Joining and Splitting Models with Markov Melding

Robert J. B. Goudie^{*}, Anne M. Presanis[†], David Lunn[‡], Daniela De Angelis[§], and Lorenz Wernisch[¶]

Abstract. Analysing multiple evidence sources is often feasible only via a modular approach, with separate submodels specified for smaller components of the available evidence. Here we introduce a generic framework that enables fully Bayesian analysis in this setting. We propose a generic method for forming a suitable joint model when *joining* submodels, and a convenient computational algorithm for fitting this joint model in stages, rather than as a single, monolithic model. The approach also enables *splitting* of large joint models into smaller submodels, allowing inference for the original joint model to be conducted via our multi-stage algorithm. We motivate and demonstrate our approach through two examples: joining components of an evidence synthesis of A/H1N1 influenza, and splitting a large ecology model.

 ${\sf Keywords:}$ model integration, Markov combination, Bayesian melding, evidence synthesis.

1 Introduction

The increasing availability of large amounts of diverse types of data in all scientific fields has prompted an explosion in applications of methods that combine multiple sources of evidence using (Bayesian) graphical models (for example, Moran and Clark, 2011; Commenges and Hejblum, 2012; Shubin et al., 2016; Birrell et al., 2016). Such evidence synthesis methods have several advantages (Ades and Sutton, 2006; Welton et al., 2012; Jackson et al., 2015): resulting estimates are typically more precise, due to the increased amount of information; they are consistent with all available knowledge; and the risk of potential biases introduced if estimation relies on a 'best quality' subset is minimised.

However, dealing with joint models of several sources of evidence, including data and expert opinion, may be inferentially imprudent, computationally challenging, or even infeasible. It is often sensible to take a modular approach, where separate submodels are specified for smaller components of the available data, facilitating computation and, importantly, allowing insight into the influence of each submodel on the joint model inference (Green et al., 2003; Liu et al., 2009). These submodels can originate in two ways: either by first specifying submodels that, in a Bayesian framework, should be

^{*}MRC Biostatistics Unit, University of Cambridge, United Kingdom, robert.goudie@mrc-bsu.cam.ac.uk

[†]MRC Biostatistics Unit, University of Cambridge, United Kingdom

[‡]MRC Biostatistics Unit, University of Cambridge, United Kingdom

[§]MRC Biostatistics Unit, University of Cambridge, United Kingdom

 $[\]P{\rm MRC}$ Biostatistics Unit, University of Cambridge, United Kingdom



Figure 1: DAG representation of a joint hierarchical model linking M submodels.

joined in a single model to allow all information and uncertainty to be fully propagated; or as a result of *splitting* an existing joint model.

Formally, consider M probability submodels $p_m(\phi, \psi_m, Y_m), m = 1, \dots, M$, for submodel-specific multivariate parameters ψ_m and observable random variables Y_m , as well as a multivariate parameter ϕ common to all submodels that acts as a 'link' between the submodels. The problem is then to *join* the submodels into a single model $p_{\text{comb}}(\phi, \psi_1, \dots, \psi_M, Y_1, \dots, Y_M)$ so that the posterior distributions for the link parameter ϕ and the submodel-specific parameters ψ_m account for all observations and uncertainty. A suitable joint model for a collection of submodels naturally arises in some contexts from standard model constructs, such as a hierarchical model (Figure 1). However, it is not immediately clear how to form such a joint model when either: the submodels are not expressed in a form conditional upon the link parameter ϕ , particularly if the link parameter is a non-invertible deterministic function of the other parameters; or the prior marginal distributions $p_m(\phi), m = 1, \ldots, M$, for the link parameter ϕ differ in the submodels. In applied research, convenient approximate two-stage approaches have been widely used, where one submodel is fitted and an approximation of the resulting posterior is provided to a second submodel (Jackson et al., 2009; Presanis et al., 2014). However, the joint model that is implied by such an approach is unclear (Eddy et al., 1992; Ades and Sutton, 2006).

Conversely, suppose a joint model $p(\phi, \psi_1, \ldots, \psi_M, Y_1, \ldots, Y_M)$ exists that needs splitting into M submodels $p_m(\phi, \psi_m, Y_m)$, $m = 1, \ldots, M$. The submodels should be faithful to the original model in the sense that joining the submodels results in the original model. In some contexts, suitable submodels arise naturally from the structure of the joint model, resulting in splitting strategies used implicitly in the context of hierarchical models (Lunn et al., 2013a; Tom et al., 2010; Liang and Weiss, 2007) and of tall data (Scott et al., 2016; Neiswanger et al., 2014). However, neither the general conditions stipulating when splitting is permissible nor a general framework for splitting a model are immediately clear.

In this paper we introduce *Markov melding*, a simple, generic approach for joining and splitting models that clarifies and generalises various proposals made in the literature under the umbrella of one theoretical framework. Markov melding builds on the theory of *Markov combination* (Dawid and Lauritzen, 1993) and super Markov combination (Massa and Lauritzen, 2010; Massa and Riccomagno, 2017) and combines it with ideas from *Bayesian melding* (Poole and Raftery, 2000), enabling evidence synthesis (Ades and Sutton, 2006; Welton et al., 2012) and model expansion (Draper, 1995) in realistic applied settings. Markov combination is a framework for combining submodels when the prior marginal distributions $p_m(\phi)$, $m = 1, \ldots, M$, are identical. If the prior marginal distributions differ, the super Markov combination family of combined models can be formed, each member of which adopts any one of the prior marginal distributions $p_m(\phi)$ in a combined model. However, in applied contexts, discarding all but one of the prior marginal distributions would usually be unjustifiable, so we propose instead to adopt a *pooled* prior marginal distribution that reflects all prior marginal distributions. Markov melding also accounts for contexts where the link parameter ϕ is a non-invertible deterministic function of other parameters in a submodel. When joining submodels, Markov melding aims to preserve the original submodels as faithfully as possible, and, in particular, always preserves the submodel-specific conditional distributions $p_m(\psi_m, Y_m \mid \phi)$ for all m. Note that, while Markov melding is defined for any collection of submodels, the results may be misleading if any evidence components (priors, submodels and data) strongly conflict (Presanis et al., 2013; Gåsemvr and Natvig, 2009). Such conflict should be investigated and resolved, for example through bias modelling (Turner et al., 2009), before proceeding with the synthesis. In terms of splitting, the Markov melding framework proposed here clarifies the conditions required and the general framework in which to conduct model splitting, facilitating the modular approach advocated above. Notably, we generalise existing tall data splitting approaches (Scott et al., 2016; Neiswanger et al., 2014) for independent, identically distributed data to other types of data.

Finally, we also develop an algorithm for fitting the Markov melded model in stages, for both joining and splitting models. This algorithm extends naturally that employed in Lunn et al. (2013a) and is closely related to those proposed in Liang and Weiss (2007) and Tom et al. (2010).

The paper is organised as follows: in Section 2 we introduce some examples motivating this work; Section 3 provides the conceptual framework underlying our approach; inferential and computational aspects of the approach are presented in Section 4; Section 5 gives details and results for the motivating examples; we conclude with a discussion and suggestions for further work in Section 6.

2 Motivating examples

We motivate and demonstrate our framework for joining and splitting models with two examples, for which we provide here a brief high-level outline. We show how Markov melding applies in each case in Section 3.3; and provide full details and results in Section 5. For both examples, as in the rest of the paper, we use directed acyclic graphs (DAGs) to represent the dependence structure between variables in a model (Figures 2 and 3). Each variable in the model is represented by a node with rectangular nodes denoting observed variables and links between the nodes indicating direct dependencies. Stochastic (distributional) dependencies are represented by solid lines and deterministic (logical) relationships by dashed lines. The joint distribution of all nodes is the product of the conditional distributions of each node given its direct parents, and conditional independence relationships can be read from the graph (Lauritzen, 1996).



Figure 2: High-level DAG representations of the influenza submodels. The double circle denotes the (highly) informative prior for χ , reflecting data from the full severity submodel that is omitted here. Detailed DAGs of these submodels are shown in Figure 6.



Figure 3: High-level DAG representations of the ecology models. Detailed DAG representations of these models are shown in Figure 9.

2.1 Joining: A/H1N1 influenza evidence synthesis

Public health responses to influenza outbreaks rely on knowledge of severity: the probability that an infection results in a severe event such as hospitalisation or death. One method to estimate severity is by combining estimates of cumulative numbers of severe events with estimates of cumulative numbers of infections obtained from synthesizing different data sources. This approach is adopted in Presanis et al. (2014) for the A/H1N1 pandemic, where information from intensive care units (ICU) is integrated with several other sources. Figure 2 provides a schematic representation of the submodels used for each evidence component. A crucial ingredient is the cumulative number of ICU admissions for the A/H1N1 strain, χ . A lower bound ϕ for χ is estimable through an immigration-death model governed by transition rates θ , from time-dependent (weekly) prevalence data y on suspected 'flu cases in ICU (Figure 2(a)). The ϕ of Figure 2(a) is a deterministic function (a sum) of latent quantities involving θ and other parameters π^{pos} . Indirect aggregate evidence on χ is also available from a severity submodel (Figure 2(b)), whose complexity is summarised here by an informative prior on χ . The lower bound ϕ is related to χ through a binomial model with probability parameter π^{det} .

The two submodels imply two different prior models for the link quantity ϕ . A further complication is that the deterministic function connecting ϕ to the ICU submodel parameters is a sum of products, which is not invertible, preventing the ICU submodel from being expressed conditional on ϕ . Presanis et al. (2014) therefore transferred information between the two separate submodels via an approximate approach (see Section 4.3 for details). We show in this paper how Markov melding can be used to join the two submodels formally into a single joint model, making all the assumptions involved explicit. We also explain the relationship between our approach and the approximate approach.

2.2 Splitting: large ecology model

As an example of splitting a large DAG model we consider a joint model (Besbeas et al., 2002) for two distinct sources of data about British Lapwings (*Vanellus vanellus*). These data sources are primarily collected to inform different aspects of studies of the birds: census-type data provide a measure of breeding population size, while mark-recapture-recovery data provide estimates of the annual survival probability of the birds via observations of the survival of uniquely marked individuals. These data are related and a joint model allows inference to account for all information available. In the joint Bayesian model of Brooks et al. (2004) (Figure 3(a)), the mark-recapture-recovery data y are modelled in terms of the recovery rate λ , and the survival rate ϕ for birds; and the census data x are modelled in terms of the survival rate ϕ , and the productivity rate γ of adult female birds. The joint model links the data sources using the common survival rate parameter ϕ .

Brooks et al. (2004) considered fitting the census and mark-recapture-recovery models both separately and jointly using standard Markov chain Monte Carlo (MCMC) algorithms, but considering the joint model simultaneously is cumbersome and MCMC convergence is slow. We describe in this paper how, through Markov melding, inference from such a joint model can be carried out in stages after splitting the model into two separate submodels (Figure 3(b)), circumventing the need to directly fit the joint model in a single MCMC procedure. The multi-stage fitting process is more computationally efficient, and gives insight into the contribution of each submodel to the joint model.

3 Conceptual framework

3.1 Joining models

To combine probabilistic models in a principled way, we propose *Markov melding* as an extension of *Markov combination*, which has been introduced by Dawid and Lauritzen (1993) and discussed extensively with generalisations and applications in Massa and Lauritzen (2010) and Massa and Riccomagno (2017).

Let p denote either a probability distribution for discrete random variables or a probability density for continuous variables (we assume such a density exists). In both cases we talk of p interchangeably as a probability or probability distribution and we express conditional probabilities as $p(\psi \mid \phi) = p(\psi, \phi)/p(\phi)$, where $p(\phi) > 0$. We will assume that when conditioning on a variable its distribution has support in the relevant region. For random variables X_1, X_2 and $X_3, X_1 \perp X_2 \mid X_3$ means that X_1 and X_2 are conditionally independent given X_3 .

Markov combination

Dawid and Lauritzen (1993) define the submodels $p_m(\phi, \psi_m, Y_m)$, $m = 1, \ldots, M$, as *consistent* in the link parameter ϕ if the prior marginal distributions $p_m(\phi) = p(\phi)$ are the same for all m. They define the *Markov combination* p_{comb} of M consistent submodels as the joint model

Joining and Splitting Models with Markov Melding

$$p_{\text{comb}}(\phi, \psi_1, \dots, \psi_M, Y_1, \dots, Y_M) = p(\phi) \prod_{m=1}^M p_m(\psi_m, Y_m \mid \phi)$$

$$= \frac{\prod_{m=1}^M p_m(\phi, \psi_m, Y_m)}{p(\phi)^{M-1}}.$$
(1)

By construction, model (1) assumes that the submodels are conditionally-independent: $(\psi_m, Y_m) \perp (\psi_{\ell}, Y_{\ell}) \mid \phi$ for $m \neq \ell$ (see Figure 1). All prior marginal distributions and submodel-specific conditional distributions, given the link parameter, are preserved: $p_{\text{comb}}(\phi, \psi_m, Y_m) = p_m(\phi, \psi_m, Y_m)$ and $p_{\text{comb}}(\psi_m, Y_m \mid \phi) = p_m(\psi_m, Y_m \mid \phi)$ for all m. Furthermore, the model has maximal entropy among the set of distributions with this marginal preservation property, and so can be viewed as the least constrained among such distributions (Massa and Lauritzen, 2010). Only prior marginals are preserved in Markov combinations: the posterior distributions of ϕ and any ψ_{ℓ} under the Markov combination model account for all data $Y_m, m = 1, \ldots, M$, rather than just the submodel-specific data Y_{ℓ} , and are not preserved.

Markov melding

If the submodels are not consistent in their link parameter ϕ , that is, if the prior marginal distributions $p_1(\phi), \ldots, p_M(\phi)$ of the link parameter differ, a Markov combination cannot be formed directly. However, the original submodels $p_m(\phi, \psi_m, Y_m), m = 1, \ldots, M$, can be altered so that the marginals $p_1(\phi), \ldots, p_M(\phi)$ for the link parameter become consistent. This is achieved by a procedure we term *marginal replacement*, where a new model $p_{\text{repl},m}(\phi, \psi_m, Y_m)$ is formed by replacing the marginal distribution $p_m(\phi)$ of ϕ in the original model $p_m(\phi, \psi_m, Y_m)$ by a new marginal distribution $p_{\text{pool}}(\phi)$:

$$p_{\text{repl},m}(\phi,\psi_m,Y_m) = p_m(\psi_m,Y_m \mid \phi) p_{\text{pool}}(\phi)$$
$$= \frac{p_m(\phi,\psi_m,Y_m)}{p_m(\phi)} p_{\text{pool}}(\phi),$$
(2)

where the pooled density $p_{\text{pool}}(\phi) = g(p_1(\phi), \dots, p_M(\phi))$ is a function g of the individual prior marginal densities. Here, and in what follows, we assume that such a pooled density exists, that g has been chosen such that $\int p_{\text{pool}}(\phi) d\phi = 1$, and that p_{pool} reflects an appropriate summary of the individual marginal distributions; we discuss options below.

Since $p_{\text{repl},m}(\phi, \psi_m, Y_m)$, $m = 1, \ldots, M$, are consistent in the link parameter ϕ (that is, they all have the same prior marginal $p_{\text{pool}}(\phi)$), we can form their Markov combination

$$p_{\text{meld}}(\phi, \psi_1, \dots, \psi_M, Y_1, \dots, Y_M) = p_{\text{pool}}(\phi) \prod_{m=1}^M p_{\text{repl},m}(\psi_m, Y_m \mid \phi)$$

$$= p_{\text{pool}}(\phi) \prod_{m=1}^M \frac{p_m(\phi, \psi_m, Y_m)}{p_m(\phi)}.$$
(3)

We term this construction Markov melding of the submodels $p_m(\phi, \psi_m, Y_m)$ with pooled density $p_{\text{pool}}(\phi) = g(p_1(\phi), \dots, p_M(\phi))$, which amounts to applying the Markov combination (1) to submodels satisfying the consistency condition after marginal replacement as in (2). The submodel-specific conditional distributions, given the link parameter, are preserved in the Markov melded model: $p_{meld}(\psi_m, Y_m \mid \phi) = p_m(\psi_m, Y_m \mid \phi)$ for all *m*. However, in contrast to Markov combination, the prior marginal distributions $p_m(\phi, \psi_m, Y_m)$ will, in general, not be preserved in the Markov melded model. Once the new model $p_{meld}(\phi, \psi_1, \ldots, \psi_M, Y_1, \ldots, Y_M)$ has been formed by Markov melding, posterior inference conditioning on the data $Y_1 = y_1, \ldots, Y_M = y_M$ can be performed (see Section 4).

If p_{pool} equals p_m for a specific submodel $m \in \{1, \ldots, M\}$ then, in the terminology of Massa and Lauritzen (2010), marginal replacement (2) forms a member of a maximally extended family and Markov melding (3) forms a member of the super Markov combination family. That is, rather than allowing a *pooled* prior marginal distribution p_{pool} that can reflect all the prior marginal distributions $p_1(\phi), \ldots, p_M(\phi)$, maximally extended families and super Markov combination require the use of a single one of the original prior marginal distributions; see Supplementary Material A (Goudie et al., 2018) for details.

By extending a similar argument used in Poole and Raftery (2000), it can be shown that marginal replacement has the attractive property that $p_{\text{repl},m}$ minimises the Kullback–Leibler divergence D_{KL} of a distribution $q(\phi, \psi, Y)$ to $p_m(\phi, \psi_m, Y_m)$ under the constraint that the marginals on ϕ agree, $q(\phi) = p_{\text{pool}}(\phi)$:

$$p_{\operatorname{repl},m}(\phi,\psi_m,Y_m) = \operatorname{argmin}_q \{ D_{\operatorname{KL}}(q \parallel p_m) \mid q(\phi) = p_{\operatorname{pool}}(\phi) \text{ for all } \phi \}.$$

Marginal replacement can also be interpreted as a generalisation of Bayesian updating in the light of new information. Details are provided in Supplementary Material B.

Markov melding with deterministic variables

Care is required when some of the dependencies in a submodel are deterministic, as in the example in Section 2.1. The considerations are identical to those for Bayesian melding (Poole and Raftery, 2000), where priors on the input and output of deterministic functions are combined. Specifically, assume the k-dimensional link parameter ϕ is deterministically related to a ℓ -dimensional parameter θ , $k \leq \ell$, in a model $p(\phi, \theta, \psi, Y)$. The probability model is effectively given by $p(\theta, \psi, Y)$ and ϕ follows an induced distribution. We assume ϕ is exclusively a deterministic function $\phi(\theta)$ of the parameter θ .

To apply Markov melding, we need to ensure that the prior marginal distribution on ϕ is well defined, and that we can apply marginal replacement to $\phi = \phi(\theta)$. We must assume that $\phi(\theta)$ is an invertible function or, in the case of $k < \ell$, that $\phi(\theta)$ can be expanded into an invertible function $\phi_e(\theta) = (\phi(\theta), t(\theta))$, with a $\ell - k$ dimensional deterministic function $t(\theta)$. We denote the inverse function by $\theta(\phi, t)$. The function ϕ_e induces a probability distribution on (ϕ, t, ψ, Y) which can be represented as

$$p(\phi, t, \psi, Y) = p(\theta(\phi, t), \psi, Y) J_{\theta}(\phi, t),$$

where $J_{\theta}(\phi, t)$ is the Jacobian determinant for the transformation $\theta(\phi, t)$. The marginal distribution on ϕ can now be obtained as $p(\phi) = \int p(\phi, t, \psi, Y) dt d\psi dY$. We show in

Supplementary Material C that $p(\phi)$ is independent of the chosen parametric extension $t(\theta)$ and so is well defined, and that we can apply marginal replacement, as defined by (2), to replace $p(\phi)$ with $p_{\text{pool}}(\phi)$:

$$p_{\rm repl}(\theta, \psi, Y) = \frac{p(\theta, \psi, Y)}{p(\phi(\theta))} p_{\rm pool}(\phi(\theta)).$$
(4)

Markov melding with $p_{repl}(\theta, \psi, Y)$ can now be applied as in (3).

Pooling marginal distributions

The pooling function g determines the prior marginal distributions $p_{meld}(\phi, \psi_m, Y_m)$, which, in general, will not match those in the original submodels. It must, therefore, be chosen subjectively, ensuring that the pooled density $p_{pool}(\phi)$ appropriately represents prior knowledge of the link parameter ϕ . Various standard pooling functions have been suggested in the multiple expert elicitation literature (see, for example, Clemen and Winkler, 1999; O'Hagan et al., 2006). The difference here is that we propose to pool prior marginal distributions of submodels, rather than directly-specified priors. A simple option is *linear pooling*,

$$p_{\text{pool}}(\phi) = \frac{1}{K_{\text{lin}}(w)} \sum_{m=1}^{M} w_m p_m(\phi), \quad K_{\text{lin}}(w) = \int \sum_{m=1}^{M} w_m p_m(\phi) \, d\phi,$$

where $w = (w_1, \ldots, w_M)^{\top}$, with $w_m \ge 0$ to weight the submodel priors. An alternative is log pooling,

$$p_{\text{pool}}(\phi) = \frac{1}{K_{\log}(w)} \prod_{m=1}^{M} p_m(\phi)^{w_m}, \quad K_{\log}(w) = \int \prod_{m=1}^{M} p_m(\phi)^{w_m} d\phi,$$

with $w_m \ge 0$, a logarithmic version of the linear pooling (for reasons why logarithmic pooling might be attractive see Supplementary Material D). A special case of log pooling is product of experts (PoE) pooling (Hinton, 2002) when $w_m = 1$ for all m

$$p_{\text{pool}}(\phi) = \frac{1}{K_{\text{poe}}} \prod_{m=1}^{M} p_m(\phi), \quad K_{\text{poe}} = \int \prod_{m=1}^{M} p_m(\phi) \, d\phi,$$

in which equal weight is given to each submodel prior. A further special case of linear or log pooling is *dictatorial pooling* $p_{\text{pool}}(\phi) = p_m(\phi)$ when one submodel $m \in \{1, \ldots, M\}$ is considered authoritative¹. We shall assume throughout this paper that the weights ware a fixed quantity, chosen subjectively, in contrast to some of the power prior literature (Neuenschwander et al., 2009), where attempts have been made to treat the weight was an unknown parameter.

¹Dictatorial pooling corresponds to left (or right) composition in the terminology of Massa and Lauritzen (2010), and for M = 2 their upper Markov combination is the family comprising two distributions namely the two possible directions for dictatorial pooling.



Figure 4: Pooled densities under PoE, log and linear pooling, with $w_1 = 0.25$, 0.5 and 0.75 (and $w_2 = 1 - w_1$), formed by pooling a N(0,1) density (---) and a N(μ_2, σ_2^2) density (---) with $\mu_2 = 1, 2$, and $\sigma_2 = 0.5, 1$.

Figure 4 shows the pooled density when combining two normal distributions under three different pooling functions with three choices of weights. The PoE approach is arguably the least intuitive pooling function due to the rather concentrated combined distribution implied. However, the required computation is greatly simplified (see Section 4), and so if (and only if) PoE adequately represents prior beliefs then PoE pooling is an attractive option. The choice of pooling function is particularly important when there is some disagreement between the priors, but if there is substantial conflict between submodel priors we do not recommend the use of Markov melding, as mentioned above.

3.2 Splitting models

We may want to split up a larger model, for example, for computational efficiency or to understand the influence of each submodel on the joint model. In this case we want to split a large joint model $p(\phi, \psi_1, \ldots, \psi_M, Y_1, \ldots, Y_M)$ into M submodels $p_m(\phi, \psi_m, Y_m)$, $m = 1, \ldots, M$, in such a way that joining the submodels using Markov melding recovers the original, joint model. If $(\psi_m, Y_m) \perp (\psi_\ell, Y_\ell) \mid \phi$ for $m \neq \ell$ in the original model, then suitable submodels are

$$p_m(\phi, \psi_m, Y_m) = p(\psi_m, Y_m \mid \phi) p_m(\phi), \qquad m = 1, \dots, M,$$

where $p_1(\phi), \ldots, p_M(\phi)$ are new prior marginal distributions. These marginal distributions and the pooling function g can chosen freely to enable efficient computation, provided that the pooled distribution $p_{\text{pool}}(\phi) = g(p_1(\phi), \ldots, p_M(\phi))$ is the same as the



Figure 5: DAG representations of stylised situations where model splitting might be desirable. Splitting the joint model is possible in (a) and (b), but not in (c).

original marginal distribution $p(\phi)$. An obvious choice is $p_m(\phi) = p(\phi)^{1/M}$ with PoE pooling, but there are many other options. For example, PoE pooling is suitable for any factorisation of $p(\phi)$ into M factors.

Note that a splitting strategy based on Markov melding is suitable *only* if $(\psi_m, Y_m) \perp (\psi_\ell, Y_\ell) \mid \phi$ for $m \neq \ell$, that is, conditioning on the link variable ϕ makes the parts that are intended for splitting conditionally independent.

Figure 5 shows a few stylised situations with M = 2 where splitting for computational purposes might be desirable. The joint distributions for all models is $p(\phi, \psi_1, \psi_2, Y_1, Y_2)$. The model in Figure 5(a) can be split into $p_1(\phi, \psi_1, Y_1) = p(\phi, Y_1 | \psi_1)p(\psi_1)$ and $p_2(\phi, \psi_2, Y_2) = p(\psi_2, Y_2 | \phi)p_2(\phi)$, with a new prior distribution $p_2(\phi)$, which could be different and computationally simpler than $p(\phi) = \int p(\phi, Y_1 | \psi_1)p(\psi_1) d\psi_1 dY_1$. Markov melding, with dictatorial pooling $p_{\text{pool}}(\phi) = p_1(\phi)$, results in

$$p_{\text{meld}}(\phi, \psi_1, \psi_2, Y_1, Y_2) = p_1(\phi) \frac{p_1(\phi, \psi_1, Y_1)}{p_1(\phi)} \frac{p_2(\phi, \psi_2, Y_2)}{p_2(\phi)}$$
$$= p(\phi, Y_1 \mid \psi_1) p(\psi_1) p_2(\psi_2, Y_2 \mid \phi) = p(\phi, \psi_1, \psi_2, Y_1, Y_2),$$

leading to the original model, regardless of the choice of $p_2(\phi)$. The case in Figure 5(b) is similar to the example in Section 2.2 (see Section 3.3 for a definition of splitting in this case). Note that in Figures 5(a) and 5(b), the dependencies between the nodes can include deterministic (logical) dependence, provided $(\psi_1, Y_1) \perp (\psi_2, Y_2) \mid \phi$, as usual. The case in Figure 5(c) cannot be split into $p_1(\phi, \psi_1, Y_1)$ and $p_2(\phi, \psi_2, Y_2)$ by Markov melding model splitting because $(\psi_1, Y_1) \not\perp (\psi_2, Y_2) \mid \phi$.

3.3 Markov melding in the motivating examples

Joining: A/H1N1 influenza evidence synthesis

Markov melding involves joining the ICU submodel (Figure 2(a)) with density $p_1(\phi, \theta, \pi^{\text{pos}}, Y)$, where ϕ is a deterministic function of θ and π^{pos} , and the severity submodel (Figure 2(b)) with density $p_2(\phi, \chi, \pi^{\text{det}})$. These submodels involve independentlycollected data from different surveillance systems, and so it is reasonable to assume conditional independence. Replacing the marginal distribution of ϕ with pooled density $p_{\text{pool}}(\phi)$ in the ICU submodel, using (4), and in the severity submodel, using (2), and then applying Markov melding, as in (3), results in

$$p_{\text{meld}}(\phi, \theta, \pi^{\text{pos}}, \chi, \pi^{\text{det}}, Y) = p_{\text{pool}}(\phi) \frac{p_1(\theta, \pi^{\text{pos}}, Y)}{p_1(\phi(\theta, \pi^{\text{pos}}))} \frac{p_2(\phi, \chi, \pi^{\text{det}})}{p_2(\phi)},$$
(5)

where $Y_1 = Y$, $\psi_1 = \{\theta, \pi^{\text{pos}}\}$, $Y_2 = \emptyset$ and $\psi_2 = \{\chi, \pi^{\text{det}}\}$ in the notation of (3).

Splitting: large ecology model

The original, joint model (Figure 3(a)), with density $p(\phi, \lambda, \gamma, Y, X)$, can be split into separate submodels (Figure 3(b)) with densities $p_1(\phi, \lambda, Y)$ and $p_2(\phi, \gamma, X)$. Provided the priors $p_1(\phi)$ and $p_2(\phi)$ in the separate submodels are such that $p_{\text{pool}}(\phi) = g(p_1(\phi), p_2(\phi))$ equals the original marginal distribution $p(\phi)$, for some choice of pooling function g, then Markov melding the submodels recovers the joint model:

$$p_{\text{meld}}(\phi, \lambda, \gamma, Y, X) = p_{\text{pool}}(\phi) p_1(\lambda, Y \mid \phi) p_2(\gamma, X \mid \phi)$$
$$= p(Y \mid \lambda, \phi) p(\lambda) p(X \mid \gamma, \phi) p(\gamma) p(\phi),$$

with $\psi_1 = \lambda$, $Y_1 = Y$, $\psi_2 = \gamma$, $Y_2 = X$ in the notation of (3).

4 Inference and computation

The joint posterior distribution, given data $Y_m = y_m$, m = 1, ..., M, under the Markov melded model in (3) is

$$p_{\text{meld}}(\phi,\psi_1,\ldots,\psi_M \mid y_1,\ldots,y_M) \propto p_{\text{pool}}(\phi) \prod_{m=1}^M \frac{p_m(\phi,\psi_m,y_m)}{p_m(\phi)}.$$
 (6)

The degree of difficulty of inference for this posterior distribution depends on the specification of the submodels. Our focus is settings in which, considered separately, each of the original collection of submodels is amenable to inference by standard Monte Carlo methods (for example, Robert and Casella, 2004).

In Section 4.1 we first consider a standard Metropolis-within-Gibbs sampler, but when the constituent submodels are complex, this sampler may be cumbersome and slow. We thus propose a multi-stage Metropolis-within-Gibbs sampler, in which inference for the full Markov melded model is generated iteratively in stages, starting with standard inference on one of the submodels. This latter sampling scheme enables a convenient modular approach to inference. Both approaches, in general, require the marginal prior densities $p_m(\phi)$ of the link parameter under each submodel, $m = 1, \ldots, M$, which will not usually be analytically tractable. In Section 4.2 we discuss approaches to estimating these densities, although there is no need to estimate them if PoE pooling is chosen. In Section 4.3 we show how approximate approaches, such as those used by Presanis et al. (2014), relate to the Markov melded model.

4.1 Metropolis-Hastings samplers

A general Metropolis-Hastings sampler for the posterior distribution (6) can be constructed in the usual way. Candidate values $(\phi^*, \psi_1^*, \ldots, \psi_M^*)$ for each parameter of the Markov melded model are drawn from a proposal distribution $q(\phi^*, \psi_1^*, \ldots, \psi_M^* \mid \phi, \psi_1, \ldots, \psi_M)$, based on the current values $(\phi, \psi_1, \ldots, \psi_M)$ of the Markov chain. The candidate values are accepted with probability min(1, r), where r is in the form

$$r = \frac{R(\phi^{\star}, \psi_1^{\star}, \dots, \psi_M^{\star}, \phi, \psi_1, \dots, \psi_M)}{R(\phi, \psi_1, \dots, \psi_M, \phi^{\star}, \psi_1^{\star}, \dots, \psi_M^{\star})}$$

where, with common normalising constants cancelled, the *target-to-proposal density ratio* is

$$R(\phi^{\star},\psi_{1}^{\star},\dots,\psi_{M}^{\star},\phi,\psi_{1},\dots,\psi_{M}) = p_{\text{pool}}(\phi^{\star}) \prod_{m=1}^{M} \frac{p_{m}(\phi^{\star},\psi_{m}^{\star},y_{m})}{p_{m}(\phi^{\star})} \times \frac{1}{q(\phi^{\star},\psi_{1}^{\star},\dots,\psi_{M}^{\star} \mid \phi,\psi_{1},\dots,\psi_{M})}.$$
(7)

Metropolis-within-Gibbs sampler

A particular form of the above general sampler is a Metropolis-within-Gibbs sampling scheme (Müller, 1991), in which samples are drawn from the full conditional distribution of each latent parameter ψ_1, \ldots, ψ_M , and then the link parameter ϕ in turn.

Latent parameter updates Markov melding does not introduce any extra complexities in sampling the parameters ψ_m in each submodel $m = 1, \ldots, M$ (conditional on the link parameter ϕ) beyond those inherent to the original submodels, considered separately. Typically, they can be sampled using standard algorithms. For instance, a Metropolis-Hastings algorithm, in which we draw a candidate value ψ_m^* from a proposal distribution $q(\psi_m^* | \psi_m)$ based upon the current value ψ_m , will be feasible whenever the corresponding algorithm is feasible for estimation of the posterior distribution of the m^{th} submodel alone. In this case, since terms involving marginal densities for the link parameter ϕ in (6) cancel, the target-to-proposal density ratio in (7) simplifies to

$$R(\phi,\psi_1,\ldots,\psi_m^\star,\ldots,\psi_M,\phi,\psi_1,\ldots,\psi_M) = p_m(\phi,\psi_m^\star,y_m) \times \frac{1}{q(\psi_m^\star \mid \psi_m)}$$

This target-to-proposal density ratio is identical to that required for a Metropolis-Hastings update for the parameter ψ_m , conditional on the link parameter ϕ , when the m^{th} submodel alone is the target distribution.

Link parameter updates To update the link parameters, a candidate value ϕ^* is drawn from an appropriate proposal distribution $q(\phi^* \mid \phi)$, based upon the current value ϕ , and is accepted according to the target-to-proposal density ratio

$$R(\phi^{\star},\psi_1,\ldots,\psi_M,\phi,\psi_1,\ldots,\psi_M) = p_{\text{pool}}(\phi^{\star})\prod_{m=1}^M \frac{p_m(\phi^{\star},\psi_m,y_m)}{p_m(\phi^{\star})} \times \frac{1}{q(\phi^{\star} \mid \phi)}.$$
 (8)

When the prior marginal distributions $p_m(\phi)$ or $p_{\text{pool}}(\phi)$ are not analytically tractable, we propose to use an approximation $\hat{p}_m(\phi)$ in their place, calculated using the methods described in Section 4.2. Note that, under PoE pooling, the terms involving the marginal distributions for the link parameter ϕ cancel in (8), leaving

$$R(\phi^{\star},\psi_1,\ldots,\psi_M,\phi,\psi_1,\ldots,\psi_M) = \prod_{m=1}^M p_m(\phi^{\star},\psi_m,y_m) \times \frac{1}{q(\phi^{\star} \mid \phi)}$$

removing the need to estimate the marginal prior distribution for the link parameter ϕ .

Multi-stage Metropolis-within-Gibbs sampler

When the constituent submodels are complex, an alternative, multi-stage approach may be computationally preferable to the Metropolis-within-Gibbs sampler. The multi-stage approach generalises the two stage approach in Lunn et al. (2013a). We assume a factorisation of the pooled prior $p_{\text{pool}}(\phi) = \prod_{m=1}^{M} p_{\text{pool},m}(\phi)$. A default factorisation for any pooling function sets $p_{\text{pool},m}(\phi) = p_{\text{pool}}(\phi)^{1/M}$, but there may be more computationallyefficient factorisations. For example, when PoE pooling is used, the factorisation with $p_{\text{pool},m}(\phi) = p_m(\phi)$ is more computationally efficient, as we describe below. The aim then is to sample, iteratively in stages $\ell = 1, \ldots, M$, from

$$p_{\text{meld},\ell}(\phi,\psi_1,\ldots,\psi_\ell \mid y_1,\ldots,y_\ell) \propto \prod_{m=1}^{\ell} \left(\frac{p_m(\phi,\psi_m,y_m)}{p_m(\phi)} p_{\text{pool},m}(\phi) \right).$$
(9)

Since $p_{\text{meld},M}(\phi, \psi_1, \ldots, \psi_M \mid y_1, \ldots, y_M) = p_{\text{meld}}(\phi, \psi_1, \ldots, \psi_M \mid y_1, \ldots, y_M)$, after M stages the samples obtained reflect the posterior distribution (6) of the full Markov melded model. Note that each $p_m(\phi)$ (and thus also $p_{\text{pool}}(\phi)$) can be estimated from the submodels in advance and independently of the following sampling scheme, as we describe in Section 4.2.

Stage 1 We obtain H_1 samples $(\phi^{(h,1)}, \psi_1^{(h,1)})$, $h = 1, \ldots, H_1$, drawn from $p_{\text{meld},1}(\phi, \psi_1 \mid y_1)$. The most appropriate method for obtaining such samples depends on the nature of the submodel $p_1(\phi, \psi_1, Y_1)$; typically, standard Monte Carlo methods, such as MCMC, will be suitable.

Stage ℓ After we have sampled up to stage $\ell - 1$ from (9), we construct a Metropoliswithin-Gibbs sampler for stage ℓ for the parameters $(\phi, \psi_1, \ldots, \psi_\ell)$ given data (y_1, \ldots, y_ℓ) . The parameter ψ_ℓ is updated, conditional on the link parameter ϕ and parameters $\psi_1, \ldots, \psi_{\ell-1}$ using a standard algorithm, such as a Metropolis-Hastings sampler, with a target-to-proposal density ratio

$$R(\phi,\psi_1,\ldots,\psi_\ell^\star,\phi,\psi_1,\ldots,\psi_\ell) = p_\ell(\phi,\psi_\ell^\star,y_\ell) \times \frac{1}{q(\psi_\ell^\star\mid\phi,\psi_\ell)}.$$

We use the samples from stage $\ell-1$ as a proposal distribution when updating the parameters $\psi_1, \ldots, \psi_{\ell-1}$ and the link parameter ϕ . Specifically, we draw an index d uniformly at random from $\{1, \ldots, H_{\ell-1}\}$, and set $(\phi^*, \psi_1^*, \ldots, \psi_{\ell-1}^*) = (\phi^{(d,\ell-1)}, \psi_1^{(d,\ell-1)}, \ldots, \psi_{\ell-1}^{(d,\ell-1)})$, so that

$$(\phi^{(d,\ell-1)},\psi_1^{(d,\ell-1)},\ldots,\psi_{\ell-1}^{(d,\ell-1)}) \sim p_{\mathrm{meld},\ell-1}(\phi,\psi_1,\ldots,\psi_{\ell-1} \mid y_1,\ldots,y_{\ell-1}).$$

The attraction of this particular proposal distribution is the resulting cancellation of likelihood terms for the first $\ell - 1$ submodels in the target-to-proposal density ratio,

$$R(\phi^{\star},\psi_{1}^{\star},\ldots,\psi_{\ell-1}^{\star},\psi_{\ell},\phi,\psi_{1},\ldots,\psi_{\ell-1},\psi_{\ell}) = \frac{p_{\ell}(\phi^{\star},\psi_{\ell},y_{\ell})}{p_{\ell}(\phi^{\star})}p_{\text{pool},\ell}(\phi^{\star}), \quad (10)$$

meaning this update step can be performed quickly. Once sampling in stage ℓ has converged, samples $(\phi^{(h,\ell)}, \psi_1^{(h,\ell)}, \ldots, \psi_\ell^{(h,\ell)})$, $h = 1, \ldots, H_\ell$, are obtained for use in stage $\ell + 1$.

The density ratio (10) does not depend on parameters $\psi_1, \ldots, \psi_{\ell-1}$, so if interest focuses entirely on the parameters (ϕ, ψ_ℓ) then $\psi_1, \ldots, \psi_{\ell-1}$ can be ignored in stage ℓ of the multi-stage sampling: they do not need to be monitored or updated by the sampling algorithm. The multi-stage sampler is nevertheless still sampling from the joint target distribution $p_{\text{meld},\ell}(\phi,\psi_1,\ldots,\psi_\ell \mid y_1,\ldots,y_\ell)$. Stage ℓ is influencing the acceptance or rejection of samples of $p_{\text{meld},\ell-1}(\phi,\psi_1,\ldots,\psi_{\ell-1} \mid y_1,\ldots,y_{\ell-1})$ from the previous stage, thus adjusting this distribution according to the requirements of the joint model.

In general, evaluation of ratio (10) requires estimates of the prior marginal distribution of the link parameter under the ℓ^{th} submodel, which can be obtained as described in Section 4.2. However, if PoE pooling is used and $p_{\text{pool},m}(\phi) = p_m(\phi), m = 1, \ldots, M$, the ratio simplifies to $R(\phi^*, \psi_1^*, \ldots, \psi_{\ell-1}^*, \psi_\ell, \phi, \psi_1, \ldots, \psi_{\ell-1}, \psi_\ell) = p_\ell(\phi^*, \psi_\ell, y_\ell)$, meaning that no estimates of the marginal distribution are required.

4.2 Estimating marginal distributions

The prior marginal densities $p_m(\phi)$ of the link parameter under each of the M submodels are central to Markov melding, and in particular are required to evaluate the acceptance probability of proposals within the MCMC samplers we proposed above. However, these marginals are not generally analytically tractable, except when the prior distribution $p_m(\phi)$ is directly-specified as a standard, tractable distribution, such as when ϕ appears as a founder node in a DAG representation of the submodel. When not available analytically, we can estimate the marginal density $p_m(\phi)$ for each submodel m by kernel density estimation (Henderson and Parmeter, 2015) with samples drawn from $p_m(\phi) = \iint p_m(\phi, \psi_m, Y_m) d\psi_m dY_m$ by standard (forward) Monte Carlo. Care is required if ϕ has high dimension because the curse of dimensionality applies to kernel density estimation; see Section 6 for further discussion.

4.3 Normal two-stage approximation method

Approximate approaches for joining submodels are widely used in applied research. In this section, we show that approximate inference for the Markov melded model formed by joining submodels (with PoE pooling) can be produced using a standard normal approximation approach.

Consider the case when the Markov melded model $p_{\text{meld}}(\phi, \psi_1, \psi_2, Y_1, Y_2)$ is formed by joining M = 2 submodels $p_1(\phi, \psi_1, Y_1)$ and $p_2(\phi, \psi_2, Y_2)$. Suppose that ψ_1 is not a parameter of interest in the posterior distribution, so that it can be integrated over

$$p_{\text{meld}}(\phi, \psi_2, Y_1, Y_2) = \int p_{\text{meld}}(\phi, \psi_1, \psi_2, Y_1, Y_2) \, d\psi_1$$

= $p_{\text{pool}}(\phi) \int p_1(\psi_1, Y_1 \mid \phi) \, d\psi_1 \, p_2(\psi_2, Y_2 \mid \phi)$
= $p_{\text{pool}}(\phi) \, p_1(Y_1 \mid \phi) \, p_2(\psi_2, Y_2 \mid \phi).$ (11)

An approximate two stage sampler that mimics the multi-stage sampler (above) can then be constructed for this marginal distribution.

Stage 1 Fit the submodel $p_1(\phi, \psi_1, Y_1)$ to obtain posterior samples from $p_1(\phi \mid Y_1)$, and approximate the posterior by a (multivariate) normal distribution with mean $\hat{\mu}$ and covariance $\hat{\Sigma}$. With p_N denoting the probability density function for a (multivariate) normal distribution,

$$p_1(\phi \mid Y_1) \approx p_N(\phi \mid \widehat{\mu}, \widehat{\Sigma}) = p_N(\widehat{\mu} \mid \phi, \widehat{\Sigma}).$$

Stage 2 Since $p_1(\phi, Y_1) \propto p_1(\phi \mid Y_1) \approx p_N(\widehat{\mu} \mid \phi, \widehat{\Sigma})$, we obtain an approximation for (11) by replacing $p_1(\phi \mid Y_1)$ by $p_N(\widehat{\mu} \mid \phi, \widehat{\Sigma})$.

$$p_{\text{meld}}(\phi, \psi_2, Y_1, Y_2) \propto p_{\text{pool}}(\phi) \frac{p_1(\phi \mid Y_1)}{p_1(\phi)} \frac{p_2(\phi, \psi_2, Y_2)}{p_2(\phi)} \\ \approx p_{\text{pool}}(\phi) \frac{p_N(\hat{\mu} \mid \phi, \hat{\Sigma})}{p_1(\phi)} \frac{p_2(\phi, \psi_2, Y_2)}{p_2(\phi)}.$$

This two-stage approximate approach is commonly used in practice (see Section 6) in the form $p_{\text{meld}}(\phi, \psi_2, Y_1, Y_2) \approx c p_N(\hat{\mu} \mid \phi, \hat{\Sigma}) p_2(\phi, \psi_2, Y_2)$, with c a data dependent constant. In this case the likelihood of the second submodel is modified by a factor $p_N(\hat{\mu} \mid \phi, \hat{\Sigma})$ (in a DAG representation a dependency of the constant $\hat{\mu}$ on ϕ and the constant $\hat{\Sigma}$ is added). This approach can be viewed as approximate Markov melding with PoE pooling, in which one submodel is represented by a normal approximation.

If, instead of PoE pooling, one wishes to regard the marginal $p_2(\phi)$ on the link variable ϕ as authoritative and thus fully retain it, dictatorial pooling $p_{\text{pool}}(\phi) = p_2(\phi)$ leads to the variant

$$p_{\text{meld}}(\phi, \psi_2, Y_1, Y_2) \propto \frac{p_1(\phi \mid Y_1)}{p_1(\phi)} p_2(\phi, \psi_2, Y_2) \approx \frac{p_N(\phi \mid \widehat{\mu}, \widehat{\Sigma})}{p_N(\phi \mid \widehat{\mu}_0, \widehat{\Sigma}_0)} p_2(\phi, \psi_2, Y_2)$$

$$\propto p_N(\phi \mid \mu_c, \Sigma_c) p_2(\phi, \psi_2, Y_2),$$



Figure 6: DAG representations of the submodels of A/H1N1 influenza. Repeated variables are enclosed by a rounded rectangle, with the label denoting the range of repetition. For simplicity the time domain is suppressed: parameters with subscripts T, U and V are collections of parameters across the time range denoted by the subscript. For example, $y_{a,U} = \{y_{a,t} : t \in U\}$.

where $\hat{\mu}_0$ and $\hat{\Sigma}_0$ are an estimate of the mean and covariance of the prior marginal $p_1(\phi)$, which can be obtained at stage one in parallel to the posterior by sampling from the prior submodel, and

$$\Sigma_c^2 = \left(\widehat{\Sigma}^{-1} - \widehat{\Sigma}_0^{-1}\right)^{-1}, \quad \mu_c = \Sigma_c \left(\widehat{\Sigma}^{-1}\widehat{\mu} - \widehat{\Sigma}_0^{-1}\widehat{\mu}_0\right).$$

Adjusting according to $\hat{\mu}_0$ and $\hat{\Sigma}_0$ removes the prior $p_1(\phi)$ from approximate joint model.

5 Results

5.1 Joining: A/H1N1 influenza evidence synthesis

Figure 6 shows DAG representations of the two submodels outlined in Section 2.1.

ICU submodel

The main data source in the ICU submodel is prevalence-type data from the Department of Health's Winter Watch scheme (Department of Health, 2011), which records the total number of patients in all ICUs in England with suspected pandemic A/H1N1 influenza infection. Weekly observations $y_{a,t}$ taken at days $t \in U = \{8, 15, 22, \ldots, 78\}$ for age group $a \in \{1, 2\}$ (children and adults respectively) are available between December 2010 and February 2011. To estimate the link parameter $\phi = (\phi_a) = (\phi_1, \phi_2)$, that is, a lower bound for the cumulative number of ICU admissions over the period of observation $t \in T = \{1, \ldots, 78\}$, from such prevalence data requires an immigration-death model for the system of admissions and exits from ICU. Assume that new ICU admissions follow an inhomogeneous Poisson process with rate $\lambda_{a,t}$ at time t, and the length of stay in ICU is exponentially distributed with rate μ_a . Then the number of patients admitted up to time t who are still present in ICU at time t follows a thinned inhomogeneous Poisson process and the observed number of prevalent patients is $y_{a,t} \sim \text{Po}(\eta_{a,t})$, $a \in \{1, 2\}$, $t \in U$, with expectation, under a discretised formulation with daily time steps, given by $\eta_{a,t} = \sum_{u=1}^{t} \lambda_{a,u} \exp\{-\mu_a(t-u)\}, t \in T$. We assume $\eta_{a,1} = 0$ to enforce the assumption that no patients with suspected 'flu were in ICU a week before observations began.

The product of the expected new admissions of suspected cases $\lambda_{a,t}$ and the proportion positive for A/H1N1 $\pi_{a,t}^{\text{pos}}$ gives the expected number of *confirmed* new admissions on day t. The link parameter ϕ_a is the non-invertible sum of these products over time:

$$\phi_a = \sum_{t \in T} \pi_{a,t}^{\text{pos}} \lambda_{a,t}, \qquad a = 1, 2.$$

We model the proportion positive $\pi_{a,t}^{\text{pos}}$ using weekly virological positivity data from the sentinel laboratory surveillance system Data Mart (Public Health England, 2014), which records the number $z_{a,v}^{\text{pos}}$ of A/H1N1-positive swabs out of the total number $n_{a,v}^{\text{pos}}$ tested during week $v \in V = \{1, \ldots, 11\}$ in age group $a \in \{1, 2\}$. We assume a uniform prior $\pi_{a,t}^{\text{pos}} \sim \text{Unif}(\omega_{a,v}, 1), t \in T$, for the true positivity, where v = 1 for $t = 1, \ldots, 14$ and $v = \lfloor (t-1)/7 \rfloor$ for $t = 15, \ldots, 78$, and where the lower bound $\omega_{a,v}$ is informed by a binomial model for the positivity data: $z_{a,v}^{\text{pos}} \sim \text{Bin}(n_{a,v}^{\text{pos}}, \omega_{a,v}), v \in V$. For the expected new admissions $\lambda_{a,t}$, we assume a random-walk prior with $\log(\lambda_{a,1}) \sim \text{Unif}(0, 250)$ and $\log(\lambda_{a,t}) \sim N(\log(\lambda_{a,t-1}), \gamma_a^{-2})$ for $t = 2, \ldots, 78$, with $\gamma_a \sim \text{Unif}(0.1, 2.7)$. For the length of ICU stays we assume constant age-group specific exit rates $\mu_1 = \exp(-\alpha)$ and $\mu_2 = \exp(-\{\alpha + \beta\})$, with $\alpha \sim N(2.7058, 0.0788^2)$ and $\beta \sim N(-0.4969, 0.2048^2)$ (Presanis et al., 2014).

Severity submodel

We consider a simplified version of the full, complex severity submodel in Presanis et al. (2014). The Winter Watch ICU data are only available for a portion of the time of the 'third wave' of the A/H1N1 pandemic, and so the cumulative number of confirmed new admissions ϕ_a from the ICU submodel is a lower bound for the true number χ_a of ICU admissions during the third wave. We thus assume $\phi_a \sim \text{Bin}(\chi_a, \pi^{\text{det}}), a \in \{1, 2\}$, where π^{det} is the age-constant detection probability, to which we assign a Beta(6, 4) prior. We incorporate the remaining evidence in the full severity submodel of Presanis et al. (2014) via informative priors $\chi_1 \sim \text{Lognormal}(4.93, 0.17^2)$ and $\chi_2 \sim \text{Lognormal}(7.71, 0.23^2)$.

Markov melded model

We joined the submodels as in (5). We considered linear and log pooling with pooling weight $w_1 = 0.25, 0.5$ and 0.75 (and $w_2 = 1 - w_1$), and PoE pooling.

We estimated the marginal priors for $\phi = (\phi_1, \phi_2)$ under the ICU and severity submodels using kernel density estimation with a bivariate *t*-distribution kernel, using 5×10^4 independent draws, sampled from the corresponding submodel by forward Monte Carlo. The marginal priors are shown in Figure 7(a). Note that the ICU submodel prior



Figure 7: Prior distributions for ϕ_a , the cumulative number of *confirmed* new admissions in age group *a*, in the A/H1N1 influenza evidence synthesis: (a) under the ICU and severity submodels; (b) pooled priors under three pooling functions with $w_1 = w_2 = 0.5$.

for ϕ is extremely flat, whereas the severity submodel prior is concentrated on a small part of the parameter space. The combined density using each of the pooling functions (with $w_1 = w_2 = 0.5$) is shown in Figure 7(b). Linear and PoE pooling in this case lead to similar densities, whereas the log pooling prior is more dispersed.

We then estimated, in stage one, the posterior distribution of the link parameter ϕ under the ICU submodel alone. We drew 5 million iterations from the ICU submodel using JAGS (Plummer, 2015b), retaining every 100th iteration, after discarding 5×10^4 iterations as burn-in. In stage two, for the Markov melded models under linear, log and PoE pooling, we drew 2×10^6 samples using the multi-stage Metropolis-within-Gibbs sampler, with the first 10^4 samples discarded as burn-in.

Figure 8 shows the results. There is a notable reduction in uncertainty in the posteriors from Markov melding compared to the ICU submodel posterior, especially in ϕ_1 , demonstrating the benefit of joining the submodels. In the adult age group (a = 2), the Markov melding results are robust to the choice of pooling function and pooling weight: the likelihood from the ICU submodel dominates over the pooled prior. There is considerable agreement between the various approaches in the child age group (a = 1)as well, although the choice of pooling weight has some influence on the upper tail under log pooling. As anticipated by Section 4.3, the normal approximation (fitted using OpenBUGS) and PoE pooling posteriors are close, due to the near normality of the ICU posterior distribution.

5.2 Splitting: large ecology model

Figure 9 is a DAG representation of the full, joint model outlined in Section 2.2.

Mark-recapture-recovery data

Mark-recapture-recovery data y_{t_1,t_2} record the number of ringed birds released before May in year $t_1 = 1, \ldots, 35$, and recovered (dead) in the 12 months up to April in year



Figure 8: Medians and 95% credible intervals for: the posterior distribution for the link parameters ϕ_1 and ϕ_2 under the ICU submodel; the posterior distribution for each parameter under the severity submodel; the posterior distribution for each parameter according to the normal approximation and the Markov melded model under each pooling function. The *x*-axis shows number of individuals, except for π^{det} , which shows probabilities.

 $t_2 = t_1 + 1, \ldots, 36$. The years correspond to observations for releases from 1963 (t = 1) to 1997 and recoveries from 1964 to 1998. The number of birds $y_{t_1,37}$ released in year t_1 and never recovered is also available. We assume

$$(y_{t_1,t_1+1},\ldots,y_{t_1,37}) \sim \operatorname{Mult}(\pi_{t_1,t_1+1},\ldots,\pi_{t_1,37}), \quad t_1 = 1,\ldots,35.$$

We model the probability π_{t_1,t_2} of recovery in year t_2 following release in year t_1 in terms of the recovery rate λ_t , and the survival rates $\eta_{C,t}$ and $\eta_{A,t}$ for immature (1 year old) and breeding (2 years or older) birds, respectively, up to April of year t:

$$\pi_{t_1,t_2} = \begin{cases} \lambda_{t_2}(1-\eta_{C,t_2}) & t_1 = 1, \dots, 35, t_2 = t_1 + 1, \\ \lambda_{t_2}\eta_{C,t_1+1}(1-\eta_{A,t_2}) & t_1 = 1, \dots, 34, t_2 = t_1 + 2, \\ \lambda_{t_2}\eta_{C,t_1+1}(1-\eta_{A,t_2}) \prod_{u=t_1+2}^{t_2-1} \eta_{A,u} & t_1 = 1, \dots, 33, t_2 = t_1 + 3, \dots, 36. \end{cases}$$



Figure 9: DAG representation of the joint ecology model. The recovery and census submodels are connected via the common parameter $\phi = (\phi_{\alpha,C}, \phi_{\alpha,A}, \phi_{\beta,C}, \phi_{\beta,A})$. For simplicity the time domain is suppressed: $y, x, \pi, \mu, \lambda, \eta_C, \eta_A, \gamma$ and z represent the collection of *all* quantities sharing the same variable name. For example, $\mu = \{\mu_{G,t} : G \in \{C, A\}, t = 3, \ldots, 36\}$.

The recovery rate is the probability that a bird that dies in year t is recovered. The probability of a bird released in year t_1 being never recovered is $\pi_{t_1,37} = 1 - \sum_{u=t_1+1}^{36} \pi_{t_1,u}$.

Census data

We assume that the observed census-type data x_t , which are available for 1965 (t = 3) to 1998, account for only breeding birds and that there is no emigration. We model the census data via the true number of breeding females $\mu_{A,t}$ and immature females $\mu_{C,t}$, and the productivity rate γ_t , the average number of female offspring per breeding female in year t, which could be greater than 1. Specifically we assume for $t = 3, \ldots, 36$

$$\begin{aligned} x_t &\sim \mathrm{N}(\mu_{A,t}, \sigma^2), \\ \mu_{C,t} &\sim \mathrm{Po}(\mu_{A,t-1}\gamma_{t-1}\eta_{C,t}), \\ \mu_{A,t} &\sim \mathrm{Bin}(\mu_{C,t-1} + \mu_{A,t-1}, \eta_{A,t}), \end{aligned}$$

with the observation variance σ^2 assumed constant.

Regression models and prior distributions

We model the parameters $\eta_{G,t}$, λ_t and γ_t with regression models, with z_t denoting the (observed) number of frost days in year t.

$$\begin{aligned} \operatorname{logit}(\eta_{G,t}) &= \phi_{\alpha,G} + \phi_{\beta,G} z_t \qquad G = C, A, \\ \operatorname{logit}(\lambda_t) &= \alpha_\lambda + \beta_\lambda t, \\ \operatorname{log}(\gamma_t) &= \alpha_\gamma + \beta_\gamma t. \end{aligned}$$

We place lognormal priors on the number of immature females $\mu_{C,2}$ and breeding females $\mu_{A,2}$ in the year prior to our data series, with scale parameter 1 and location param-



Figure 10: Histograms of the posterior densities of the link parameters $\phi_{\alpha,C}$, $\phi_{\alpha,A}$, $\phi_{\beta,C}$ and $\phi_{\beta,A}$ under the recovery submodel (Stage 1), and under the full joint model, as estimated by Stage 2 of the two stage sampler and by a standard MCMC sampler for the joint model.

eters $\mu_C = 200$ and $\mu_A = 1000$ respectively. We assume $\sigma^2 \sim \text{Inv-Gam}(0.001, 0.001)$ a priori, and independent N(0, 10²) prior distributions for all 8 regression parameters $(\phi_{\alpha,C}, \phi_{\alpha,A}, \alpha_{\lambda}, \alpha_{\gamma}, \phi_{\beta,C}, \phi_{\beta,A}, \beta_{\lambda}, \beta_{\gamma}).$

Results

We split the joint model, as described in Section 3.3, into two components: the mark-recapture-recovery submodel and the census submodel. Denote by $\Omega_0 = (\eta_C, \eta_A, \phi_{\alpha,C}, \phi_{\alpha,A}, \phi_{\beta,C}, \phi_{\beta,A})$ the parameters shared by both submodels and by $\Omega_1 = (\pi, \lambda, \alpha_\lambda, \beta_\lambda)$ the parameters specific to the recovery submodel. Under both the mark-recapture-recovery submodel (stage one) and the census submodel (stage two), we use independent normal priors, with mean 0 and standard deviation $\sqrt{200}$, for each component $\phi_{\alpha,C}$, $\phi_{\alpha,A}, \phi_{\beta,C}$ and $\phi_{\beta,A}$ of the link parameter. These priors were chosen so that PoE pooling of these priors results in the original prior for the link parameters under the joint model.

In stage one we drew samples from the posterior distribution $p_1(\Omega_1, \Omega_0 \mid y)$ under the recovery submodel, and retained these samples for use as a proposal distribution in stage two, in which we drew samples under the full joint model. In stage one, we drew 2.5×10^5 MCMC iterations from the posterior distribution of the mark-recapturerecovery submodel, taking 7 hours on a single core of an Intel Xeon E5-2620 2.0GHz CPU. In stage two, we discarded all but every 100th iteration, leaving 2.5×10^5 MCMC iterations for inference. This took $6\frac{1}{2}$ hours.

Figure 10 shows the results. We compare the two-stage estimates to the estimates of the joint distribution based upon 6×10^5 MCMC iterations (retaining every 10^{th} iteration) drawn using a standard (one stage) MCMC sampler, which took 22 hours to run in OpenBUGS. We regard these results as the 'gold standard' that we aim to match with the two stage sampling approach. The components of the link parameters $\phi_{\alpha,C}$ and $\phi_{\beta,C}$ corresponding to the immature birds have posterior distributions that closely agree under the joint model and mark-recapture-recovery submodel alone, but there are differences in the parameters corresponding to mature birds. In particular there is a sizeable difference for the regression parameter $\phi_{\beta,A}$, which is estimated to be notably higher under the joint model than under the mark-recapture-recovery submodel alone. The two-stage approach accurately captures this shift (Figure 10, right-hand panel). The similarity of $\phi_{\alpha,C}$, $\phi_{\alpha,A}$ and $\phi_{\beta,C}$ in the stage 1 posterior (recovery model) and the stage 2 posterior implies that the census submodel contains little information about these parameters. In contrast, the census submodel does contain information about the regression parameter $\phi_{\beta,A}$ describing the relationship between the survival rate of adult birds and the number of frost days. The census information suggests that $\phi_{\beta,A}$ should be less negative than implied by the recovery information, implying that adult survival rate decreases only slightly in harsher winters.

6 Further work and discussion

We have presented a unifying view and a generic method for joining and splitting probabilistic submodels that share a common variable. We have extended the notions of Markov combination and super Markov combination to permit pooling of prior marginal distributions, enabling a principled approach to joining models in realistic applied settings, assuming that there is not strong conflict between evidence components and that it is reasonable to assume that the submodels are conditionally independent. We also introduced a computational algorithm that allows inference for submodels to be efficiently conducted in stages, when considering either joining or splitting models. The remainder of this section discusses related work, computational issues and alternative approaches.

6.1 Related work

The key idea for a melding approach can be attributed to Poole and Raftery (2000), but their presentation focuses on a limited set of models and is tied up with a deterministic link parameter ϕ . This slightly obscures the key issues that we present more generally in Section 3, where we clearly separate issues relating to marginal replacement from issues related to deterministic transformation of random variables. A further influence is Markov combination of submodels with consistent marginal distributions (Dawid and Lauritzen, 1993), and the concept of super Markov combination (Massa and Lauritzen, 2010) of a pair of families of submodels, in which submodel-specific conditionals and link-parameter marginal distributions (which must be marginals of members of one of the original families) are "mixed-and-matched" to form a family of possible combined models. Markov melding extends these approaches by permitting, via the pooling function, the marginal distribution of the link parameter to reflect the prior marginal distributions of all submodels, rather than just one. Markov combination forms part of the literature on decomposable graphical models, where a key concept is the separator, a subset of variables that splits the model into two parts that are independent conditional on the separator. Separators correspond to link variables in Markov melding. The rich literature on decomposable graphs and corresponding algorithms, such as junction tree algorithms (Lauritzen, 1996), suggests extensions of Markov melding to a series of link variables (separators) for joining several submodels into chain or tree formations.

Evidence synthesis models (Eddy et al., 1992; Jackson et al., 2009; Albert et al., 2011; Commenges and Hejblum, 2012) often employ the approximate approach of summarising the results of a first-stage submodel via a Gaussian or other distribution, for use in a second-stage submodel as a likelihood term. We demonstrated in Section 4.3that this approach is an approximation to Markov melding under PoE pooling, therefore justifying the approximation. Similar approximations are widely used in standard and network meta-analysis (for example, Hasselblad et al., 1992; Ades and Sutton, 2006; Welton et al., 2008). Similarly, in more general hierarchical models, splitting models to make inference faster or easier has previously been considered (Liang and Weiss, 2007; Tom et al., 2010; Lunn et al., 2013a). In this setting, posterior inference is first obtained from independent unit-specific submodels, with flat, independent priors replacing all hierarchical priors in the joint model. Inference for the joint model is recovered in stage two through Markov melding of these unit-specific submodels with dictatorial pooling, so that only the hierarchical prior is reflected in the final results. This can make crossvalidation more convenient (Goudie et al., 2015). Splitting models into conditionally independent components at a set of separator or link parameters is also a key aspect of cross-validatory posterior predictive methods, including "node-splitting", for assessing conflict across subsets of evidence (Presanis et al., 2013; Gåsemyr and Natvig, 2009). Markov melding may provide a natural, computationally-efficient approach for systematic conflict assessment (Presanis et al., 2016).

Our framework can also be viewed as encapsulating a range of approaches proposed in the big data literature for handling a large number of observations ('tall data'). With tall data it may be infeasible even to store all of the data on a single computer, nevermind evaluate functions depending on the whole dataset thousands of times, as needed in MCMC. Instead, a divide-and-conquer approach can be taken, in which the original exchangeable data y are partitioned into B batches y_1, \ldots, y_B , each of which contains few enough observations that standard statistical methods can be applied without undue trouble. The key observation is that the full posterior distribution $p(\phi \mid y)$ can be split into a number of submodel posteriors $p_b(\phi \mid y_b) \propto p(y_b \mid \phi)p(\phi)^{1/B}$, $b = 1, \ldots, B$. This is a form of model splitting (Section 3.2), with PoE pooling and the original prior apportioned equally among the batches. Various approaches for integrating the batchspecific posteriors to approximate the overall posterior have been proposed (Huang and Gelman, 2005; Scott et al., 2016; Neiswanger et al., 2014; Wang and Dunson, 2013; Bardenet et al., 2017; Minsker et al., 2017). However, this literature has so far only considered independent, identically distributed data, whereas we have considered more general models and data.

6.2 Computational challenges

In our examples, the link variable ϕ is comparatively low dimensional and simple kernel density estimation using a multivariate *t*-distribution kernel proved sufficient. Moreover, the results were robust with respect to the choice of kernel and kernel bandwidth. For higher-dimensional link variables more care in the choice of kernel estimation method

might be required (Henderson and Parmeter, 2015), or, alternatively, we might wish to estimate the ratio of densities directly to improve stability (Sugiyama et al., 2012).

The multi-stage sampler (Section 4.1) broadly falls into the category of a sequential Monte Carlo sampler (Doucet et al., 2013), as described in Supplementary Material E. While the Markov melding model is invariant to the ordering of the submodels used (assuming the pooling function is also), the efficiency of the multi-stage algorithm may not be in practice, due to the need to obtain sufficient stage one samples in the appropriate region. If two submodels contain an approximately equal amount of non-conflicting information, then the ordering is unlikely to be important. In other settings, more care may be required. For example, suppose submodel M_1 contains considerably more information than M_2 . If stage one uses M_1 , then the stage one posterior may be so precise that it is unable to be adjusted for the extra information in M_2 . In contrast, if M_2 is used first, then the estimate of the posterior distribution may be very coarse, due to a lack of samples in the central part of the posterior distribution. Further research will be needed to identify the best ordering to adopt in general.

6.3 Alternative approaches

We obtained consistency in the link parameter ϕ , as required by Markov combination, through marginal replacement (Section 3.1). This approach assumes the priors differ in substance across submodels. Alternatively, as we outline in Supplementary Material F, we could assume that the priors differ only due to different scalings in each submodel, and so can be made consistent through rescaling, similar to when deriving multivariate distributions from copulas (Durante and Sempi, 2010). Yet another approach is a supra-Bayesian approach (Lindley et al., 1979; Leonelli, 2015), in which the decision maker models the experts' opinions.

The prior pooling approach considered within our framework includes a judgement as to how to weight the different submodels. Various other methods have been proposed for weighting evidence, including the cut operator (Lunn et al., 2013b; Plummer, 2015a); the power prior approach in clinical trials (for example, Neuenschwander et al., 2009); and modularisation in the computer models literature (Liu et al., 2009). Further research is required to investigate the relationship of Markov melding to other weighting approaches.

Supplementary Material

Supplementary Material for: "Joining and splitting models with Markov melding" (DOI: 10.1214/18-BA1104SUPP; .pdf).

References

Ades, A. E. and Sutton, A. J. (2006). "Multiparameter evidence synthesis in epidemiology and medical decision-making: current approaches." *Journal of the Royal* Statistical Society: Series A (Statistics in Society), 169(1): 5-35. MR2222010. doi: https://doi.org/10.1111/j.1467-985X.2005.00377.x. 1, 2, 23

- Albert, I., Espié, E., de Valk, H., and Denis, J.-B. (2011). "A Bayesian evidence synthesis for estimating Campylobacteriosis prevalence." *Risk Analysis*, 31(7): 1141–1155. 23
- Bardenet, R., Doucet, A., and Holmes, C. (2017). "On Markov chain Monte Carlo methods for tall data." Journal of Machine Learning Research, 18(47): 1–43. MR3670492. 23
- Besbeas, P., Freeman, S. N., Morgan, B. J. T., and Catchpole, E. A. (2002). "Integrating mark-recapture-recovery and census data to estimate animal abundance and demographic parameters." *Biometrics*, 58(3): 540–547. MR1933532. doi: https://doi.org/10.1111/j.0006-341X.2002.00540.x. 5
- Birrell, P. J., De Angelis, D., Wernisch, L., Tom, B. D., Roberts, G. O., and Pebody, R. G. (2016). "Efficient real-time monitoring of an emerging influenza epidemic: how feasible?" arXiv:1608.05292. 1
- Brooks, S. P., King, R., and Morgan, B. J. T. (2004). "A Bayesian approach to combining animal abundance and demographic data." Animal Biodiversity and Conservation, 27(1): 515–529. 5
- Clemen, R. T. and Winkler, R. L. (1999). "Combining probability distributions from experts in risk analysis." *Risk Analysis*, 19(2): 187–203.
- Commenges, D. and Hejblum, B. P. (2012). "Evidence synthesis through a degradation model applied to myocardial infarction." *Lifetime Data Analysis*, 19(1): 1–18. MR3016946. 1, 23
- Dawid, A. P. and Lauritzen, S. L. (1993). "Hyper Markov laws in the statistical analysis of decomposable graphical models." Annals of Statistics, 21(3): 1272–1317. MR1241267. doi: https://doi.org/10.1214/aos/1176349260. 2, 5, 22
- Department of Health (2011). "Department of Health Winter Watch." http://winterwatch.dh.gov.uk. 16
- Doucet, A., de Freitas, N., and Gordon, N. (eds.) (2013). Sequential Monte Carlo Methods in Practice. New York: Springer Science & Business Media. MR1847783. doi: https://doi.org/10.1007/978-1-4757-3437-9. 24
- Draper, D. (1995). "Assessment and propagation of model uncertainty." Journal of the Royal Statistical Society: Series B (Methodological), 57(1): 45–97. MR1325378.
- Durante, F. and Sempi, C. (2010). "Copula theory: an introduction." In Jaworski, P., Durante, F., Härdle, W. K., and Rychlik, T. (eds.), Copula Theory and its Applications, 3–31. Berlin: Springer-Verlag. MR3075361. doi: https://doi.org/10.1007/ 978-3-642-12465-5. 24
- Eddy, D. M., Hasselblad, V., and Shachter, R. (1992). Meta-Analysis by the Confidence Profile Method. London: Academic Press. 2, 23

- Gåsemyr, J. and Natvig, B. (2009). "Extensions of a conflict measure of inconsistencies in Bayesian hierarchical models." *Scandinavian Journal of Statistics*, 36(4): 822–838. MR2573310. doi: https://doi.org/10.1111/j.1467-9469.2009.00659.x. 3, 23
- Goudie, R. J. B., Hovorka, R., Murphy, H. R., and Lunn, D. (2015). "Rapid model exploration for complex hierarchical data: application to pharmacokinetics of insulin aspart." *Statistics in Medicine*, 34(23): 3144–3158. MR3402581. doi: https://doi.org/10.1002/sim.6536. 23
- Goudie, R. J. B., Presanis, A. M., Lunn, D., De Angelis, D., and Wernisch, L. (2018).
 "Supplementary Material for: "Joining and splitting models with Markov melding"." Bayesian Analysis. doi: https://doi.org/10.1214/18-BA1104SUPP. 7
- Green, P. J., Hjort, N. L., and Richardson, S. (2003). "Introducing highly structured stochastic systems." In Green, P. J., Hjort, N. L., and Richardson, S. (eds.), *Highly Structured Stochastic Systems*, 1–12. Oxford: Oxford University Press. MR2082404. 1
- Hasselblad, V., Eddy, D. M., and Kotchmar, D. J. (1992). "Synthesis of environmental evidence: nitrogen dioxide epidemiology studies." Journal of the Air & Waste Management Association, 42(5): 662–671. 23
- Henderson, D. J. and Parmeter, C. F. (2015). Applied Nonparametric Econometrics. New York: Cambridge University Press. 14, 24
- Hinton, G. E. (2002). "Training products of experts by minimizing contrastive divergence." Neural Computation, 14(8): 1771–1800.
- Huang, Z. and Gelman, A. (2005). "Sampling for Bayesian computation with large datasets." Working paper. 23
- Jackson, C. H., Best, N. G., and Richardson, S. (2009). "Bayesian graphical models for regression on multiple data sets with different variables." *Biostatistics*, 10(2): 335–351. 2, 23
- Jackson, C. H., Jit, M., Sharples, L. D., and De Angelis, D. (2015). "Calibration of complex models through Bayesian evidence synthesis." *Medical Decision Making*, 35(2): 148–161. 1
- Lauritzen, S. L. (1996). *Graphical Models*. Oxford: Clarendon Press. MR1419991. 3, 23
- Leonelli, M. (2015). "Bayesian decision support in complex modular systems: an algebraic and graphical approach." Ph.D. thesis, University of Warwick, UK. 24
- Liang, L.-J. and Weiss, R. E. (2007). "A hierarchical semiparametric regression model for combining HIV-1 phylogenetic analyses using iterative reweighting algorithms." *Biometrics*, 63(3): 733-741. MR2395710. doi: https://doi.org/10.1111/ j.1541-0420.2007.00753.x. 2, 3, 23
- Lindley, D. V., Tversky, A., and Brown, R. V. (1979). "On the reconciliation of probability assessments." *Journal of the Royal Statistical Society: Series A (General)*, 142(2): 146–180. MR0547236. doi: https://doi.org/10.2307/2345078. 24

- Liu, F., Bayarri, M. J., and Berger, J. O. (2009). "Modularization in Bayesian analysis, with emphasis on analysis of computer models." *Bayesian Analysis*, 4(1): 119–150. MR2486241. doi: https://doi.org/10.1214/09-BA404. 1, 24
- Lunn, D., Barrett, J., Sweeting, M., and Thompson, S. (2013a). "Fully Bayesian hierarchical modelling in two stages, with application to meta-analysis." *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 62(4): 551–572. MR3083911. doi: https://doi.org/10.1111/rssc.12007. 2, 3, 13, 23
- Lunn, D., Jackson, C., Best, N., Thomas, A., and Spiegelhalter, D. (2013b). The BUGS Book: A Practical Introduction to Bayesian Analysis. Boca Raton: CRC Press. 24
- Massa, M. S. and Lauritzen, S. L. (2010). "Combining statistical models." In Viana, M. A. G. and Wynn, H. P. (eds.), Contemporary Mathematics: Algebraic Methods in Statistics and Probability II, 239–260. MR2730753. doi: https://doi.org/10.1090/ conm/516/10179. 2, 5, 6, 7, 8, 22
- Massa, M. S. and Riccomagno, E. (2017). "Algebraic representations of Gaussian Markov combinations." *Bernoulli*, 23(1): 626–644. MR3556787. doi: https://doi.org/10.3150/15-BEJ759. 2, 5
- Minsker, S., Srivastava, S., Lin, L., and Dunson, D. B. (2017). "Robust and scalable Bayes via a median of subset posterior measures." *Journal of Machine Learning Re*search, 18(124): 1–40. 23
- Moran, E. V. and Clark, J. S. (2011). "Estimating seed and pollen movement in a monoecious plant: a hierarchical Bayesian approach integrating genetic and ecological data." *Molecular Ecology*, 20(6): 1248–1262. 1
- Müller, P. (1991). "A generic approach to posterior integration and Gibbs sampling." Technical Report 91-09, Purdue University. 12
- Neiswanger, W., Wang, C., and Xing, E. P. (2014). "Asymptotically exact, embarrassingly parallel MCMC." In Proceedings of the Thirtieth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-14), 623–632. Corvallis, Oregon: AUAI Press. 2, 3, 23
- Neuenschwander, B., Branson, M., and Spiegelhalter, D. J. (2009). "A note on the power prior." *Statistics in Medicine*, 28(28): 3562-3566. MR2744383. doi: https://doi.org/10.1002/sim.3722. 8, 24
- O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E., and Rakow, T. (2006). Uncertain Judgements: Eliciting Experts Probabilities. Chichester: John Wiley & Sons. 8
- Plummer, M. (2015a). "Cuts in Bayesian graphical models." *Statistics and Computing*, 25(1): 37–43. MR3304902. doi: https://doi.org/10.1007/s11222-014-9503-z. 24
- Plummer, M. (2015b). "JAGS Version 4.0.1 user manual." 18
- Poole, D. and Raftery, A. E. (2000). "Inference for deterministic simulation models: The Bayesian melding approach." *Journal of the American Statistical Association*,

95(452): 1244-1255. MR1804247. doi: https://doi.org/10.2307/2669764. 2, 7, 22

- Presanis, A. M., Ohlssen, D., Cui, K., Rosinska, M., and De Angelis, D. (2016). "Conflict diagnostics for evidence synthesis in a multiple testing framework." arXiv:1702.07304. 23
- Presanis, A. M., Ohlssen, D., Spiegelhalter, D. J., and De Angelis, D. (2013). "Conflict diagnostics in directed acyclic graphs, with applications in Bayesian evidence synthesis." *Statistical Science*, 28(3): 376–397. MR3135538. doi: https://doi.org/ 10.1214/13-STS426. 3, 23
- Presanis, A. M., Pebody, R. G., Birrell, P. J., Tom, B. D. M., Green, H. K., Durnall, H., Fleming, D., and De Angelis, D. (2014). "Synthesising evidence to estimate pandemic (2009) A/H1N1 influenza severity in 2009–2011." Annals of Applied Statistics, 8(4): 2378–2403. MR3292502. doi: https://doi.org/10.1214/14-A0AS775. 2, 4, 11, 17
- Public Health England (2014). "Sources of UK flu data: influenza surveillance in the UK." https://www.gov.uk/guidance/sources-of-uk-flu-data-influenza-surveillance-in-the-uk. 17
- Robert, C. P. and Casella, G. (2004). Monte Carlo Statistical Methods. New York: Springer. MR2080278. doi: https://doi.org/10.1007/978-1-4757-4145-2. 11
- Scott, S. L., Blocker, A. W., Bonassi, F. V., Chipman, H. A., George, E. I., and Mc-Culloch, R. E. (2016). "Bayes and big data: The consensus Monte Carlo algorithm." *International Journal of Management Science and Engineering Management*, 11(2): 78–88. 2, 3, 23
- Shubin, M., Lebedev, A., Lyytikäinen, O., and Auranen, K. (2016). "Revealing the true incidence of pandemic A(H1N1)pdm09 influenza in Finland during the first two seasons—An analysis based on a dynamic transmission model." *PLOS Computational Biology*, 12(3): e1004803.
- Sugiyama, M., Suzuki, T., and Kanamori, T. (2012). Density Ratio Estimation in Machine Learning. New York: Cambridge University Press. MR2895762. doi: https://doi.org/10.1017/CB09781139035613. 24
- Tom, J. A., Sinsheimer, J. S., and Suchard, M. A. (2010). "Reuse, recycle, reweigh: combating influenza through efficient sequential Bayesian computation for massive data." *Annals of Applied Statistics*, 4(4): 1722–1748. MR2829934. doi: https://doi.org/ 10.1214/10-AOAS349. 2, 3, 23
- Turner, R. M., Spiegelhalter, D. J., Smith, G. C. S., and Thompson, S. G. (2009). "Bias modelling in evidence synthesis." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 172(1): 21–47. MR2655603. doi: https://doi.org/ 10.1111/j.1467-985X.2008.00547.x. 3
- Wang, X. and Dunson, D. B. (2013). "Parallel MCMC via Weierstrass Sampler." arXiv:1312.4605. 23

- Welton, N. J., Cooper, N. J., Ades, A. E., Lu, G., and Sutton, A. J. (2008). "Mixed treatment comparison with multiple outcomes reported inconsistently across trials: Evaluation of antivirals for treatment of influenza A and B." *Statistics in Medicine*, 27(27): 5620–5639. MR2573773. doi: https://doi.org/10.1002/sim.3377. 23
- Welton, N. J., Sutton, A. J., Cooper, N. J., Abrams, K. R., and Ades., A. (2012). Evidence Synthesis for Decision Making in Healthcare. Chichester: John Wiley & Sons. 1, 2

Acknowledgments

This work was supported by the UK Medical Research Council [programme codes MC_UU_00002/2, MC_UU_00002/11 and MC_UU_00002/1]. We are grateful to Ian White, Sylvia Richardson, Brian Tom, Michael Sweeting, Paul Kirk, Adrian Raftery, and the 2015 Armitage lecturers (Leonhard Held and Michael Höhle) for helpful discussions of this work; and to the referees and editors for insightful comments that led to an improved manuscript. We also thank colleagues at Public Health England for providing data.