

Bayesian Inference in Incomplete Multi-Way Tables

Adrian Dobra Claudia Tebaldi and Mike West

Technical Report #2003-2
February 27, 2003

Statistical and Applied Mathematical Sciences Institute
PO Box 14006
Research Triangle Park, NC 27709-4006
www.samsi.info

Bayesian inference in incomplete multi-way tables

ADRIAN DOBRA*, CLAUDIA TEBALDI† & MIKE WEST*

adobra@isds.duke.edu tebaldi@ucar.edu mw@isds.duke.edu

February 2003

**Institute of Statistics and Decision Sciences, Duke University, Durham, NC 27708-0251, USA.*

†*National Center for Atmospheric Research, Boulder, CO 80305, USA.*

SUMMARY

We describe and illustrate approaches to Bayesian inference in multi-way contingency tables for which partial information, in the form of subsets of marginal totals, is available. In such problems, interest lies in questions of inference about the parameters of models underlying the table together with imputation for the individual cell entries. We discuss questions of structure related to the implications for inference on cell counts arising from assumptions about log-linear model forms, and a class of simple and useful prior distributions on the parameters of log-linear models. We then discuss “local move” and “global move” Metropolis-Hastings simulation methods for exploring the posterior distributions for parameters and cell counts, focusing particularly on higher-dimensional problems. As a by-product, we note potential uses of the “global move” approach for inference about numbers of tables consistent with a prescribed subset of marginal counts. Illustration and comparison of MCMC approaches is given, and we conclude with discussion of areas for further developments and current open issues.

Some key words: Bayesian inference; Disclosure limitation; Fixed margins problem; Imputation; Log-linear models; Markov basis; Markov chain Monte Carlo; Missing data.

1 Introduction

The general problem of inference in contingency tables based on partial information in terms of observed counts on a set of margins has become of increasing interest in recent years. We address this problem here in a framework involving inference on parameters of statistical models underlying multi-way tables together with inference about missing cell entries.

Some of our initial motivating interest in this area came from socio-economic and demographic studies that involve and rely on survey and census data representing differing levels of aggregation (hence marginalisation) of population characteristics. Much of the activity in these latter areas has been referred to as micro-simulation, and the development of micro-simulation methods in areas such as transportation policy and planning rely heavily on an ability to impute individual household level counts, for example, from more highly aggregated data from local or population census data.

More recently, the last several years has seen a very significant upsurge of interest in development of methods to aid in the creation, dissemination and use of public data sets from governmental sources, related to serious societal and legal concerns about data confidentiality and security. A range of issues then arise about the potential to infer individual level data from sets of interlinked aggregate-level data, and this can be focused on the problem of inferring cell counts in multi-way tables based on observation of some sets of marginal totals. Here there are questions of the extent to which sets of observed margins can inform on cell counts under varying assumptions about the structure of candidate statistical models (such as log-linear models) (Dobra, 2002; Dobra et al., 2003), with related questions about the role and impact of specific prior distributions on parameters of such models.

Historically, strong interest in the general problem area has focused on problems of counting tables with prescribed marginal totals, a goal of interest in both data confidentiality studies and in the traditional context of estimating significance levels in testing approaches (Agresti, 1992; Diaconis & Efron, 1985; Smith et al., 1996).

Bayesian inference in this context is in principle straightforward: we aim to compute posterior distributions for the unobserved cell counts and parameters underlying models for cell probabilities, jointly. In practice this may be addressed using Markov chain Monte Carlo (MCMC) simulations; this requires creativity in dealing, in particular, with simulations of the missing cell counts from appropriate conditional posterior distributions. The contributions of this article address a number of questions and needs in support of the practical development of such methods. First, we describe the general problem and context, and develop some theoretical insights into the nature and role of assumptions about model structure in its relation to the problem of inference on individual cell counts based on observation of sets of marginal totals. Then, we discuss MCMC approaches to joint analysis of parameters and missing cell counts. This uses a simple but flexible class of prior distributions on parameters of log-linear models at the parametric level. In imputing cell counts, we discuss, implement and evaluate “local move” algorithms that rely on Markov basis construction, together with a new class of “global move” approaches that have some relative attractions, especially as problems increase in dimension and complexity. A detailed discussion of an example demonstrates both the implementation and the efficacy of the approaches, and we conclude with some discussion of current open issues and challenges. In addition to innovations in modelling and computation, the work represents a broad overview of the recent and current issues in this important field.

2 Analysis Framework and Goals

2.1 Definitions and Notation

Consider a k -way contingency table of counts over a k -vector of discrete random variables $X = \{X_1, X_2, \dots, X_k\}$. Let $K = \{1, 2, \dots, k\}$, and for each $j \in K$, suppose X_j takes values in $\mathcal{I}_j = \{1, \dots, I_j\}$. Write $\mathcal{I} = \mathcal{I}_1 \times \dots \times \mathcal{I}_k$ and denote an element of \mathcal{I} by $i = (i_1, \dots, i_k)$.

The contingency table $n = \{n(i)\}_{i \in \mathcal{I}}$ is a k -dimensional array of non-negative integer numbers, with cell entries $n(i) = \#\{X = i\}$, ($i \in \mathcal{I}$), and a total of $m = I_1 \cdot \dots \cdot I_k$ cells. Any set of marginal counts is obtained by summation over one or more of the X variables. For any target subset of variables $D \subset K$, the D -marginal n_D of n has cells $i_D \in \mathcal{I}_D = \times_{j \in D} \mathcal{I}_j$, with cell entries

$$n_D(i_D) = \sum_{j \in \mathcal{I}_{K \setminus D}} n(i_D, j).$$

If $D = \emptyset$ then n_\emptyset is the grand total over n .

Cells are ordered lexicographically with the index of the k -th variable varying fastest, so that $\mathcal{I} = \{i^1, \dots, i^m\}$ where $i^1 = (1, \dots, 1)$ is the first cell and $i^m = (I_1, \dots, I_k)$ is the last cell.

2.2 Tables with Sets of Fixed Margins

Our interest lies in problems in which we observe only subsets of margins from the full table. Suppose the l margins $\mathcal{D} = \{n^1, \dots, n^l\}$ are recorded. Additional information may be available, such as upper and lower bounds for some (or all) the cells in the full table, or cases of structural zeroes (that can also be represented by fixed upper and lower bounds, in this case each at zero). In such cases, the constraints can be added to \mathcal{D} without changing the development below.

Observing the margins \mathcal{D} induces constraints that imply upper and lower bounds $U(i) = \max\{n(i) : n \in \mathcal{T}\}$ and $L(i) = \min\{n(i) : n \in \mathcal{T}\}$ on each cell entry $n(i)$, $i \in \mathcal{I}$. This limits attention to tables satisfying these constraints. Denote the set of such tables by \mathcal{T} – thus \mathcal{T} is the set of k -way tables $n = \{n(i)\}_{i \in \mathcal{I}}$ strictly compatible with \mathcal{D} . Write $\mathcal{M}(\mathcal{T})$ for the number of such tables.

2.3 Statistical Models over Tables and Observed Margins

Focus on the independent Poisson sampling model that underlies the canonical multinomial distribution for n . That is, cell counts are conditionally independent Poisson, $n(i) \sim Po(\lambda(i))$, for each $i \in \mathcal{I}$, with positive means $\lambda = \{\lambda(i)\}$.

The implied probability of the observed set of margins \mathcal{D} is then, theoretically, simply

$$\text{pr}(\mathcal{D}|\lambda) = \sum_{n' \in \mathcal{T}} \text{pr}(n'|\lambda). \quad (1)$$

In problems of inference on λ , this defines the likelihood function, but, evidently, in other than trivial, low-dimensional problems, direct evaluation is impossible. The essential role of n as (in part) missing or latent data underlies simulation-based methods using MCMC approaches that, following Tebaldi & West (1998a,b), iteratively resimulate the “missing” components of n and the parameters λ from relevant conditional distributions. For the former, note that

$$\text{pr}(n, \mathcal{D}|\lambda) = \begin{cases} \text{pr}(n|\lambda), & \text{if } n \in \mathcal{T}, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

Then, conditional on the observed margins and model parameters, inference on the missing components of the table are derived directly from the implied conditional posterior

$$\text{pr}(n|\mathcal{D}, \lambda) = \text{pr}(n|\lambda) / \sum_{n' \in \mathcal{T}} \text{pr}(n'|\lambda) \quad (3)$$

if $n \in \mathcal{T}$, being zero otherwise. When embedded in a simulation-based analysis, a key technical issue is that of developing methods to simulate this distribution – proportional to the product of Poisson components conditioned by the complicated set of constraints $n(i) \in [L(i), L(i) + 1, \dots, U(i) - 1, U(i)]$ over cells $i \in \mathcal{I}$ defined by the observed margins \mathcal{D} .

2.4 Population Models and Parameters

Inference about the underlying population structure is based on choices of model, such as traditional log-linear models, for λ . Any model \mathcal{A} has parameters θ that define $\lambda = \lambda(\theta)$. Then priors are specified on θ , so on λ indirectly and the latter will incorporate deterministic constraints imposed by the model.

Thus we work interchangeably between λ and θ in notation. Sampling distributions are conditional on λ , thus implicitly θ . For the given model \mathcal{A} , denote a prior density for the implied model parameters θ by $\text{pr}(\theta|\mathcal{A})$. Posterior inference is theoretically defined through the intractable likelihood function (1). When elaborated to include the missing data to fill out the full table, the more tractable conditional posterior is simply the prior modified by the product of Poisson likelihood components $\text{pr}(n|\lambda)$.

2.5 MCMC Framework

From the above components of the posterior distribution, we may implement MCMC methods to generate, ultimately, samples from $\text{pr}(n, \theta|\mathcal{D}, \mathcal{A})$ by iteratively re-simulating values of θ (and hence, by direct evaluation, λ), and the missing cell counts in n . Start with $n^{(0)} \in \mathcal{T}$. At the s^{th} step of the algorithm, do:

- Simulate $\theta^{(s+1)}$ from $\text{pr}(\theta|\mathcal{A}, n^{(s)}) \propto \text{pr}(\theta|\mathcal{A})\text{pr}(n^{(s)}|\lambda(\theta))$, and compute the implied new value of $\lambda^{(s+1)} = \lambda(\theta^{(s+1)})$.
- Simulate $n^{(s+1)}$ from $\text{pr}(n|\mathcal{D}, \lambda^{(s+1)})$.

Section 3 discusses aspects of prior specification and model structure in their impact on $\text{pr}(n|\mathcal{D}, \lambda)$. Section 4 then considers specific priors and the resulting analysis in log-linear models. Section 5 discusses algorithms to sampling the critical conditional posterior for the missing elements of n under the relevant distribution (3).

3 Log-linear Models and Structural Information

Some initial theoretical results describe structural aspects of the conditional posterior $\text{pr}(n|\mathcal{D}, \lambda)$ relevant in consideration of classes of log-linear models. We first note that the conditional distribution (3) of course has precisely the same form if we assume multinomial sampling for n .

First, consider the model \mathcal{A}_0 with unrelated, free parameters $\lambda(i)$. Evidently, $\text{pr}(n|\mathcal{D}, \lambda)$ then depends on all the $\lambda(i)$ and does not further simplify. Hence a prior specification across all the Poisson means has an influence in defining resulting inferences on the full table based on any set of observed margins \mathcal{D} .

Now, suppose \mathcal{A} is a specified log-linear model defining λ in terms of underlying log-linear parameters θ . This may be defined in terms of a specification of the minimal sufficient statistics of \mathcal{A} ; suppose these to be defined by the index sets $\{C_1, C_2, \dots, C_q\}$, where $C_j \subset K$. Writing $\mathcal{C} = \{C : \emptyset \neq C \subset C_j \text{ for some } j \in \{1, 2, \dots, q\}\}$, the log-linear model has the form

$$\lambda(i) = \mu \prod_{C \in \mathcal{C}} \psi_C(i_C), \quad (4)$$

where ψ_C depends on $i \in \mathcal{I}$ only through the indices in C . To make the parameters identifiable, the aliasing constraints are to set each $\psi_C(i_C) = 1$ whenever there exists $p \in C$ with $i_p = 1$ – see, for example, Whittaker (1990). One immediate consequence of these aliasing constraints is that $\lambda(i^1) = \mu$. The parameter set is then

$$\theta = \{\mu\} \cup \{\psi_C(i_C) : C \in \mathcal{C}, i_C \in \mathcal{I}_C \text{ with } i_p \neq 1 \text{ for } p \in C\}. \quad (5)$$

The following theorem now shows that the posterior $\text{pr}(n|\mathcal{D}, \lambda)$ from (3) simplifies and does not depend on all the parameters θ from (5). This represents an extension of an earlier result by Haberman (1974).

Theorem 1. *If the marginal n_C of the full table is determined from \mathcal{D} , then, under model \mathcal{A} , the posterior distribution $\text{pr}(n|\mathcal{D}, \lambda)$ does not depend on the parameters $\psi_C(i_C)$ for all $i_C \in \mathcal{I}_C$.*

Proof. It is not hard to see that, for any table n' , we have

$$\prod_{i \in \mathcal{I}} \lambda(i)^{n'(i)} = \mu^{n'_\emptyset} \prod_{C \in \mathcal{C}} \prod_{i_C \in \mathcal{I}_C} \psi_C(i_C)^{n'_C(i_C)}.$$

If $n' \in \mathcal{T}$, it follows that the grand totals of n' and n are equal, i.e., $n'_\emptyset = n_\emptyset$. Moreover, if the marginal n_C is known from \mathcal{D} , then the marginals n'_C and n_C coincide. Thus

$$\prod_{i_C \in \mathcal{I}_C} \psi_C(i_C)^{n'_C(i_C)} = \prod_{i_C \in \mathcal{I}_C} \psi_C(i_C)^{n_C(i_C)},$$

for every table $n' \in \mathcal{T}$. It follows that the terms involving $\psi_C(i_C)$, $i_C \in \mathcal{I}_C$, cancel in the denominator and the numerator of (3) and hence $\text{pr}(n|\mathcal{D}, \lambda)$ does not depend on these parameters. Note that this posterior also does not depend on the grand mean parameter μ . \square

Theorem 1 implies that the model of independence \mathcal{A}_0 and the saturated log-linear model induce different posterior distributions for $n \in \mathcal{T}$ since the posterior associated with the latter model depends on fewer parameters. This is somewhat surprising since these two models are usually considered equivalent since the saturated log-linear model is just a re-parameterization of \mathcal{A}_0 . Theorem 1 also shows how the form of an assumed log-linear model influences the conditional inference on n given \mathcal{D} . We explore this further.

If E is a set of subsets of K , let $\mathcal{P}(E)$ be those elements of E that are maximal with respect to inclusion, i.e., $\mathcal{P}(E) = \{C : C \in E \text{ and there is no } C' \in E \setminus \{C\} \text{ with } C \subset C'\}$. Write $\mathcal{E}(\mathcal{D}) = \{C : C \subseteq K \text{ and } n_C \text{ is known from } \mathcal{D}\}$ and $\mathcal{E}'(\mathcal{D}) = \{C : C \subseteq K \text{ and } C \notin \mathcal{E}(\mathcal{D})\}$. For any model \mathcal{A} , for the following we make explicit the dependence of the Poisson means on the model, via the notation $\lambda_{\mathcal{A}}$.

Theorem 2. *Consider the family $\mathcal{Q}(\mathcal{D})$ of all the log-linear models \mathcal{A} with minimal sufficient statistics induced by the index sets $\mathcal{P}(\mathcal{E}(\mathcal{D}) \cup E')$, where E' is a subset of $\mathcal{E}'(\mathcal{D})$. For any log-linear model \mathcal{A}' , there exists a model $\mathcal{A} \in \mathcal{Q}(\mathcal{D})$ such that $\text{pr}(n|\mathcal{D}, \lambda_{\mathcal{A}'})$ and $\text{pr}(n|\mathcal{D}, \lambda_{\mathcal{A}})$ depend on the same subset of parameters.*

Proof. An arbitrary log-linear model \mathcal{A}' has minimal sufficient statistics induced by some index sets $\mathcal{P}(E_0 \cup E'_0)$, where $E_0 \subset \mathcal{E}(\mathcal{D})$ and $E'_0 \subset \mathcal{E}'(\mathcal{D})$. Take the model $\mathcal{A} \in \mathcal{Q}(\mathcal{D})$ defined by the index sets $\mathcal{P}(\mathcal{E}(\mathcal{D}) \cup E'_0)$. From Theorem 1 we see that $\text{pr}(n|\mathcal{D}, \lambda_{\mathcal{A}'})$ and $\text{pr}(n|\mathcal{D}, \lambda_{\mathcal{A}})$ depend on the same subset of parameters. \square

Finally, we note the key case of the hypergeometric distribution for n given \mathcal{D} implied by lack of dependence on parameters. This is based on the set $\mathcal{B}(\mathcal{D})$ of log-linear models \mathcal{A} having minimal sufficient statistics $\mathcal{P}(E)$ for some $E \subset \mathcal{E}(\mathcal{D})$, i.e., cases in which the log-linear models have marginals embedded in \mathcal{D} . From Theorem 1 we see that, if $\mathcal{A} \in \mathcal{B}(\mathcal{D})$, then the posterior $\text{pr}(n|\mathcal{D}, \mathcal{A}, \lambda)$ is constant as a function of λ and it reduces to the hypergeometric distribution,

$$\text{pr}(n|\mathcal{D}, \lambda) \equiv \text{pr}(n|\mathcal{D}) = \left[\prod_{i \in \mathcal{I}} n(i)! \right]^{-1} / \sum_{n' \in \mathcal{T}} \left[\prod_{i \in \mathcal{I}} n'(i)! \right]^{-1} \quad (6)$$

Therefore the hypergeometric distribution is a special case of the posterior in (3) obtained by conditioning on a log-linear model whose minimal sufficient statistics are fully determined by the available data \mathcal{D} . As an aside, we note that Sundberg (1975) shows that the normalizing constant in (6) can be directly evaluated if the log-linear model \mathcal{A} is decomposable (Whittaker, 1990; Lauritzen, 1996); otherwise, this normalizing constant can be computed only if the set of tables \mathcal{T} can be exhaustively enumerated.

4 Prior Specification and Posterior Sampling of Model Parameters

We consider classes of priors for the unknown Poisson rates λ that are consistent with the models discussed in Section 3. We start with the model of independence, \mathcal{A}_0 , under which $\text{pr}(\lambda|\mathcal{D}, n) \equiv \text{pr}(\lambda|n) = \prod_i \text{pr}(\lambda(i)|n(i))$ where, for specified, independent priors, the $\lambda(i)$ have independent posteriors via the update through independent Poisson-based likelihoods. Gamma priors are of course conjugate.

Now consider a general log-linear model \mathcal{A} with minimal sufficient statistics $\{C_1, C_2, \dots, C_q\}$. Thus, if $\lambda = \{\lambda(i)\}_{i \in \mathcal{I}}$ is consistent with \mathcal{A} , then $\lambda(i)$ is represented as in (4). As for the model of independence \mathcal{A}_0 , a relatively unstructured class of priors λ is induced by exchangeable/independence assumptions for the core parameters $\psi(\cdot)$ in this representation. As developed in West (1997), and then extended

in Tebaldi & West (1998a) and Tebaldi & West (1998b), independent gamma priors on the multiplicative parameters in the log-linear model representation imply that all complete, univariate conditional posteriors are also of gamma form. Thus drawing new values for the model parameters, and hence the values of the $\lambda(i)$, is immediately accessible using Gibbs sampling. Specifically, we note that:

- $\text{pr}(\mu|\mathcal{A}, n, \theta \setminus \mu) \propto \text{pr}(\mu)\mu^{n_0} \exp \left\{ -\mu \sum_{i \in \mathcal{I}} \prod_{C \in \mathcal{C}} \psi_C(i_C) \right\}$.
- for each $\alpha := \psi_{C_0}(i_{C_0}^0) \in \Theta \setminus \mu$, the complete conditional posterior for α is proportional to

$$\text{pr}(\alpha)\alpha^{n_{C_0}(i_{C_0}^0)} \exp \left\{ -\alpha \sum_{\{i \in \mathcal{I}: i_{C_0} = i_{C_0}^0\}} \prod_{C \in \mathcal{C}} \psi_C(i_C) \right\}. \quad (7)$$

Now, when \mathcal{A} is the saturated log-linear model, \mathcal{A} is specified by one minimal sufficient statistic given by the complete index set K . In this case \mathcal{C} comprises all the non-empty subsets of K . For $i \in \mathcal{I}$ the conditional posterior for $\beta := \psi_K(i)$ is then proportional to

$$\text{pr}(\beta)\beta^{n(i)} \exp \left\{ -\beta \prod_{\{C: C \subset K, C \neq K\}} \psi_C(i_C) \right\}.$$

Hence, as noted, adopting independent gamma priors for each of the parameters in θ induces gamma complete conditionals in each case. Alternatively, finite uniform priors, $U(0, L)$ for an appropriately large upper bound L , serve as at least initial objective priors. Thus, conditional on a complete table n , we can simulate new parameter values for Poisson rates $\lambda(i)$, $i \in \mathcal{I}$, via sets of independent draws from gamma distributions, or truncated gamma distributions.

5 Imputing Cell Counts

The major component of MCMC analyses relates to the sampling algorithms to generate complete tables of counts n that are consistent with the constraints defined by the marginal count information \mathcal{D} .

Direct sampling is computationally infeasible because the normalizing constant in (3) would have to be

evaluated at each iteration, and this evaluation cannot be done quickly, if at all, due to the existence of a huge number of tables consistent with \mathcal{D} . Hence, some form of embedded Metropolis-Hastings method is required within the overall MCMC that also samples λ .

Given a current state $(n^{(s)}, \lambda^{(s+1)})$, a candidate table n^* is generated from a specified proposal distribution $q(n^{(s)}, n^*)$, and accepted with probability

$$\min \left[1, \frac{\text{pr}(n^* | \lambda^{(s+1)}) q(n^*, n^{(s)})}{\text{pr}(n^{(s)} | \lambda^{(s+1)}) q(n^{(s)}, n^*)} \right]. \quad (8)$$

The only requirement the proposal distribution has to satisfy is that $q(n^{(s)}, n^*) > 0$ if and only if $q(n^*, n^{(s)}) > 0$. In contrast with the direct sampling approaches, it is neither necessary to identify the support \mathcal{T} of $\text{pr}(n | \mathcal{D}, \lambda)$, nor to evaluate $\text{pr}(n | \lambda)$ completely across the support. If a proposal distribution generates candidate tables outside \mathcal{T} , they will be rejected as they lead to zero acceptance probabilities.

Two approaches are considered: “local” and “global” moves methods.

5.1 Local Moves

Diaconis & Sturmfels (1998) proposed generating a candidate table $n^* \in \mathcal{T}$ using Markov bases of “local moves”. A local move $g = \{g(i)\}_{i \in \mathcal{I}}$ is a multi-way array containing integer entries $g(i) \in \{\dots, -2, -1, 0, 1, 2, \dots\}$. A Markov basis $\text{MB}(\mathcal{T})$ associated with \mathcal{T} allows any two tables n_1, n_2 in \mathcal{T} to be connected by a series of local moves g^1, g^2, \dots, g^r , i.e.

$$n_1 - n_2 = \sum_{j=1}^r g^j.$$

If the chain is currently at $n^{(s)} \in \mathcal{T}$, a new candidate n^* is generated by uniformly choosing a move $g \in \text{MB}(\mathcal{T})$. The candidate $n^* = n^{(s)} + g$ belongs to \mathcal{T} if and only if $n^*(i) \geq 0$ for all $i \in \mathcal{I}$. Such a

move g is said to be permissible for the current table $n^{(s)}$. If the selected move is not permissible, the chain stays at $n^{(s)}$. Otherwise, the chain moves to $n^{(s+1)} = n^*$ with probability $\min\{1, \rho\}$, where

$$\rho := \frac{\text{pr}(n^*|\lambda^{(s+1)})}{\text{pr}(n^{(s)}|\lambda^{(s+1)})} = \prod_{\{i \in \mathcal{I}: n^*(i) \neq n^{(s)}(i)\}} \frac{n^{(s)}(i)!}{n^*(i)!} \exp\{g(i) \log \lambda^{(s+1)}(i)\}.$$

Note that the proposal distribution $q(\cdot, \cdot)$ induced by a Markov basis is symmetric, i.e., $q(n^*, n^{(s)}) = q(n^{(s)}, n^*)$.

The Markov basis required by the “local move” algorithm constitutes both the strength and the weakness of this sampling procedure. The basis has to be generated before the actual simulation begins. This extra step is likely to involve long and tedious computations in algebra systems (such as Macaulay or Cocoa) following the approach for computing a Markov basis suggested by Diaconis & Sturmfels (1998). They showed that a Markov basis for a set of tables \mathcal{T} can be determined from a Gröbner basis of a well-specified polynomial ideal.

An alternative to this algebraic approach was proposed by Dobra (2002), who gave direct formulæ for dynamically generating a Markov basis in the special case when the information \mathcal{D} available about the original table n consists of a set of marginals that define a decomposable log-linear model. Although the use of Dobra’s formulæ require minimal computational effort, they cannot be extended to more general cases (non-decomposable models). Nevertheless, once a Markov basis is computed, the “local move” method can be very fast since generating a new candidate table is done only by additions and subtractions.

Another possible disadvantage of the “local move” method is that the current table and the candidate table can be very similar, and this is critical if \mathcal{T} is large. The Markov bases associated with such large spaces of tables contain moves that change only few table entries and the change can be as small as ± 1 . For example, in the case of two-way tables with fixed row and column totals, the moves have counts of zero everywhere except four cells that contain two counts of 1 and two counts of -1 . This type of

moves are called primitive. Actually, Dobra (2002) proved that primitive moves are the only moves needed to connect tables in the decomposable case mentioned above. Clearly, changing only four cells in a high-dimensional contingency table that might contain millions of cells will undoubtedly lead to high dependencies between consecutive sample tables and the corresponding parameter values.

Therefore, with the “local move” method, one cannot control how far the chain “jumps” in the space of feasible tables because the jumps are pre-specified by the Markov basis employed. We present an alternative to the “local move” method, which we call the “global move” method, that allows one to control and adjust the distance between the current and candidate tables.

5.2 Global Moves

The idea behind the “global move” method is straightforward, and utilizes compositional sampling: a new table in \mathcal{T} can be generated by sequentially drawing a value for each cell in the table from the set of possible values for that cell while updating the corresponding upper and lower bounds for the rest of the cells. This strategy is similar to the approach and algorithm of Tebaldi & West (1998a,b) that was modified to a sequential form as utilized by Liu (2001) and Chen (2001). Moreover, the same idea constitutes the core of the generalized shuttle algorithm (Dobra, 2002) that calculates sharp upper and lower bounds for cells in a contingency tables only by efficiently exploiting the unique structure of the categorical data.

Using the chain rule, we re-write the target distribution $\text{pr}(n|\mathcal{D}, \lambda) \propto \text{pr}(n|\lambda)$ from (3) as

$$\begin{aligned} \text{pr}(n|\lambda) &= \text{pr}(n(i^1), \dots, n(i^m)|\lambda), \\ &= \text{pr}(n(i^1)|\lambda) \prod_{a=2}^m \text{pr}(n(i^a)|n(i^1), \dots, n(i^{a-1}), \lambda). \end{aligned} \tag{9}$$

The support of $\text{pr}(n(i^1)|\lambda)$ is a subset of the set of integers \mathcal{H}_1 defined by the upper and lower bounds induced by \mathcal{D} on the cell i^1 , i.e., $\mathcal{H}_1 := \{L(i^1), L(i^1) + 1, \dots, U(i^1) - 1, U(i^1)\}$. Similarly, the sup-

port of $\text{pr}(n(i^a)|n(i^1), \dots, n(i^{a-1}), \lambda)$ is a subset of the set of integers \mathcal{H}_a defined by the upper and lower bounds induced on the cell i^a by \mathcal{D} and by the additional constraints resulted from fixing the counts in the cells i^1, \dots, i^{a-1} , i.e., $\mathcal{H}_a := \{L(n; i^a), L(n; i^a) + 1, \dots, U(n; i^a) - 1, U(n; i^a)\}$, where $L(n; i^a) = \min\{n'(i^a) : n' \in \mathcal{T}, n'(i^1) = n(i^1), \dots, n'(i^{a-1}) = n(i^{a-1})\}$ and $U(n; i^a) = \max\{n'(i^a) : n' \in \mathcal{T}, n'(i^1) = n(i^1), \dots, n'(i^{a-1}) = n(i^{a-1})\}$.

If a Markov basis associated with the set of feasible tables \mathcal{T} can be constructed such that this basis contains only primitive moves, then the supports of $\text{pr}(n(i^1)|\lambda)$ and $\text{pr}(n(i^a)|n(i^1), \dots, n(i^{a-1}), \lambda)$ will be exactly \mathcal{H}_1 and \mathcal{H}_a , respectively. Otherwise, to the best of our knowledge, there is no theoretical result which shows that the supports of these two distributions should coincide with \mathcal{H}_1 and \mathcal{H}_a , nor do we know of an example when this property does not hold.

A candidate table $n^* \in \mathcal{T}$ is generated with the “global move” method as follows:

- Draw a value $n^*(i^1)$ from \mathcal{H}_1 .
- for $a = 2, \dots, m$ do
 1. Calculate $L(n^*; i^a)$ and $U(n^*; i^a)$.
 2. Draw a value $n^*(i^a)$ from \mathcal{H}_a .

end for

We still need a way to draw cell values from $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_m$ such that the resulting candidate tables n^* will be neither “too different” nor “too similar” to the current state $n^{(s)} \in \mathcal{T}$. Candidate tables that are “too different” are very likely to be rejected by the Metropolis-Hastings step, while candidates that are “too similar” will be likely to be accepted, but the chain will not advance fast enough in the target space, so inducing very correlated sample path. One approach to balancing these issues uses an annealing idea.

Consider the scaling factors v_1, v_2, \dots, v_m with $v_a \in (0, 1)$. For each $a \in \{1, 2, \dots, m\}$, draw a value $n^*(i^a)$ from the proposal distribution with probabilities

$$q_a(n^{(s)}(i^a), n^*(i^a)) \propto v_a^{|n^*(i^a) - n^{(s)}(i^a)|}, \quad (10)$$

Note that the current value $n^{(s)}(i^a)$ of cell i^a does not have to belong to the support \mathcal{H}_a of $q_a(n^{(s)}(i^a), \cdot)$. However, candidate cell values in \mathcal{H}_a that are closer to $n^{(s)}(i^a)$ receive a higher probability to be selected. This probability increases as the scaling factor v_a decreases towards 0. The full unnormalized proposal distribution can then be written as:

$$q(n^{(s)}, n^*) = \prod_{a=1}^m q_a(n^{(s)}(i^a), n^*(i^a)), \quad (11)$$

over contingency tables $n^* \in \mathcal{T}$. Any feasible table $n^* \in \mathcal{T}$ has a strictly positive probability of being sampled given any current state $n^{(s)}$, hence the Markov chain obtained by employing the proposal distribution $q(\cdot, \cdot)$ from (11) will be irreducible.

6 Example – Czech Autoworkers Data

The data in Table 1 come from a prospective epidemiological study of 1,841 workers in a Czechoslovakian car factory, as part of an investigation of potential risk factors for coronary thrombosis (see Edwards & Havranek (1985)). In the left-hand panel of Table 1, A indicates whether the worker smokes or not, B corresponds to “strenuous mental work”, C corresponds to “strenuous physical work”, D corresponds to “systolic blood pressure”, E corresponds to “ratio of β and α lipoproteins” and F represents “family anamnesis of coronary heart disease”. Since all the workers in this factory participated in the survey, Table 1 represents population data, hence the identity of individuals with uncommon characteristics in this population which translate into small cell counts of 1 or 2 might be at risk if the dataset would be made public—see Dobra (2002) for additional details and explanations about disclosure limitation issues

related to this dataset. Table 1 has three such small counts: one count of 1 and two counts of 2. In this paper we focus only on the cell $(1, 2, 2, 1, 1, 2)$ that contains the population unique.

We assume that the information we have about this six-way table n consists of the set of marginals

$$\mathcal{R}_1 = \{n_{\{A,C,D,E,F\}}, n_{\{A,B,D,E,F\}}, n_{\{A,B,C,D,E\}}, n_{\{B,C,D,F\}}, n_{\{A,B,C,F\}}, n_{\{B,C,E,F\}}\}.$$

Note that \mathcal{R}_1 contains most of the marginals of the Czech Autoworkers data – the omissions are the three five-way marginals. Let \mathcal{T}_1 be the set of dichotomous six-way tables consistent with \mathcal{R}_1 . Using the generalized shuttle algorithm of Dobra (2002) we find that \mathcal{T}_1 contains 810 tables. The upper and lower bounds on cell entries induced by \mathcal{R}_1 are given in the right-hand panel of Table 1. Given \mathcal{R}_1 , every cell in this table can take 10 or 11 possible values.

In the following analyses, priors on log-linear model parameters are all taken as uniform on a fixed, large range.

6.1 Independence Model

We start by assuming the model of independence \mathcal{A}_0 for the parameters λ . We generated a sample of 25,000 from the joint distribution $\text{pr}(n, \lambda | \mathcal{R}_1, \mathcal{A}_0)$ using the “global move” method and the “local move” method. In both cases, to reduce the correlation between two consecutive draws, we discarded 25 pairs (n, λ) before selecting another pair in the sample. The burn-in time was 2,500 which should be appropriate given the small number of tables in \mathcal{T}_1 .

The scaling factors v_a from the “global move” method were taken to be equal to 0.5 for cell $(1, 1, 1, 1, 1, 1)$ in the table and 0.05 for the rest of the cells. We used the CoCoa system to generate a Markov basis that preserve the marginals \mathcal{R}_1 of dichotomous six-way tables. This basis has 20 moves, 12 moves of “type 1” and 8 moves of “type 2”. The moves of “type 1” contain 16 counts of -1 , 32 counts of 0 and 16 counts of 1, while the moves of “type 2” have 4 count of -2 , 16 counts of -1 , 24 counts of 0, 16 counts of 1 and 4 counts of 2. When employing the “local move” algorithm, we also

need to take into account the negatives of these 20 moves. Because the space of tables induced by \mathcal{R}_1 is very constrained, each local move changes many cells when applied to a candidate table which leads to sequences of values with relatively low correlation between iterates (see Figure 1), hence the “local move” method is very similar to the “global move” method in this particular case. The rejection rates of the Metropolis-Hastings step are close when using the two methods: 48% for “local move” and 63% for “global move”.

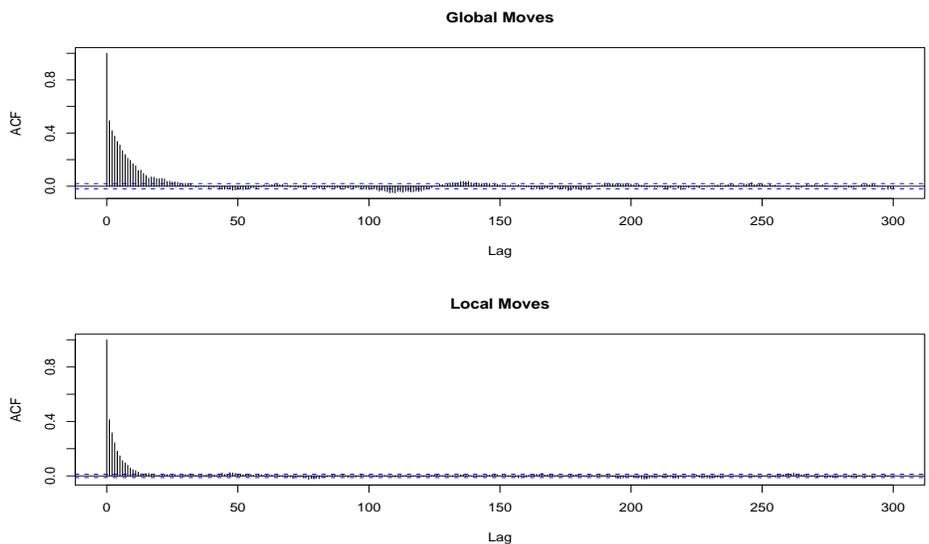


Figure 1: Autocorrelation plots for the Poisson rate associated with the cell $(1, 2, 2, 1, 1, 2)$ of the Czech Autoworkers data if the “global move” method was employed (top) and if the “local move” method was employed (bottom).

Figures 2 and 3 display approximate posterior distributions for the Poisson rate and for the possible values of cell $(1, 2, 2, 1, 1, 2)$. Note that these posterior distributions calculated from samples generated with the “global move” method and with the “local move” method are very similar. The quantile-quantile plots we give show that the samples come from the same marginal distributions, hence the two methods yield similar results as it was expected.

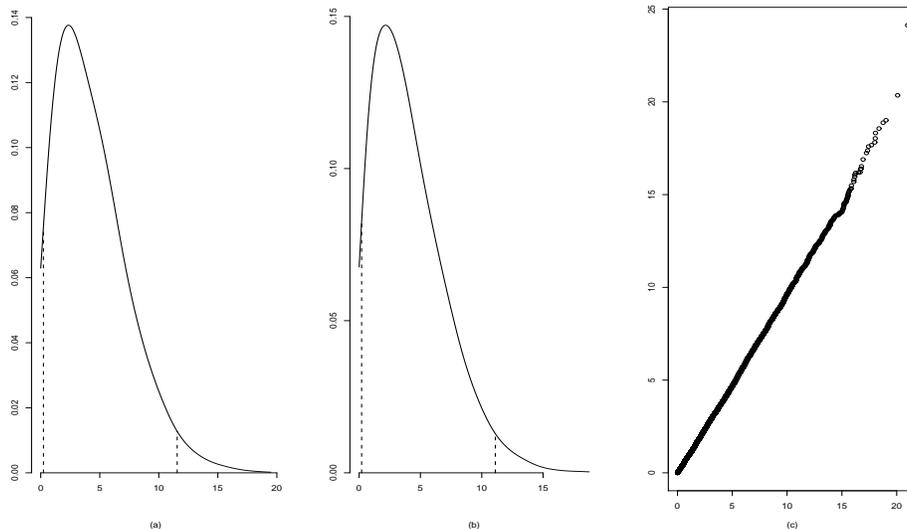


Figure 2: Marginal posterior distributions under the model of independence \mathcal{A}_0 for the Poisson rate associated with cell $(1, 2, 2, 1, 1, 2)$ of the Czech Autoworkers data obtained using (a) Global moves and (b) Local moves. A quantile-quantile plot of the samples obtained using the two methods is presented at (c): the vertical axis represents the “global move” sample, while the horizontal axis represents the “local move” sample.

6.2 The Hypergeometric Distribution

Since \mathcal{T}_1 is small enough to enumerate all its elements, we were able to compute exactly the hypergeometric distribution induced by fixing the marginals \mathcal{R}_1 —see (6). In the first row of Table 2 we present the exact marginal distribution for cell $(1, 2, 2, 1, 1, 2)$ obtained from the hypergeometric. We can see that this marginal distribution is unimodal and that the cell values 0, 1, 2, 3, 4 and 5 receive almost all the probability mass. The marginal probabilities for the other five possible values of this cell are strictly positive, but they are very small.

We approximated the exact marginal of hypergeometric distribution for cell $(1, 2, 2, 1, 1, 2)$ by generating samples of size 10,000 using three methods: (a) direct sampling from hypergeometric; (b) the

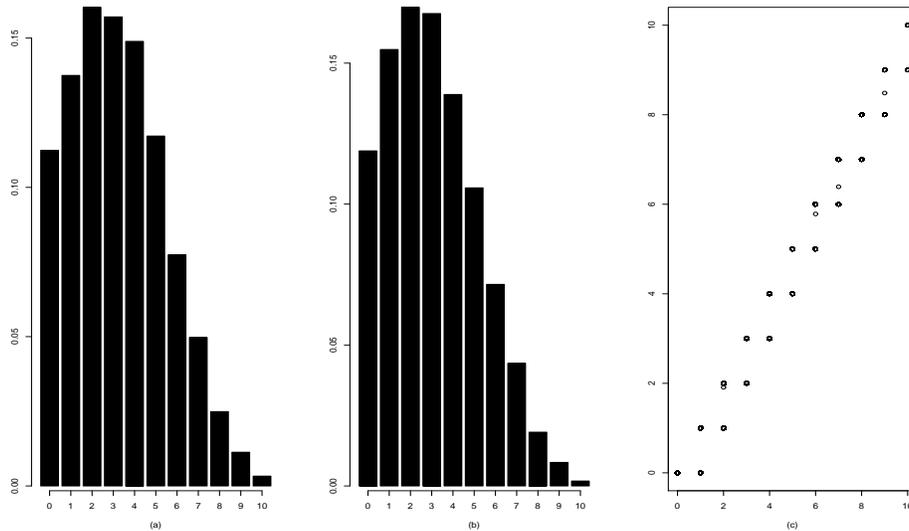


Figure 3: Marginal posterior distributions under the model of independence \mathcal{A}_0 for the possible values of cell (1, 2, 2, 1, 1, 2) of the Czech Autoworkers data obtained using (a) Global moves and (b) Local moves. A quantile-quantile plot of the samples of cell values obtained using the two methods is presented at (c): the vertical axis represents the “global move” sample, while the horizontal axis represents the “local move” sample.

“global move” method and (c) the “local move” method. The results are displayed in the second, third and fourth rows of Table 2. We see that all three methods yield reasonable good approximations of the “true” exact probabilities. The values 6, 7, 8, 9 and 10 were never obtained when sampling directly from the hypergeometric and when the “global move” method was used. When using the “local move” method, the value 6 was generated once.

The scaling factors v_a for the “global move” method were all equal to 0.25 which induce a rejection rate of 69%. The same Markov basis as before was used for the “local move” method. In this case the rejection rate was 64%. For both imputation methods the burn-in time was 2,500 and the gap between two consecutive draws selected in the samples was 25.

6.3 Log-linear Models

In Section 3 we argued that the log-linear models \mathcal{A} we need to consider in order to obtain different posterior distributions $\text{pr}(n|\mathcal{D}, \mathcal{A}, \cdot)$ are those log-linear models whose minimal sufficient statistics embed the set of fixed marginals which constitute part or all the available information \mathcal{D} . If the marginals \mathcal{R}_1 are fixed, the set of relevant log-linear models is given in Table 3. Again, the hypergeometric distribution is obtained by conditioning on a log-linear model whose marginals are embedded in the set of fixed marginals \mathcal{R}_1 .

Using the “global move” method, we simulated a samples of size 25,000 from the joint distributions $\text{pr}(n, \lambda|\mathcal{R}_1, \mathcal{A}_j)$ for each $j = 1, 2, \dots, 8$. The simulation details (e.g., burn-in time, scaling factors, rejection rates, etc) were very similar as in the case of independence model \mathcal{A}_0 .

In Table 4 we summarized the approximate marginal posterior distributions for cell $(1, 2, 2, 1, 1, 2)$ obtained by conditioning on the model of independence \mathcal{A}_0 (first row in the table), on any log-linear model whose marginals are known from \mathcal{R}_1 which leads to the hypergeometric distribution (second row) and on the “relevant” log-linear models $\mathcal{A}_1, \dots, \mathcal{A}_8$ (third row up to tenth row). We note that these marginal distributions can be very different. In some instances, these distributions appear to be unimodal (with a mode at 0 or 2), while in some other instances, the posteriors seem to be bimodal (with a mode at 0 and another mode at 10).

7 Counting Tables

The “global move” method for sampling tables consistent with a set of constraints can be employed to estimate the total number of tables consistent with these constraints. Write $\mathcal{M}(\mathcal{T})$ for the number of tables in the constrained set \mathcal{T} . Chen (2001) suggested that $\mathcal{M}(\mathcal{T})$ can be estimated by sampling from the uniform distribution on \mathcal{T} , i.e., $\text{pr}(n|\mathcal{D}) = 1/\mathcal{M}(\mathcal{T})$, but this is infeasible in any but trivial problems

since it would be equivalent to knowing $\mathcal{M}(\mathcal{T})$ exactly. Instead, we simulate tables from the proposal distribution defined in (11) with scaling factors v_a equal (at least eventually) to 1 for all the cells in the table. This particular choice of the scaling factor makes this proposal distribution independent of the table previously sampled, i.e.,

$$q(n) \propto \prod_{a=1}^m \frac{1}{U(n; i^a) - L(n; i^a) + 1}.$$

Then, we can write

$$1 = \sum_{n \in \mathcal{T}} \frac{\text{pr}(n|\mathcal{D})}{q(n)} q(n) = \frac{1}{\mathcal{M}(\mathcal{T})} \sum_{n \in \mathcal{T}} \frac{1}{q(n)} q(n).$$

This suggests an estimate for $\mathcal{M}(\mathcal{T})$ given by

$$\frac{1}{S} \sum_{s=1}^S \frac{1}{q(n^{(s)})} = S^{-1} \sum_{s=1}^S \prod_{a=1}^m [U(n^{(s)}; i^a) - L(n^{(s)}; i^a) + 1], \quad (12)$$

where $n^{(1)}, n^{(2)}, \dots, n^{(S)}$ are sampled independently from $q(\cdot)$. As an example, consider two sets of marginals of the Czech Autoworkers data. For the first set we can count the exact number of tables with a modified version of the generalized shuttle algorithm (Dobra, 2002); for the second, there are so many feasible tables and the generalized shuttle algorithm is infeasible.

Let \mathcal{R}_2 be the 15 four-way marginals of the Czech Autoworkers data. The upper and lower bounds induced by \mathcal{R}_2 are given in the left-hand panel of Table 5. We generated 100 samples of 5,000 tables each from \mathcal{T}_2 , the set of tables consistent with \mathcal{R}_2 —see Figure 4. The true number of tables in \mathcal{T}_2 is 705,884. The mean of our estimates of $\mathcal{M}(\mathcal{T}_2)$ is 703,126 with a coefficient of variation of 0.0012, suggesting a very efficient sampling method. A 95% confidence interval for $\mathcal{M}(\mathcal{T}_2)$ is 651,919–750,234.

The second example assesses the number of dichotomous six-way tables that have a count of 1 in cell $(1, 2, 2, 1, 1, 2)$ (which identifies the population unique in the Czech Autoworkers data n) and that are consistent with the marginals \mathcal{R}_3 given in (13),

$$\mathcal{R}_3 := \{n_{\{B,F\}}, n_{\{A,B,C,E\}}, n_{\{A,D,E\}}\}. \quad (13)$$

Let \mathcal{T}_3 denote this set of tables. The upper and lower bounds associated with \mathcal{T}_3 are given in the right-hand panel of Table 5. We generated 1,000 samples of 35,000 tables each—see Figure 4. We work on the \log_{10} scale as $\mathcal{M}(\mathcal{T}_3)$ is very large. The mean of $\log_{10}\{\mathcal{M}(\mathcal{T}_3)\}$ is 58 with a coefficient of variation of 0.00005, and a 95% confidence interval of 57.60–59.31. The true number of tables $\mathcal{M}(\mathcal{T}_3)$ is, of course, unknown, but we have confidence in the estimate of 57 – 59.

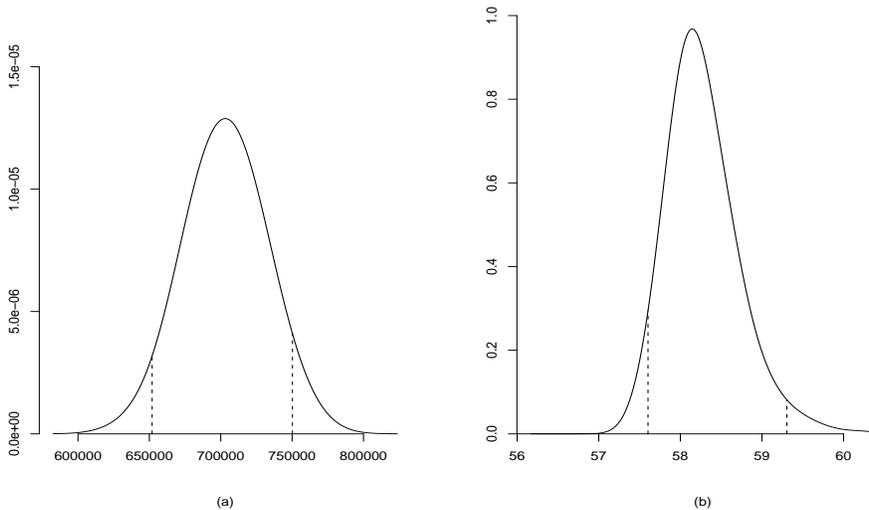


Figure 4: Estimates for (a) $\mathcal{M}(\mathcal{T}_2)$ and (b) $\log_{10}\{\mathcal{M}(\mathcal{T}_3)\}$.

8 Concluding Comments

Combined parameter inference and missing-data imputation in contingency tables is a very broad reaching problem. The specific context of inference based on data in terms of subsets of marginal counts is pervasive, and is central in problems of data disclosure, dissemination and confidentiality. Issues of prior specification are relevant both for direct analysis, and in connection with questions about potential uses of data released to multiple consumers who will each bring their own priors, or classes of priors, to bear on interpreting the margins. Our work describes some of the basic theoretical and structural

issues in the context of log-linear models, and presents a detailed development of MCMC approaches, with examples.

Central to the work reported is the development and implementation of Bayesian inference under a specific class of structured priors for log-linear model parameters, and the introduction and development of a “global move” simulation approach for imputing missing elements of contingency tables subject to observed margins. Our examples illustrate the use of both “local” and “global” move algorithms for sampling tables. Each has advantages and disadvantages, though we generally prefer the “global move” approach as it is relatively easily set up and implemented. The speed of the above procedure is directly influenced by the number of cells that need to be fixed at a certain value before a full table consistent with the data \mathcal{D} is determined. Every time we assign a value to a cell, we need to update the upper and lower bounds for the rest of the cells in the table. Consequently, the smaller the number of cells we need to fix, the faster the algorithm is. This number is a function of the pattern of constraints induced by the full information \mathcal{D} .

One of the requirements of the “global move” approach sampling algorithm is that the bounds defining the values a cell count can take given data \mathcal{D} and given that some other cells have been fixed at a certain value have to be sharp. Gross bounds approximating the corresponding sharp bounds will frequently lead to combinations of cell values that do not correspond to tables in \mathcal{T} , and hence the use of gross bounds will significantly decrease the efficiency of this sampling procedure. Unfortunately, computing sharp bounds to determine the admissible values of every cell in the target table could become a serious computational burden in the case of large high-dimensional tables having thousands or possibly millions of cells. Therefore, one very difficult computational problem (the generation of a Markov basis) from the “local move” algorithm is replaced, in the “global move” algorithm, with another very challenging problem—the calculation of sharp integer bounds. As we pointed out before, even if sharp bounds are calculated and used at each step, there is no theory which shows that infeasible

combinations of cell values cannot be generated. However, this has never happened in a range of examples studied so far. Note that the “local move” method generates candidate tables that are outside \mathcal{T} if a move that is not permissible for the current state of the chain is selected. In this respect, the “global move” method seems to perform better than the “local move” method.

In addition, a simulation algorithm employing the “global move” method can be started right away without any possibly tedious preliminary computations required in most cases by the “local move” method. There are situations when the generation of a Markov basis could take too long to complete (hence no samples can be drawn with the “local move” method), while the “global move” method might still be applied and samples will still be generated even if these samples might be expensive in term of computing time.

As opposed to the “local move” method that modifies only some small subset of cells affected by the chosen local move, the “global move” method potentially changes the entire table at each step, which leads to sequences of tables with very low correlation from one step to the next, and clearly facilitates an effective exploration of the multidimensional support of the random variables in exam.

Throughout the paper we employed the generalized shuttle algorithm (Dobra, 2002) to compute sharp upper and lower bounds. This is a very flexible algorithm that allows one to specify different types of constraints on the cell entries of a table. These constraints can be, but are not limited to, fixed marginals or cells fixed at a certain value. Moreover, the generalized shuttle algorithm can be modified to exhaustively enumerate the tables consistent with a set of constraints.

We have illustrated how posterior inferences on cell counts can vary substantially based on assumed forms of log-linear models. It may in future be of interest to consider extensions of this work to build analyses across models, utilising model-mixing and model-averaging ideas (Kass & Raftery, 1995; Madigan & York, 1995), and some further developments of the proposed MCMC methods here that extend to reversible-jump methods will be relevant there.

Acknowledgments

We thank Stephen Fienberg for comments and useful discussions, and Ian Dinwoodie for assistance and advice on the generation of Markov bases. Elements of this work formed part of the Stochastic Computation program (2002-03) at the Statistical and Applied Mathematical Sciences Institute, RTP, USA. Partial support was also provided under NSF grant DMS-0102227.

References

- AGRESTI, A. (1992). A survey of exact inference for contingency tables. *Statist. Sci.* **7**, 131–177.
- CHEN, Y. (2001). *Sequential importance sampling and its applications*. Ph.D. thesis, Stanford University.
- DIACONIS, P. & EFRON, B. (1985). Testing for independence in a two-way table: new interpretations of the chi-square statistic. *The Annals of Statistics* **13**, 845–874.
- DIACONIS, P. & STURMFELS, B. (1998). Algebraic algorithms for sampling from conditional distributions. *The Annals of Statistics* **26**, 363–397.
- DOBRA, A. (2002). *Statistical tools for disclosure limitation in multi-way contingency tables*. Ph.D. thesis, Department of Statistics, Carnegie Mellon University.
- DOBRA, A., FIENBERG, S. E. & TROTTINI, M. (2003). Assessing the risk of disclosure of confidential categorical data (with discussion). In *Bayesian Statistics 7*, J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith & M. West, eds. Oxford University Press.
- EDWARDS, D. E. & HAVRANEK, T. (1985). A fast procedure for model search in multidimensional contingency tables. *Biometrika* **72**, 339–351.
- HABERMAN, S. J. (1974). *The Analysis of Frequency Data*. Chicago: University of Chicago Press.

- KASS, R. E. & RAFTERY, A. E. (1995). Bayes factors. *Journal of the American Statistical Association* **90**, 773–795.
- LAURITZEN, S. L. (1996). *Graphical Models*. Oxford: Clarendon Press.
- LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. New York: Springer.
- MADIGAN, D. & YORK, J. (1995). Bayesian graphical models for discrete data. *International Statistical Review* **63**, 215–232.
- SMITH, P., FORSTER, J. & McDONALD, J. (1996). Monte carlo exact tests for square contingency tables. *Journal of the American Statistical Association* **2**, 309–321.
- SUNDBERG, R. (1975). Some results about decomposable (or markov-type) models for multidimensional contingency tables: distribution of marginals and partitioning of tests. *Scandinavian Journal of Statistics* **2**, 71–79.
- TEBALDI, C. & WEST, M. (1998a). Bayesian inference on network traffic using link count data (with discussion). *Journal of the American Statistical Association* **93**, 557–576.
- TEBALDI, C. & WEST, M. (1998b). Reconstruction of contingency tables with missing data. ISDS Discussion Paper #98-01, Duke University.
- WEST, M. (1997). Statistical inference for gravity models in transportation flow forecasting. Technical Report #60, National Institute of Statistical Sciences.
- WHITTAKER, J. (1990). *Graphical Models in Applied Multivariate Statistics*. New York: John Wiley & Sons.

Table 1: Czech Autoworkers data from Edwards & Havranek (1985). The left-hand panel contains the cell counts and the right-hand panel contains the bounds given the margins \mathcal{R}_1 . The cell containing the population unique and its corresponding bounds are marked with a box.

F	E	D	C	B				B			
				A	no	yes		A	no	yes	
				no	yes	no	yes	no	yes	no	yes
neg	< 3	< 140	no	44	40	112	67	[35, 45]	[35, 44]	[111, 121]	[63, 72]
			yes	129	145	12	23	[128, 138]	[141, 150]	[3, 13]	[18, 27]
	≥ 140	no	35	12	80	33	[29, 39]	[5, 14]	[76, 86]	[31, 40]	
		yes	109	67	7	9	[105, 115]	[65, 74]	[1, 11]	[2, 11]	
	≥ 3	< 140	no	23	32	70	66	[16, 25]	[26, 35]	[68, 77]	[63, 72]
			yes	50	80	7	13	[48, 57]	[77, 86]	[0, 9]	[7, 16]
		≥ 140	no	24	25	73	57	[19, 28]	[16, 25]	[69, 78]	[57, 66]
			yes	51	63	7	16	[47, 56]	[63, 72]	[2, 11]	[7, 16]
pos	< 3	< 140	no	5	7	21	9	[4, 14]	[3, 12]	[12, 22]	[4, 13]
			yes	9	17	1	4	[0, 10]	[12, 21]	[0, 10]	[0, 9]
	≥ 140	no	4	3	11	8	[0, 10]	[1, 10]	[5, 15]	[1, 10]	
		yes	14	17	5	2	[8, 18]	[10, 19]	[1, 11]	[0, 9]	
	≥ 3	< 140	no	7	3	14	14	[5, 14]	[0, 9]	[7, 16]	[8, 17]
			yes	9	16	2	3	[2, 11]	[10, 19]	[0, 9]	[0, 9]
		≥ 140	no	4	0	13	11	[0, 9]	[0, 9]	[8, 17]	[2, 11]
			yes	5	14	4	4	[0, 9]	[5, 14]	[0, 9]	[4, 13]

Table 2: Marginal associated with cell $(1, 2, 2, 1, 1, 2)$ of the hypergeometric distribution induced by fixing the marginals \mathcal{R}_1 . The exact as well as the estimated cell probabilities were rounded to the fourth decimal. We marked with a “–” the cell values that did not appear in a sample. The values $7, \dots, 10$ receive a probability weight very close to 0 and do not appear in the table.

	Possible Cell Values						
	0	1	2	3	4	5	6
Exact Distribution	0.0028	0.1022	0.4663	0.3697	0.0570	0.0016	0.0000
Direct Sampling	0.0035	0.1051	0.4712	0.3637	0.0547	0.0018	–
Global Moves	0.0032	0.1066	0.4553	0.3759	0.0561	0.0029	–
Local Moves	0.0029	0.1085	0.4595	0.3734	0.0538	0.0018	0.0001

Table 3: Relevant log-linear models for \mathcal{R}_1 .

Log-linear Model	Minimal Sufficient Statistics
\mathcal{A}_1	$\mathcal{R}_1 \cup \{n_{\{B,C,D,E,F\}}\}$
\mathcal{A}_2	$\mathcal{R}_1 \cup \{n_{\{A,B,C,E,F\}}\}$
\mathcal{A}_3	$\mathcal{R}_1 \cup \{n_{\{A,B,C,D,F\}}\}$
\mathcal{A}_4	$\mathcal{R}_1 \cup \{n_{\{B,C,D,E,F\}}, n_{\{A,B,C,E,F\}}\}$
\mathcal{A}_5	$\mathcal{R}_1 \cup \{n_{\{B,C,D,E,F\}}, n_{\{A,B,C,D,F\}}\}$
\mathcal{A}_6	$\mathcal{R}_1 \cup \{n_{\{A,B,C,E,F\}}, n_{\{A,B,C,D,F\}}\}$
\mathcal{A}_7	$\mathcal{R}_1 \cup \{n_{\{B,C,D,E,F\}}, n_{\{A,B,C,E,F\}}, n_{\{A,B,C,D,F\}}\}$
\mathcal{A}_8	Saturated

Table 4: Marginal posterior distributions under several models for the possible values of cell (1, 2, 2, 1, 1, 2) of the Czech Autoworkers data obtained by employing the “global move” method.

Model	Possible Cell Values										
	0	1	2	3	4	5	6	7	8	9	10
\mathcal{A}_0	0.1124	0.1375	0.1603	0.1571	0.1489	0.1172	0.0774	0.0497	0.0249	0.0113	0.0033
\mathcal{A}_1	0.3274	0.2564	0.1737	0.1060	0.0672	0.0354	0.0184	0.0093	0.0040	0.0015	0.0008
\mathcal{A}_2	0.2408	0.1764	0.1291	0.0948	0.0762	0.0576	0.0521	0.0442	0.0456	0.0394	0.0437
\mathcal{A}_3	0.1749	0.1553	0.1604	0.1449	0.1243	0.0929	0.0662	0.0404	0.0252	0.0110	0.0046
\mathcal{A}_4	0.3373	0.1566	0.1219	0.0907	0.0744	0.0545	0.0496	0.0393	0.0316	0.0248	0.0194
\mathcal{A}_5	0.2623	0.1674	0.1474	0.1188	0.0944	0.0704	0.0532	0.0328	0.0231	0.0153	0.0149
\mathcal{A}_6	0.2289	0.08423	0.0672	0.0690	0.0645	0.0619	0.0749	0.0776	0.0918	0.0831	0.0969
\mathcal{A}_7	0.2735	0.0915	0.0717	0.0590	0.0598	0.0638	0.0668	0.0709	0.0778	0.0724	0.0929
\mathcal{A}_8	0.2412	0.1030	0.0930	0.0873	0.0651	0.0632	0.0627	0.0538	0.0595	0.0519	0.1193

Table 5: Bounds for the Czech Autoworkers data given the marginals \mathcal{R}_2 (left-hand panel) and given the marginals \mathcal{R}_3 (right-hand panel). The bounds for the cell (1, 2, 2, 1, 1, 2) are marked with a box. These bounds were calculated using the generalized shuttle algorithm.

				B				B			
				no		yes		no		yes	
F	E	D	C	A	no	yes	no	yes	A	no	yes
neg	< 3	< 140	no	[27, 58]	[25, 56]	[96, 134]	[44, 82]	[0, 88]	[0, 62]	[0, 224]	[0, 117]
			yes	[108, 149]	[123, 168]	[0, 22]	[9, 37]	[0, 261]	[0, 246]	[0, 24]	[0, 38]
	≥ 140	no	[22, 49]	[0, 24]	[60, 96]	[16, 52]	[0, 88]	[0, 62]	[0, 224]	[0, 117]	
		yes	[91, 127]	[45, 85]	[0, 18]	[0, 20]	[0, 261]	[0, 151]	[0, 24]	[0, 38]	
	≥ 3	< 140	no	[10, 37]	[17, 44]	[48, 86]	[49, 89]	[0, 58]	[0, 60]	[0, 170]	[0, 148]
			yes	[30, 68]	[58, 102]	[0, 19]	[0, 25]	[0, 115]	[0, 173]	[0, 20]	[0, 36]
	≥ 140	no	[13, 37]	[8, 36]	[55, 90]	[38, 76]	[0, 58]	[0, 60]	[0, 170]	[0, 148]	
		yes	[30, 67]	[45, 86]	[0, 19]	[0, 27]	[0, 115]	[0, 173]	[0, 20]	[0, 36]	
pos	< 3	< 140	no	[0, 15]	[0, 13]	[4, 31]	[0, 23]	[0, 88]	[0, 62]	[0, 125]	[0, 117]
			yes	[0, 21]	[3, 30]	[0, 10]	[0, 9]	[0, 134]	[0, 134]	[1, 1]	[0, 38]
	≥ 140	no	[0, 11]	[0, 10]	[0, 24]	[0, 18]	[0, 88]	[0, 62]	[0, 125]	[0, 117]	
		yes	[0, 26]	[2, 30]	[0, 11]	[0, 9]	[0, 134]	[0, 134]	[0, 24]	[0, 38]	
	≥ 3	< 140	no	[1, 14]	[0, 9]	[0, 26]	[0, 26]	[0, 58]	[0, 60]	[0, 125]	[0, 125]
			yes	[0, 19]	[4, 29]	[0, 9]	[0, 9]	[0, 115]	[0, 134]	[0, 20]	[0, 36]
	≥ 140	no	[0, 9]	[0, 9]	[0, 26]	[0, 22]	[0, 58]	[0, 60]	[0, 125]	[0, 125]	
		yes	[0, 19]	[0, 23]	[0, 9]	[0, 13]	[0, 115]	[0, 134]	[0, 20]	[0, 36]	