13

Some explorations into Bayesian modelling of risks due to pesticide intake from food

Hilko van der Voet[#] *and M. João Paulo*[#]

Abstract

This paper presents some common types of data and models in pesticide exposure assessment. The problems of traditional methods are discussed in connection with possibilities to address them in a Bayesian framework. We present simple Bayesian models for consumption of food and for residue monitoring data.
Keywords: Bayesian model; Monte Carlo; MCMC; exposure assessment; food safety; risk analysis; dietary consumption; residue monitoring

Introduction

Food safety has become a major focus of attention for consumers, and therefore also for producers and other parties in the agro-food chain. In this paper we will consider the risks associated with the possible presence of pesticides in vegetables and fruits.

Traditionally the analysis of such risks by food-safety authorities has been in terms of deterministic estimates of exposure (*Submission and evaluation of pesticide residues data for the estimation of maximum residue levels in food and feed* 2002), but currently there is a shift in thinking from deterministic to probabilistic approaches (Ferrier et al. 2002). However, currently available probabilistic approaches require data of sufficient quality and quantity, which are often not available. Bayesian models hold a promise to allow an efficient use of prior knowledge or an efficient pooling of parameters.

The objective of this paper is to explore the possibilities of Bayesian models for exposure assessment in order to prepare more realistic risk-analysis methods in the field of food safety. The methods found in current Bayesian literature cannot be applied directly to the type of data which are commonly available, and therefore new models are developed.

Data for exposure assessment

For exposure assessment it is necessary to have data on food consumption (x) and residue levels in food (c). Also, it is necessary to have knowledge of the population of individuals for whom the exposure assessment will be assumed to be relevant (including knowledge of their body weights w). We will now successively consider consumption data and concentration data.

[#] Biometris, Wageningen University and Research centre, P.O. Box 100, 6700 AC Wageningen, Netherlands. E-mail: hilko.vandervoet@wur.nl and joao.paulo@wur.nl

The most common type of data on food consumption arise from National Food Consumption Surveys (Verger et al. 2002). In the European Union it is currently attempted to harmonize such national surveys in the interest of international comparability (Brussaard et al. 2002). In a Food Consumption Survey (FCS) a stratified random sample of consumers is asked to complete a diary writing down their consumption for a number of days. For example, from the Dutch National FCS of 1997/1998 (*Zo eet Nederland: resultaten van de voedselconsumptie-peiling 1997-1998* 1998) the food consumption of 6250 Dutch individuals in 2564 households is available for two consecutive days. Here we can already note that there is much structure in these data (stratification, households, consecutive days) which needs attention in statistical modelling. In current applications this data structure is often ignored.

A common feature in food consumption data sets is the sparseness, that is the fact that most products on any day are *not* consumed. In statistical parametric modelling it will be necessary to account explicitly for the spike of zero consumptions.

A disadvantage of usual FCS data for the assessment of chronic risks is that hardly any information about the long-term use of food products is included. An alternative type of data about consumption is provided by a food-frequency questionnaire (FFQ). An example is the Dutch VEG questionnaire on the consumption of vegetables and fruit (Van Dooren-Flipsen, Van Donkersgoed and Van Klaveren 1999). Here a typical question is: How often (in the last summer) did you eat apples? with 7 possible categorized answers ranging from *never or less than 1 day per month* to *6-7 days per week*. In addition it is then asked how many apples are typically consumed on such a day. Such data can be converted to rough estimates of the average consumption of apples for this person, but no information about the variability of consumption and about the between-product correlations is obtained.

In order to make food consumption data x compatible with concentration data c it will almost always be necessary to apply food conversion tables to transform the amount of products as eaten (e.g. pizza) to raw agricultural commodities (e.g. tomato). This conversion, currently often applied without regard for variability and uncertainty, is clearly another stage where probability models may be incorporated in a Bayesian context.

Data on concentrations of chemicals in raw agricultural commodities may be available from supervised trials or, for substances which are already in use for some time, from monitoring programmes. In this paper we consider the latter type of data from the Dutch Quality Programme for Agricultural Products (Van Klaveren 1998). Under this programme a database is filled with residue- and contaminant-monitoring results in many types of food. Here we will consider the example of the pesticide iprodione in vegetables and fruit. On a period of 5 years (1998-2002) iprodione was found on 53 vegetable and fruit commodities. In the monitoring database we had between 12 and 1367 measurements per commodity (average 232), but only between 1 and 293 of those were positive values (average 27). Among the 53 commodities there were only 14 (26 %) with at least 20 positive values, and for as much as 31 commodities (58 %) there were less than 10 positive values available. It can be remarked that this will present a problem in parametric modelling, especially considering the fact that in practice an exposure assessment is often based on data for a shorter time period than 5 years.

Often, a pesticide is found in higher concentrations on only a few commodities. For example in the iprodione example the highest concentrations were found on endive and kiwi fruit. But note that we cannot be certain that other commodities are

not relevant for acute exposure, because incidental high values may occur for other commodities, and also high consumptions with relatively low concentration levels may still contribute significantly to the exposure.

Models for exposure assessment

The basic model of exposure assessment is

$$y_{ij} = \frac{\sum_{k=1}^{p} x_{ijk}\ c_{ijk}}{w_i}$$

where y_{ij} is the intake by individual i on day j (in µg chemical substance per kg body weight), x_{ijk} is the consumption by individual i on day j of food commodity k (in g), c_{ijk} is the concentration of the chemical substance in commodity k eaten by individual i on day j (in mg/kg), and w_i is the body weight of individual i (in kg). Finally, p is the number of food commodities accounted for in the model.

In this section we discuss four types of model for exposure assessment:

1. deterministic models as currently in use by official organizations;
2. probabilistic models using simulation
   a) based on resampling empirical data;
   b) based on parametric modelling of data;
   c) Bayesian models (parametric modelling of data and prior information).

Types 2a and 2b are also known as Monte Carlo (MC) models, whereas type 2c can be implemented as a Markov Chain Monte Carlo (MCMC) model.

Traditionally, exposure assessment has been done without regard for variability, that is, by multiplying *average* consumption by *average* concentration for all commodities. Considering that the risks of acute exposure could be underestimated due to the big variability of both consumptions and concentrations, a natural reaction has been to apply worst-case values, leading to exposure estimates of the type

$$y = \frac{x_{97.5} \cdot c_{\max}}{w_{mean}}$$

where $x_{97.5}$ is the 97.5th percentile in the consumption distribution of the eaters of the chosen commodity, $c_{\max}$ is the largest observed residue concentration, and w_{mean} is the mean body weight in a population. This exposure estimate has been further developed into the International Estimate of Short-Term Intake (IESTI) which is now adopted by the Joint FAO/WHO Meeting on Pesticide Residues (JMPR) for the evaluation of residue data when maximum residue levels (MRLs) in food and feed have to be estimated (*Submission and evaluation of pesticide residues data for the estimation of maximum residue levels in food and feed* 2002). It is clear that the actual intake is overestimated by the combination of two conservative (worst case) values from distributions that are independent. It may be noted that this type of estimate can only be applied to single food commodities because summation over commodities makes no sense as this would correspond to a simultaneous worst case situation for all products.

In a Monte Carlo analysis using empirical data (type 2a) the empirical intake distribution is estimated by combining, in a large number of iterations, the consumption pattern and body weight of a randomly selected consumer with concentration values randomly selected from the monitoring data (allowing for the observed proportion of non-detects). See Figure 1 for an example of such an empirical distribution, obtained with the programme MCRA (Van Der Voet et al. 2003). The

risk assessment is usually made by comparing toxicological limits, such as for example the acute reference dose, with percentiles of this distribution. Figure 2 shows selected percentile estimates, together with bootstrap estimates of the uncertainty in each percentile estimate due to sampling errors of the consumption and concentration data (the bootstrap distributions were obtained by calculating the selected percentiles in 1000 intake distributions based on 1000 Monte Carlo iterations each).

Although this type of Monte Carlo analysis performs well in cases where sufficient regular data are available, there are often problems which prompted us to investigate the possibilities of parametric modelling in a frequentist or fully Bayesian framework:

1. The quantity of data is often small. For example, winter carrot was only consumed on 5 out of 12,500 days, and over a period of five years only 12 measurements of iprodione in kohlrabi have been made (one of which was positive). In such cases resampling the data can only generate a strictly limited set of values. Fitting a parametric distribution to the data allows the generation of a more smooth distribution of future values.
2. With empirical modelling each commodity is separate, and in that sense empirical models are maximally complex. Parametric modelling allows models with common parameters, e.g. coefficients of variation that are equal for all commodities in a certain class. Bayesian modelling allows the specification of priors to estimate parameters which are intermediate between ‘all separate’ and ‘all equal’.
3. Sometimes no primary data are available, but only data in summary form, for example as means, standard deviations, percentiles or maxima. This type of information can only be used in connection with a distributional assumption to generate new values.
4. Much of the structure in the data sets (e.g. days within persons within households) is commonly ignored in an empirical analysis. Hierarchical modelling techniques seem appropriate to take such structure into account.
5. There are analytical limitations to detect low residue levels. This means that the failure to find a positive concentration (non-detect) actually can be due to two causes: either the pesticide is really absent (pesticide was never used) or it is present in a low concentration below the limit of reporting. In parametric analysis this can be modelled as a mixture of distributions.

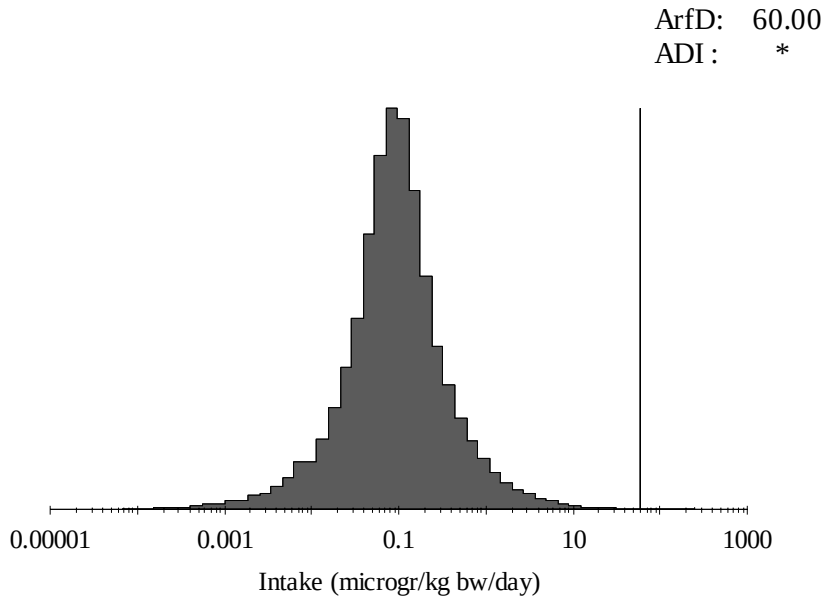

Figure 1. Empirical intake distribution of iprodione from vegetables and fruit based on 1,000,000 iterations. Non-detects were replaced by the limit of reporting. ArfD is the acute reference dose, which is the toxicological limit-value set for acute risks

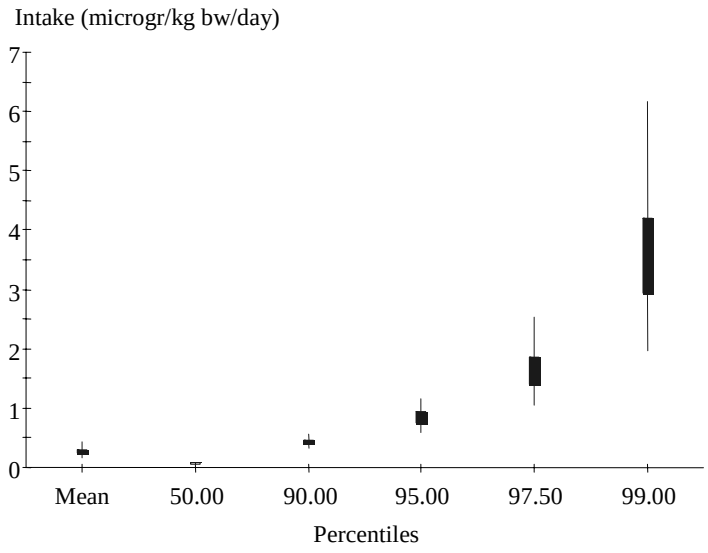

Figure 2. Estimated mean and selected percentiles of empirical intake distribution. Bars and line segments limit the quartiles and 2.5% and 97.5 % points of the uncertainty distribution based on bootstrapping

A Bayesian model for consumption

As an exercise in Bayesian modelling we developed a simple model for the example data set. Given the sparseness of the data it was considered necessary to model the binary result (commodity consumed yes or no) separate from the quantitative result (amount eaten). For one commodity the model can be written as:

$$p(B_j) = Bernoulli(\pi_j); \quad p(x_j \mid B_j) = \begin{cases} \delta(0) & \text{if } B_j = 0 \\ LogNormal(\mu_j, \sigma_j^{\,2}) & \text{if } B_j = 1 \end{cases}$$

where B_j is an indicator function for consumption, and x_j is the amount of product j consumed. $\delta(0)$ represents a spike at $x_j = 0$. There are three parameters (the probability of consumption π_j and the mean μ_j and variance σ_j^2 of the ln-transformed consumptions) for which we need to specify prior distributions. We chose the conventional prior distributions for the parameters of the (log)normal distribution, that is a normal prior distribution for μ_j and an prior inverse-chi square distribution for σ_j^2. In a multivariate context (with p commodities) this is easily generalized to a multivariate normal distribution and an inverse-Wishart distribution, respectively (see e.g. Gelman et al. 1995). Therefore a straightforward Bayesian model for the consumed amount is:

$$\ln(x) \sim N_p(\mu_x, \Sigma_x); \quad \mu_x \sim N_p(\mu_0, \Sigma_0); \quad \Sigma_x \sim W_p(R, df).$$

Typically, in an analysis where there are sufficient data, the prior distributions will be taken vague by specifying covariance matrices Σ_0 and R equal to e.g. 1000 times the identity matrix, and means μ_0 equal to 0. The degrees of freedom df should be at least p+2 in order for the mean of R^{-1} to exist. This choice therefore gives the maximum amount of vagueness attainable in the inverse Wishart distribution.
For the Bernoulli parameter π_j, the conventional conjugate prior is the beta distribution. However, in a multivariate context this model is less attractive because there is no straightforward generalization to the multivariate case. We investigated two possibilities: a logit-normal model and a loglinear model.
In the logit-normal model we introduce an auxiliary variable z_j, defined as the logit transform of the Bernoulli parameter π_j, and then assume a normal prior distribution for z_j. This can be easily generalized to a multivariate normal distribution in a model for more food commodities. The multivariate Bayesian model for binary food consumption specified by a binary vector B is then specified by

$$B \sim Bern(\pi); \qquad \text{logit}(\pi) \sim N_p(\mu_\pi, \Sigma_\pi).$$

In the loglinear model the binary consumption patterns are tabulated in a 2^p contingency table. The counts n_B for each possible consumption pattern B are then modelled as Poisson variables with a parameter given by a loglinear model:

$$n_B \sim Poisson(\lambda_B); \qquad \lambda_B = \exp\left(u_0 + \sum_j B_j u_{1j} + \sum_j \sum_{k<j} B_j B_k u_{2jk} + ... \right)$$

where u_0, u_{1j} and u_{2jk} are parameters describing the mean, the main effects of the components of *B*, and the two-factor interaction effects, respectively. The model is easily extended with higher-order interactions. All *u* parameters are given vague normal prior distributions with mean 0 (exception: u_0 has a higher mean, e.g. 5) and a large variance (*e.g.* 1000):

$$u_0 \sim N(\mu_0, \sigma^2); \quad u_{1j} \sim N(0, \sigma^2); j = 1,..., p;$$
$$u_{2jk} \sim N(0, \sigma^2); j = 1,..., p; k = 1...j-1.$$

Although both models for binary patterns could be used, we ultimately preferred the loglinear model, because it is the most general model allowing to model any kind of interaction, and because in the logit-normal model correlations between commodities could only be modelled via the prior distribution.

Example: consumption of relatively-high-risk commodities

The consumption models were implemented in Winbugs 1.4 (Spiegelhalter et al. 2003), and could be used to generate posterior distributions for the parameters, and predictive distributions for consumption *p*(*x*). However, the models could only be applied to a limited set of food commodities. For the full set of 53 commodities in our example data set there are far too many parameters that would have to be estimated. Therefore in this example we first ran the empirical Monte Carlo model to identify the four commodities which contributed most to the upper tail (upper 2 %) of the exposure distribution. These four commodities turned out to be endive, cabbage lettuce, grape and kiwi fruit (relative contributions to the tail 21, 15, 12 and 10 %). The frequencies of occurrence are shown in Table 1.

Table 1. Occurrence of four relatively-high-risk commodities in the diet of 6250 consumers (fractions in parentheses)

	3. kiwi fruit	0		1	
	4. grape	0	1	0	1
1. endive	2. cabbage lettuce				
0	0	3088	1938	83	116
		(0.494)	(0.310)	(0.013)	(0.019)
	1	282	195	9	10
		(0.045)	(0.031)	(0.001)	(0.002)
1	0	230	138	10	11
		(0.037)	(0.022)	(0.002)	(0.002)
	1	76	54	4	6
		(0.012)	(0.009)	(0.001)	(0.001)

Table 2. Posterior statistics (15000 samples in three Markov chains, after a burn-in run of 5000 samples each) for cell probabilities (*prob*), mean log amounts (*mu.x*), and amounts consumed (*yy*). Indices of *prob* correspond with the row-wise order of cells in Table 1

node	mean	sd	median	90.0%	95.0%	97.5%
prob[1]	0.4942	0.006166	0.4942	0.5022	0.5044	0.5063
prob[2]	0.3103	0.005729	0.3102	0.3177	0.3198	0.3216
prob[3]	0.01352	0.001383	0.01349	0.01531	0.01587	0.01636
prob[4]	0.0181	0.001606	0.01804	0.02019	0.02082	0.02137
prob[5]	0.04491	0.002522	0.04487	0.04818	0.0491	0.04989
prob[6]	0.03123	0.002045	0.03116	0.03386	0.03465	0.03546
prob[7]	0.001309	2.835E-4	0.001291	0.001684	0.00181	0.001922
prob[8]	0.001938	3.987E-4	0.001916	0.002454	0.002632	0.002801
prob[9]	0.03644	0.002239	0.0364	0.03932	0.04021	0.04094
prob[10]	0.02212	0.001699	0.02208	0.02433	0.02499	0.02557
prob[11]	0.001564	3.248E-4	0.001542	0.001993	0.002129	0.002241
prob[12]	0.002022	4.08E-4	0.001995	0.002559	0.002722	0.002882
prob[13]	0.01253	0.00118	0.01251	0.01407	0.01454	0.01493
prob[14]	0.008429	9.069E-4	0.008395	0.009611	0.009993	0.0103
prob[15]	5.701E-4	1.529E-4	5.535E-4	7.814E-4	8.542E-4	9.126E-4
prob[16]	8.166E-4	2.186E-4	7.882E-4	0.00111	0.001222	0.001319
mu.x[endive]	2.972	0.1073	2.972	3.11	3.15	3.185
mu.x [cabbage lettuce]	3.028	0.07216	3.027	3.12	3.149	3.17
mu.x [kiwi fruit]	3.673	0.1574	3.675	3.873	3.931	3.981
mu.x [grape]	2.024	0.04342	2.024	2.08	2.094	2.109
yy[endive]	72.23	4596.0	0.0	0.0	12.92	93.55
yy[cabbage lettuce]	10.16	130.4	0.0	0.0	19.9	64.0
yy[kiwi fruit]	51.51	2556.0	0.0	0.0	0.0	19.46
yy[grape]	40.47	855.8	0.0	32.44	87.55	192.6

The posterior statistics in Table 2 show that a loglinear model with only main effects and two-factor interactions could well estimate the original cell probabilities in Table 1 (for example, prob[1] nicely estimates 3088/6250 = 0.494). More importantly, the posterior predictive consumption distributions, integrated over consumers and non-consumers, show that more than 50% of the intakes are 0 for all 4 products (median is zero). The 97.5th percentiles are estimated at 94 g of endive, 64 g of cabbage lettuce, 19 g of kiwi fruit, and 193 g of grapes. It may be noted that among these four products, grape has the lowest mean daily portion size (among grape consumers the geometric mean consumption is $e^{mu.x[grape]} = e^{2.024} = 8$ g against *e.g.* $e^{mu.x[kiwi\ fruit]} = e^{3.673} = 39$ g). However, the 97.5th percentile in the distribution of consumption of all individuals (both consumers and non-consumers) is 193 g, which is higher than for the other products (e.g. 19 g for kiwi fruit). These differences can be explained by the proportion of consumers and the different variability of consumption sizes (both higher for grape than for kiwi fruit).

In a further analysis (outside the scope of this paper) it will now be straightforward to combine these consumption distributions with the distributions for concentration derived in the second part of this paper.

Chemical residue present in food

The residue data come from monitoring programmes run at several stages of the agro-food chain. Monitoring programmes are run to ensure that the residue concentration present in the food commodities is below the safe limit as set by the government. Data from monitoring programmes are also useful to assess the risk posed to consumers. Monitoring data essentially are a collection of measurements performed on a number of food samples on the concentration of the residue present in the sample. The residue concentration present in food strongly depends on the sort

and origin of the crop, and in the case of pesticides, it also depends on whether use of that particular pesticide is allowed on the crop or not. Apart from that, the number of samples taken varies very much from one food commodity to another. To make things worse, the measurement devices are not sensitive to very low concentrations, and only results above a certain value are reported, called the Limit of Reporting (LOR). A sample of food is either found to be contaminated with a concentration *c* above the LOR, or the residue is not detected (*ND*). If no residue is detected in a sample it obviously does not necessarily imply that residue is not present; it may very well be that residue is present at a low concentration, below the LOR. Formally, the residue concentration present in a food sample, c_j^*, will be either zero or positive, with a given probability:

$$p(c_j^* \mid I_j) = \begin{cases} \delta(0) \text{ if } I_j = 0 \\ f(\theta) \text{ if } I_j = 1 \end{cases}$$

where we further define I_j as a variable taking the value 0 if the chemical in question was not used on the population from where sample *j* was collected, and taking value 1 otherwise. The observed/measured concentration c_j on sample *j* is then

$$c_j = \begin{cases} ND \text{ if } c_j^* < \text{LOR} \\ c_j^* \text{ if } c_j^* \geq \text{LOR} \end{cases}$$

The lognormal distribution has been found to describe the positive concentrations very well, and therefore is usually used as $f(\theta)$. The lognormal density function is given by

$$f(c) = \frac{1}{\sqrt{2\pi}\sigma c} \exp\left\{ -\frac{(\ln c - \mu)^2}{2\sigma^2} \right\}, \quad c > 0$$

with parameters μ and σ. It has the nice property that if $c \sim LN(\mu, \sigma)$ then $\ln(c) \sim N(\mu, \sigma)$. Figure 3 shows the histogram of iprodione residue found in strawberry (the y-axis is set in two scales to accommodate the 1096 non-detects, shown below the LOR, and 271 detects), and Figure 4 shows the same data, ln-transformed (y-axis shown in two scales, point LOR in the x-axis is actually ln(LOR)).

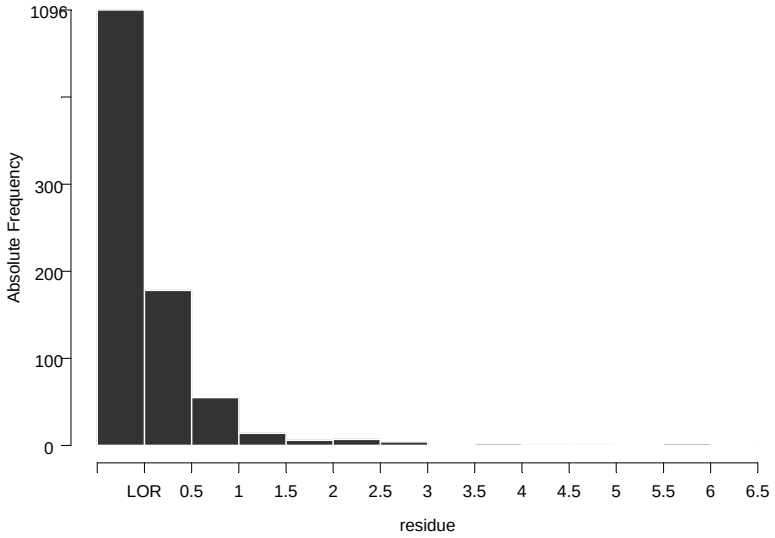

Figure 3. Histogram of iprodione residue found in strawberry. The limit of reporting is 0.02. The y-axis was set in two scales to accommodate the 1096 non-detects, shown below the LOR, and 271 detects

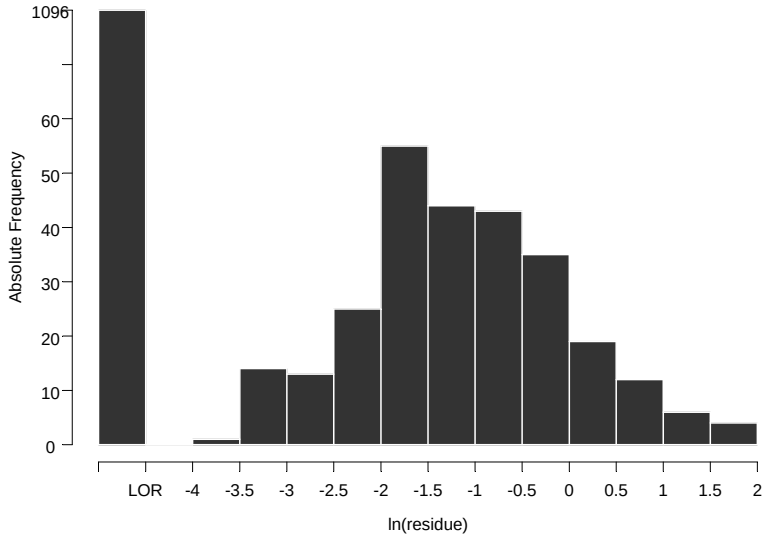

Figure 4. Histogram of the *ln-transformed* iprodione residue found in strawberry. The y-axis was set in two scales to accommodate the 1096 non-detects and 271 detects. LOR in the x-axis is actually ln(LOR)

Model for ln(residue)

We model the number of observed zeros (i.e. censored concentrations and real zeros) separately from the sample of detects:

$$\begin{cases} N_0 \sim Bin\left(N\ ,\ p_0 + (1-p_0)\Phi(\alpha) \right) \\ p(c) = \dfrac{1}{(1-\Phi(\alpha))\sqrt{2\pi}\sigma} \exp\left(-\dfrac{(\ln(c)-\mu)^2}{2\sigma^2} \right), \quad c \geq \mathrm{LOR} \end{cases} \tag{1}$$

The number of observed zeros (N_0) has a binomial distribution with probability that an observed zero is either a real zero or a non-detect. The detects (c) are a truncated sample from a lognormal distribution, $\alpha = (\ln(\mathrm{LOR}) - \mu)/\sigma$, and $\Phi(\alpha)$ is the cumulative distribution function of the standard normal at α.

Example: iprodione concentrations in vegetables and fruit

Iprodione concentrations were measured in 53 sorts of vegetables and fruit. A different number of samples were taken in each food commodity, the total number of samples varied between 12 and 1367. The number of detects varied between 1 and 293, since the LOR for iprodione was 0.02 mg/kg (in practice the LOR is likely to vary from one laboratory to another, but here it is assumed to be constant).
The ln-transformed observed positive concentrations were tested as to whether their underlying distribution was Normal, using the Anderson-Darling goodness-of-fit test. The Anderson-Darling test (Anderson and Darling 1954) compares the empirical distribution $F_n(c)$ of a sample of size n with a given distribution function $F(c)$, through the weighing of the squared difference of the two functions with $F(c)(1-F(c))$. The difference is then integrated over the domain of $F(c)$:

$$A^2 = n \int \left(F_n(c) - F(c)\right)^2 / \left[F(c)(1-F(c))\right]\, dF(c) .$$

This test has similarities with the Kolmogorov-Smirnov test, but has some advantages over the latter, namely that it allows parameter estimation from the data, and it is more sensitive at the tails of the distribution. A discrete version of the test is obtained for the ordered observations $\{c_1, c_2, \ldots, c_n\}$ (Stephens 1974), and it is given by

$$A^2 = -n - \sum_{i=1}^{n} \frac{2i-1}{n} \left[\ln F(c_i) + \ln\left(1 - F(c_{n+1-i})\right)\right]$$

The test rejects normality at the 95% confidence level for test values above 0.752. The test was performed on the products with a number of 4 or more detects. For each product two normality tests were carried, using the sample mean and sample standard deviation, respectively (ignoring the fact that the observations are truncated at the LOR), and using the maximum likelihood parameters of the truncated normal distribution. Figure 5 shows the resulting test statistics for the second test (maximum likelihood parameters) versus sample size.

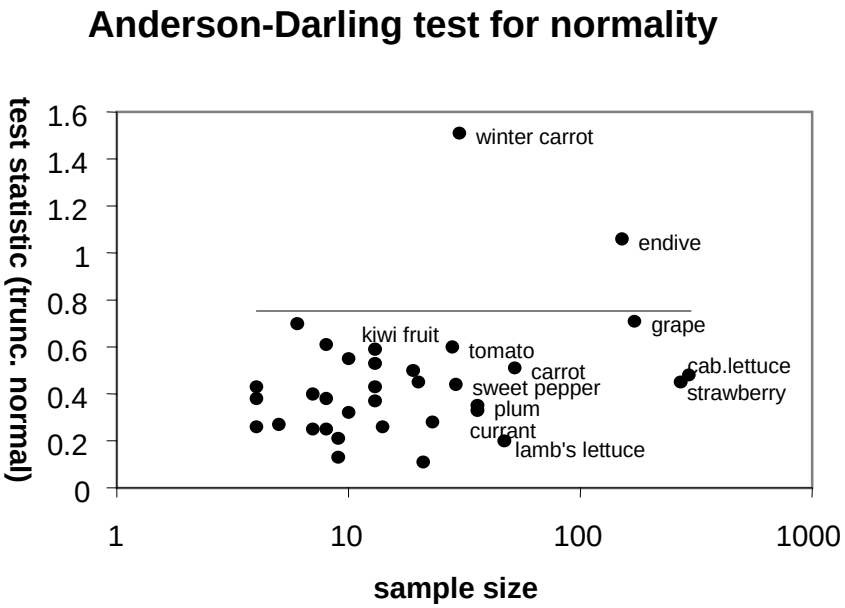

Figure 5. Results of normality test for 34 products with sample size n ≥ 4. The line represents the critical value at the 95% confidence level

The two tests produced very similar results, except for one commodity, kiwi fruit, where the first distribution was rejected, and the second was not. For the 34 products, the test rejected normality in 2 samples, endive and winter carrot, even after accounting for truncation. The histograms of the ln-transformed positive concentrations in endive and winter carrot are shown in Figure 6 and Figure 7, respectively. Although we recommend caution when estimating the risks associated to these two products, we find it reasonable to assume a lognormal distribution for the residue for all other products.

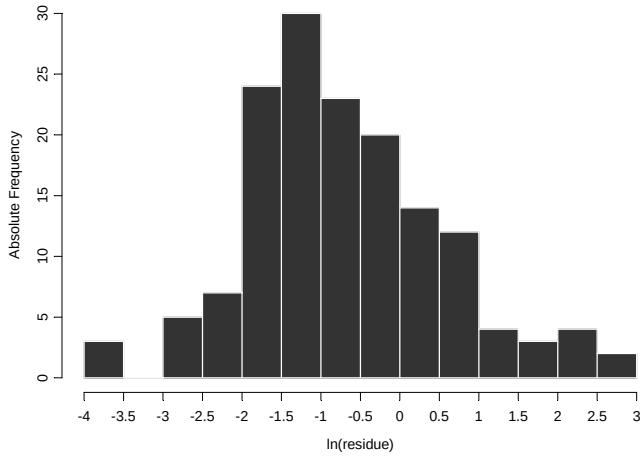

Figure 6. Histogram of the *ln-transformed* positive residue found in endive. The Anderson-Darling goodness-of-fit test rejected normality for this sample, even after accounting for left truncation

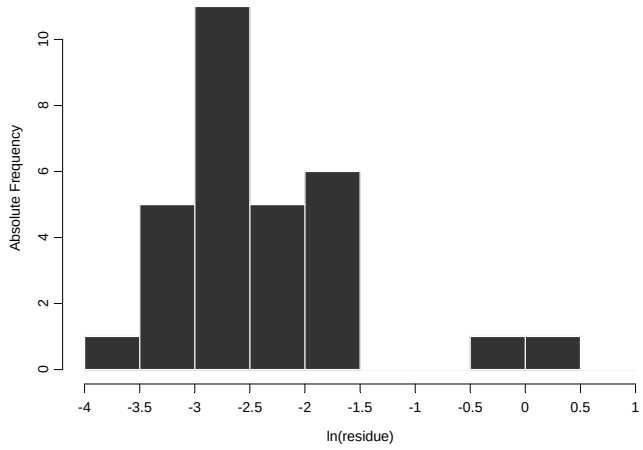

Figure 7. Histogram of the *ln-transformed* positive residue found in winter carrot. The Anderson-Darling goodness-of-fit test rejected normality for this sample, even after accounting for left truncation

Model (1) was implemented in WinBUGS 1.4 (Spiegelhalter et al. 2003), the current version of BUGS. BUGS (**B**ayesian inference **U**sing **G**ibbs **S**ampling) is a software for Bayesian analysis that uses Markov chain Monte Carlo (MCMC) methods.
Model (1) was coded in BUGS as follows:

```
model truncated{

  low   <- log(LOR)
  sigma <- sqrt(1/prec)
  pcens <- phi((low - mu)/sigma)
  pzm   <- p0+(1-p0)*pcens
  big   <- 1.0E6

  N0     ~ dbin(pzm,N)
  for(i in 1:Npos){
     z[i]    <- (ln.res[i] - mu)/sigma
     arg[i]  <- -pow(z[i],2)/2
     dens[i] <- sqrt(prec/(2*3.14159))*exp(arg[i])/(1-pcens)
     like[i] <- dens[i]/big
     ones[i] <- 1
     ones[i] ~ dbern(like[i])
  }

# priors
  mu   ~ dnorm(0,0.01)
  prec ~ dgamma(0.1,0.1)
  p0   ~ dbeta(1,1)

}# end-of-model
```

We show the results of two applications of the model, namely to residue found in strawberry and in kiwi fruit.

Residue in strawberry

A positive residue was detected in 271 samples of strawberry, out of a total of 1367 samples. The histogram of the ln-transformed detected residue with non-detects (residue below the LOR) is shown in Figure 4. The ln-detected residue in the data has a sample mean equal to -1.2 and sample standard deviation 0.9. Since we observe a truncated sample, we are more interested in obtaining the parameters of the original distribution. We used vague priors for the model parameters to express our lack of knowledge in this situation. The resulting statistics of the stationary posterior distributions are shown in Table 3. The posterior kernel densities are shown in Figure 8. The posterior distributions synthesize how much information we have about the parameters from the prior distributions and from the data combined. The mean posterior μ and σ (mu and sigma) are -1.2 and 1.2, respectively (slightly higher posterior σ than in the data). The posterior probability of observing a real zero is between 0.78 and 0.82 with a 95% chance, with mean at 0.8, i.e. about 20% of the samples are likely to be positive. In our data set we could expect to have 273 positive samples, i.e. 2 censored samples.

Table 3. Posterior statistics in strawberry

node	mean	sd	MC error	2.50%	median	97.50%
p0	0.799	0.011	0.000	0.777	0.799	0.820
Pcens	0.011	0.004	0.000	0.005	0.011	0.021
pzm	0.801	0.011	0.000	0.780	0.802	0.822
mu	-1.198	0.076	0.001	-1.348	-1.197	-1.052
sigma	1.183	0.058	0.000	1.076	1.180	1.305

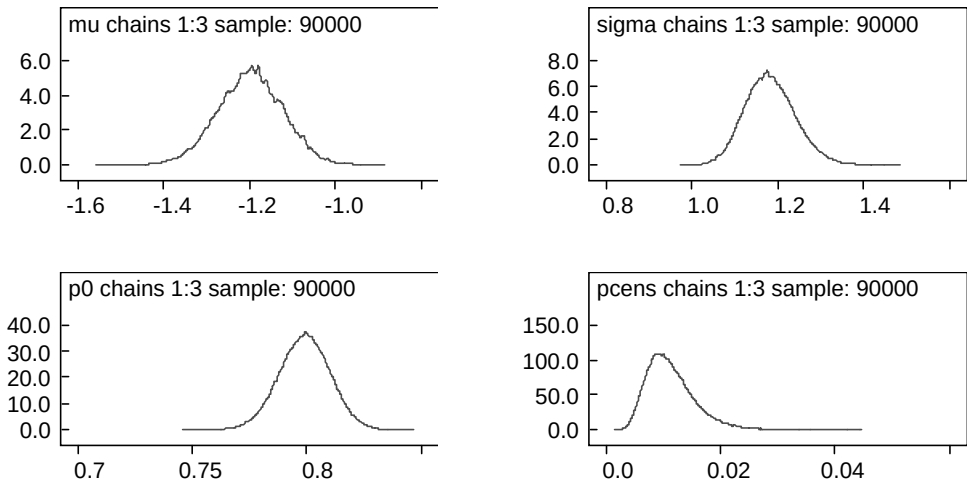

Figure 8. Posterior kernel densities in strawberry

Residue in kiwi fruit

In kiwi fruit there were 115 samples taken, residue was detected in 13 samples. The histogram of the residue in kiwi fruit can be seen in Figure 9, and the histogram of the ln-transformed residue is shown in Figure 10.

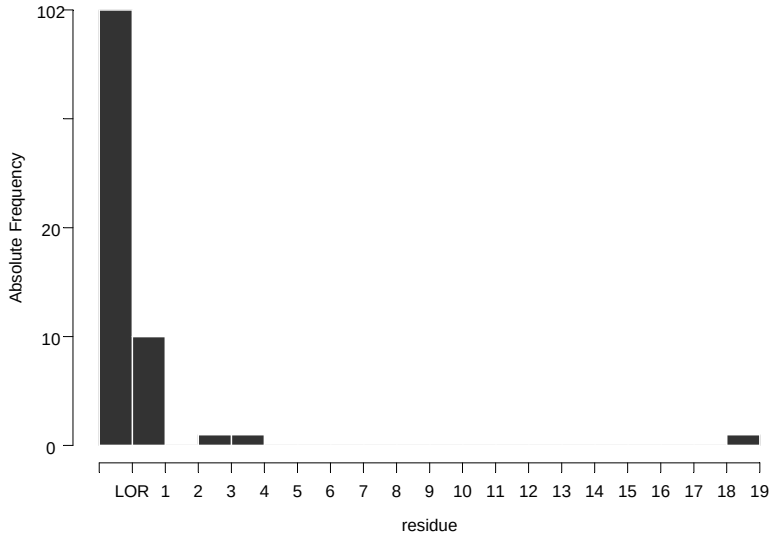

Figure 9. Histogram of iprodione residue found in kiwi fruit. The limit of reporting is 0.02. The y-axis was set in two scales to accommodate the 102 non-detects, shown below the LOR, and 13 detects

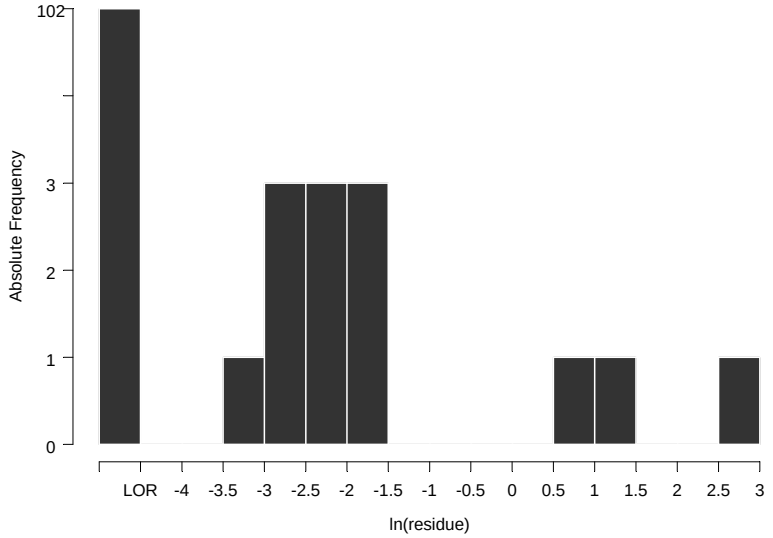

Figure 10. Histogram of the *ln-transformed* iprodione residue found in kiwi fruit. The y-axis was set to accommodate the 102 non-detects and 13 detects. LOR in the x-axis is actually ln(LOR))

The ln-detected residue has sample mean equal to -1.4 and sample standard deviation 5.0, meaning that the detected residue has a very large dispersion (in fact a very high concentration, 18.4 mg/kg, was measured in kiwi fruit). The Bayesian model was solved as before, using vague priors, and the resulting posterior statistics and kernel densities are displayed in Table 4 and Figure 11, respectively. The posterior distributions have a lot of variation. The mean posterior p_0 is 0.53, but in fact $p_0 \in [0.04, 0.88]$ with a 95% probability. In our data set we could expect 47% of

the samples with a positive concentration (54 positive samples), meaning that a substantial part of the concentrations would be below the LOR and was not detected. The mean posterior μ (-6.4) is also much lower than the sample mean, as one would expect if there is strong censoring in the data. These values need to be handled with caution though, since variation in the posterior distributions is so large.

Table 4. Posterior statistics for kiwi fruit

node	mean	sd	MC error	2.50%	median	97.50%
p0	0.531	0.253	0.005	0.036	0.576	0.881
Pcens	0.662	0.211	0.005	0.152	0.730	0.905
pzm	0.885	0.029	0.000	0.822	0.888	0.936
mu	-6.413	2.817	0.066	-12.030	-6.409	-1.637
Sigma	4.150	1.369	0.026	1.994	3.987	7.297

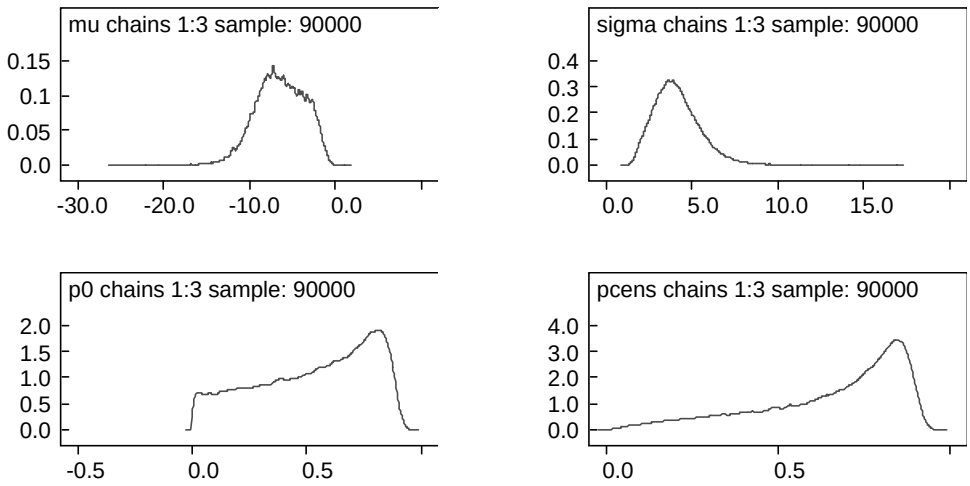

Figure 11. Posterior kernel densities in kiwi fruit

Discussion

In this paper we analysed some uses of Bayesian models for risks due to pesticide intake from food. Such models can in principle be applied to the type of data which is available from food consumption surveys and residue-monitoring programmes. We presented separate models for consumption and concentration. A next (and trivial) step will be to combine these two submodels into a model for pesticide intake. It will be interesting to incorporate more of the detailed structure of the data, such as hinted at in this paper, into such a model, and to devise a proper weighting scheme of prior information and available data.

Bayesian methods have several advantages. First, in principle all available information can be combined to obtain posterior distributions. This is certainly important in food safety, where information often comes from several sources and data can be very scarce. For example, Bayesian approaches seem to offer possibilities to combine the information of both FCS data and FFQ data in one model. For example, in a situation where the FFQ data have been obtained for a subset of the FCS

sample, a sketch of a useful hierarchical model (restricting ourselves here to one product) may be

$$xFCS_{ij} = mean + person_i + day_{ij} \quad (i = 1,...,6000)$$
$$xFFQ_i = mean + person_i + error_i \quad (i = 1,...,1500)$$

where, with *i* and *j* indices for person and day, *xFCS* and *xFFQ* refer to (the logarithm of) daily consumption as estimated from FCS and FFQ respectively, *mean* and *person* are mean consumption and person-specific deviation (equal for both submodels), *day* is the day-specific deviation in the FCS data (including any error component), and *error* is the measurement error in the FFQ data. Parameters to estimate in this simple model would be the mean consumption, and three variance components for between-person variation, between-day variation and questionnaire error variation. Note that the last two components can be separately estimated thanks to the combination with the FCS model. Expanding such a hierarchical model to a full Bayesian model is standard (Gelman et al. 1995), at least for simple cases.

A second advantage of Bayesian modelling is the fact that non-informative priors can be used where no information is available. We showed two examples for which no prior information was available. In the first example, iprodione residue in strawberry, there was a substantial number of detects, yielding posterior distributions with small variances. In the second example, iprodione residue in kiwi, there was a small number of detects and as a result the posterior distributions had large variances.

Third, hierarchical models are more naturally specified in a Bayesian framework. They allow us to combine data from different sources. Hierarchical models can be applied, for example, to the consumption of several commodities by individuals grouped in households. They can also be of great use to combine information from similar products, when modelling the amount of residue for commodities with a very small number of detects. Fourth, there is no limit for Bayesian inference either in the amount of prior information or in the amount of data, so that Bayesian inference is possible even with small samples.

Conclusions

Bayesian modelling of the multivariate consumption pattern was possible in simple cases (few commodities), but needs to be elaborated for higher dimensions. The modelling of chemical-residue data was easier because this could be done univariately. The lognormal distribution proved to be a sensible choice for the positive iprodione concentrations found in almost all commodities.

Acknowledgements

We like to thank our colleagues Michiel Jansen and Cajo ter Braak for their great help in formulating Bayesian models for this project, and Gerda van Donkersgoed (RIKILT Institute of Food Safety) for making available the data.

References

Anderson, T.W. and Darling, D.A., 1954. A test of goodness of fit. *Journal of the American Statistical Association,* 49, 765-769.

Brussaard, J.H., Lowik, M.R.H., Steingrimsdottir, L., et al., 2002. A European food consumption survey method: conclusions and recommendations. *European Journal of Clinical Nutrition,* 56 (Supplement 2), S89-S94.

Ferrier, H., Nieuwenhuijsen, M., Boobis, A., et al., 2002. Current knowledge and recent developments in consumer exposure assessment of pesticides: a UK perpective. *Food Additives and Contaminants,* 19 (9), 837-852.

Gelman, A., Carlin, J.B., Stern, H.S., et al., 1995. *Bayesian data analysis*. Chapman and Hall, London. Texts in statistical science series.

Spiegelhalter, D., Thomas, A., Best, N., et al., 2003. *WinBUGS user manual, version 1.4*. Available: [http://www.mrc-bsu.cam.ac.uk/bugs] (1 Sep 2003).

Stephens, M.A., 1974. EDF statistics for goodness of fit and some comparisons. *Journal of the American Statistical Association,* 69, 730-737.

Submission and evaluation of pesticide residues data for the estimation of maximum residue levels in food and feed, 2002. Food and Agriculture Organization of the United Nations, Rome. [http://www.fao.org/ag/agp/agpp/pesticid/jmpr/Download/faom2002.doc]

Van Der Voet, H., De Boer, W.J., Boon, P.E., et al., 2003. *MCRA: a Genstat program for Monte Carlo Risk Assessment. Release 2: reference guide*. Available: [http://www2.rikilt.dlo.nl/mcra/Referencemanual.pdf] (1 Sep 2003).

Van Dooren-Flipsen, M.M.H., Van Donkersgoed, G. and Van Klaveren, J.D., 1999. *VEG-voedselfrequentievragenlijst 97/98 : 1. Blootstelling aan nitraat via de voeding : 2. Rapportage consumptiegegevens*. Rikilt-Dlo, Wageningen. Rapport / RIKILT-DLO 99.005.

Van Klaveren, J.D., 1998. *Programme for the quality of agricultural products (KAP): results residue monitoring in the Netherlands*. Agricultural Research Department (DLO), Wageningen.

Verger, P., Ireland, J., Moller, A., et al., 2002. Improvement of comparability of dietary intake assessment using currently available individual food consumption surveys. *European Journal of Clinical Nutrition,* 56 (Supplement 2), S18-S24.

Zo eet Nederland: resultaten van de voedselconsumptie-peiling 1997-1998, 1998. Voedingscentrum, Den Haag.