

# BUGS in Bayesian stock assessments

Renate Meyer and Russell B. Millar

**Abstract:** This paper illustrates the ease with which Bayesian nonlinear state–space models can now be used for practical fisheries stock assessment. Sampling from the joint posterior density is accomplished using Gibbs sampling via BUGS, a freely available software package. By taking advantage of the model representation as a directed acyclic graph, BUGS automates the hitherto tedious calculation of the full conditional posterior distributions. Moreover, the output from BUGS can be read directly into the software CODA for convergence diagnostics and statistical summary. We illustrate the BUGS implementation of a nonlinear nonnormal state–space model using a Schaefer surplus production model as a basic example. This approach extends to other assessment methodologies, including delay difference and age-structured models.

**Résumé :** Le présent article montre la facilité avec laquelle des modèles d'espace d'états Bayésiens non linéaires peuvent maintenant être utilisés dans l'évaluation des stocks de pêche. On échantillonne des unités dans la fonction de densité a posteriori conjointe en utilisant la méthode d'échantillonnage de Gibbs avec BUGS, un logiciel facilement disponible. En tirant avantage de la représentation du modèle sous la forme d'un graphe acyclique orienté, BUGS automatise le calcul jusqu'ici fastidieux des distributions conditionnelles a posteriori complètes. De plus, le produit de BUGS peut être lu directement dans le logiciel CODA pour établir un diagnostic de convergence et un sommaire statistique. Nous illustrons l'application d'un modèle d'espace d'états non linéaire et non normal en utilisant, à titre d'exemple, un modèle de production excédentaire de Schaefer. Cette approche peut être étendue à d'autres méthodes d'évaluation, notamment aux modèles à différences retardées et aux modèles structurés selon l'âge.

[Traduit par la Rédaction]

## Introduction

State–space models are among the most powerful tools for dynamic modeling and forecasting (Fahrmeir and Tutz 1994). They have started to enjoy an increasing popularity in fisheries stock assessment (Sullivan 1992; Pella 1993; Gudmundsson 1994; Schnute 1994; Freeman and Kirkwood 1995; Kinas 1996; Reed and Simons 1996; Meyer and Millar 1999; Millar and Meyer 1999a) because they can realistically account for both measurement and process error. However, unrealistic assumptions such as linearity of state transitions and Gaussian error distributions are imperative for maximum likelihood estimation via Kalman filtering and have limited the number of fisheries models that can be fit within the classical/frequentist paradigm.

Meyer and Millar (1999) explained how delay difference and surplus production models can be cast into the framework of state–space modeling. They demonstrated a fully Bayesian approach using Gibbs sampling for posterior computation following Carlin et al. (1992). In contrast with the classical approach, the Bayesian approach can easily handle realistic distributional assumptions as well as nonlinearities in state and observation equations. The papers by Millar and Meyer (1999a) and Meyer and Millar (1999) provide the un-

derlying theory of fitting Bayesian state–space surplus production and delay difference models, respectively.

Here, we report on significant progress made in facilitating the routine implementation that may have a revolutionary effect on Bayesian stock assessment in everyday practice. This is achieved through BUGS (Bayesian inference using Gibbs sampling), a recently developed software package (Spiegelhalter et al. 1996) by the Medical Research Council Biostatistics Unit, Institute of Public Health, Cambridge, England. BUGS samples from the joint posterior distribution by using the Gibbs sampler (Gilks et al. 1996), i.e., by cyclically sampling from each of the full conditionals. For reviews on BUGS the reader is referred to Thomas et al. (1992), Gilks et al. (1994), and Gentleman (1997). BUGS is available free of charge from <http://www.mrc-bsu.cam.ac.uk/bugs/Welcome.html> for the operating systems UNIX, LINUX, and Windows, among others. It comes with complete documentation and two example volumes. These examples demonstrate the variety of complex models (random effects, generalized linear, proportional hazards, latent variable, and frailty models) amenable to a Bayesian analysis via BUGS. We will show that nonlinear non-Gaussian state–space models can be added to this list.

In nonlinear non-Gaussian state–space models the full conditional distributions required for Gibbs sampling are typically not log-concave. Both Meyer and Millar (1999) and Millar and Meyer (1999a) implemented the Gibbs sampler in C code, using the C subroutines ARS (Gilks and Wild 1992) and ARMS (Gilks et al. 1995) to sample from univariate log-concave and non-log-concave full conditional distributions, respectively. This, however, required the explicit derivation “by hand” of the full conditional distribution of each parameter in the model, a nontrivial, substantial,

Received September 18, 1998. Accepted February 19, 1999.  
J14789

**R. Meyer<sup>1</sup> and R.B. Millar.** Department of Statistics,  
University of Auckland, Private Bag 92019, Auckland,  
New Zealand.

<sup>1</sup>Author to whom all correspondence should be addressed.  
e-mail: [meyer@stat.auckland.ac.nz](mailto:meyer@stat.auckland.ac.nz)

**Table 1.** Catch (1000's t) and CPUE (kg/100 hooks) data for South Atlantic albacore from Polacheck et al. (1993).

Year	Catch	CPUE
1967	15.9	61.89
1968	25.7	78.98
1969	28.5	55.59
1970	23.7	44.61
1971	25.0	56.89
1972	33.3	38.27
1973	28.2	33.84
1974	19.7	36.13
1975	17.5	41.95
1976	19.3	36.63
1977	21.6	36.33
1978	23.1	38.82
1979	22.5	34.32
1980	22.5	37.64
1981	23.6	34.01
1982	29.1	32.16
1983	14.4	26.88
1984	13.2	36.61
1985	28.4	30.07
1986	34.6	30.75
1987	37.5	23.36
1988	25.9	22.36
1989	25.3	21.91

and tedious task that may deter the practical stock assessment scientist from fitting these models. We explain how BUGS alleviates this chore by making use of conditional independence assumptions in the model that are graphically represented by a directed acyclic graph (DAG). Moreover, the new version 0.6 of BUGS contains a method to sample from any not necessarily log-concave full conditional density. This makes it now possible to fit nonlinear non-Gaussian state-space models.

The paper is organized as follows. In the first section, we briefly describe a biomass dynamics model, the Schaefer surplus production model, that is the most commonly used non-age-structured model in practical stock assessments. This simple, parsimonious model was chosen for illustrative purposes. We use the same catch-effort data set on South Atlantic albacore (*Thunnus alalunga*) as in Millar and Meyer (1999a) and in Polacheck et al. (1993). In the second section, we put this surplus production model into the context of nonlinear, non-Gaussian state-space methodology. This model is then represented as a DAG in the third section. In the subsequent sections, we explain how this DAG assists the model's implementation in BUGS and how to perform convergence diagnostics using CODA (Best et al. 1995).

## The problem

The data available for stock assessment purposes quite often consist of a time series of annual catches  $C_t$ ,  $t = 1, \dots, N$ , and relative abundance indices  $I_t$ ,  $t = 1, \dots, N$ , such as research survey catch rates or catch-per-unit-effort (CPUE) in-

stances from commercial fisheries. For example, Table 1 gives a historical data set of catch-effort data of South Atlantic albacore from 1967 to 1989. Age composition data are not available for this stock. This data set has previously been analysed by Polacheck et al. (1993) and Yeh et al. (1991). Objectives include the estimation of the size of the stock at the end of 1989 and management parameters such as the maximum surplus production (MSP), the biomass at which MSP occurs ( $B_{MSP}$ ), and the optimal effort ( $E_{MSP}$ ), the level of commercial fishing effort required to harvest MSP when the stock is at  $B_{MSP}$ .

When only catch-effort data are available, biomass dynamics models are the primary assessment tools for many fisheries (Hilborn and Walters 1992). They relate the current biomass to previous biomass plus terms for growth and recruitment minus terms for natural mortality and catch. Surplus production models aggregate terms for recruitment, growth, and natural mortality into one term for "surplus production" so that the biomass dynamics equations can be written in the form (Polacheck et al. 1993)

$$(1) \quad B_t = B_{t-1} + g(B_{t-1}) - C_{t-1}$$

where  $B_t$ ,  $C_t$ , and  $g(B_t)$  denote biomass at the start of year  $t$ , catch during year  $t$ , and the surplus production function, respectively. The surplus production function is usually assumed to be nonnegative with  $g(0) = g(K) = 0$ , where  $K$  is the carrying capacity (corresponding to the level of the stock biomass at equilibrium prior to commencement of the fishery). The Schaefer (1954) form of the surplus production function is

$$(2) \quad g(B_{t-1}) = rB_{t-1} \left( 1 - \frac{B_{t-1}}{K} \right)$$

Substituting eq. 2 into eq. 1 gives a parsimonious model describing the annual biomass dynamics transitions with just the two parameters  $r$ , the intrinsic growth rate, and  $K$ :

$$(3) \quad B_t = B_{t-1} + rB_{t-1} \left( 1 - \frac{B_{t-1}}{K} \right) - C_{t-1}.$$

Note that the annual catch is treated as a fixed constant. A common, although simplifying assumption is that the relative abundance index is directly proportional to the biomass, i.e.:

$$(4) \quad I_t = qB_t$$

with catchability parameter  $q$ .

For the Schaefer surplus production model, the maximum surplus production  $MSP = rK/4$  occurs at  $B_{MSP} = K/2$ . When the biomass indices are CPUE's from commercial fishing, then eq. 4 gives  $MSP/E_{MSP} = qK/2$  and thereby the optimal effort is  $E_{MSP} = r/2q$ .

For a more detailed discussion on surplus production models, the interested reader is referred to Hilborn and Walters (1992, chap. 8), Polacheck et al. (1993), Millar and Meyer (1999a), and references therein.

## Bayesian nonlinear state-space model

Polacheck et al. (1993) compared three commonly used statistical techniques for fitting the model defined by eqs. 3

and 4: process error models, observation error models, and equilibrium models. None of these is capable of incorporating uncertainty present in both equations: natural variability underlying the annual biomass dynamics transitions (process error) and uncertainty in the observed abundance indices due to measurement and sampling error (observation error). This is possible, however, using a state-space model, as shown in Meyer and Millar (1999) and Millar and Meyer (1999a). State-space models (see Meyer and Millar 1999 and references therein) relate time series observations  $\{I_t\}$  to unobserved "states"  $\{B_t\}$  through a stochastic observation model for  $I_t$  given  $B_t$ . The states are assumed to follow a stochastic transition model. Equations 3 and 4 are the deterministic versions of the stochastic state and observation equations. We assumed lognormal error structures and used a reparameterization ( $P_t = B_t/K$ ) by expressing the annual biomass as a proportion of carrying capacity as in Millar and Meyer (1999a) to speed mixing (i.e., sampling over the support of the posterior distribution) of the Gibbs sampler. The state equations are rewritten as

$$(5) \quad P_t | \sigma^2 = e^{u_t}$$

$$P_t | P_{t-1}, K, r, \sigma^2 = (P_{t-1} + rP_{t-1}(1 - P_{t-1}) - C_{t-1}/K)e^{u_t}, \quad t = 2, \dots, N$$

and the observation equations are

$$(6) \quad I_t | P_t, q, \tau^2 = qKP_t e^{v_t}, \quad t = 1, \dots, N$$

where  $u_t$  are iid normal with mean 0 and variance  $\sigma^2$  and  $v_t$  are iid normal with mean 0 and variance  $\tau^2$ . Here,  $X|Y$  denotes the conditional distribution of  $X$  given  $Y$ .

A fully Bayesian model consists of the joint prior distribution of all unobservables, here the five parameters  $K, r, q, \sigma^2$ , and  $\tau^2$  and the unknown states  $P_1, \dots, P_N$ , and the joint distribution of the observables, here the relative abundance indices  $I_1, \dots, I_N$ . Bayesian inference is then based on the posterior distribution of the unobservables given the data. In the sequel, we will denote the probability density function of a parameter  $\theta$  by  $p(\theta)$ . We assume that the parameters  $K, r, q, \sigma^2$ , and  $\tau^2$  are independent a priori. By a successive application of Bayes' theorem and conditional independence of subsequent states, the joint prior density is given by

$$(7) \quad p(K, r, q, \sigma^2, \tau^2, P_1, \dots, P_N)$$

$$= p(K)p(r)p(q)p(\sigma^2)p(\tau^2)p(P_1 | \sigma^2)$$

$$\times \prod_{i=2}^N p(P_i | P_{i-1}, K, r, \sigma^2).$$

A noninformative prior is chosen for  $q$ . Prior distributions for  $K, r, \sigma^2$ , and  $\tau^2$  are specified using biological knowledge and inferences from related species and stocks as discussed in Millar and Meyer (1999a):

$$K \sim \text{lognormal}(\mu_K = 5.04, \sigma_K = 0.5162)$$

$$r \sim \text{lognormal}(\mu_r = -1.38, \sigma_r = 0.51)$$

$$p(q) \propto 1/q$$

$$\sigma^2 \sim \text{inverse-gamma}(3.79, 0.0102)$$

$$\tau^2 \sim \text{inverse-gamma}(1.71, 0.0086).$$

For general guidelines on the choice of prior distributions for parameters in stock assessment models, the reader is referred to Punt and Hilborn (1997).

Because of the conditional independence assumption of the relative abundance indices given the unobserved states, the sampling distribution is

$$(8) \quad p(I_1, \dots, I_N | K, r, q, \sigma^2, \tau^2, P_1, \dots, P_N)$$

$$= \prod_{t=1}^N p(I_t | P_t, q, \tau^2).$$

Then, by Bayes' theorem, the joint posterior distribution of the unobservables given the data,  $p(K, r, q, \sigma^2, \tau^2, P_1, \dots, P_N | I_1, \dots, I_N)$ , is proportional to the joint posterior distribution of all unobservables and observables:

$$(9) \quad p(K, r, q, \sigma^2, \tau^2, P_1, \dots, P_N, I_1, \dots, I_N)$$

$$= p(K)p(r)p(q)p(\sigma^2)p(\tau^2)p(P_1 | \sigma^2)$$

$$\times \prod_{i=2}^N p(P_i | P_{i-1}, K, r, \sigma^2) \prod_{t=1}^N p(I_t | P_t, q, \tau^2)$$

i.e., the product of eqs. 7 and 8, the product of prior and sampling distribution. Thus, the joint distribution (eq. 9) is the crucial one for subsequent Bayesian inference.

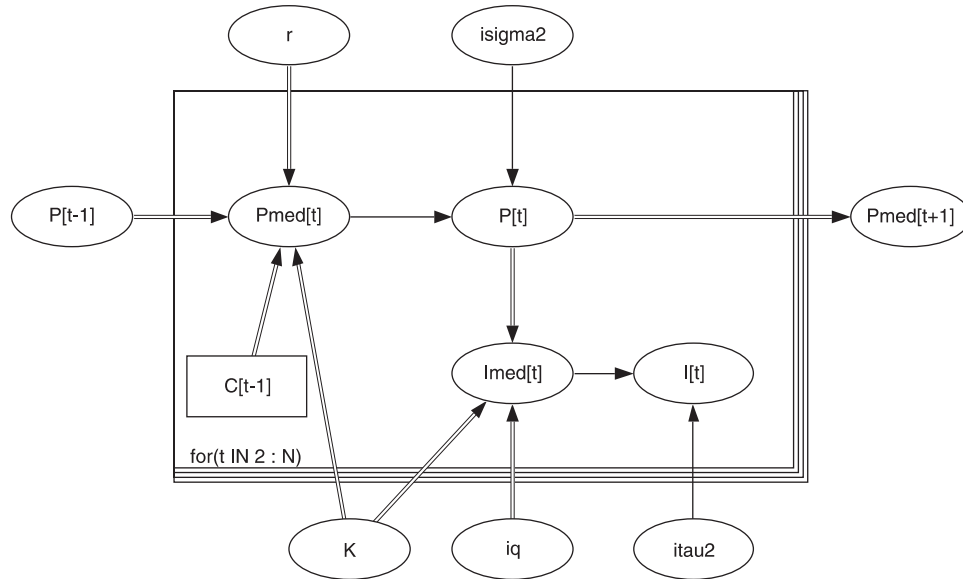
### Model representation as DAG

Now let us step back and look at a graphical representation of the fully Bayesian model. This abstraction has the advantage that we can concentrate on the essential model structure without getting bogged down by all the details about densities. For any year  $t$ , let us represent all unobservables ( $K, r, q, \sigma^2, \tau^2, P_t$ ) and observables ( $I_t$ ) as ellipses and constants ( $C_t$ ) as rectangles. A way to express the conditional independence assumptions is by drawing solid arrows between nodes (see Fig. 1). (By analogy, a graphical representation can be drawn for the node  $P_1$  whose state equation is different from that of the nodes  $P_t, t = 2, \dots, N$ ; see eq. 5.) Hollow arrows go to deterministic nodes, which are logical functions of other nodes. The logarithm of the conditional median of  $P_t$ ,  $\text{Pmed}[t]$ , is an example of a deterministic node, as it is a function of the nodes  $K, r, C_{t-1}$ , and  $P_{t-1}$ .

This renders a model representation as a DAG, as all edges in the graph are directed and there are no cycles because of the conditional independence assumptions. Let  $V$  denote the set of all nodes in the graph. Direct predecessors of a node  $v \in V$  are called "parents" and direct offspring the "children". The solid arrows indicate that *given its parent nodes*, each node  $v$  is independent of all other nodes except descendants of  $v$ . For instance, if we are in year  $t$  and know the biomass in year  $t - 1$  and the values of the parameters  $r, K$ , and  $\sigma^2$ , then our belief in  $P_t$  is independent of the biomass in previous years 1 to  $t - 2$  and the data of all other years except the current relative abundance index  $I_t$ .

It is then easy to construct the joint probability distribution of all stochastic nodes using the graphical description of the conditional independence assumptions:

Fig. 1. Representation of the surplus production model as a DAG.



$$(10) \quad p(V) = \prod_{v \in V} p(v | \text{parents}(v)).$$

For our specific surplus production model, eq. 10 is the graph-theoretical version of eq. 9. In this way, the DAG (Fig. 1) assists in constructing the full Bayesian model. For further reading on conditional independence graphs and graphical chain models, the interested reader is referred to Wermuth and Lauritzen (1990).

### Bayesian inference using BUGS

Let  $V_u$  denote the subset of unobservable nodes and  $V_o$  the subset of observable nodes. Once  $p(V)$  has been obtained from eq. 10, a general technical difficulty encountered in any application of Bayesian inference is calculating the high-dimensional integral necessary to find the normalization constant in the posterior distribution of the unobservables given the data:

$$(11) \quad p(V_u | V_o) = \frac{p(V_u, V_o)}{\int p(V_u, V_o) dV_u}.$$

In our specific example, this would require an  $(N + 5)$ -dimensional integration, as we have to integrate over the unobservables  $K, r, q, \sigma^2, \tau^2, P_1, \dots, P_N$ . Calculating the marginal posterior distribution of any variable would require a subsequent  $(N + 4)$ -dimensional integration. High-dimensional integration problems can be solved via Markov chain Monte Carlo (MCMC) as reviewed in Gilks et al. (1996). The Gibbs sampler, a special MCMC algorithm, generates a sample from the posterior eq. 11 by iteratively sampling from each of the univariate full conditional posterior distributions as explained in Meyer and Millar (1999). These univariate full conditional posterior distributions  $p(v | V \setminus v)$ , for  $v \in V_u$ , can be easily constructed from the joint posterior distribution  $p(V)$  in eq. 10 by picking out those terms that depend on  $v$ :

$$(12) \quad p(v | V \setminus v) \propto p(v | \text{parents}(v)) \times \prod_{w \in \text{parents}(w)} p(w | \text{parents}(w)).$$

This is facilitated by the graphical representation (Fig. 1), as the full conditional posterior distribution of any node  $v$  depends only on its parents, children, and coparents. For instance, if  $v = P_t$ , then the full conditional posterior distribution of  $P_t$ ,  $p(P_t | K, r, q, \sigma^2, \tau^2, P_1, \dots, P_{t-1}, P_{t+1}, P_N, I_1, \dots, I_N)$ , is proportional to

$$p(P_t | P_{t-1}, K, r, \sigma^2) \times p(P_{t+1} | P_t, K, r, \sigma^2) \times p(I_t | P_t, q, \tau^2).$$

Here, the dependence of the deterministic nodes  $Pmed[t]$  and  $lmed[t]$  as logical functions of  $P_t, r, K$ , and  $q$  has been resolved. In this way, BUGS exploits the representation of the model as a DAG for constructing these full conditional posterior distributions for all unobservable nodes. Once this is accomplished, it uses certain sophisticated sampling methods to sample from these univariate densities. BUGS contains a small expert system for choosing the best sampling method. The first choice is to identify conjugacy, where the full conditional reduces analytically to a well-known distribution, and to sample accordingly. If the density is not conjugate but turns out to be log-concave, it employs the adaptive rejection sampling (ARS) algorithm (Gilks and Wild 1992). If the density is not log-concave, BUGS uses a Metropolis–Hastings (MH) step. The MH algorithms differ across the various BUGS versions and platforms. The current UNIX version 0.6 uses the Griddy–Gibbs sampler as developed by Ritter and Tanner (1992). More efficient MH implementations currently under development include slice sampling (Neal 1997) for variables with a restricted range and adaptive techniques (Gilks et al. 1998) for variables with unrestricted range. A first version has been released under WinBUGS, the BUGS version for the Windows 95 operating system.

### Model implementation in BUGS

For a typical BUGS run using the UNIX version 0.6, four different files have to be specified: (i) the data are entered in a file with extension `.dat`, here called `tuna.dat`, (ii) initial val-

ues for all unobservables needed to start the Gibbs sampler are in a .in file, here surplus.in, (iii) the file that contains the model descriptions has the .bug extension, here surplus.bug, and (iv) the commands for running BUGS that control the number of sampled values, which parameters to monitor, and so on go into the surplus.cmd file; in WinBUGS, these commands are options in various menus. These four files are listed in the Appendix.

The data file can be either in rectangular format like tuna.dat or in SPLUS format as in tuna\_S.dat.

The format of the initial value file follows that of the data file, here, for instance, surplus.in in SPLUS format. If initial values for parameters are not given in the initial value file, then values will be generated by forward sampling from the prior distributions specified in the model.

The DAG representation of our model not only serves BUGS-internal purposes but can assist us in specifying our model in the BUGS language. The file containing the model specification, surplus.bug, consists of two sections: the declarations and the model description. The declaration part specifies the name of the model, which nodes are constants (rectangles), and which nodes are stochastic (ellipses) and gives the names of the files containing the data and initial values. The model description forms a declarative representation of the fully Bayesian model. It is enclosed in curly brackets {...}. This is a translation of the graphical model into BUGS syntax. Each statement consists of a relationship of two kinds:  $\sim$ , which means "is distributed as" and substitutes the solid arrows, and  $\leftarrow$ , which means "is to be replaced by" and substitutes the hollow arrows. Quantities on the left of a  $\sim$  are stochastic; those on the left of  $\leftarrow$  are deterministic. In general, each quantity should appear once and only once on the left-hand side of a statement. The order of the expressions within a pair of braces is irrelevant.

WinBUGS even has an option Doodle that allows the user to specify the model graphically by drawing a DAG. It uses a hyperdiagram approach to add extra information to the graph to give a complete model specification. DoodleBUGS then writes the corresponding model in BUGS syntax to a file.

Note that BUGS uses a nonstandard parametrization of distributions in terms of the precision (1/variance) instead of the variance. For example, a lognormal distribution denoted by  $\text{dlnorm}(v, \phi^2)$  has density function

$$\int (x|v, \phi^2) = \sqrt{\frac{\tau}{2\pi}} \frac{1}{x} e^{-\frac{1}{2}(\log x - v)^2 \phi^2}.$$

Thus, we need the nodes isigma2 and itau2. BUGS does not permit an expression to be used as a parameter of a distribution, and hence, we need the deterministic nodes Pmed[t] and Imed[t]. BUGS only allows use of proper prior distributions. The noninformative distribution for  $q$  is improper, however. Therefore, we use a  $\text{gamma}(0.001, 0.001)$  prior for the inverse of  $q$ , which is practically equivalent to  $p(q) \propto 1/q$ . Furthermore, we have to restrict the ranges of those nodes with non-log-concave full conditional distributions by specifying lower and upper bounds using the I(lower, upper) function. This is not necessary in the WinBUGS version because of the adaptive MH algorithm.

It is very easy to obtain posterior samples of the manage-

ment parameters MSP and  $E_{\text{MSP}}$ , which are functions of the parameters  $K$ ,  $r$ , and  $q$ , simply by adding these deterministic nodes MSP and EMSP to the BUGS code (see file surplus.bug in the Appendix). The same holds for predictions like the depletion  $P_{1990}$  and total biomass  $B_{1990}$  in the year 1990.

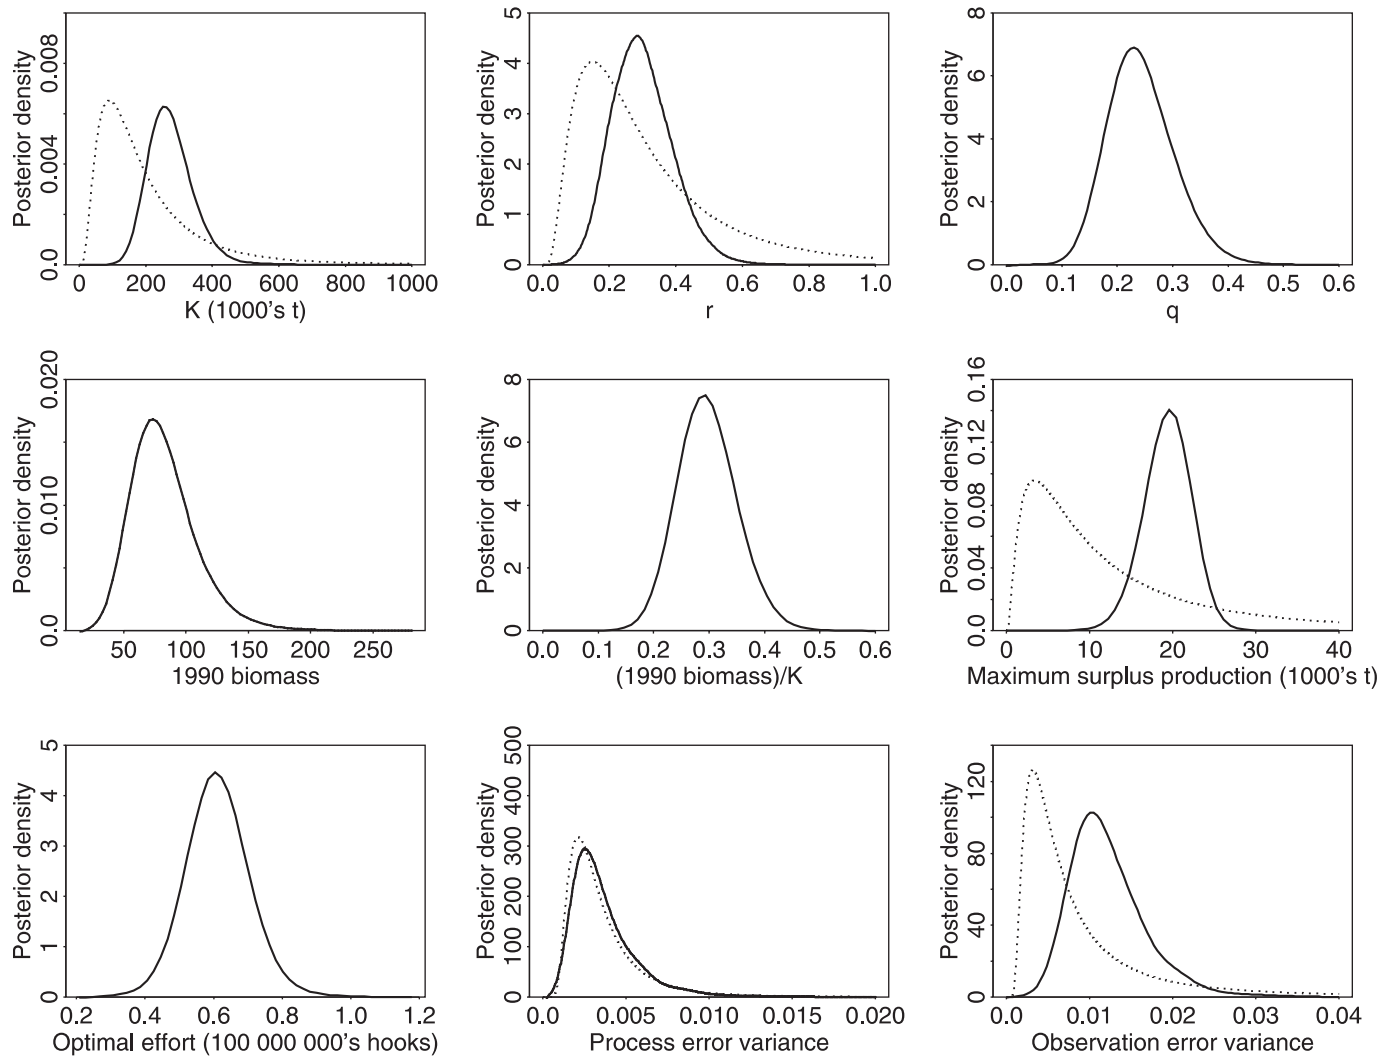
Finally, the file surplus.cmd compiles the BUGS commands in surplus.bug, generates an initial 25 000 iterations (the so-called "burn-in" period), monitors every specified parameter for the next 225 000 iterations, stores every 25th value, and calculates summary statistics of the sampled values for each specified parameter. Our preference is to submit these commands as a batch job by using the command backbugs "surplus.cmd" with session output automatically directed to the bugs.log file. Alternatively, BUGS can be run interactively by using the command bugs. This run took 90 min on a 233-MHz laptop.

BUGS generates five files after completion as follows.

- (i) The file bugs.log contains the BUGS code (surplus.bug) that was used, any error messages, the running time, and the requested summary statistics of the marginal posterior distribution of each parameter. The posterior summaries from this file are listed in the Appendix. Note that the results from fitting this surplus production model are presented and discussed in detail in Millar and Meyer (1999a).
- (ii) The file bugs.out contains two columns. The first column gives the iteration number and the second column gives the corresponding sampled value.
- (iii) The file bugs.ind tells you which line of the bugs.out file corresponds to which monitored variable. Here, lines 1–9000 of bugs.out are samples from variable  $K$ , lines 9001 – 18 000 are samples from variable  $r$ , and so on.
- (iv) The file bugs1.out contains the results of the stats command in a rectangular format suitable for reading into statistical packages for producing graphs, tables, etc. The 10 columns contain the summary statistics mean, SD, observed lower percentile being reported (default 2.5%), observed lower percentile, observed upper percentile being reported, observed upper percentile (default 97.5%), median, number of iterations, start iteration, and finish iteration.
- (v) The file bugs1.ind contains the node names for the variables listed and the corresponding row number in the bugs1.out file.

Whenever one employs MCMC methods to sample from the posterior distribution, a question of extreme importance is "has the chain converged to its target distribution?" Convergence diagnostics is still an area of active research. A recent review on methods used for establishing whether an MCMC algorithm has converged, i.e., whether its output can be regarded as samples from the target distribution of the Markov chain, has been done by Cowles and Carlin (1996). Some of these methods are implemented in the software CODA (Best et al. 1995), a menu-driven collection of SPLUS functions for analysing the output obtained by BUGS. Besides trace plots and the usual tests for convergence, CODA calculates statistical summaries of the posterior distributions and kernel density estimates. CODA is

**Fig. 2.** Kernel density estimates (solid lines) of the posterior distribution of various model and management parameters. Proper prior densities are given by the broken lines.



being maintained and distributed by the same research group responsible for BUGS.

CODA can be used to automatically generate trace plots and kernel density plots of model parameters based on the samples stored in the file bugs.out and the information provided in bugs.ind. Kernel estimates of the posterior densities of the parameters  $K$ ,  $r$ ,  $q$ ,  $B_{1990}$ ,  $P_{1990}$ , MSP,  $E_{MSP}$ ,  $\sigma^2$ , and  $\tau^2$  are shown in Fig. 2, together with their respective prior densities. From the kernel density estimates, one can see that the posterior modes of  $K$ ,  $r$ , and  $q$  are similar to the maximum likelihood estimates obtained by Polacheck et al. (1993) under the observation error model. This may hold in general when priors are diffuse and the observation error variance ( $\tau^2$  has posterior mean of 0.012) is substantially larger than the process error variance ( $\sigma^2$  has posterior mean of 0.0031). Table 2 compares the posterior means of the parameters  $K$ ,  $r$ ,  $q$ ,  $B_{1990}$ ,  $P_{1990}$ , MSP, and  $E_{MSP}$  with the observation error and process error estimates of Polacheck et al. (1993).

Extensive convergence diagnostics for this chain were calculated using the CODA software of Best et al. (1995). All

parameters passed the Heidelberger and Welch stationarity and halfwidth tests. Geweke's  $Z$  scores do not fall within the extreme tails of a standard normal distribution, suggesting that the chain fully converged. The Raftery and Lewis convergence diagnostics (estimating the 2.5th percentile up to an accuracy of 0.02 with probability 0.9) confirmed that the thinning, burn-in period, and sample size were sufficient.

We noted a strong dependency of the mixing behaviour of the chains on the specification of bounds for each parameter with non-log-concave full conditional posterior. The tighter those bounds, the faster the convergence due to the Griddy-Gibbs sampler used in the implementation of the MH step that is necessary to sample from non-log-concave full conditional posteriors. We therefore conclude that slow convergence of the MCMC chains generated by BUGS is an initial impediment that will be overcome with more efficient MH implementations like those based on slice sampling and adaptive techniques that are currently being developed and implemented.

As to the practical limitations in the number of parameters, an age-structured sequential population analysis applied

**Table 2.** Comparison of estimates of a number of management parameters: maximum likelihood estimates obtained by Polacheck et al. (1993) using observation error and process error models, respectively, and posterior means obtained by applying a Bayesian state-space model.

Parameter	Observation error model	Process error model	State-space model
$K$ (1000's t)	239.6	153.4	279.8
$r$	0.328	0.620	0.293
$q \times 10^4$	26.71	43.72	23.89
$B_{1990}$	75.51	50.04	83.97
$P_{1990}$	0.315	0.326	0.297
MSP (1000's t)	19.65	23.78	19.26
$E_{MSP}$	61.4	70.9	60.86

to northern Atlantic cod (*Gadus morhua*) data that included about 200 unknowns is the most sophisticated model that we have yet fitted using BUGS (Millar and Meyer 1999b). This is basically at the current limit of computational practicability. Runs were prohibitively slow, taking 3 h on a 233-MHz laptop to compute 15 000 iterations of the Gibbs sampler. This, however, should only be a transient curb, as continuous algorithmic improvements in the BUGS software are enhancing its performance and capabilities.

## Discussion

BUGS provides a means for a Bayesian analysis of the most complex fisheries models such as fully age-structured models (Millar and Meyer 1999b) and fully Bayesian hierarchical random effects models for stock-recruitment data (Liermann and Hilborn 1997) and for tackling the important Bayesian hierarchic metaanalysis problems identified in the "future directions" of Hilborn and Liermann (1998). BUGS will have considerable implications for the day-to-day practice of the fisheries stock assessment scientist. He/she can concentrate on the essential realistic modeling of data and of prior knowledge, handing over the mathematical intricacies of model fitting to a software package. Because prior and sampling distributions can be easily modified without having to recalculate full conditional distributions, the scientist's time is freed up to experiment with different scenarios — a hitherto far too time-consuming procedure.

The recent evolution in Bayesian computation and software within the last few years parallels and outperforms in some ways the development seen over the last decades of *frequentist* software packages based on maximizing the likelihood function. Once stable numerical routines for nonlinear optimization (like the Newton-Raphson algorithm) had been developed, these were implemented as modules in software packages (like SAS, GLIM, or SPLUS) for maximizing the likelihood function. Fitting a complicated model became possible for the statistical layperson just by running a certain DATA and PROC step in SAS, for instance. Knowledge about mathematical differentiation and optimization was no longer required.

The inevitable difficulty in calculating the high-dimensional integrals necessary for the determination of posterior probability distributions has hindered routine *Bayesian* data analysis. For scientists applying Bayesian inference, a good grasp of numerical and Monte Carlo integration techniques has been essential. Until recently, all

applications have involved writing one-off computer code in low- or intermediate-level languages such as C or Fortran. Even with reliable MCMC subroutines like the MH algorithm and the Gibbs sampler to overcome the multidimensional integrations, writing and debugging a specific program would have taken anything from a few days to several weeks. A subsequent modification of the program, perhaps an extension of the model or an application to a different data set, might well have taken several hours. So Bayesian computations were previously impractical when compared with model fitting in SAS or SPLUS for which model specification may take just a few minutes and modifications may take just a few seconds. Therefore, the Bayesian software package BUGS has been developed with the Gibbs sampler and MH algorithms as building blocks. Instead of having to choose the right PROC in SAS for the problem at hand, with different commands and options for each procedure, the user of BUGS only has to specify prior and sampling distributions for his/her model. In this respect, the Bayesian software package BUGS is considerably more flexible, general, and user friendly. Moreover, as demonstrated in the context of state-space models, the Bayesian approach in conjunction with the Gibbs sampler is capable of handling far more complex models than the classical approach.

The progress made in Bayesian computation via MCMC methods has already revolutionized many scientific disciplines and provided solutions to problems that were hitherto considered computationally intractable. The routine implementation of Bayesian inference that is now possible will "almost surely" have an impact on fisheries stock assessment.

Despite the immense promise of this new Bayesian approach via Gibbs sampling, the reader should also be aware of potential pitfalls with the Bayesian paradigm and the use of MCMC techniques for posterior computation. Playing devil's advocates, we will point some of them out. One of the most controversial issues in Bayesian inference is the specification of prior distributions. An overview of selection of prior distributions by formal rules is given in Kass and Wasserman (1996) and guidelines on prior specification in fisheries stock assessment models are given in Punt and Hilborn (1997). A sensitivity analysis that explores how much the posterior inferences change when other prior distributions are assumed should always accompany a Bayesian data analysis. Basic techniques for model checking by comparing data with posterior predictive distributions and using

Bayes factors are given in Gelman et al. (1996). Despite the ease with which MCMC techniques can be employed using BUGS, a number of subtleties remain, such as assessing convergence, using parameter transformations to reduce posterior correlations and thereby enhance the mixing of the chains, the choice of starting values, and computing simulation standard errors. These are topics of active current research. An article giving advice and guidance to novice users of MCMC that addresses these topics is Kass et al. (1998). We caution against the noninformed use of BUGS, as against any other statistical software package, whether Bayesian or frequentist, and end with the warning issued on the title page of the BUGS manual, "Beware — Gibbs sampling can be dangerous!"

## References

- Best, N.G., Cowles, M.K., and Vines, S.K. 1995. CODA manual version 0.30. Medical Research Council Biostatistics Unit, Institute of Public Health, Cambridge, England.
- Carlin, B.P., Polson, N.G., and Stoffer, D.S. 1992. A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *J. Am. Stat. Assoc.* **87**: 493–500.
- Cowles, M.K., and Carlin, B.P. 1996. Markov chain Monte Carlo convergence diagnostics: a comparative review. *J. Am. Stat. Assoc.* **91**: 883–904.
- Fahrmeir, L., and Tutz, G. 1994. Multivariate statistical modelling based on generalized linear models. Springer, New York.
- Freeman, S.N., and Kirkwood, G.P. 1995. On a structural time series method for estimating stock biomass and recruitment from catch and effort data. *Fish. Res.* **22**: 77–98.
- Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, R.B. 1996. Bayesian data analysis. Chapman and Hall, London, U.K.
- Gentleman, R. 1997. A review of BUGS: Bayesian inference using Gibbs sampling. *Chance*, **10**: 48–51.
- Gilks, W.R., and Wild, P. 1992. Adaptive rejection sampling for Gibbs sampling. *Appl. Stat.* **41**: 337–348.
- Gilks, W.R., Thomas, A., and Spiegelhalter, D.J. 1994. A language and program for complex Bayesian modelling. *Statistician*, **43**: 169–178.
- Gilks, W.R., Best, N.G., and Chan, K.K.C. 1995. Adaptive rejection Metropolis sampling within Gibbs sampling. *Appl. Stat.* **44**: 455–472.
- Gilks, W.R., Richardson, S., and Spiegelhalter, D.J. 1996. Markov chain Monte Carlo in Practice. Chapman and Hall, London, U.K.
- Gilks, W.R., Roberts, G.O., and Sahu, S.K. 1998. Adaptive Markov chain Monte Carlo through regeneration. MCMC preprint server (<http://www.stats.bris.ac.uk/MCMC/>).
- Gudmundsson, G. 1994. Time series analysis of catch-at-age observations. *Appl. Stat.* **43**: 117–126.
- Hilborn, R., and Liermann, M. 1998. Standing on the shoulders of giants: learning from experience in fisheries. *Rev. Fish Biol. Fish.* **8**: 1–11.
- Hilborn, R., and Walters, C.J. 1992. Quantitative fisheries stock assessment: choice, dynamics and uncertainty. Chapman and Hall, New York.
- Kass, R.E., and Wasserman, L. 1996. The selection of prior distributions by formal rules. *J. Am. Stat. Assoc.* **91**: 1343–1370.
- Kass, R.E., Carlin, B.P., Gelman, A., and Neal, R.M. 1998. Markov chain Monte Carlo in practice: a roundtable discussion. *Am. Stat.* **52**: 93–100.
- Kinas, P.G. 1996. Bayesian fishery stock assessment and decision making using adaptive importance sampling. *Can. J. Fish. Aquat. Sci.* **53**: 414–423.
- Liermann, M., and Hilborn, R. 1997. Depensation in fish stocks: a hierarchic Bayesian meta-analysis. *Can. J. Fish. Aquat. Sci.* **54**: 1976–1983.
- Meyer, R., and Millar, R.B. 1999. Bayesian stock assessment using a state-space implementation of the delay difference model. *Can. J. Fish. Aquat. Sci.* **56**: 37–52.
- Millar, R.B., and Meyer, R. 1999a. Nonlinear state-space modeling of fisheries biomass dynamics using Metropolis-Hastings within Gibbs sampling. Tech. Rep. STAT9901. Department of Statistics, University of Auckland, Auckland, N.Z.
- Millar, R.B., and Meyer, R. 1999b. Bayesian state-space modeling of age-structured data: fitting the model is just the beginning. Tech. Rep. STAT9902. Department of Statistics, University of Auckland, Auckland, N.Z.
- Neal, R.M. 1997. Markov chain Monte Carlo methods based on "slicing" the density function. Tech. Rep. No. 9722. Department of Statistics, University of Toronto, Toronto, Ont.
- Pella, J.J. 1993. Utility of structural time series models and the Kalman filter for predicting consequences of fishery actions. *In* Proceedings of the International Symposium on Management Strategies for Exploited Fish Populations. Edited by G. Kruse, D.M. Eggers, R.J. Marasco, C. Pautzke, and T.J. Quinn II. Alaska Sea Grant College Program Rep. 93-02. University of Alaska, Fairbanks, Alaska. pp. 571–593.
- Polacheck, T., Hilborn, R., and Punt, A.E. 1993. Fitting surplus production models: comparing methods and measuring uncertainty. *Can. J. Fish. Aquat. Sci.* **50**: 2597–2607.
- Punt, A.E., and Hilborn, R. 1997. Fisheries stock assessment and decision analysis: the Bayesian approach. *Rev. Fish Biol. Fish.* **7**: 35–63.
- Reed, W.J., and Simons, C.M. 1996. Analyzing catch-effort data by means of the Kalman filter. *Can. J. Fish. Aquat. Sci.* **53**: 2157–2166.
- Ritter, C., and Tanner, M.A. 1992. Facilitating the Gibbs sampler: the Gibbs stopper and the Griddy-Gibbs sampler. *J. R. Stat. Soc. Ser. B (Methodol.)*, **59**: 291–317.
- Schaefer, M.B. 1954. Some aspects of the dynamics of populations important to the management of commercial marine fisheries. *Inter-Am. Trop. Tuna Comm. Bull.* **1**: 27–56.
- Schnute, J.T. 1994. A general framework for developing sequential fisheries models. *Can. J. Fish. Aquat. Sci.* **51**: 1676–1688.
- Spiegelhalter, D.J., Thomas, A., Best, N., and Gilks, W.R. 1996. BUGS 0.5. Bayesian inference using Gibbs sampling. Manual version ii. Medical Research Council Biostatistics Unit, Institute of Public Health, Cambridge, England.
- Sullivan, P.J. 1992. A Kalman filter approach to catch-at-length analysis. *Biometrics*, **48**: 237–257.
- Thomas, A., Spiegelhalter, D.J., and Gilks, W.R. 1992. BUGS: a program to perform Bayesian inference using Gibbs sampling. *In* Bayesian statistics 4. Edited by J.M. Bernardo, J.O. Berger, A.P. Dawid, and A.F.M. Smith. Oxford University Press, Oxford, U.K. pp. 837–842.
- Wermuth, N., and Lauritzen, S.L. 1990. On substantive research hypothesis, conditional independence graphs and graphical chain models (with discussion). *J. R. Stat. Soc. Ser. B (Methodol.)*, **52**: 21–72.
- Yeh, S.-Y., Tsou, T.-S., and Liu, H.-C. 1991. Assessment of the South Atlantic albacore resource by using surplus production models, 1976–1988. *Collect. Vol. Sci. Pap. ICCAT*, **39**: 166–170.



**Appendix**tuna.dat

15.9 61.89  
 25.7 78.98  
 28.5 55.59  
 23.7 44.61  
 25.0 56.89  
 33.3 38.27  
 28.2 33.84  
 19.7 36.13  
 17.5 41.95  
 19.3 36.63  
 21.6 36.33  
 23.1 38.82  
 22.5 34.32  
 22.5 37.64  
 23.6 34.01  
 29.1 32.16  
 14.4 26.88  
 13.2 36.61  
 28.4 30.07  
 34.6 30.75  
 37.5 23.36  
 25.9 22.36  
 25.3 21.91

tuna\_S.dat

list(C=c(15.9,25.7,...,25.3), I=c(61.89,78.98,...,21.91))

surplus.in

list(P=c(0.99,0.98,0.96,0.94,0.92,0.90,0.88,0.86,0.84,0.82,  
 0.80,0.78,0.76,0.74,0.72,0.70,0.68,0.66,0.64,0.62,0.60,0.5  
 8,0.56),  
 r=0.8, K=200, iq=5, isigma2=100, itau2=100)

surplus.bug

```
model surplusproduction;
const N=23;
var C[N], I[N], Imed[N], P[N], Pmed[N],
r, K, q, iq, sigma2, isigma2, tau2, itau2, MSP, EMSP,
P1990, B1990;
data C, I in "tuna.dat";
inits in "surplus.in";
{
```

# prior distribution of K: lognormal with 10% and 90%  
 quantile at 80 and 300  
 K ~ dlnorm(5.042905,3.7603664)I(10,1000);

# prior distribution of r: lognormal with 10% and 90%  
 quantile at 0.13 and 0.48  
 r ~ dlnorm(-1.38,3.845)I(0.01,1.2);

# prior distribution of q: instead of improper (prop. to 1/q)  
 use just proper IG  
 iq ~ dgamma(0.001,0.001)I(0.5,100);  
 q <- 1/iq;

# prior distribution of sigma2: inv. gamma with 10% and  
 90% qu. at 0.04 and 0.08  
 isigma2 ~ dgamma(3.785518,0.010223);  
 sigma2 <- 1/isigma2;

# prior distribution of tau2: inv. gamma with 10% and  
 90% qu. at 0.05 and 0.15  
 itau2 ~ dgamma(1.708603,0.008613854);  
 tau2 <- 1/itau2;

# (conditional) prior distribution of Ps (from state equa-  
 tions):

```
Pmed[1] <- 0;
P[1] ~ dlnorm(Pmed[1],isigma2) I(0.001,2.0)
for (t in 2:N) { Pmed[t] <- log(P[t-1] + r*P[t-1]*(1-P[t-1])
- C[t-1]/K);
P[t] ~ dlnorm(Pmed[t],isigma2)I(0.001,2.0) }
```

# sampling distribution:

```
for (t in 1:N) { Imed[t] <- log(q*K*P[t]);
I[t] ~ dlnorm(Imed[t],itau2);
}
```

# further management parameters and predictions:

```
MSP <- r*K/4;
EMSP <- r/(2*q);
P1990 <- P[N]+r*P[N]*(1-P[N]) -C[N]/K;
B1990 <- P1990*K;
}
```

surplus.cmd

```
compile("surplus.bug")
update(25000)
monitor(K,25)
monitor(r,25)
monitor(q,25)
monitor(sigma2,25)
monitor(tau2,25)
monitor(P[],25)
monitor(P1990,25)
monitor(B1990,25)
monitor(MSP,25)
monitor(EMSP,25)
update(225000)
stats(K)
stats(r)
stats(q)
stats(sigma2)
stats(tau2)
stats(P[])
stats(P1990)
stats(B1990)
stats(MSP)
stats(EMSP)
q()
```

bugs.log

Bugs &gt; stats(K)

mean	sd	2.5% :97.5% CI	median	sample
2.798E+2	6.307E+1	1.826E+2 4.251E+2	2.709E+2	9000

Bugs &gt; stats(r)

mean	sd	2.5% : 97.5% CI	median	sample
2.928E-1	8.663E-2	1.436E-1 4.793E-1	2.869E-1	9000

Bugs &gt; stats(q)

mean	sd	2.5% : 97.5% CI	median	sample
2.389E-1	5.694E-2	1.434E-1 3.627E-1	2.344E-1	9000

Bugs &gt; stats(sigma2)

mean	sd	2.5% : 97.5% CI	median	sample
3.100E-3	1.894E-3	1.133E-3 7.826E-3	2.601E-3	9000

Bugs &gt; stats(tau2)

mean	sd	2.5% : 97.5% CI	median	sample
1.215E-2	4.422E-3	5.731E-3 2.270E-2	1.137E-2	9000

Bugs &gt; stats(P[])

mean	sd	2.5% : 97.5% CI	median	sample
[1] 1.018E+0	5.432E-2	9.165E-1 1.133E+0	1.015E+0	9000

[2] 9.959E-1	7.436E-2	8.741E-1 1.171E+0	9.874E-1	9000
--------------	----------	-------------------	----------	------

[3] 8.807E-1	6.579E-2	7.646E-1 1.025E+0	8.755E-1	9000
--------------	----------	-------------------	----------	------

.

.

.

[22] 3.543E-1	3.603E-2	2.910E-1 4.324E-1	3.519E-1	9000
---------------	----------	-------------------	----------	------

[23] 3.282E-1	4.014E-2	2.557E-1 4.136E-1	3.255E-1	9000
---------------	----------	-------------------	----------	------

Bugs &gt; stats(P1990)

mean	sd	2.5% : 97.5% CI	median	sample
2.973E-1	4.654E-2	2.123E-1 3.955E-1	2.946E-1	9000

Bugs &gt; stats(B1990)

mean	sd	2.5% : 97.5% CI	median	sample
8.397E+1	2.579E+1	4.457E+1 1.441E+2	8.017E+1	9000

Bugs &gt; stats(MSP)

mean	sd	2.5% : 97.5% CI	median	sample
1.926E+1	2.412E+0	1.401E+1 2.354E+1	1.944E+1	9000

Bugs &gt; stats(EMSP)

mean	sd	2.5% : 97.5% CI	median	sample
6.086E-1	8.770E-2	4.382E-1 7.903E-1	6.079E-1	9000