levels that had been used in previous Home Office reports (Singleton *et al.*, 2006). For each of the nine regional contingency tables at the Government Office level, there are four unknown cell entries, corresponding to the number of individuals who were not observed by any of the sources for each gender $\times$ age group classification.

For a given region, we let $\mathbf{n}_{\text{obs}}$ and $\mathbf{n}_{\text{unobs}}$ denote the set of observed and unobserved cell entries respectively, and $\mathbf{n} = \{\mathbf{n}_{\text{obs}}, \mathbf{n}_{\text{unobs}}\}$. The total number of observed individuals is denoted by n. Further, for each individual region, we let $n_{(g,a)}$ denote the observed number of individuals of gender g in age group a and $n_{(g,a):\mathbf{k}}$ the number of individuals of gender g in age group a that belong to cell $\mathbf{k} \in \{0, 1\}^4$ for $g \in \{\text{M}, \text{F}\}$ ($\text{M} \equiv$ male; $\text{F} \equiv$ female) and $a \in \{15\text{–}34, 35\text{–}64\}$. Thus, $n_{(g,a):\mathbf{0}} = n_{(g,a):\{0,0,0,0\}}$ denotes the number of individuals of gender g in age group a who are not observed (i.e. the missing cell for the given cross-classification). Additionally we let $n_{\mathbf{0}} = \Sigma_{g,a}\, n_{(g,a):\{0,0,0,0\}}$ denote the total number of unobserved individuals.

We let $N_{(g,a)}$ denote the total number of individuals of gender g in age group a for $g \in \{\text{M}, \text{F}\}$ and $a \in \{15\text{–}34, 35\text{–}64\}$; and $\mathbf{N} = \{N_{(g,a)} : g \in \{\text{M}, \text{F}\}; a \in \{15\text{–}34, 35\text{–}64\}\}$, so that

Table 1. Number of unique IDUs observed in each region and each cross-classification of gender and age

Region	*Total*	*Males, 15–34 years*	*Females, 15–34 years*	*Males, 35–64 years*	*Females, 35–64 years*
East of England	3408	1574	605	962	267
East Midlands	5717	3365	963	1117	272
London	8198	2687	1062	3492	957
North East	4585	2944	858	643	140
North West	11309	4678	1756	3904	971
South East	5444	2605	940	1498	401
South West	8767	4091	1580	2405	691
West Midlands	6627	3886	1081	1332	328
Yorkshire and Humber	11221	6413	2221	2089	498
England	65276	32243	11066	17442	4525

$$N_{(g,a)} = n_{(g,a)} + n_{(g,a):\mathbf{0}} = \sum_{\boldsymbol{k} \in \{0,1\}^4} n_{(g,a):\mathbf{k}}.$$

We let $N_{\text{tot}} = n + n_{\mathbf{0}} = \Sigma_{g,a} N_{(g,a)}$ denote the total number of IDUs in the given region. To provide a brief summary of the data, we present the observed number of unique individuals identified in each region in Table 1 along with the corresponding number observed for each combination of gender and age (i.e. n and $n_{(g,a)}$ for $g \in \{\text{M}, \text{F}\}$ and $a \in \{15\text{–}34, 35\text{–}64\}$) for each region. Appendix A provides the corresponding contingency tables for each region, but where cells entries between 1 and 4 have been omitted.

3. Analysis

The observed contingency table for each region is analysed independently of all other regions. We consider log-linear models that were initially introduced by Fienberg (1972), where the logarithms of the contingency table cell probabilities are a linear sum of main effects and interaction terms between the sources and/or covariates (and normalized so that the sum of the cell probabilities equals 1). We let $\theta_x^{S_i}$ denote the main effect log-linear terms for source S_i, $i \in \{1, 2, 3, 4\}$, at level $x \in \{0, 1\}$ and θ_b^B the main effect log-linear terms for covariate $B \in \{G, A\}$ ($G \equiv$ gender; $A \equiv$ age) for the different levels (i.e. $b \in \{\text{M}, \text{F}\}$ for $B = G$ and $b \in \{15\text{–}34, 35\text{–}64\}$ for $B = A$). We restrict the set of possible interactions to that of two-way interactions corresponding to source $\times$ source (six in total), source $\times$ covariate (eight in total) and covariate $\times$ covariate (only one) interactions. These interactions remove the independence assumption between the different sources. For example, a two-way interaction between sources S_1 and S_2 implies that being observed by source S_1 (probation) increases or decreases the probability of also being observed by source S_2 (DIP prison assessment), and similarly for all other interactions between sources and/or covariates. Notationally, we let $\theta_{x,y}^{S_i,S_j}$ denote the source $\times$ source interaction between sources S_i and S_j ($i, j \in \{1, 2, 3, 4\}$, $i \neq j$, and $x, y \in \{0, 1\}$); $\theta_{x,b}^{S_i,B}$ the source $\times$ covariate interaction between source $i \in \{1, 2, 3, 4\}$ and covariate $B \in \{G, A\}$ for $x \in \{0, 1\}$ and $b \in \{\text{M}, \text{F}\}$ if $B = A$ and $b \in \{15\text{–}34, 35\text{–}64\}$ if $B = A$; and $\theta_{b,c}^{G,A}$ the covariate $\times$ covariate interaction for $b \in \{\text{M}, \text{F}\}$ and $c \in \{15\text{–}34, 35\text{–}64\}$. For identifiability (and prior consistency; see for example King and Brooks (2001b)), we specify sum-to-zero constraints over the levels of each source or covariate on each of the log-linear terms. For example, we specify $\theta_0^{S_1} + \theta_1^{S_1} = 0$, and similarly

for all other source and covariate main effect terms. Similar constraints are specified on the interaction terms, e.g. $\theta^{S_1,S_2}_{0,x} + \theta^{S_1,S_2}_{1,x} = 0$ for $x = 0, 1$.

We let $p_{(g,a):\mathbf{k}}$ denote the probability that an individual is of gender $g \in \{\mathrm{M}, \mathrm{F}\}$ in age group $a \in \{15\text{–}34,\ 35\text{–}64\}$ and lies in the contingency table cell $\mathbf{k} \in \{0, 1\}^4$ relating to the four data sources. The saturated model (in terms of the presence of all main effect and two-way interaction terms) has log-linear cell probabilities of the form

$$p_{(g,a):\mathbf{k}} \propto \exp\left(\sum_{i=1}^{4} \theta^{S_i}_{k(i)} + \theta^{G}_{g} + \theta^{A}_{a} + \sum_{i=1}^{3} \sum_{j=i+1}^{4} \theta^{S_i,S_j}_{k(i),k(j)} + \sum_{i=1}^{4} \theta^{S_i,G}_{k(i),g} + \sum_{i=1}^{4} \theta^{S_i,A}_{k(i),a} + \theta^{G,A}_{g,a} \right),$$

where $k(i)$ is the ith element of $\mathbf{k}$ (i.e. the value of $\mathbf{k}$ corresponding to source S_i). Notationally, we let the probability of not being observed by any source be denoted by $p_{\mathbf{0}} = \Sigma_{g,a}\, p_{(g,a):\mathbf{0}}$. Submodels are obtained by setting the two-way interactions terms to be equal to 0 for all levels. We let the set of log-linear parameters be denoted by $\boldsymbol{\theta}$. Finally, we note that

$$(\mathbf{n}_{\mathrm{obs}}, n_{\mathbf{0}}) | N_{\mathrm{tot}}, \boldsymbol{\theta} \sim \text{ multinomial}(N_{\mathrm{tot}}, \mathbf{q}_{\mathrm{obs}}),$$

where $\mathbf{q}_{\mathrm{obs}}$ denotes the set of probabilities of the observed cells (i.e. $\mathbf{k} \in \{0, 1\}^4 \backslash \mathbf{0}$ corresponding to being observed by each combination of sources, excluding not being observed by any source) for each gender $\times$ age group cross-classification and probability of not being observed (i.e. $p_{\mathbf{0}}$). For further discussion see, for example, King *et al.* (2005, 2009), who considered similar models in relation to IDUs in Scotland, with region as an additional two-level factor.

3.1. Bayesian approach

We consider a Bayesian approach and analyse the data from each region independently of all other regions so that, without loss of generality, we condition on a given region. For a given log-linear model m (in terms of the log-linear parameters in the model), we let the corresponding set of log-linear parameters be denoted by $\boldsymbol{\theta}_m$. We then form the joint posterior distribution over the set of log-linear parameters and total number of individuals in each gender $\times$ age group cross-classification,

$$\begin{aligned} \pi(N_{\mathrm{tot}}, \boldsymbol{\theta}_m | \mathbf{n}_{\mathrm{obs}}) &\propto f(\mathbf{n}_{\mathrm{obs}}, n_{\mathbf{0}} | N_{\mathrm{tot}}, \boldsymbol{\theta}_m)\, p(N_{\mathrm{tot}}, \boldsymbol{\theta}_m) \\ &\propto \frac{N_{\mathrm{tot}}!}{(N_{\mathrm{tot}} - n)!} \prod_{g \in \{\mathrm{M},\mathrm{F}\}} \prod_{a \in \{15-34,\, 35-64\}} \prod_{\mathbf{k} \in \{0,1\}^4 \backslash \mathbf{0}} p_{(g,a):\mathbf{k}}^{n_{(g,a):\mathbf{k}}} \times p_{\mathbf{0}}^{n_{\mathbf{0}}}\, p(N_{\mathrm{tot}}, \boldsymbol{\theta}_m). \end{aligned}$$

The first terms in the posterior distribution correspond to the multinomial joint probability mass function of the observed cell entries given the total population count and log-linear parameters (and hence cell probabilities) and $p(N_{\mathrm{tot}}, \boldsymbol{\theta}_m) = p(N_{\mathrm{tot}})\, p(\boldsymbol{\theta}_m)$, the prior on the total population count and log-linear parameters that are assumed to be independent of each other. We present an alternative parameterization in Appendix B which may be of particular interest if there is prior information on the gender $\times$ age group total population counts. In particular, the gender $\times$ age group total population counts are specified as model parameters with an associated prior distribution. However, this alternative parameterization does not permit the estimation of the covariate-only log-linear parameters. Within our analysis, the log-linear interaction terms are of particular interest (including the presence or absence of such interactions and if present the sign of the interaction) so that we retain the parameterization that was presented above, but in Appendix B we discuss the implications of using the alternative parameterization.

We do not specify the log-linear model *a priori*, in terms of the log-linear interaction terms that are present in the model, but consider a model discrimination approach. We follow the approach of Madigan and York (1997) and King and Brooks (2001a) and extend the posterior distribution to include the model space. In other words, we treat the model itself as a discrete parameter, given the observed data, and form the joint posterior distribution over both the model and the parameter space, denoted by $\pi(N_{\text{tot}}, \boldsymbol{\theta}_m, m|\mathbf{n}_{\text{obs}})$. The (marginal) posterior model probability for model m, given the data, can be expressed in the form

$$\pi(m|\mathbf{n}_{\text{obs}}) \propto \int_{\boldsymbol{\theta}_m} \sum_{N_{\text{tot}}} \pi(N_{\text{tot}}, \boldsymbol{\theta}_m, m|\mathbf{n}_{\text{obs}})\, \mathrm{d}\boldsymbol{\theta}_m,$$

where the denominator once again ensures that the sum of the posterior distribution over admissible models is equal to 1. In addition, we can also calculate the posterior (model-averaged) distribution of the population sizes, accounting for both parameter and model uncertainty. For example, the posterior model-averaged distribution for the number of IDUs for the total population size is given by

$$\pi(N_{\text{tot}}|\mathbf{n}_{\text{obs}}) = \sum_m \pi(N_{\text{tot}}|\mathbf{n}_{\text{obs}}, m)\pi(m|\mathbf{n}_{\text{obs}}),$$

where $\pi(N_{\text{tot}}|\mathbf{n}_{\text{obs}}, m)$ denotes the marginal posterior distribution for N_{tot} under model m.

Model averaging can also be performed within a classical framework (Buckland *et al.*, 1997; Hook and Regal, 1997). However, identifying the set of models with reasonable support to include can be difficult, particularly for large model spaces. In addition, irrespectively of using a Bayesian or classical approach, care should be taken when providing model-averaged point estimates. In particular, in the case where competing models with large support provide very different estimates of the parameter(s) of interest, the corresponding model-averaged point estimate could lie in an area of little or no support. In a Bayesian framework such circumstances can usually be identified in terms of a multimodal marginal posterior density of the parameter(s), and this may be of interest in itself (see for example King, Morgan, Gimenen and Brooks (2009)). For further general discussion of model averaging, see for example chapter 6.5 of King, Morgan, Gimenen and Brooks (2009) and Burnham and Anderson (2002).

3.2. Prior expert information

External information is available that can be combined with expert prior beliefs to provide an informative prior in relation to the total number of IDUs, N_{tot}. In particular, we have independent data relating to the number of DRDs for each region between 2004 and 2007 and prior beliefs relating to the annual DRD rate for injectors. The totality of DRDs includes those with any combination of heroin or morphine, methadone, cocaine, benzodiazapines and alcohol in their systems at the time of death. We make the following decisions regarding the classification of DRDs as pertaining to current injectors to obtain an estimate of the proportion of injecting DRDs in each region. We assume that current IDUs are only those with any heroin or morphine in their system (irrespectively of any other drugs identified). Note that we do not include methadone-only deaths (i.e. no heroin or morphine identified) in our definition of injecting DRDs since methadone-only DRDs may occur preferentially to those enrolled in opiate substitution therapy or to those for whom methadone was not prescribed. Individuals taking a mixture of methadone and heroin or morphine are already identified within the HRDs. The corresponding mean annual number of HRDs from 2004 to 2007 recorded by year of death (see

Table 2. Mean number of HRDs per year of death for each region by using data from the four calendar years 2004–2007

Region	*Total*	*Males, 15–34 years*	*Females, 15–34 years*	*Males, 35–64 years*	*Females, 35–64 years*
East of England	62.75	21.0	3.75	33.0	5.0
East Midlands	64.75	28.5	4.75	27.0	4.5
London	59.5	22.5	4.25	27.25	5.5
North East	39.25	24.5	3.5	9.5	1.75
North West	124.25	49.0	6.25	56.5	12.5
South East	124.0	52.5	10.0	49.75	11.75
South West	111.75	47.75	10.0	46.75	7.25
West Midlands	79.25	34.25	5.25	33.5	6.25
Yorkshire and Humber	107.25	61.75	7.0	33.25	5.25
England	772.75	341.75	54.75	316.5	59.75

`www.rss.org.uk/policy` 'Registration of deaths in England and Wales') for each region is provided in Table 2.

To form the prior on the total IDU population size, we couple this information with the prior beliefs relating to the annual injecting DRD rate. We specify a symmetric 90% interval for IDUs' annual injecting DRD rate of (0.3%, 1.2%) with a median of 0.6% (this prior was informed by the analysis of Merrall *et al.* (2012) of DRD rate for drug treatment clients in Scotland from 1996–2006 and by Scotland's injectors as analysed by King, Bird, Hay and Hutchinson (2009)). We note that injecting DRD rates are generally higher for older individuals and for males (Merrall *et al.*, 2012; Cornish *et al.*, 2010), and it is possible to consider different prior intervals for different cross-classified groups (using an alternative model reparameterization, as discussed in Appendix B). However, to avoid increased variability of annual estimates, and explicitly to model the gender $\times$ age group interaction, we specify a relatively wide interval for the overall injecting DRD rate.

3.3. Prior distributions

We initially specify priors on the log-linear parameters where we do not have any prior information, before we consider the parameter on which there is some expert prior belief: the IDU population size. We complete the prior specification with the prior model probabilities in terms of the interactions that we present in each model. For each individual region and each possible log-linear model we follow King and Brooks (2001a) and specify a hierarchical $N(\mathbf{0}, \sigma^2 I)$ distribution on the set of log-linear parameters in the model and use the non-informative prior $\sigma^2 \sim \Gamma^{-1}(0.001, 0.001)$. This implies that, given that a two-way interaction is present in the model, there is an equal prior probability that the interaction is positive or negative. See King *et al.* (2005) for an alternative prior specification if there is expert prior belief for a positive or negative interaction.

To represent the expert prior information on the total population size N_{tot}, we specify a log-normal prior independently over models, because the prior information is specified in multiplicative form, which results in a skewed distribution. For example, suppose that for a given region the mean annual number of HRDs is X. We specify a prior on the logarithm of the total number of IDUs for the region to be normally distributed with mean $\log(X/0.6\%)$ (so that the prior median is accurately reflected) and variance 0.1776 (to reflect the specified prior 90% interval).

Finally, we specify a prior over the model space. We define the set of possible models to be those models with a maximum of second-order interaction terms (essentially specifying a prior probability of 0 for all higher order interaction terms). This allows interactions between two sources, two covariates and one covariate and one source. For example, this permits a two-way interaction between DIP community assessment and drug treatment data, so that being observed by DIP community assessment makes it more or less likely to be observed within the drug treatment data. Clearly, the aim of DIP community assessment work is to increase the number of individuals receiving treatment, so a positive interaction is desired. However, as stated above, we specify an equal prior probability on each two-way interaction being positive or negative, given that the interaction is present in the model. Similarly, a two-way interaction between gender and drug treatment data would be interpreted as male IDUs being more or less likely (than females) to be receiving treatment for drug addiction.

Considering two-way interactions only significantly reduces the number of possible hierarchical log-linear models and aims to focus on the most important direct interactions between the different sources and/or covariates and to retain epidemiologically interpretable models without data dredging. Without any strong prior information relating to the two-way interactions that may be present we specify a prior probability of 0.5 that each interaction is present in the model, but we note that the interactions that are identified within the analysis are of direct interest particularly in terms of any relationships between the different criminal justice sources and/or drug treatment agencies. The specified prior induces an equal prior probability for each possible model in the set of plausible models.

To assess the sensitivity of the posterior estimates of the number of IDUs on the above prior specification, we conduct a prior sensitivity analysis (see Section 4.4) and compare the results that are obtained. Firstly, we consider the sensitivity of the posterior with respect to the priors specified on the model parameters, using an uninformative prior specification in the form of uniform priors on the total population count and standard deviation of the log-linear variance term. Secondly, we remove the restriction of considering only two-way interaction terms and allow all possible hierarchical log-linear models, including for example three-way interactions, with each model equally likely (and note that this increases the prior probability of two-way interactions being present in the model).

3.4. (Reversible jump) Markov chain Monte Carlo algorithm

The posterior distribution is defined over both parameter and model space, so we implement a reversible jump Markov chain Monte Carlo (MCMC) algorithm (Green, 1995) since the posterior distribution is multi-dimensional (as the number of parameters differs between models). The advantage of the reversible jump algorithm is that the Markov chain simultaneously explores the parameter and model space. This means that we do not need to fit each possible model individually. Irrespectively of the number of possible models, only a single chain is necessary (though typically, as the model space increases, so does the length of the Markov chain that is needed). Within the algorithm, we use a two-step procedure.

Step 1: conditional on the model, we cycle through each individual parameter in turn and propose to update the parameter by using a Gibbs or Metropolis–Hastings step (note that we also simulate population counts for each gender $\times$ age group cross-classification from the posterior conditional distribution).
Step 2: update the model by using a reversible jump step by adding or removing a log-linear interaction term from the model.

We consider each step in turn.

3.4.1. Step 1: updating the parameters

We update σ^2 by using a Gibbs step, since the posterior conditional distribution is of standard form (i.e. inverse gamma) and a single-update random-walk Metropolis–Hastings algorithm is used for all the other log-linear parameters and total population size. See Brooks (1998) for a general description of these algorithms and King and Brooks (2001a) for the specific application to the log-linear parameters. We note that, not only is the total population size of interest, but also the population sizes for each gender × age group cross-classification. These can be easily obtained within the MCMC algorithm by simply simulating these population sizes from their posterior conditional predictive distribution at each iteration of the Markov chain. In particular, we have that

$$\mathbf{n}_{\text{unobs}} | N_{\text{tot}}, \boldsymbol{\theta}, \mathbf{n}_{\text{obs}} \sim \text{multinomial}(N_{\text{tot}} - n, \mathbf{q}_{\text{unobs}}),$$

where $\mathbf{q}_{\text{unobs}} = \{q_{(g,a):\text{unobs}} : g \in \{\text{M}, \text{F}\}, a \in \{15\text{–}34, 34\text{–}65\}\}$ and

$$q_{(g,a):\text{unobs}} = p_{(g,a):\mathbf{0}} \Big/ \sum_{g,a} p_{(g,a):\mathbf{0}}.$$

In other words $q_{(g,a):\text{unobs}}$ denotes the probability that an individual is of gender g and in age group a given that they are not observed within the study.

3.4.2. Step 2: updating the model

To update the log-linear interaction terms within the model we use a reversible jump step (Green, 1995). For a single reversible jump step, we propose to add or remove a single two-way interaction term (since we consider only models with two-way interactions). We choose each log-linear interaction with equal probability. If the parameter is in the model, we propose to remove the parameter; if it is not in the model, we propose to add the parameter. Suppose that we propose to add a given two-way interaction parameter. We propose a candidate value from a proposal distribution q, which in this case is a normal distribution. The corresponding proposal mean is obtained by using the posterior mean of the given parameter from a pilot MCMC run in the model containing all two-way interactions. The proposal variance is chosen via pilot tuning. The corresponding acceptance probability reduces to the ratio of the likelihood function of the proposed and current parameter values respectively, multiplied by the ratio of the prior density function to the proposal density function for the newly proposed log-linear parameter (the Jacobian is equal to 1). See King and Brooks (2001a) for further details by using an analogous approach and Forster *et al.* (2012) and Papathomas *et al.* (2011) for alternative reversible jump implementations.

For each region, the reversible jump MCMC algorithm is run for a total of 10 million iterations with the first 10% discarded as burn-in. For memory storage purposes the observations are thinned every five iterations. Three independent replications using overdispersed starting points obtained similar results (all with the same interpretation) so we conclude that the algorithm has sufficiently converged. Additionally, using the Brooks–Gelman–Rubin statistic on the missing cell entries provided no evidence for lack of convergence. The mean acceptance probabilities for adding or removing the log-linear terms lay around 1.5% for each of the regions (the mean values ranged from 0.7% to 2.5%). The mean acceptance values are not high, but this is partially explained by many of the log-linear terms having either a very high probability of being present or not being present (Table 3), so removing or adding such terms respectively was largely rejected in the reversible jump chain. In other words, taking into account the number of possible models, there was relatively little uncertainty in the parameters that were in the model.

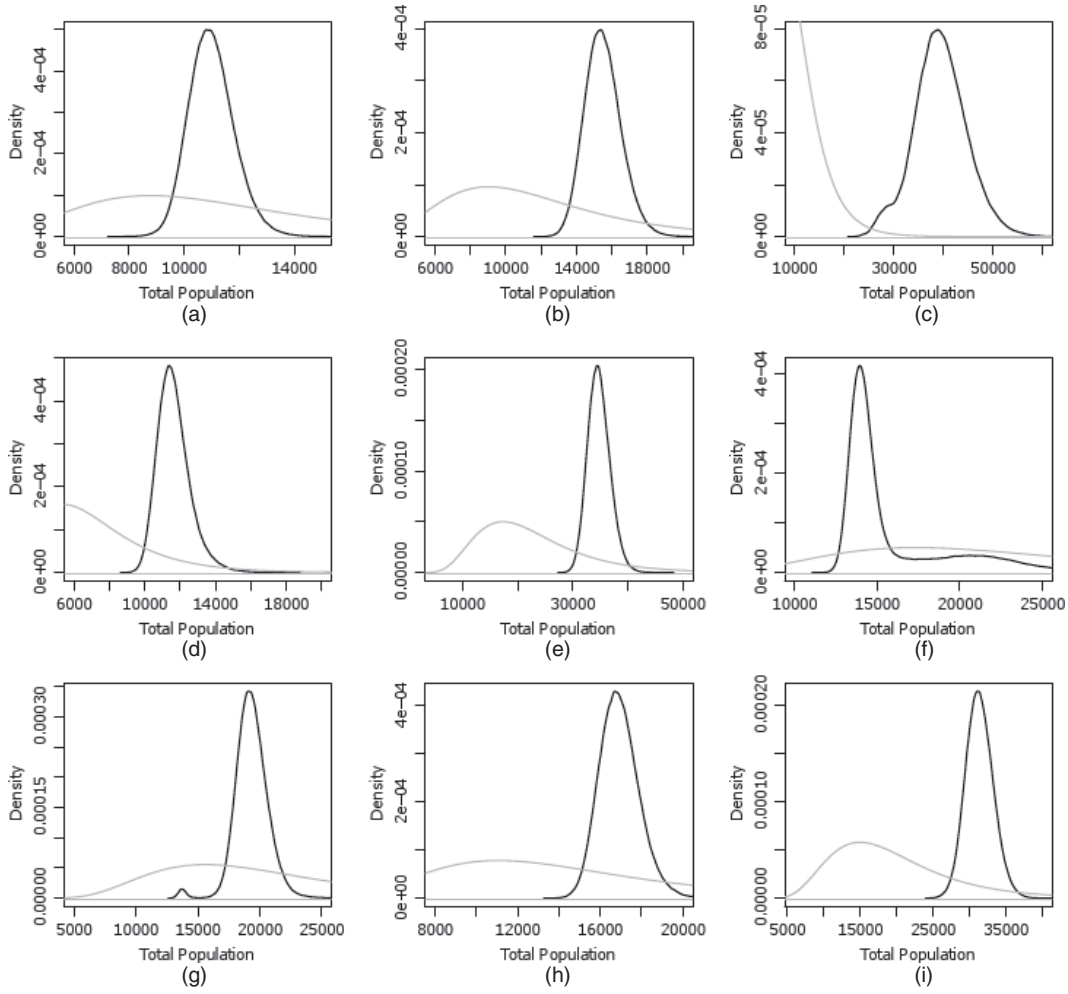

Fig. 1. Posterior distribution for the total population size for each region (———) and the corresponding prior distribution (———): (a) East of England; (b) East Midlands; (c) London; (d) North East; (e) North West; (f) South East; (g) South West; (h) West Midlands; (i) Yorkshire and Humber

4. Results

4.1. Estimating the number of injecting drug users

Fig. 1 provides plots of the prior and (model-averaged) marginal posterior distributions for the number of IDUs in each region. These model-averaged density estimates appear to be unimodal, so models with reasonable posterior support appear to provide similar estimates of population size. For regions the East Midlands, London, the North East, North West, West Midlands and Yorkshire and Humber, the priors generally appear to underestimate the number of IDUs in the regions. The most significant difference between the prior and posterior distributions is clearly for London with virtually no overlap between the prior and posterior distributions. This would potentially suggest, for these regions, and London particularly, that

Table 3. Posterior mean and 95% symmetric credible interval (in parentheses) for the total number of IDUs in each region and each cross-classification of gender and age and aggregated to the England level by using a Monte Carlo approach (rounded to the nearest 10)

Region	*Total*	*Results for the following groups:*			
		Males, 15–34 years	*Females, 15–34 years*	*Males, 35–64 years*	*Females, 35–64 years*
East of England	11000 (9450, 12680)	5120 (4340, 5950)	1680 (1370, 2000)	3420 (2790, 4050)	780 (610, 950)
East Midlands	15490 (13540, 17540)	9030 (7860, 10230)	2280 (1860, 2760)	3460 (2950, 4000)	720 (560, 900)
London	39390 (27870, 50060)	14430 (10050, 18730)	4630 (3090, 6200)	16770 (11770, 21520)	3570 (2430, 4750)
North East	11650 (9940, 13540)	7350 (6250, 8570)	2220 (1810, 2650)	1680 (1290, 2050)	390 (280, 490)
North West	34770 (30920, 38780)	13250 (11810, 14740)	5560 (4890, 6260)	12580 (11120, 14090)	3370 (2930, 3820)
South East	15930 (12550, 23720)	7230 (5690, 10670)	2760 (2120, 4120)	4660 (3570, 7020)	1290 (960, 1960)
South West	19320 (16980, 22040)	8680 (7610, 9860)	3550 (2970, 4140)	5470 (4740, 6310)	1620 (1330, 1910)
West Midlands	16930 (15100, 18850)	9480 (8460, 10540)	2890 (2510, 3270)	3580 (3130, 4040)	990 (820, 1150)
Yorkshire and Humber	31360 (27710, 35110)	17040 (15100, 19060)	6530 (5710, 7370)	6190 (5400, 7000)	1590 (1360, 1830)
England	195840 (181700, 210480)	91610 (85610, 97950)	32100 (29760, 34550)	57810 (52260, 63420)	14320 (12910, 15800)

(a) the number of injecting DRDs is an underestimate and/or
(b) the injecting DRD rate is lower than the prior expert beliefs.

We return to this issue below when discussing the posterior injecting DRD rates.

Table 3 provides the posterior estimates for the total population size and each combination of gender × age group cross-classifications for each of the regions, in addition to the corresponding population sizes for England (i.e. posterior estimates summed over each region). The posterior mean of the total current injector population for England can be easily calculated as the sum of the posterior means of the estimates for each region. However, the corresponding credible intervals at the England level cannot be obtained directly from the credible intervals for each individual region. For example, summing the 2.5% quantiles (which are used for the lower bound of the 95% credible interval) over all regions will not give the corresponding 2.5% quantile for England (the value obtained would be for a much lower quantile for the total population size for England). We can obtain the 95% credible interval at the national level by considering a Monte Carlo approach. Recall that the regional data sets are analysed independently of each other, so the posterior (marginal) distributions of the population sizes are independent across regions. To obtain a sample observation from the posterior distribution of the population size for England, we simply take a sample observation of the number of IDUs from each region and sum these values. By repeatedly sampling from the set of regional posterior distributions for IDU population sizes, we can obtain, for England, a Monte Carlo estimate for the credible intervals of interest.

From Table 3 we see that three regions (London, the North West and Yorkshire and Humber)

appear to have significantly higher absolute numbers of IDUs. In addition, consistently, there is a larger estimated number of males than females in each region for each age group considered. Overall, the posterior mean ratio of males to females (aggregated to the England level) is 3.22 with corresponding 95% symmetric credible interval (3.02, 3.38). The posterior mean male-to-female ratio over the regions ranges from 2.74 (South West) to 4.19 (East Midlands). Capture propensities also appear to differ between regions, in terms of the proportion of individuals who are observed by at least one source. Injectors in London have the least propensity of being observed (posterior mean 0.21 with 95% credible interval (0.16, 0.28)), IDUs in the South West the highest propensity (posterior mean 0.46 with 95% credible interval (0.40, 0.51)). The posterior means for all other regions lay within the range (0.31–0.40).

For comparison with the estimate of the number of IDUs in England in Table 3 by aggregating the posterior regional estimates, we perform a further analysis where we aggregate the raw data across the Government Office regions and analyse the resulting contingency table by using the same Bayesian approach. To analyse these data, we use the same prior beliefs as before, which provide a prior median for the total population size of 128792 with 90% interval (64396, 257583). This lower bound is actually less than the number of observed IDUs (see Table 1). The corresponding posterior mean (rounded to the nearest 10) of the total population size is 209820 with 95% symmetric credible interval (197930, 222200). Thus, the regionally derived England estimate (i.e. obtained by aggregating the posterior regional estimates) is generally lower than that obtained when analysing the data without heed to the regional component (although there is some overlap between the credible intervals). If we consider the corresponding estimates for the cross-classifications when aggregating at the data level we obtain posterior means and 95% symmetric credible intervals (rounded to the nearest 10) for males aged 15–34 years of 96440 (90960, 102100), for females aged 15–34 years of 36940 (34480, 39430), for males aged 35–64 years of 59830 (56190, 63620) and for females aged 35–64 years of 16610 (15380, 17810). The posterior estimates for males are fairly consistent with the regionally derived England estimates in Table 3 (with significant overlap between the credible intervals), but estimates were higher for females. In other words, allowing for heterogeneity at the regional level results in lower estimates for female IDUs.

A previous estimate for England (using the same capture–recapture data but analysed at the DAT level) obtained by Hay *et al.* (2009) is significantly lower, with a point estimate of 129980 and 95% confidence interval (125790, 137030), rounded to the nearest 10. Estimates aggregated at the Government Office regions are also generally smaller (with the exception of the South East). We return to possible reasons for this apparent discrepancy in Section 4.3 when we discuss in detail the interactions that were identified for each of the Government Office regions. Alternatively, using a Bayesian evidence synthesis approach to estimating the prevalence of hepatitis C virus infections, de Angelis *et al.* (2009) provided a posterior median for the current IDU population in 2003 for England and Wales of 217000 with 95% credible interval (157000, 309000), which is broadly consistent with our estimate when taking their inclusion of Wales into account. The same analysis also provided estimates for London and the North West, with posterior means of 38000 and 23000 with 95% credible intervals (30000, 48000) and (14000, 38000) respectively, which again appear to be largely consistent with the estimates that were obtained by using only the capture–recapture data here if only because of wide uncertainty. For example, the estimate by de Angelis *et al.* (2009) for the North West has a relatively much wider credible interval than does ours. Finally, we note that Hickman *et al.* (2004) provided a capture–recapture estimate for London of 34400, for the slightly earlier year of 2001 for those aged 15–44 years.

Table 4 relates the centrally estimated number of current injectors to regions' mid-2005 population aged 15–64 years, since the regions differ in population size. England has an estimated

Table 4. Current injector totals set in context by regions' mid-2005 population aged 15–64 years and estimated ratio of young to old (i.e. 15–34 to 35–64 years) for each gender in each region†

Region	*Mid-2005 population ($\times 10^3$) aged 15–64 years*	*Posterior mean of current injectors (per 1000) population aged 15–64 years*	*Posterior mean of current injectors to nearest 50*	*Posterior mean of male injector ratio by age group (15–34/35–64 years)*	*Posterior mean of female injector ratio by age group (15–34/35–64 years)*
East of England	3604.0	3.1 (2.6, 3.5)	11000 (9450, 12700)	1.50 (1.32, 1.76)	2.15 (1.83, 2.52)
East Midlands	2839.0	5.5 (4.8, 6.2)	15500 (13550, 17550)	2.61 (2.41, 2.82)	3.17 (2.72, 3.58)
London	5269.0	7.5 (5.3, 9.5)	39400 (27850, 50050)	0.86 (0.76, 0.97)	1.30 (1.13, 1.48)
North East	1686.1	6.9 (5.9, 8.0)	11650 (9950, 13550)	4.40 (3.80, 5.47)	5.85 (4.46, 7.50)
North West	4497.0	7.7 (6.9, 8.6)	34750 (30900, 38800)	1.05 (1.01, 1.10)	1.65 (1.54, 1.77)
South East	5338.0	3.0 (2.4, 4.4)	15950 (12550, 23700)	1.56 (1.45, 1.66)	2.15 (1.92, 2.37)
South West	3252.7	5.9 (5.2, 6.8)	19300 (17000, 22050)	1.59 (1.51, 1.67)	2.20 (2.03, 2.38)
West Midlands	3499.9	4.8 (4.3, 5.4)	16950 (15100, 18850)	2.65 (2.46, 2.83)	2.92 (2.58, 3.32)
Yorkshire and Humber	3325.7	9.4 (8.3, 10.6)	31350 (27700, 35100)	2.75 (2.61, 2.90)	4.11 (3.76, 4.48)
England	33311.4	5.9 (5.5, 6.3)	195850 (181700, 210500)	1.59 (1.51, 1.67)	2.24 (2.12, 2.36)

†95% credible intervals are given in parentheses.

5.9 current injectors per 1000 of the population aged 15–64 years (with 95% symmetric credible interval 5.5–6.3). The estimated injector prevalence is low (posterior mean around 3) in the East of England and the South East, high (posterior mean around 7.5) in London, the North East and the North West and very high (posterior mean around 9) for Yorkshire and Humber. However, it is an encouraging sign for London and the North West (with high prevalence rates) that their injector age group ratios (15–34 to 35–64 years) are relatively low compared with England as a whole (posterior mean 1.59 for males and 2.24 for females; and see Millar *et al.* (2006) for further detailed discussion of problem drug use in the North West up to 2001). Regions with high injector ratios by age group may have experienced later diffusion with younger injectors predominating. These regions include the East and West Midlands, North East and Yorkshire and Humber, the last of which is also beset by the largest overall injector prevalence per 1000 of the population aged 15–64 years.

4.2. Injecting drug-related death rates

We obtain a sample from the posterior distribution for the injecting DRD rates by taking the ratio of the mean annual number of HRDs (as provided in Table 2) with the total number of IDUs for each gender × age group cross-classification at each iteration of the Markov chain. The corresponding posterior mean and symmetric 95% credible interval of the injecting DRD rates are provided in Table 5. Recall that the prior 90% interval on the injecting DRD rates was (0.3%, 1.2%). We comment first at the England level and observe that the posterior injecting DRD rate

Table 5. Posterior mean and 95% symmetric credible interval (in parentheses) for the injecting DRD rate, in each region and each cross-classification of gender and age

Region	*Total (%)*	*Results (%) for the following groups:*			
		Males, 15–34 years	*Females, 15–34 years*	*Males, 35–64 years*	*Females, 35–64 years*
East of England	0.57 (0.49, 0.66)	0.41 (0.35, 0.48)	0.23 (0.18, 0.27)	0.97 (0.80, 1.16)	0.65 (0.51, 0.79)
East Midlands	0.42 (0.37, 0.47)	0.32 (0.28, 0.36)	0.21 (0.17, 0.25)	0.79 (0.67, 0.90)	0.63 (0.48, 0.77)
London	0.15 (0.11, 0.20)	0.16 (0.11, 0.21)	0.09 (0.06, 0.13)	0.17 (0.12, 0.22)	0.17 (0.11, 0.21)
North East	0.34 (0.29, 0.39)	0.34 (0.28, 0.39)	0.16 (0.13, 0.19)	0.57 (0.45, 0.72)	0.46 (0.34, 0.60)
North West	0.36 (0.32, 0.40)	0.37 (0.33, 0.41)	0.11 (0.10, 0.13)	0.45 (0.40, 0.50)	0.37 (0.32, 0.42)
South East	0.81 (0.51, 0.97)	0.75 (0.48, 0.90)	0.38 (0.24, 0.46)	1.11 (0.69, 1.35)	0.95 (0.58, 1.18)
South West	0.58 (0.50, 0.65)	0.55 (0.48, 0.62)	0.28 (0.24, 0.33)	0.86 (0.73, 0.97)	0.45 (0.37, 0.53)
West Midlands	0.47 (0.42, 0.52)	0.36 (0.32, 0.40)	0.18 (0.16, 0.21)	0.94 (0.82, 1.06)	0.64 (0.53, 0.75)
Yorkshire and Humber	0.34 (0.30, 0.38)	0.36 (0.32, 0.41)	0.11 (0.09, 0.12)	0.54 (0.47, 0.61)	0.34 (0.28, 0.38)
England	0.40 (0.37, 0.42)	0.37 (0.35, 0.40)	0.17 (0.16, 0.18)	0.55 (0.50, 0.60)	0.42 (0.38, 0.46)

is at the lower end of the prior distribution informed by the Scottish analyses. We note that the overall posterior estimate for the injecting DRD rate is lower than that presented by Bloor *et al.* (2008) who investigated the 'Scottish effect' of higher DRD rates in Scotland compared with England and offered an estimate for Scotland of 0.8% (with 95% uncertainty interval 0.5–1.2% by using data from 2001–2005); our estimated injecting DRD rates for England indeed fall below the lower end of their uncertainty interval. In addition, the injecting DRD rate in England appears to be significantly lower for younger than older injectors: for males, a posterior mean of 0.37% for the younger age group compared with 0.55% for the older age group with non-overlapping credible intervals and likewise for females, with posterior means of 0.17–0.42% for the younger and older age groups. We note that the previous Scottish analysis of King, Bird, Hay and Hutchinson (2009), using data from 2003–2005, estimated significantly higher injecting DRD rates for the cross-classified groups in Scotland but, unlike this analysis, identified a lower female injecting DRD rate only for young injectors with no gender differential for older injectors. For England, more definitively than for Scotland, we observe that older females' injecting DRD rate is also significantly lower than for older males (posterior mean 0.42% *versus* 0.55% with non-overlapping credible intervals). See King, Bird, Hay and Hutchinson (2009) for further details and results relating to the analyses of the Scottish data. Finally, we note that England's overall injecting DRD rate, as defined by us, appears to be similar to the DRD rate of 0.36% that was reported by Merrall *et al.* (2012) for all Scotland's drug treatment clients in the five years to the end of March 2006, although this estimate for Scotland relates to problem drug users who had sought treatment and who included non-injectors.

We now consider the results at the regional level. Comparing the results in Table 5 with the 90% prior interval for injecting DRD rate, it is clear that the London result appears to be the most at odds with these prior beliefs, with the upper 97.5% posterior quantiles of injectors' DRD rates lower than 0.3% (the lower 5% prior quantile) for each gender × age group. Comparing the prior and posterior distributions of numbers of IDUs in Fig. 1 we see very little overlap between these distributions. The significantly higher posterior estimate of the population size (compared with the prior specification) consequently produces the lower estimates of the injecting DRD rates.

For all regions, the lowest injecting DRD rates are for females in the younger age group (15–34 years), with many regions having an injecting DRD rate in the lower 5% quantile of the prior interval. Overall, the females injecting DRD rates are generally lower than for the males. The older age group (35–64 years) has a higher injecting DRD rate for both males and females, relative to the younger age group (15–34 years). This appears to be broadly consistent with other studies showing increased mortality rates for older individuals and males (Cornish *et al.*, 2010; Merrall *et al.*, 2012).

The difference in the injecting DRD rates across the different regions could be a result of

(a) a genuine artefact across the regions,
(b) misestimation of the number of IDUs (i.e. the denominator),
(c) misclassification of the number of injecting DRDs (i.e. the numerator) or
(d) misestimation of the number of IDUs and misclassification of the number of injecting DRDs.

It is not possible to rule out either misestimation or misclassification. As discussed in Section 1 there may be some heterogeneity with respect to the misclassification of injecting DRDs, e.g. in the recording of the presence or absence of heroin or morphine and/or methadone, based on the presence of toxicology or based on whether it was implicated in the death. In addition (except for London as we discuss in Section 4.4) the regional estimates obtained are insensitive to the priors that were specified on the parameters and models, which suggests that significant misestimation is unlikely. Thus, assuming that there are some genuine regional differences in the risk of mortality for IDUs, three regions in Table 5 (the East of England, South East and South West) had particularly high injecting DRD rates, the first two of which (the East of England and South East) can be seen from Table 4 as regions with the lowest prevalence of current injectors per 1000 of the population aged 15–64 years.

4.3. Marginal log-linear probabilities

The corresponding marginal posterior probability that each covariate is present in the model for each separate region is provided in Table 6. Note that we identify 'positive' evidence for a Bayes factor of 3 or greater, for the presence of an interaction, corresponding to a posterior model probability 0.75 or greater, and 'strong' evidence for a Bayes factor of 20 or greater, or posterior probabilities 0.95 or greater (Kass and Raftery, 1995). There are several points of interest. Multiple interactions are clearly important across all (or the majority of) regions, namely $S_1 \times S_2$ (probation × DIP prison data), $S_1 \times S_3$ (probation × drug treatment data), $S_2 \times G$ (DIP prison data × gender) and $S_4 \times A$ (DIP community assessment data × age). For all these interactions, the sign of the interaction is consistent across regions, in particular, a decreased probability of being observed by DIP prison data for females, a decreased probability of being observed by DIP community assessment data for the older age group and positive interactions for $S_1 \times S_2$ and $S_1 \times S_3$, indicating, as we would perhaps expect, an increased probability of being

Table 6. Marginal posterior probability for each two-way interaction being present in the model for each region†

Interaction	*Results for the following regions:*								
	East of England	*East Midlands*	*London*	*North East*	*North West*	*South East*	*South West*	*West Midlands*	*Yorkshire and Humber*
Source × source									
$S_1 \times S_2$	0.996	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
$S_1 \times S_3$	0.998	1.000	0.951	1.000	1.000	1.000	0.988	1.000	1.000
$S_2 \times S_3$	0.068	0.049	0.994	0.160	0.055	0.281	0.076	0.086	0.981
$S_1 \times S_4$	1.000	1.000	1.000	1.000	1.000	0.486	1.000	1.000	1.000
$S_2 \times S_4$	1.000	0.999	1.000	1.000	1.000	0.095	0.690	1.000	1.000
$S_3 \times S_4$	0.060	1.000	1.000	1.000	1.000	0.330	0.988	1.000	1.000
Source × covariate									
$S_1 \times G$	0.050	0.168	0.226	0.061	0.997	0.115	0.997	0.995	0.843
$S_2 \times G$	1.000	1.000	0.833	0.985	1.000	0.998	1.000	1.000	1.000
$S_3 \times G$	0.971	0.959	0.953	0.171	0.033	0.039	0.117	0.060	0.031
$S_4 \times G$	0.119	0.470	0.379	0.957	1.000	0.916	1.000	0.955	1.000
$S_1 \times A$	0.657	1.000	0.109	0.816	1.000	0.981	1.000	0.967	0.986
$S_2 \times A$	0.929	1.000	1.000	0.537	1.000	1.000	1.000	1.000	1.000
$S_3 \times A$	0.117	0.054	0.999	0.262	0.026	0.045	0.062	0.061	0.055
$S_4 \times A$	0.993	1.000	1.000	0.998	1.000	1.000	1.000	0.999	1.000
Covariate × covariate									
$G \times A$	0.997	0.787	1.000	0.867	1.000	0.999	1.000	0.489	1.000

†Recall that $S_1 \equiv$ probation, $S_2 \equiv$ DIP prison assessment data, $S_3 \equiv$ drug treatment data, $S_4 \equiv$ DIP community assessment data, $G \equiv$ gender and $A \equiv$ age.

observed within DIP prison data and drug treatment data if an individual is observed within probation (individuals released from prison are often placed on probation and drug treatment can be a requirement of probation).

Similarly there is a set of interactions, each of which is identified in a majority of regions. These are $S_1 \times S_4$ (probation × DIP community assessment data for all regions except the South East), $S_2 \times S_4$ (DIP prison data × DIP community assessment data for all regions except the South East and South West, though there is posterior uncertainty in the South West), $S_3 \times S_4$ (drug treatment data × DIP community assessment data for all regions except the East of England and the South East), $S_1 \times A$ (probation × age for all regions except the East of England and London), $S_2 \times A$ (DIP prison data × age for all regions except the North East), $G \times A$ (gender × age identified in all regions except the West Midlands). Once more, for the regions where the interaction is identified the sign of the interaction is consistent. We note that the positive interaction $S_3 \times S_4$ (treatment data × DIP community assessment data), as identified in all regions except the East of England and the South East, is a highly desired cross-linkage via increased uptake of drug treatment for individuals in DIP community assessment programmes.

There are some further discrepancies over the different regions regarding the presence of particular interactions. These include the following.

(a) London and Yorkshire and Humber are the only regions to identify the interaction $S_2 \times S_3$ (DIP prison assessment data × drug treatment data), despite large investment in the DIP initiative to lead to increased drug treatment. As we would expect, when this interaction is identified, it is positive. The lack of identification of this interaction is disappointing for

other areas, as there does not appear to be the prison–drug treatment centre links made that are intended.

(b) London is the only region that identifies the interaction $S_3 \times A$, with older individuals more likely to be observed by the treatment data. However, for this region, the interaction $S_1 \times A$ is not identified whereby, in other regions, fewer younger individuals are identified via source S_1 (probation).

(c) The East of England, East Midlands and London identify an interaction $S_3 \times G$ (drug treatment data $\times$ gender) but no interaction between $S_4 \times G$ (DIP community assessment data $\times$ gender). For these regions, there is an increased probability for females to be observed within drug treatment agencies but no support for the interaction $S_4 \times G$, identified by all other regions, wherein there is a decreased probability of females being observed in DIP community assessment programmes.

(d) The North West, South West, West Midlands and Yorkshire and Humber are the only regions to provide positive support for the interaction $S_1 \times G$ (probation $\times$ gender) with a decreased probability of being observed within probation for females.

Finally, we return to the comparison of results that were obtained within this analysis and those of Hay *et al.* (2009), who considered the same data but analysed at the DAT area level. Within their analyses, they did not include the covariate information and considered only the set of log-linear models with a maximum of two source $\times$ source interactions (a total of 22 models). Typically, the model with lowest Akaike information criterion value was chosen (although see Hay *et al.* (2009) for more specific details) and the corresponding estimate for the total population was as given by the chosen log-linear model. For all Government Office regions, except the South East, the number of source $\times$ source interactions that were identified in our models typically lies between 4 and 6. Further, all of the source $\times$ source interactions that are identified with large posterior support for each region have a posterior mean that is positive. Thus, not including such interactions (as for eight of the nine regions) results in the decreased estimate of population size obtained by the previous analysis of Hay *et al.* (2009), rather than differences being due to the use of the lower DAT area level data or ignoring the gender and age group covariate information. Conversely, for the South East (where only two source $\times$ source interactions are identified with positive support), Hay *et al.* (2009) provided an overall estimate and 95% confidence interval of 13 270 (10 290, 16 380), which is reasonably consistent with the estimate that is provided in Table 2, with both point estimates contained in the alternative analysis's uncertainty interval but this is not so for any other region.

4.4. Sensitivity analyses

We present two sensitivity analyses. The first considers the prior specification on the parameters, whereas the second considers the prior on the set of possible models. We initially consider the prior specification on the parameters, and the set of models allowing only two-way interaction terms with each possible model equally likely, as in the previous analysis. We specify a uniform prior on the total population size and set the standard deviation of the log-linear terms to be uniform, with a suitable large upper limit (Gelman, 2006). In particular, we set $\sigma \sim U[0, 100]$. The reversible jump MCMC algorithm is run for each of the different regional contingency tables, and the aggregated England data set. The posterior distributions that were obtained for the majority of analyses (all except London) are very similar for those obtained for the previous informative priors (for example, estimated posterior means for the total number of IDUs within 4% of each other and the same interactions are identified as before), suggesting that the posterior distribution in these cases are data driven. For London, larger estimates are

obtained (approximately 19% higher), which suggests that the informative prior specification had some influence on the posterior estimates in this region.

Secondly, we remove the restriction on the models considered within the analysis and allow higher order interactions. We specify an equal prior probability for each possible hierarchical log-linear model, retaining the previous informative priors, and run a reversible jump MCMC algorithm for each of the Government Office regions and for the aggregated England data. We initially discuss the estimates that were obtained at the England level for both the results aggregated by using the regional level data and the results obtained by aggregating at the data level.

The estimated posterior mean of the total number of IDUs (to the nearest 1000) obtained by aggregating the estimates obtained from the analyses of the regional data is 191 000 with 95% posterior credible interval (176000, 210000), which is slightly lower than the previously obtained estimate by using only two-way interactions, but with large overlap between the credible intervals. Analysis of the data aggregated to the England level before model fitting obtains a posterior mean total estimate (to the nearest 1000) of 211000 with 95% symmetric credible interval (198 000, 224 000). This is slightly higher than the estimate that was obtained when considering the regional data but with overlapping credible intervals. We note that there is positive posterior support for only one three-way interaction, $S_1 \times S_2 \times G$ (probation $\times$ DIP prison data $\times$ gender), with females having an increased probability of being observed by both sources.

We now consider the results that were obtained at the regional level. For all the regional level data sets, except for London, the estimated posterior means for the cross-classified (and total) number of IDUs all lie within 10% of the results that were obtained by using only two-way interaction terms (with the majority lying within 5%), although there are some differences with regard to interactions observed to be present. Unsurprisingly, the 95% posterior credible intervals are generally slightly wider, representing the additional model uncertainty. For London, the estimated total number of IDUs is significantly lower, with a posterior mean of 26430 with 95% symmetric credible interval (17 490, 42 830), which still has significant overlap with the previous estimate. The decrease in estimate (approximately 30%) is consistent across the gender $\times$ age cross-classifications and directly leads to the reduced overall estimate for England. The reason for this lower estimate in London appears to be related to the identification of a three-way interaction term between sources $S_2 \times S_3 \times S_4$ (DIP prison $\times$ drug treatment $\times$ DIP community assessment data), but note that this interaction is not identified within any other region. In particular, if an individual is identified by both forms of DIP data (or not observed by either of these sources), they have an increased probability of being observed in drug treatment.

Allowing this interaction to be present removes the presence of the two-way interaction between probation and drug treatment data within London (which is clearly identified as being present in all other regions). Alternatively, for the North West, South West, West Midlands and Yorkshire and Humber regions, the three-way interaction between criminal justice sources $S_1 \times S_2 \times S_4$ (probation $\times$ DIP prison $\times$ DIP community assessment data) was identified, with an increased probability of being observed by all three sources. Further three-way interactions that were identified were $S_1 \times S_4 \times A$ (DIP prison data $\times$ DIP community data $\times$ age) for the South East (for the older age group a reduced probability of being observed by both sources), $S_1 \times S_4 \times G$ (probation $\times$ DIP community data $\times$ gender) for the South West, $S_2 \times S_4 \times G$ (DIP prison data $\times$ DIP community data $\times$ gender) for the West Midlands and $S_1 \times S_2 \times G$ (probation $\times$ DIP prison data $\times$ gender) for Yorkshire and Humber. For all interactions identified, females had an increased probability of being observed by both of the given sources. However, identifying additional higher order interactions for the regions (except London) did not appear

Table 7. Counts for the East of England

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	184	58	89	29
0	1	0	0	85	6	36	*
1	1	0	0	9	*	*	0
0	0	1	0	912	422	673	209
1	0	1	0	98	31	46	13
0	1	1	0	19	5	6	*
1	1	1	0	5	*	*	0
0	0	0	1	166	43	67	8
1	0	0	1	24	7	13	*
0	1	0	1	6	*	*	*
1	1	0	1	*	*	*	0
0	0	1	1	41	21	16	*
1	0	1	1	10	6	6	*
0	1	1	1	9	0	0	0
1	1	1	1	*	0	0	0
Total				1574	605	962	267

Table 8. Counts for the East Midlands

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	299	66	81	14
0	1	0	0	205	8	35	0
1	1	0	0	31	*	0	0
0	0	1	0	1769	651	749	220
1	0	1	0	226	64	51	9
0	1	1	0	60	*	13	0
1	1	1	0	26	*	6	0
0	0	0	1	308	72	80	9
1	0	0	1	37	7	6	0
0	1	0	1	21	0	*	0
1	1	0	1	8	*	*	0
0	0	1	1	267	64	73	16
1	0	1	1	80	21	16	*
0	1	1	1	21	*	*	0
1	1	1	1	7	*	0	*
Total				3365	963	1117	272

to have a significant effect on the estimated population sizes, as noted above, yet did generally improve the goodness of fit to the observed data.

5. Discussion

Estimating the number of IDUs and the injecting DRD rate is an inherently difficult problem as injecting is an ostracized behaviour and yet injectors have a clear social and economic effect within society. The use of data from capture–recapture studies for estimating such hidden populations has a long history. The use of log-linear models is appealing because of their direct

Table 9. Counts for London

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	121	30	125	20
0	1	0	0	127	33	90	6
1	1	0	0	7	*	5	0
0	0	1	0	1554	752	2582	789
1	0	1	0	28	14	44	5
0	1	1	0	33	6	23	8
1	1	1	0	8	*	0	0
0	0	0	1	557	144	397	84
1	0	0	1	16	5	12	*
0	1	0	1	21	5	11	*
1	1	0	1	*	*	*	0
0	0	1	1	184	62	171	37
1	0	1	1	8	*	11	*
0	1	1	1	20	*	19	*
1	1	1	1	*	*	*	*
Total				2687	1062	3492	957

Table 10. Counts for the North East

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	228	74	47	7
0	1	0	0	135	17	21	*
1	1	0	0	18	*	*	0
0	0	1	0	1778	584	465	122
1	0	1	0	242	63	35	5
0	1	1	0	55	12	9	0
1	1	1	0	30	9	*	0
0	0	0	1	189	35	17	0
1	0	0	1	24	*	*	0
0	1	0	1	10	*	0	0
1	1	0	1	5	*	*	0
0	0	1	1	145	35	21	*
1	0	1	1	60	17	12	*
0	1	1	1	13	*	5	0
1	1	1	1	12	*	0	0
Total				2944	858	643	140

modelling (and interpretation) of interactions between the different data sources and/or covariates which are likely to be present within such complex systems. The corresponding estimates of IDU prevalence are model dependent. We implement a model averaging approach to take into account both parameter and model uncertainty within the estimation of population size, although there can still be dependence on the set of possible models considered.

The estimated total IDU population size in England of approximately 200 000 in 2005–2006 is broadly consistent with the previous estimate that was obtained by de Angelis *et al.* (2009) when investigating the prevalence of hepatitis C but has considerably less uncertainty associated with

Table 11. Counts for the North West

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	390	113	238	58
0	1	0	0	167	25	73	11
1	1	0	0	21	*	6	0
0	0	1	0	2736	1231	2797	775
1	0	1	0	229	85	166	25
0	1	1	0	52	7	21	*
1	1	1	0	20	5	*	*
0	0	0	1	457	144	272	54
1	0	0	1	63	13	43	9
0	1	0	1	31	5	9	*
1	1	0	1	9	*	*	0
0	0	1	1	323	86	215	25
1	0	1	1	139	32	47	8
0	1	1	1	33	5	7	*
1	1	1	1	8	*	*	0
Total				4678	1756	3904	971

Table 12. Counts for the South East

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	246	99	123	30
0	1	0	0	96	24	31	*
1	1	0	0	18	*	*	*
0	0	1	0	1609	648	1101	322
1	0	1	0	160	45	83	14
0	1	1	0	43	*	13	*
1	1	1	0	17	*	*	0
0	0	0	1	267	69	94	17
1	0	0	1	11	7	5	*
0	1	0	1	*	*	0	0
1	1	0	1	*	0	0	0
0	0	1	1	107	33	36	8
1	0	1	1	18	5	8	*
0	1	1	1	5	*	0	0
1	1	1	1	*	0	0	0
Total				2605	940	1498	401

it. Our analysis also provides a regional dimension, offering new insights into injecting DRD rates regionally, and to regional interactions between sources.

Providing regional cross-classified estimates of IDUs and injecting DRD rates gives more detailed information that may be useful in assessing the regional effect of opiate substitution therapy in reducing the risk of mortality. In addition, the risk of transmission for blood-borne viruses may be better assessed for different cross-categories at the regional level, by providing estimates of potential carriers. Besides providing regional estimates of IDUs, regional differences in terms of the underlying interactions may be of interest because they provide insight into cross-linkages between the different sources of data and/or covariates.

Table 13. Counts for the South West

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	146	35	61	6
0	1	0	0	120	26	46	*
1	1	0	0	16	0	*	0
0	0	1	0	3151	1377	2075	659
1	0	1	0	206	63	66	11
0	1	1	0	90	13	18	0
1	1	1	0	24	*	5	0
0	0	0	1	120	21	45	*
1	0	0	1	12	*	5	0
0	1	0	1	*	0	*	0
1	1	0	1	0	0	0	0
0	0	1	1	159	30	73	6
1	0	1	1	29	9	5	*
0	1	1	1	13	*	*	*
1	1	1	1	*	0	0	0
Total				4091	1580	2405	691

Table 14. Counts for the West Midlands

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	277	69	77	22
0	1	0	0	225	24	43	5
1	1	0	0	34	*	7	0
0	0	1	0	2252	780	958	259
1	0	1	0	239	36	58	11
0	1	1	0	92	9	10	*
1	1	1	0	53	6	6	*
0	0	0	1	312	67	68	11
1	0	0	1	44	7	15	0
0	1	0	1	37	5	*	*
1	1	0	1	10	*	*	*
0	0	1	1	199	45	59	13
1	0	1	1	68	15	21	*
0	1	1	1	23	8	*	0
1	1	1	1	21	*	*	0
Total				3886	1081	1332	328

Owing to the structure of the data, which allows for age group and gender as covariates, population estimates can be obtained at these lower cross-classification levels within each region and permit the identification of more complex underlying structure and/or patterns. For example, for both male and female IDUs in the North East, an unusually high proportion are younger individuals (15–34 years) with posterior means for the ratio of younger to older IDUs greater than 4 (see Table 4). We also note that, consistently within each region, and aggregated to the England level, the younger-to-older ratio is higher for females than for males, indicating that a larger proportion of younger IDUs are female than of older IDUs. The higher proportion of younger female IDUs was also observed by King, Bird, Hay and Hutchinson (2009) within

Table 15. Counts for Yorkshire and Humber

S_1	S_2	S_3	S_4	*Males, young*	*Females, young*	*Males, old*	*Females, old*
1	0	0	0	372	139	133	24
0	1	0	0	196	23	43	*
1	1	0	0	18	9	*	0
0	0	1	0	3604	1481	1388	378
1	0	1	0	364	109	95	20
0	1	1	0	99	13	14	*
1	1	1	0	36	7	*	*
0	0	0	1	676	193	174	28
1	0	0	1	100	16	27	6
0	1	0	1	41	9	5	*
1	1	0	1	10	*	0	0
0	0	1	1	538	150	144	28
1	0	1	1	281	58	46	5
0	1	1	1	52	*	8	*
1	1	1	1	26	9	7	0
Total				6413	2221	2089	498

Scotland. We note that, in this analysis, the capture–recapture data for each region were assumed to be independent of each other. It is possible to consider a single integrated analysis with the region itself as a categorical covariate within the analysis, with each level of the covariate corresponding to each region, and once more allowing interactions between the different sources or covariates and region. This may potentially allow the borrowing of information across the regions and is an area of current research.

The estimates of IDU prevalence can be combined with the number of injecting DRDs to obtain the injectors' drug-related risk of mortality. Within our analysis, we take the number of injecting DRDs to be the average annual number of HRDs in each region which occurred over the 4-year period 2004–2007. There is additional potential heterogeneity in terms of identifying and reporting illicit drugs in post-mortem examinations. The numbers of HRDs are used both for constructing a prior for the total population size and in calculating the injecting DRD rate, by combining the number of HRDs with the estimated number of IDUs. It is possible to consider adding a further level of uncertainty to the number of HRDs per region. This would widen the prior interval specified on the total population size but would have little effect on the posterior estimates of prevalence of IDUs since the posterior distributions are largely data driven (although doing so may create greater overlaps between the prior and posterior estimates of population size). Additionally, for each region, assuming a Poisson or negative binomial distribution, say, for the annual number of HRDs (with mean equal to the annual mean number of HRDs) would result in essentially the same posterior mean for the injecting DRD rate (assuming that the posterior distributions for the total population sizes are unchanged), but with a wider credible interval to reflect the additional level of uncertainty that is incorporated.

The generally lower injecting DRD rates for England than in Scotland (King, Bird, Hay and Hutchinson, 2009) suggests that Scotland may have something to learn from the cross-linkages that England has put in place. Discussion of regional source × source interactions with regions' criminal justice or drug treatment practitioners may shed further light on regional implications when local expertise is brought to bear on their interpretation. This analysis appears to offer a broad reassurance that criminal justice and drug treatment interventions are working together.

However, there are concerns also—particularly for those regions in which injector ratios by age group (15–34 to 35–64 years) are high and thereby suggest an unwelcome preponderance of younger injectors, which means that greater resistance to injecting needs to be engendered in their young people.

Acknowledgements

We thank the Home Office and the research team at the Centre for Drug Misuse Research, University of Glasgow, for providing the requested cross-count contingency table data, Claudia Wells at the Office for National Statistics for the requested number of DRDs by year of death, toxicology and cross-classified by region, sex and age group, and Tim Millar for some useful discussions. We also thank two referees and the Associate Editor for providing very insightful comments on a previous version of the manuscript, including the suggestion to investigate higher order interactions. SMB is funded by the Medical Research Council (programme: MC_U_A030_0007/01), AMO was funded by, and RK partly funded by, the Medical Research Council Addiction Research Cluster, 'Nationally integrated quantitative understanding of addiction harm').

Appendix A: Data

The observed contingency tables for each region, cross-classified by gender and age, are provided in Tables 7–15. The four sources correspond to S_1, probation data, S_2, DIP prison assessments, S_3, drug treatment data, and S_4, DIP community assessments.

Small observed cell sizes (i.e. cell entries with values 1–4) have been replaced by asterisks to comply with the Home Office request relating to avoiding potential deductive disclosure.

Appendx B: Alternative model parameterization

An alternative parameterization of the model specifies the total cell counts for each gender × age group as explicit model parameters. Such a parameterization may be desirable if expert prior information is available at this level and the data themselves are not sufficiently informative. The corresponding log-linear parameters are the main effect terms for each source and two-way interactions for source × source and source × covariate combinations (the covariate-only log-linear parameters relating to main effect covariate terms and covariate × covariate interaction term are no longer strictly estimable). We let $\mathbf{N} = \{N_{(g,\,a)} : g \in \{\mathrm{M}, \mathrm{F}\}, a \in \{15\text{–}34, 35\text{–}64\}\}$ and $p_{\mathbf{k}|(g,\,a)}$ denote the probability that an individual is observed in cell $\mathbf{k} \in \{0, 1\}^4$, conditional on being of gender g in age group a. The saturated log-linear model (up to two-way interactions) for the conditional cell probabilities is given by

$$p_{\mathbf{k}|(g,\,a)} \propto \exp\left(\sum_{i=1}^{4} \theta_{k(i)}^{S_i} + \sum_{i=1}^{3} \sum_{j=i+1}^{4} \theta_{k(i),\,k(j)}^{S_i,\,S_j} + \sum_{i=1}^{4} \theta_{k(i),\,g}^{S_i,\,G} + \sum_{i=1}^{4} \theta_{k(i),\,a}^{S_i,\,A} \right).$$

In addition, letting $\mathbf{n}_{(g,\,a)} = \{\mathbf{n}_{(g,\,a):\mathbf{k}} : \mathbf{k} \in \{0, 1\}^4\}$, for each combination of gender g and age group a,

$$\mathbf{n}_{(g,\,a)} | N_{(g,\,a)}, \boldsymbol{\theta}_m \sim \text{ multinomial}(N_{(g,\,a)}, \mathbf{p}_{(g,\,a)}),$$

where $\mathbf{p}_{(g,\,a)} = \{p_{\mathbf{k}|(g,\,a)} : \mathbf{k} \in \{0, 1\}^4\}$. The posterior distribution of the model parameters is given by

$$\begin{aligned}
\pi(\mathbf{N}, \boldsymbol{\theta}_m | \mathbf{n}_{\mathrm{obs}}) &\propto f(\mathbf{n}_{\mathrm{obs}} | \mathbf{N}, \boldsymbol{\theta}_m)\, p(\mathbf{N}, \boldsymbol{\theta}_m) \\
&\propto \left\{ \prod_{g \in \{\mathrm{M},\,\mathrm{F}\}} \prod_{a \in \{15-34,\,35-64\}} f(\mathbf{n}_{(g,\,a)} | N_{(g,\,a)}, \boldsymbol{\theta}_m) \right\} p(\mathbf{N}, \boldsymbol{\theta}_m) \\
&\propto \left(\prod_{g \in \{\mathrm{M},\,\mathrm{F}\}} \prod_{a \in \{15-34,\,35-64\}} \frac{N_{(g,\,a)}!}{n_{(g,\,a):\mathbf{0}}!} \prod_{\mathbf{k} \in \{0,\,1\}^4} p_{\mathbf{k}|(g,\,a)}^{n_{(g,\,a):\mathbf{k}}} \right) p(\mathbf{N}, \boldsymbol{\theta}_m).
\end{aligned}$$

We note that the priors are specified on the total population size for each gender × age group cross-classification, i.e. $\mathbf{N}$ (and typically independently of $\boldsymbol{\theta}_m$). The corresponding MCMC algorithm would, for example, update each $N_{(g,a)}$ in turn by using a Metropolis–Hastings step (analogous to that for N_{tot} in Section 3.4.1).

We note that prior information may not always be of the form of the total population size of each gender × age group cross-classification, but functions of these. For example, prior information may be available on the total population size, male-to-female ratio (denoted R) and the proportion of males (and females) that are young (denoted by P_1 and P_2 respectively). Prior information of this form can be incorporated in this model parameterization by specifying a prior distribution on N_{tot}, R, P_1 and P_2, denoted by $p(N_{\text{tot}}, R, P_1, P_2)$, and calculating the corresponding prior on the total population counts for each gender × age group cross-classification, denoted by $p(\mathbf{N})$, using a transformation-of-variables argument. For this example,

$$R = N_{(\text{M})}/N_{(\text{F})},$$
$$P_1 = N_{(\text{M, 15–34})}/N_{(\text{M})}$$

and

$$P_2 = N_{(\text{F, 15–34})}/N_{(\text{F})}.$$

Then, we can write

$$p(\mathbf{N}) = p(N_{\text{tot}}, R, P_1, P_2) \left| \frac{\text{d}(N_{\text{tot}}, R, P_1, P_2)}{\text{d}\mathbf{N}} \right|,$$

where the final term corresponds to the determinant of the Jacobian. It is straightforward to show that

$$\left| \frac{\text{d}(N_{\text{tot}}, R, P_1, P_2)}{\text{d}\mathbf{N}} \right| = \begin{vmatrix} 1 & 1 & 1 & 1 \\ \dfrac{1}{N_{(\text{M})}} - \dfrac{N_{(\text{M, 15–34})}}{N_{(\text{M})}^2} & -\dfrac{N_{(\text{M, 15–34})}}{N_{(\text{M})}^2} & 0 & 0 \\ \dfrac{1}{N_{(\text{F})}} & \dfrac{1}{N_{(\text{F})}} & -\dfrac{N_{(\text{M})}}{N_{(\text{F})}^2} & -\dfrac{N_{(\text{M})}}{N_{(\text{F})}^2} \\ 0 & 0 & \dfrac{1}{N_{(\text{F})}} - \dfrac{N_{(\text{F, 15–34})}}{N_{(\text{F})}^2} & -\dfrac{N_{(\text{F, 15–34})}}{N_{(\text{F})}^2} \end{vmatrix} = \frac{N_{\text{tot}}}{N_{(\text{M})} N_{(\text{F})}^3}.$$

In general, such results are easily obtained by using an algebraic computer package, such as MAPLE.

Finally, we note that further parameterizations are possible which may be suitable for different prior information. For example, following on from the prior specification above, if there is expert prior information on only N_{tot} and R, it is possible to express the likelihood of observed data given the total number of males and total number of females by using an analogous approach, conditioning on gender only (instead of gender and age group as above). The log-linear parameter corresponding to the main effect term for gender is no longer estimable, but the rest of the covariate (and source) log-linear parameters are.

References

de Angelis, D., Hickman, M. and Yang, S. Y. (2004) Estimating long-term trends in the incidence and prevalence of opiate use/injecting drug use and the number of former users: back-calculation methods and opiate overdose deaths. *Am. J. Epidem.*, **160**, 994–1004.

de Angelis, D., Sweeting, M., Ades, A., Hickman, M., Hope, V. and Ramsay, M. (2009) An evidence synthesis approach to estimating Hepatitis C prevalence in England and Wales. *Statist. Meth. Med. Res.*, **18**, 361–379.

Beynon, C., Bellis, M. A., Millar, T., Meier, P., Thomson, R. and Jones, K. M. (2001) Hidden need for drug treatment services: measuring levels of problematic drug use in the North West of England. *J. Publ. Hlth Med.*, **23**, 286–291.

Bloor, M., Gannon, M., Hay, G., Jackson, G., Leyland, A. H. and McKeganey, N. (2008) Contribution of problem drug users' deaths to excess mortality in Scotland: secondary analysis of cohort study. *Br. Med. J.*, **337**, article a478.

Brooks, S. P. (1998) Markov chain Monte Carlo method and its application. *Statistician*, **47**, 69–100.

Buckland, S. T., Burnham, K. P. and Augustin, N. H. (1997) Model selection: an integral part of inference. *Biometrics*, **53**, 603–618.

Burnham, K. P. and Anderson, D. R. (2002) *Model Selection and Multimodel Inference*. New York: Springer.

Chao, A., Tsay, P. K., Lin, S., Shau, W. and Chao, D. (2001) The application of capture-recapture models to epidemiological data. *Statist. Med.*, **20**, 2123–2157.

Cornish, R., Macleod, J., Strang, J., Vickerman, P. and Hickman, M. (2010) Risk of death during and after opiate substitution treatment in primary care: prospective observational study in UK General Practice Research Database. *Br. Med. J.*, **341**, article c5475.

Fienberg, S. E. (1972) The multiple recapture census for closed populations and incomplete 2^k contingency tables. *Biometrika*, **59**, 591–603.

Forster, J. J., Gill, R. C. and Overstall, A. M. (2012) Reversible jump methods for generalised linear models and generalised linear mixed models. *Statist. Comput.*, **22**, 107–120.

Frischer, M., Leyland, A., Cormack, R., Goldberg, D. J., Bloor, M., Green, S. T., Taylor, A., Covell, R., McKeganey, N. and Platt, S. (1993) Estimating the population prevalence of injection drug use and infection with human immunodeficiency virus among injection drug users in Glasgow, Scotland. *Am. J. Epidem.*, **138**, 170–181.

Gelman, A. (2006) Prior distributions for variance parameters in hierarchical models. *Baysn Anal.*, **1**, 515–534.

Green, P. J. (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, **82**, 711–732.

Hay, G., Gannon, M., MacDougall, J., Eastwood, C., Williams, K. and Millar, T. (2009) Capture-recapture and anchored prevalence estimation of injecting drug users in England: national and regional estimates. *Statist. Meth. Med. Res.*, **18**, 323–339.

Hickman, M., Higgins, V., Hope, V., Bellis, M., Tilling, K., Walker, A. and Henry, J. (2004) Injecting drug use in Brighton, Liverpool, and London: best estimates of prevalence and coverage of public health indicators. *J. Epidem. Commty Hlth*, **58**, 766–771.

Hook, E. B., Albright, S. G. and Cross, P. K. (1980) Use of Bernoulli census and log-linear methods for estimating the prevalence of spina bifida in livebirths and the completeness of vital record reports in New York state. *Am. J. Epidem.*, **112**, 750–758.

Hook, E. B. and Regal, R. R. (1995) Capture-recapture methods in epidemiology: methods and limitations. *Epidem. Rev.*, **17**, 243–264.

Hook, E. B. and Regal, R. R. (1997) Validity of methods for model selection, weighting for model uncertainty, and small sample adjustment in capture-recapture estimation. *Am. J. Epidem.*, **145**, 1138–1144.

Hook, E. B. and Regal, R. R. (1999) Recommendations for presentation and evaluation of capture-recapture estimates in epidemiology. *J. Clin. Epidem.*, **52**, 917–926.

Hook, E. B. and Regal, R. R. (2000) On the need for a 16th and 17th recommendation for capture-recapture analysis. *J. Clin. Epidem.*, **52**, 1275–1277.

Hutchinson, S. J., Taylor, A., Gruer, L., Barr, C., Mills, C., Elliott, L., Goldberg, D. J., Scott, R. and Gilchrist, G. (2000) One year follow-up of opiate injectors treated with oral methadone in a GP-centred programme. *Addiction*, **95**, 1055–1068.

Kass, R. E. and Raftery, A. E. (1995) Bayes factors. *J. Am. Statist. Ass.*, **90**, 773–793.

King, R., Bird, S. M., Brooks, S. P., Hutchinson, S. J. and Hay, G. (2005) Prior information in behavioural capture-recapture methods: demographic influences on drug injectings' propensity to be listed in data sources and their drug-related mortality. *Am. J. Epidem.*, **162**, 694–703.

King, R., Bird, S. M., Hay, G. and Hutchinson, S. J. (2009) An update on the estimation of the prevalence of injector-drug users in Scotland via capture-recapture methods. *Statist. Meth. Med. Res.*, **18**, 341–359.

King, R. and Brooks, S. P. (2001a) On the Bayesian analysis of population size. *Biometrika*, **88**, 317–336.

King, R. and Brooks, S. P. (2001b) Prior induction in log-linear models for general contingency table analysis. *Ann. Statist.*, **29**, 715–747.

King, R., Morgan, B. J. T., Gimenez, O. and Brooks, S. P. (2009) *Bayesian Analysis for Population Ecology*. Boca Raton: CRC Press.

Madigan, D. and York, J. C. (1997) Bayesian methods for estimation of the size of a closed population. *Biometrika*, **84**, 19–31.

Mastro, T. D., Kitayaporn, D. and Weniger, B. G. (1994) Estimating the number of HIV-infected injection drug users in Bangkok: a capture-recapture method. *Am. J. Publ. Hlth*, **84**, 1094–1099.

Merrall, E. L. C., Bird, S. M. and Hutchinson, S. J. (2012) Mortality of those who attended drug services in Scotland 1996-2006: record linkage study. *Int. J. Drug Poly*, **23**, 24–32.

Millar, T., Gemmel, I., Hay, G., Heller, R. F. and Donmall, M. (2006) How well do trends in incidence of heroin use reflect hypothesised trends in prevalence of problem drug use in the North West of England? *Addn Res. Theor.*, **14**, 537–549.

Otis, D. L., Burnham, K. P., White, G. C. and Anderson, D. R. (1978) Statistical inference from capture data on closed animal populations. *Wldlf Monogr.*, **62**, 1–135.

Papathomas, M., Dellaportas, P. and Vasdekis, V. G. S. (2011) A novel reversible jump algorithm for generalized linear models. *Biometrika*, **98**, 231–236.

Singleton, N., Murray, R. and Tinsley, L. (eds) (2006) Measuring different aspects of problem drug use: methodological developments. *Online Report 16/06*. Home Office, London.

Skodbo, S., Brown, G., Deacon, S., Cooper, A., Hall, A., Millar, T., Smith, J. and Whitham, K. (2007) The Drug Intervention Programme (DIP): addressing drug use and offending through 'Tough Choices'. *Research Report*. Home Office, London.

Strang, J., Hall, W., Hickman, M. and Bird, S. M. (2010) Impact of supervision of methadone consumption on deaths related to methadone overdose (1993-2008): analyses using OD4 index in England and Scotland. *Br. Med. J.*, **341**, article c4851.

Strang, J., Manning, V., Mayet, S., Ridge, G., Best, D. and Sheridan, J. (2007) Does prescribing for opiate addiction change after national guidelines?: methadone and buprenorphine prescribing to opiate addicts by general practitioners and hospital doctors in England, 1995-2005. *Addiction*, **102**, 761–770.

Tilling, K. and Sterne, J. A. C. (1999) Capture-recapture models including covariate effects. *Am. J. Epidem.*, **149**, 392–400.

Tilling, K., Sterne, J. A. C. and Wolfe, C. D. (2001) Estimation of the incidence of stroke using a capture-recapture model including covariates. *Int. J. Epidem.*, **30**, 1351–1359.