INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Bayesian Inference Using Gibbs Sampling in Applications and Curricula of Decision Analysis

Mauricio Diaz, Daniel M. Frances

Vol. 14, No. 2, January 2014, pp. 86–95
ISSN 1532-0545 (online)

INFORMS Transactions on Education

http://dx.doi.org/10.1287/ited.2013.0120

Bayesian Inference Using Gibbs Sampling in Applications and Curricula of Decision Analysis

Mauricio Diaz, Daniel M. Frances
Department of Mechanical and Industrial Engineering, University of Toronto, Ontario M5S 3G8, Canada
{mauricio.diaz@mail.utoronto.ca, frances@mie.utoronto.ca}

Applications and curricula of decision analysis currently do not include methods to compute Bayes' rule and obtain posteriors for nonconjugate prior distributions. The current convention is to force the decision maker's belief to take the form of a conjugate distribution, leading to a suboptimal decision. Bayesian inference using Gibbs sampling (BUGS) software, which uses Markov chain Monte Carlo methods, numerically obtains posteriors for nonconjugate priors. By using the decision maker's true nonconjugate belief, the problems explored suggest that BUGS can produce a posterior distribution that leads to optimal decision making. Other methods exist that can use nonconjugate priors, but they must be implemented ad hoc because they do not have any supporting software. BUGS offers the distinct advantage of being implemented in existing software, and with simple coding can solve a wide range of decision analysis problems. BUGS is useful in making optimal decisions, and it is easy to learn and implement; therefore, including BUGS in decision analysis curricula is valuable.

Keywords: decision analysis; Bayesian inference; BUGS; BRugs; Gibbs sampling; nonconjugate prior; decision analysis curricula; Bayesian decision analysis; Bayesian updating; Bayesian inference using Gibbs sampling
History: Received: July 2013; accepted: September 2013.

Introduction

As operations research (OR) educators, we often have to remind students that mathematical models are only an abstraction of reality, with just enough detail, we hope, to yield improved solutions to the real problem at hand. Thus when we teach linear programming, we remind students to ensure that all relations are approximately linear, that the variables can be modeled as continuous, and that uncertainty can be ignored. Furthermore we make them aware that should these assumptions be unrealistic, there are more advanced methods, such as nonlinear programming for nonlinear relations, integer programming for discrete variables, and stochastic programming for handling uncertainty. And we always encourage them to use the simplest possible methods and to reserve the more advanced toolkits until there is evidence that they are indeed necessary. Similarly, in stochastic OR we give priority to analytical methods that assume the Markov property, but we make sure we also cover simulation methods to handle more complex situations.

This paper identifies an area of OR for which we regularly tell our students that the methods require a simplifying assumption, yet we fail to mention readily available tools for dealing with situations when this assumption is clearly not satisfied. We hope this paper will raise the awareness about such tools and thereby increase their application to yield better decisions.

The Decision Analysis Setting

Decision analysis can be loosely defined as the science of making good decisions under uncertainty. It is widely taught (see Table 1) and practiced. The uncertainty is modeled by probability distributions of uncertain events and parameters, often based initially on the subjective belief of the decision maker. As additional information about the events and parameters becomes available, Bayes' rule can be used to objectively update the decision maker's subjective prior distribution to a posterior distribution that incorporates the new information.

When the distributions are discrete, Bayes' rule can be readily implemented, but when continuous distributions are involved, Bayes' rule becomes problematic (see Chap. 9 of Lee 2012). In such cases, to make Bayes' rule computationally feasible, it is common to force the decision maker's subjective belief to take the form of a specific type of distribution (conjugate distribution; see Section 2.1 of Lee 2012). Although the resulting distribution may not accurately represent the decision maker's true beliefs, this method is widely taught and used.

In the literature, there exist methods such as Bayesian inference using Gibbs sampling (BUGS) to compute posteriors for prior distributions of any form (nonconjugate distributions). However, these methods are currently absent from applications and curricula of decision analysis. This paper advocates that such methods should be taught extensively and that their use should be encouraged in decision analysis practice.

The term BUGS has several meanings:

- The BUGS method: a Markov chain Monte Carlo method (MCMC) that uses Gibbs sampling and the Metropolis-Hastings algorithm to numerically perform Bayesian inference (Gilks et al. 1994).
- The BUGS Project: the project that spawned the BUGS method and led to the development of BUGS software (BUGS Project 2012a).
- BUGS software: software implementing the BUGS method, e.g., WinBUGS (Lunn et al. 2000), OpenBUGS (Lunn et al. 2009), and JAGS (Plummer 2003).
- The BUGS language: the language used to code models within the various BUGS software packages.

In the context of this paper, unless stated otherwise, BUGS refers to OpenBUGS software.

Whereas currently the decision maker's subjective belief is forced to take the form of a conjugate distribution, if BUGS were adopted it would allow the decision maker's subjective belief to take any form. This would produce a posterior distribution that more accurately reflects the decision maker's true beliefs, thus allowing the decision maker to make a more informed decision and thereby facilitate better decision making.

The Current Situation

Curricula

Current decision analysis course descriptions posted on the Internet from 10 engineering institutions (Table 1) reveal that the majority (7 of 10) include conjugate priors but none mentions BUGS software or any method to compute Bayes' rule for nonconjugate priors.

We can infer that many of the decision analysis courses examined in Table 1 are Bayesian since they teach conjugate distributions. BUGS is recognized as a valuable Bayesian tool, as evidenced by its inclusion in several courses that teach Bayesian methods (see BUGS Project 2012b for a list of courses teaching BUGS). Despite this, BUGS appears to be absent from decision analysis courses.

Applications

The BUGS method was created by and is typically used by statisticians (see BUGS Project 2012c for a sample of relevant publications). Despite exhaustive searching, we do not find a single paper discussing the use of BUGS in the context of decision analysis practice.

Miller and Rice (1983) discuss the discretization of probability densities to obtain posteriors. This converts the integral in Bayes' theorem to a Riemann sum. However, for this to be a good approximation, the discretization requires a large number of bins, which then makes the problem computationally infeasible. In the context of hybrid influence diagrams, i.e., those that include continuous distributions, methods exist that address the shortcomings of conjugate distributions. Poland and Shachter (1993) describe mixtures of Gaussian influence diagrams, though this requires several assumptions about the problem.

Table 1 Current Decision Analysis Curricula

Institution	Course	Type of document	Reference	Influence diagrams	Conjugate distributions	BUGS/other nonconjugate methods
Stanford University	MS&E 353 Frontiers of Decision Analysis/MS&E 355 Influence Diagrams and Probabilistic Networks	Syllabus	Howard (2011)	•	•	
Princeton University	WWS5940 Risk Analysis	Syllabus	Craft (2012)	•		
UC Berkeley	IEOR 166 Decision Analysis	Course description	Oren (2012)	•		
University of Michigan	IOE 460 Decision Analysis	Course description	Bordley (2004)		•	
University of Singapore	IE5203 Decision Analysis	Syllabus	Leng (2012)	•		
University of Minnesota	IE 5545 Decision Analysis	Syllabus	Gupta (2007)	•	•	
University of Illinois	GE 550 AA—Decision Analysis II	Syllabus	Abbas (2013)		•	
Hong Kong Polytechnic University	AMA484 Decision Analysis	Syllabus	Hoi-Lun (2008)	•	•	
Auburn University	INSY 5630 Decision Analysis and Real Options	Syllabus	Park (2005)		•	
Purdue University	IE 546 Economic Decisions in Engineering	Syllabus	Liu (2010)	•	•	
Total	10			7	7	0

To simplify the integral in Bayes' theorem, Moral et al. (2001) approximate nonconjugate priors with mixtures of truncated exponentials; similarly, Shenoy and West (2011) approximate the probability densities with mixtures of polynomials. All of these methods must be implemented ad hoc because they do not have any supporting software.

There appear to be no references to BUGS, mixtures of polynomials, mixtures of truncated exponentials, or any other method of computing posteriors for nonconjugate priors being used to solve specific, real-world problems.

In contrast, there are recent papers that use conjugate priors in decision making. For example, Greenland (2001) uses conjugate priors to model epidemiologic risk and Huang (2008) uses conjugate priors to determine an optimal cost sharing warranty policy. The published evidence indicates that decision analysis practitioners currently use conjugate priors. Both Greenland and Huang state that conjugate priors are used for simplified computation. This indicates a perceived difficulty with using nonconjugate priors, which leads them to use conjugate priors as taught in curricula.

Better Decisions with BUGS in an Insurance Context

This setting is an extension of Hesselager (1993), who describes a class of conjugate priors with applications to excess-of-loss reinsurance, and Martínez-Miranda et al. (2012), who mentions the need for insurance companies to meet their claims liability cash flows.

An auto insurance company had 1,000 collision claims exceeding \$5,000 in the previous year; these claims totalled \$8 million. The auto insurance company decides to purchase reinsurance on its collision claims. They decide to buy an excess-of-loss cover for a \$10,000 layer in excess of \$5,000; i.e., for a collision with severity $s \geq \$5{,}000$, the auto insurance company can claim $c = \min(s - \$5{,}000, \$10{,}000)$ from the reinsurer. The reinsurer will pay the auto insurance company a lump sum at the end of each year.

The reinsurance manager assumes that the number of collision claims exceeding \$5,000, n, is Poisson distributed with unknown rate λ. The manager also assumes that the severity of the collision claims exceeding \$5,000, s, is Pareto distributed with minimum value \$5,000 and unknown shape parameter ψ.

The manager is interested in forecasting the total claim cost, $\sum_{i=1}^{n} c_i$, so the reinsurance company can build the monetary reserves needed to meet its liability at the end of the year. The manager would like to build a monetary reserve of size M equal to the expected liability, $\mathbb{E}(\sum_{i=1}^{n} c_i) = \mathbb{E}(liability)$; i.e., the manager's goal is to minimize the deviation of M from the expected liability, $\min_M D$, where $D = |M - \mathbb{E}(liability)|$. Because the number of claims and their severity are assumed to be independent, the expected liability is equal to the expected number of claims multiplied by the expected claim amount, $\mathbb{E}(liability) = \mathbb{E}(n)\,\mathbb{E}(c)$.

To solve the problem, the manager must choose prior distributions for the unknown parameters λ and ψ using expert knowledge. Suppose the manager agrees that the following gamma distribution (which is conjugate to the Pareto sampling distribution of s) is appropriate for ψ.

$$\psi \sim \Gamma(\alpha_\psi = 1{,}000, \beta_\psi = 375)$$

Suppose that the manager does not believe λ follows a gamma distribution (which is conjugate to Poisson) but instead follows a lightly left- or right-skewed distribution, a heavily left- or right-skewed distribution, or a bimodal distribution (seen in Table 2). Through a Monte Carlo simulation, we can find the expected liability for each of the priors. Suppose the manager builds a monetary reserve equal to the expected liability.

Over the following year, the auto insurance company has $N = 988$ claims exceeding \$5,000, each with known severity S_i (in the absence of real claims data, these values were generated from a Pareto distribution, $S_i \sim \mathrm{Par}(x_m = \$5{,}000, \psi = 2.4)$). Using data from the first year, the manager must forecast the total claims cost and decide on the size of the monetary reserve to be built for the end of the second year. This requires the manager's prior of ψ to be updated using the 988 generated samples of severity, which is easily done using the formulas for conjugacy. The manager's nonconjugate prior of λ must also be updated using the sample data $n_1 = 988$. Current teaching would suggest forcing the manager's prior to be a gamma distribution and updating this using the formulas for conjugacy. With BUGS we can update the manager's true belief, the nonconjugate prior, and obtain the expected liability through a Monte Carlo simulation.

We find that for the nonconjugate priors of λ, the true expected liability will range from \$2.64 million to \$2.67 million, and by building a reserve of size M equal to the expected liability the manager can achieve $D = 0$. Had we used the currently accepted assumption of conjugacy, in the worst case when the manager believes λ follows a heavily left-skewed prior and builds a monetary reserve according to the solution from the conjugate approximation ($M =$ \$2.65 million), the company would incur an expected \$2.67 million liability, resulting in a deviation $D =$ \$18,700. For other priors, the approximate solution leads to smaller deviations, summarized in Table 2. In all cases BUGS can be used to make a better decision.

Table 2 **Nonconjugate Priors/Posteriors and Impact of Conjugate Approximations—Insurance Setting**

Prior	Posterior	Outcome using solution from best conjugate approximation
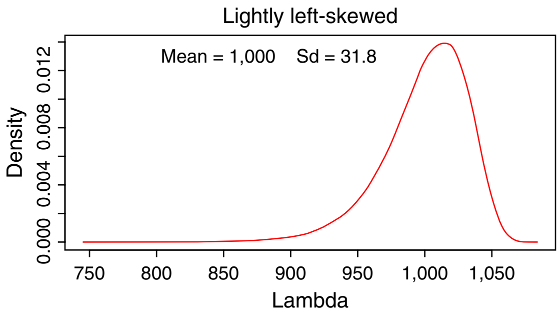	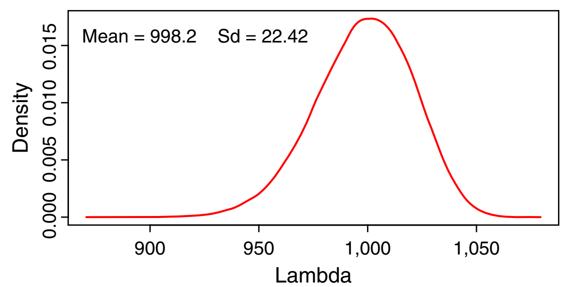	$D = \$11,200$
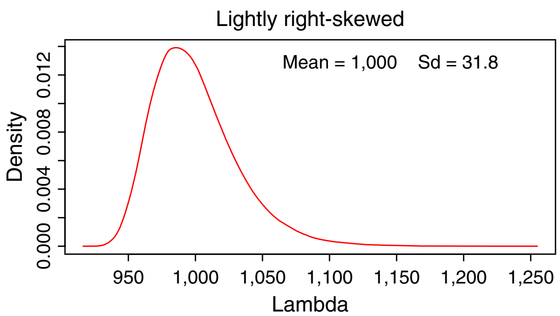	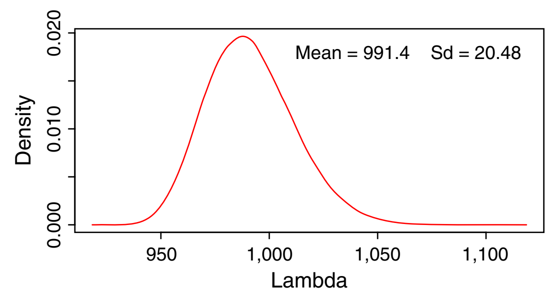	$D = \$6,900$
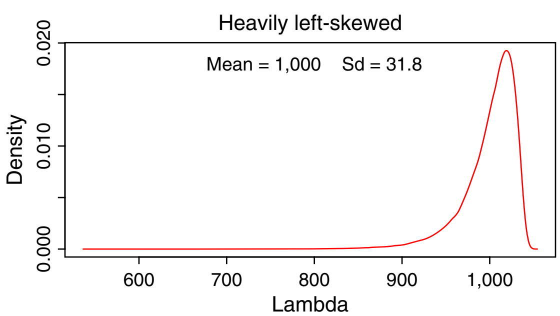	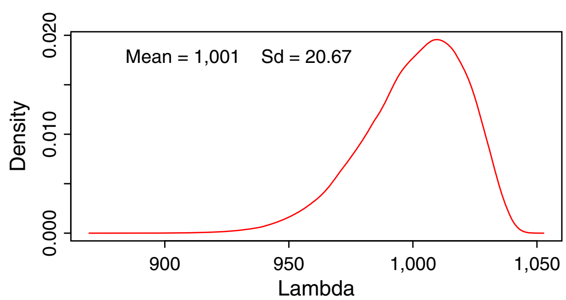	$D = \$18,700$
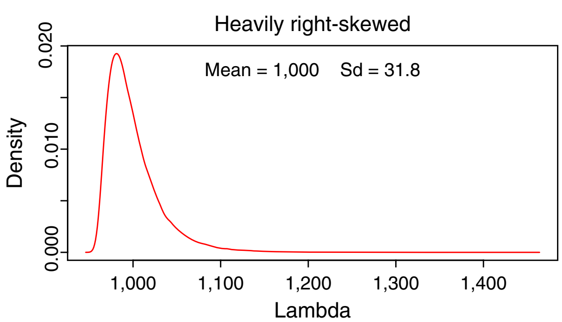	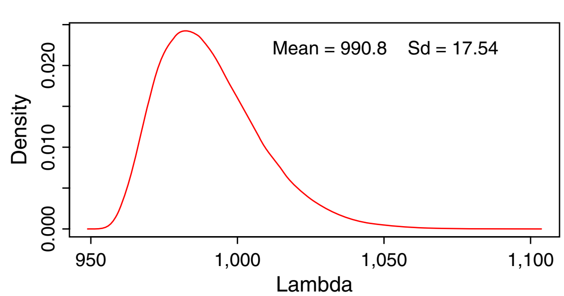	$D = \$8,500$
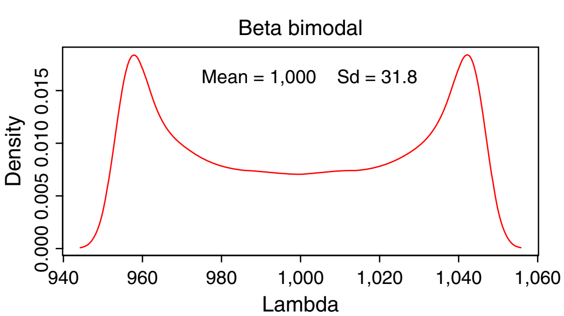	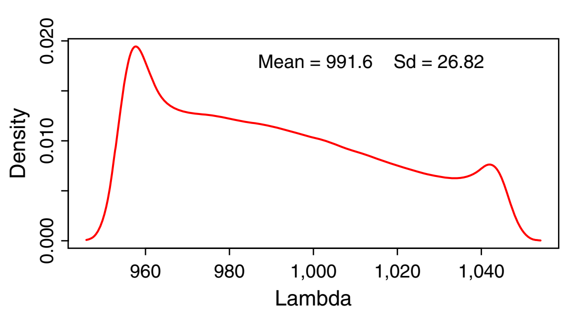	$D = \$6,400$

Better Decisions with BUGS in a Fishing Context

This setting is an extension of Clark et al. (1985), who use conjugate priors in deciding the optimal fishing capacity for a developing fishery. Here we will assume instead nonconjugate priors.

In the fishing industry, the proportion of fish stock to be caught is commonly referred to as the fishing capacity. A fisheries manager is responsible for setting the optimal fishing capacity (F) for a developing fishery of prawn for the next nine years to maximize the expected net present value (NPV) subject to expected sustainability of the prawn population. The following is a summarized and slightly simplified version of the mathematical model in Clark et al. (1985).

$$\max_{F} \ NPV = \mathbb{E}\left(\sum_{t=1}^{9}(C_t p \alpha^t)\right) \tag{1}$$
$$\text{subject to:} \ \ \mathbb{E}(N_9) \geq N_0,$$

where

- $p = \$6.5\ \text{kg}^{-1}$ is the profit per kg of prawn fished.
- $\alpha = 0.96$ is a discount factor.
- C_t is the total catch of prawn in year t in kg, given by

$$C_t = \frac{F}{F+M}(1 - e^{-F-M})N_t w.$$

- M is the natural mortality of the prawns, estimated to be 0.1.
- F is the fishing capacity, to be set by the manager.
- w is the average weight of the prawns, estimated to be 0.025 kg per prawn.
- N_t is the number of prawns after t years, with the current population N_0 estimated to be 2.00×10^{10} prawns, which is assumed to follow the popular fishery dynamics model:

$$N_t = N_{t-1}e^{-F-M} + \frac{R_t \phi(N_{t-1})}{w}. \tag{2}$$

- $\phi(N_t)$ is the stock-recruitment function, estimated to be equal to $\min(1, N_t/(1.50 \times 10^{10}))$.
- R_t is a sample annual recruitment of the prawns in year t in kg. The random variable recruitment (R) follows a lognormal distribution where $\ln R$ is assumed to have a known standard deviation $\sigma = 0.5$, and $\ln R$ has unknown mean μ.

The fisheries manager is assumed to have expert knowledge that can be quantified into a prior distribution of μ. Based on experience, the fisheries manager does not believe that a conjugate normal distribution is appropriate for μ, the mean of $\ln R$. Suppose instead that the manager believes μ follows the right-skewed distribution in Figure 1. This prior distribution can be modeled as a transformed log normally distributed variable: $L \sim \ln \mathcal{N}$ (mean $= 2$, standard deviation $= 0.324$), $\mu = L/9 + 17.13$.

Figure 1 The Manager's Prior of μ

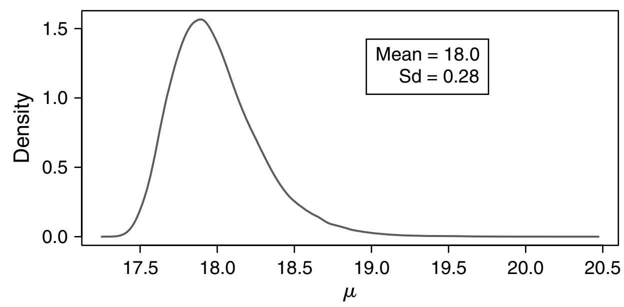

For a given capacity F, we can now compute the objective and determine the feasibility of the model in (1), i.e., compute the expected NPV and the expected prawn population after 9 years. This is accomplished through an ordinary Monte Carlo simulation. We can vary the value of F to find the fishing capacity that produces the largest expected NPV and is expected to be sustainable.

OpenBUGS was used to perform the Monte Carlo simulation. Note that this does not require computation of Bayes' rule, and therefore the BUGS method was not used; OpenBUGS was simply used to perform a basic Monte Carlo simulation.

We find that optimally $F = 0.067$. The manager sets the fishing capacity to this value. Three years later, the manager finds that the actual recruitment of the prawns was the following:

$$R_1^1 = 9.80 \times 10^7 \ \text{kg}, \quad \ln R_1^1 = 18.4$$
$$R_2^1 = 5.94 \times 10^7 \ \text{kg}, \quad \ln R_2^1 = 17.9$$
$$R_3^1 = 8.86 \times 10^7 \ \text{kg}, \quad \ln R_3^1 = 18.3$$

for years 1, 2, and 3, respectively. This leads, via Equation (2), to an estimated prawn population of 2.05×10^{10} at the present time.

Now the prior distribution must be updated using the three new data points for recruitment in order to solve for the optimal fishing capacity for years 3 through 12. Because this prior is nonconjugate, it cannot be updated analytically and thus decision analysts using current curricula would not have learned the methodology for dealing with this situation.

With BUGS, it is possible to obtain the posterior for the manager's true belief, the nonconjugate right-skewed prior distribution of μ in Figure 1. Through Monte Carlo simulation we find that the optimal $F = 0.072$, resulting in an expected NPV of \$1.63 billion and a sustainable expected prawn population of 2.05×10^{10} after nine years. Had we used

Table 3 Nonconjugate Priors/Posteriors and Impact of Conjugate Approximations—Fisheries Setting

Prior	Posterior	Outcome using solution from best conjugate approximation
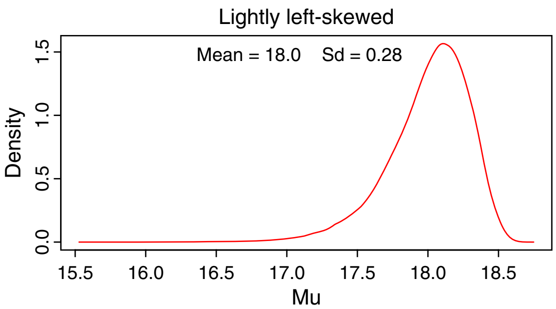	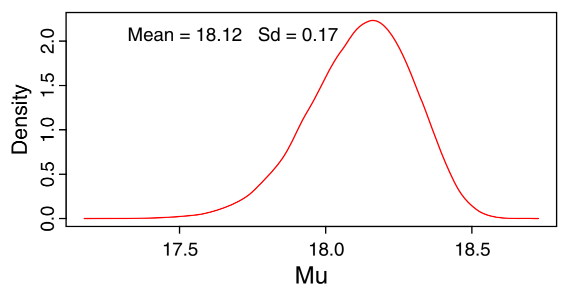	$\Delta NPV = -\$170$ million, suboptimal $\mathbb{E}(N_{[9]}) = 2.08 \times 10^{10}$ prawns
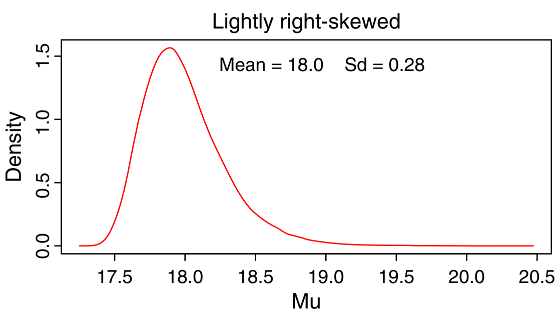	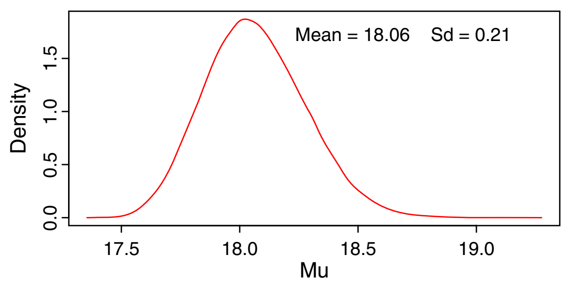	$\Delta NPV = \$110$ million $\mathbb{E}(N_{[9]}) = 2.00 \times 10^{10}$ prawns, infeasible
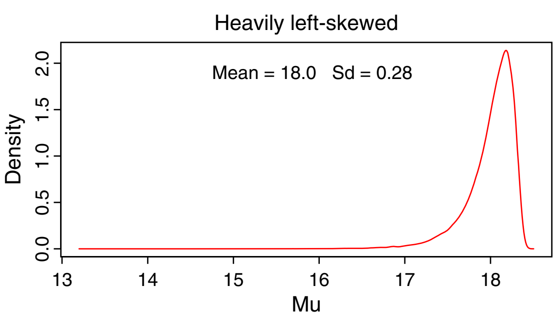	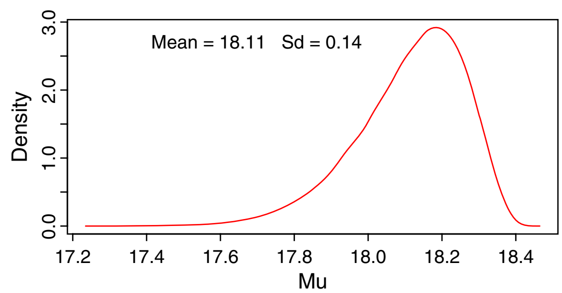	$\Delta NPV = -\$20$ million, suboptimal $\mathbb{E}(N_{[9]}) = 2.06 \times 10^{10}$ prawns
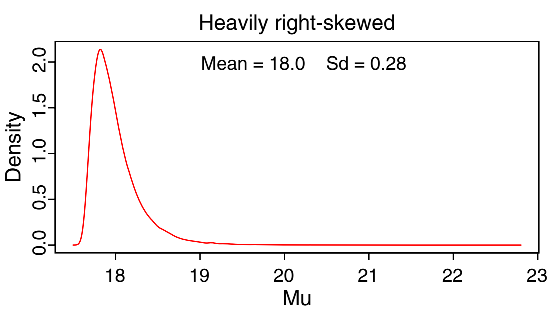	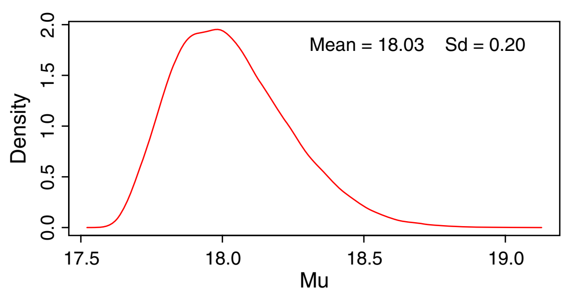	$\Delta NPV = \$200$ million $\mathbb{E}(N_{[9]}) = 1.95 \times 10^{10}$ prawns, infeasible
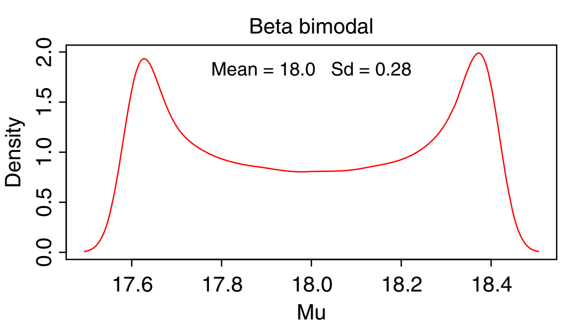	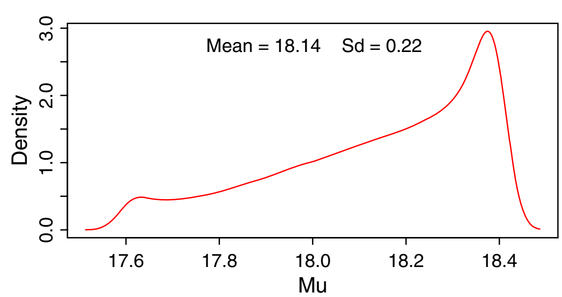	$\Delta NPV = -\$140$ million, suboptimal $\mathbb{E}(N_{[9]}) = 2.12 \times 10^{10}$ prawns

Figure 2 **Hybrid ID of the Fisheries Problem, Three Year Time Horizon**

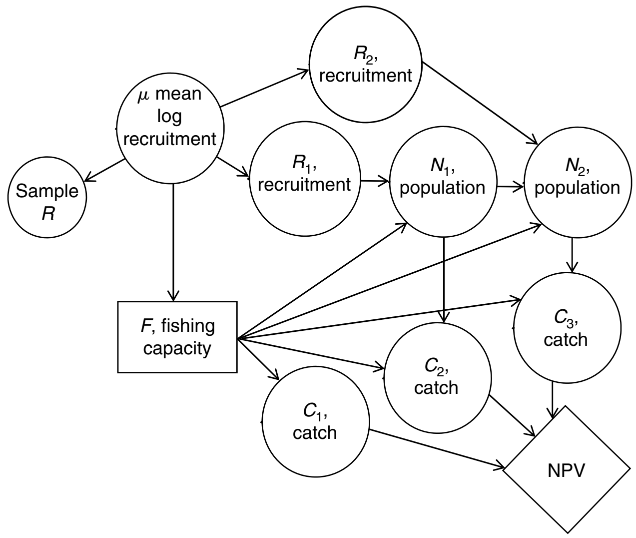

the best approximate normal conjugate prior as currently taught and applied in decision analysis, we would find the optimal $F = 0.078$. An evaluation of this decision with the manager's true prior distribution yields an expected NPV of $1.74 billion and an expected prawn population of 2.00×10^{10} after nine years. This represents a ΔNPV of $110 million but is infeasible because it leads to expected overfishing and depletion of the prawn.

Other nonconjugate priors tested were the following: a light left skew, a heavy left skew, a heavy right skew, and a bimodal beta distribution. In all cases, the solutions suggested by using the best possible conjugate approximations were found to be either suboptimal or infeasible. Results are summarized in Table 3.

Using BUGS in Decision Analysis Courses

Influence diagrams are widely taught in decision analysis. A hybrid influence diagram is one where chance and/or decision variables are continuous. We can think of the fisheries problem as a hybrid influence diagram, as shown in Figure 2.

BUGS software can inherently accept influence diagrams without decision nodes to calculate the posterior distribution of the NPV for a given fishing capacity F.

At its most intuitive level, OpenBUGS allows the user to graphically create models, referred to by BUGS as Doodles, as shown in Figure 3.

Through OpenBUGS's drop-down menus, students can perform various actions, including loading the model, updating posterior distributions, and viewing statistics of variables of interest.

OpenBUGS converts the graphical model into BUGS code, which the students can then compare to the graphical model to help them understand the code. Eventually students may wish to code models instead of using the graphical interface. This is much more efficient; each node or variable is essentially a single line of code.

Except for the shorter time horizon, the Doodle in Figure 3 produces code similar to Figure 4, which contains the code used to solve the fisheries problem.

OpenBUGS has distinctly simpler and more advanced ways of performing the same tasks, which helps students learn the software. We first included this tool in a decision analysis course to a class of about 50 senior undergraduate industrial engineering students in 2012. We noted that once a student has achieved some proficiency, complex models can quickly be coded and solved. We are again using it in

Figure 3 **BUGS Doodle of the Fisheries Problem, Three Year Time Horizon**

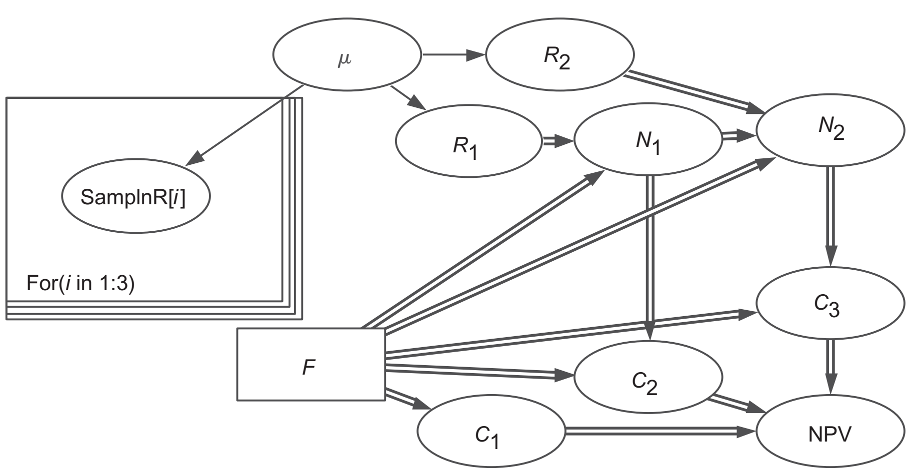

Figure 4 Sample BUGS Code, Fisheries Setting

```
list(tau=4, M=0.1, F=0.073, N0=2.05E+10, p=6.5, ...
alpha=0.96, samplnR=c(18.4,17.9,18.3))

model{
     L ~ dlnorm(2, 9.5)
     #mu = L/9+18.87                 #left skewed prior of mu
     mu = L/9+17.13                       #right skewed prior of mu

     #H ~ dlnorm(1.1, 2.9)
     #mu <- -H/8+18.45               #heavily left skewed prior of mu
     #mu <- H/8+17.55                #heavily right skewed prior of mu

     #B ~ dbeta(0.5, 0.5)
     #mu <- 17.6+0.8 * B             #bimodel prior of mu

     for(i in 1:3){
          samplnR[i] ~ dnorm(mu, tau)     #samples of log recruitment
     }

     for(i in 1:9){
          logR[i] ~ dnorm(mu, tau)   #logR is a variable name
          log(R[i])<-logR[i]         #transformation of logR to R
     }

     N[1]<-exp(-F-M) * N0+(R[1]/0.025) * min(N0/1.5E+10,1)
     for(i in 2:9){
          N[i]<-exp(-F-M) * N[i-1]+(R[i]/0.025) * min(N[i-1]/1.5E+10,1)
     }                               #forecasting the prawn population

     C[1]<-(F/(F+M)) * (1-exp(-F-M)) * N0 * 0.025
     for(i in 2:9){
          C[i]<-(F/(F+M)) * (1-exp(-F-M)) * N[i-1] * 0.025
     }                               #computing the annual catch

     NPV[1]<-C[1] * p * pow(alpha, 1)
     for(i in 2:9){
          NPV[i]<-NPV[i-1]+C[i] * p * pow(alpha, i)
     }               #computing the net present value of the annual catches
}
```

2013 to a similar group of similar size. We believe that OpenBUGS can be considered a practical, functional, and easy to learn tool for all Bayesian updating.

It is important to note that BUGS is a tool primarily used by statisticians and does not solve for optimal decisions. For example, it does not automatically find the value of the decision variable F that maximizes the expected NPV in the fisheries setting. Students could manually change the value of F, rerun BUGS, and find the optimal F in this way. Alternatively, students can use the BRugs package by Thomas et al. (2006), which provides a comprehensive R interface to OpenBUGS. BRugs allows students to code iterative procedures to replace the extensive typing and clicking required by OpenBUGS. For example, to solve our influence diagram, we need only load up the model, inject a value of F into the data set, compute the expected NPV, and loop this procedure for different values of F while storing the F value that produces the highest expected NPV. Sample BRugs code to perform this task is shown in Figure 5. Further BRugs examples can be found in Kruschke (2010).

Conclusion

Conjugate distributions continue to be taught and applied in decision analysis. There are cases where a conjugate approximation of a nonconjugate prior will lead to significantly suboptimal decision making. BUGS provides value by allowing more realistic nonconjugate priors and producing an overall optimal decision.

BUGS has a graphical interface in the form of Doodles as well as a more advanced coding mode. The transition from graphical models to coding allows the student to start with a more intuitive method before advancing to more efficient but abstract methods. For a wide range of problems, BRugs can solve hybrid influence diagrams and find the optimal decisions with a simple exhaustive search.

We conclude that there is significant value in including BUGS in decision analysis curricula and

Figure 5 Sample BRUGS Code, Fisheries Setting

```
library("BRugs")

NPVopt = 0
FTEMP = 0.05

for (i in 1:50){
FTEMP = FTEMP+i/1000            # F value to be used
                                #inject F value into the data set
bugsData(list(mu0 = 18, tau = 4, M = 0.1, F = FTEMP, N0 = 2.05E+10, ...
p = 6.5, alpha = 0.96, samplnR = c(18.4,17.9,18.3)), "fishdata.txt")
modelCheck("fishmodel.txt")             #load the model from Figure 4
modelData("fishdata.txt")               #load the data
modelCompile(numChains = 1)             #number of markov chains = 1
modelGenInits( )                        #generate initial chain values
modelUpdate(20000)                      #posterior computation burn in
samplesSet(c("N[9]", "NPV[9]"))         #set a watch on variables of interest
modelUpdate(20000)                      #compute posteriors

     susttest = samplesStats("N[9]")[1,1]      #pull population value
     NPVtest = samplesStats("NPV[9]")[1,1]     #pull NPV value

     if (susttest> = 20500000000){             #check if sustainable
          if (NPVtest>NPVopt){                 #check if NPV is optimal
               FOPT = FTEMP                    #store F value if optimal
               sustopt = susttest
          NPVopt = NPVtest
          }
     }
samplesClear("*")          #clear sample collection before next simulation
}                          #loop
```

practice. With this tool, students will know a practical way to proceed when conjugate priors are not acceptable to the decision maker.

References

Abbas A (2013) GE: 550 AA—Decision Analysis II, University of Illinois. Accessed December 2013, http://www.bayesiannetworks.com/index.php?option=com_content&view=article&id=64&Itemid=82.

Bordley R (2004) IOE 460 F 2004: Decision Analysis, University of Michigan. Accessed January 2013, http://sitemaker.umich.edu/dec.btr/files/bordley_decisionanalysis_industrialandoperationsengineering_undergrad.pdf.

BUGS Project (2012a) Bayesian inference using Gibbs sampling. Accessed June 2013, http://www.mrc-bsu.cam.ac.uk/bugs/welcome.shtml.

BUGS Project (2012b) BUGS on the Web. Accessed August 2013, http://www.mrc-bsu.cam.ac.uk/bugs/weblinks/webresource.shtml#courses.

BUGS Project (2012c) BUGS user papers and publications. Accessed June 2013, http://www.mrc-bsu.cam.ac.uk/bugs/references/bugs-uses-papers.shtml.

Clark CW, Charles A, Beddington JR, Mangel M (1985) Optimal capacity decisions in a developing fishery. *Marine Resource Econom.* 2:25–53.

Craft A (2012) WWS594O: Risk Analysis, Princeton University. Accessed January 2013, http://wws.princeton.edu/grad/course_offerings/spring-2012/syllabi/WWS-594o.pdf.

Gilks WR, Clayton DG, Spiegelhalter DJ, Best NG, McNeil AJ (1994) Modelling complexity: Applications of Gibbs sampling in medicine. *J. Roy. Statist. Soc.* 55:39–52.

Greenland S (2001) Putting background information about relative risks into conjugate prior distributions. *Biometrics* 57: 663–670.

Gupta D (2007) IE 5545 Decision Analysis, University of Minnesota. Accessed January 2013, http://www.me.umn.edu/courses/pdf/IE%;205545.pdf.

Hesselager O (1993) A class of conjugate priors with applications to excess-of-loss reinsurance. *Astin Bull.* 23:77–93.

Hoi-Lun A (2008) AMA484: Decision Analysis, Hong Kong Polytechnic University. Accessed January 2013, http://www.polyu.edu.hk/~ama/subject/63021/ama484_08.pdf.

Howard R (2011) MS&E 353 Frontiers of Decision Analysis, Stanford University. Accessed January 2013, http://main.da.stanford.edu/classes-1/ms-e-353-frontiers-of-decision-analysis.

Huang YS (2008) A cost sharing warranty policy for products with deterioration. *IEEE Trans. Engrg. Management* 55:617–627.

Kruschke JK (2010) *Doing Bayesian Data Analysis: A Tutorial with R and BUGS* (Academic Press, Burlington, MA).

Lee P (2012) *Bayesian Statistics: An Introduction* (John Wiley & Sons Ltd., Chichester, UK).

Leng P (2012) IE5203 Decision Analysis, National University of Singapore. Accessed January 2013, http://www.ise.nus.edu.sg/course/ie5203/.

Liu AL (2010) IE 546 Economic Decisions in Engineering, Purdue University. Accessed December 2013, http://web.ics.purdue.edu/~liu334/IE546_Syllabus_Spr10.pdf.

Lunn DJ, Spiegelhalter D, Thomas A, Best N (2009) The BUGS project: Evolution, critique, and future directions. *Statist. Medicine* 28:3049–3067.

Lunn DJ, Thomas A, Best N, Spiegelhalter D (2000) WinBUGS—A Bayesian modelling framework: Concepts, structure, and extensibility. *Statist. Comput.* 10:325–337.

Martínez-Miranda MD, Nielsen JP, Wüthrich MV (2012) Statistical modelling and forecasting in nonlife insurance. *Statist. Oper. Res. Trans.* 36:195–218.

Miller III AC, Rice T (1983) Discrete approximations of probability distributions. *Management Sci.* 29:352–362.

Moral S, Rumi R, Salmeron A (2001) Mixtures of truncated exponentials in hybrid Bayesian networks. *Symbolic Quant. Approaches Reasoning Uncertainty* 2143:156–167.

Oren S (2012) IEOR 166: Decision Analysis, University of California, Berkeley. Accessed December 2013, http://www.scribd.com/doc/98773466/166-intro2012.

Park C (2005) INSY 5630 Real Options and Decision Analysis, Auburn University. Accessed January 2013, https://fp.auburn.edu/provost/ucc/INSY5630Syllabus.pdf.

Plummer M (2003) JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. *Proc. 3rd Internat. Workshop on Distributed Statist. Comput., Vienna, Austria*.

Poland WB, Shachter RD (1993) Mixtures of Gaussians and minimum relative entropy techniques for modeling continuous uncertainties. Heckerman D, Mamdani E, eds. *Proc. Ninth Internat. Conf. Uncertainty in Artificial Intelligence* (Morgan Kaufmann, San Francisco), 183–190.

Shenoy PP, West JC (2011) Inference in hybrid Bayesian networks using mixtures of polynomials. *Internat. J. Approximate Reasoning* 52:641–657.

Thomas A, O'Hara B, Ligges U, Sturtz S (2006) Making BUGS open. *R News* 6:12–17.