



# Getting started with WinBUGS

James B. Elsner and Thomas H. Jagger  
Department of Geography, Florida State University



# The BUGS Project

## The BUGS Project

*welcome*



<http://www.mrc-bsu.cam.ac.uk/bugs/>

1. Click on [WinBUGS](#) and follow instructions for downloading the software.

2. Click on [Overview](#) and view the movie.

[Welcome Page](#)

[Latest News](#)

[Overview & Demo](#)

[Contact us/BUGS list](#)

[WinBUGS](#)

[New WinBUGS examples](#)

[GeoBUGS](#)

[PKBUGS](#)

[Classic BUGS](#)

[CODA](#)

## About BUGS

Bayesian inference Using Gibbs Sampling is a piece of computer software for the Bayesian analysis of complex statistical models using Markov chain Monte Carlo (MCMC) methods. It grew from a statistical research project at the MRC Biostatistics Unit, but now is developed jointly with the Imperial College School of Medicine at St Mary's, London.

**New users** should see the [Overview](#) page for a summary of how the programs work, and what sort of problems they're useful for.

The main program, [WinBUGS](#), has a graphical user interface and on-line monitoring and convergence diagnostics. A restricted version is available free, and users who [register](#) receive a 'key' providing full functionality (also free).

The original, 'Classic' BUGS program uses text-based model description and a command-line interface, and versions are available for major

[FAQ Pages](#)

[BUGS resources online](#)

[References](#)

[Events](#)

[Quick Links](#)

[WinBUGS development site](#)

computer platforms. It is not being developed further but is still freely available.

[CODA](#) is a suite of S-plus/R functions for convergence diagnostics, suitable for use with either version of BUGS.

The programs are reasonably easy to use and come with a range of examples. Considerable caution is, however, needed in their use, since the software is not perfect and MCMC is inherently less robust than analytic statistical methods. There is no in-built protection against misuse.



## Obtaining the File

Please read the [LICENSE AGREEMENT](#) for WinBUGS, before downloading anything.

The single file you need is [WinBUGS14.exe](#), and selecting the link will yield a dialog box inviting you to save the file onto your system.

[WinBUGS 1.1](#) and [WinBUGS 1.2](#) are still available for those with historical interests. [WinBUGS 1.3](#) and its [key](#) are also available and may still be useful if problems are obtained with WinBUGS 1.4 and for running PKBUGS (see below).

## Installing WinBUGS 1.4 in Windows

For installation, run the file WinBUGS14.exe. One way to do this is as follows:

1. Exit all other programs currently running (particularly if using Windows XP)
2. Copy WinBUGS14.exe to your computer
3. Go into Explore and double click on WinBUGS14.exe
4. Follow the instructions in the dialog box
5. You should have a new directory called WinBUGS14 within Program Files
6. Inside the WinBUGS14 directory is a program called WinBUGS14.exe
7. Right-click on the pretty WinBUGS icon, select 'create shortcut', then drag this shortcut to the desktop.
8. Double click on WinBUGS14.exe to run WinBUGS14.

If you have problems after installation, check again that you did not have programs running. You could also try disabling any virus checker.

## Obtaining the key for unrestricted use

WinBUGS is free, but has absorbed a lot of time and grant money over the last 12 years or so. We need to keep our employers happy with our time spent on the project, and it helps if we can keep track of how many people are downloading it and where they come from. It is also very helpful to have an idea of what it is being used for.

Please fill in the [registration form](#) and return it to us. We will then e-mail you the key which will remove the restrictions in WinBUGS 1.4. After following the instructions given in the key, check that the **Keys.ocf** file in `..WinBUGS14\Bugs\Code\` has been updated. (Some people have found they need to re-boot the machine to complete installation of the key.)

Please note that the key currently being issued is valid until June 30th 2004. You will be sent a new key before this one expires.



Click on WinBUGS icon to start a session.

Register to get full version access.

WinBUGS14

File Tools Edit Attributes Info Model Inference Options Doodle Map Text Window Help

WinBUGS Licence

**BUGS**

## Licence Agreement

Introduction

This software and any associated documentation whether electronic or printed (hereinafter called "WinBUGS PACKAGE") is made available under a licence agreement and may be used only in accordance with the terms of that agreement. This is a legal agreement between you (the Licensee), and MRC and Imperial College of Science, Technology and Medicine (the Licensor). The terms of the licence are provided in the following pages.

Users are required to register and to pay a fee for the use of the WinBUGS PACKAGE. Details of fees and the procedure for registration and acceptance of the licence terms is provided here. There is no fee payable for the use of the demonstration (Internet) version of the WinBUGS package. Users of the demonstration version of the WinBUGS package can upgrade to the full version on payment of a fee.

The current fee is zero dollars (\$0).

By completing and sending the registration you demonstrate your agreement to the terms of this licence and will become legally bound to the terms therein.

It should be emphasised that the statistical tools provided in the WinBUGS PACKAGE are by their very nature partly subjective. The Licensor cannot offer advice on interpretation of results obtained using the WinBUGS PACKAGE. Any assistance will be strictly limited to attempting to help if there are problems formulating the statistical problem with the WinBUGS PACKAGE.

The Licensor makes no representation or warranties with respect to the WinBUGS PACKAGE and specifically disclaims any implied warranties of merchantability and fitness for a particular purpose. The Licensor reserves the right to revise WinBUGS PACKAGE and to make changes therein from time to time without obligation to notify any person or organisation of such revision or changes. While Licensor will make every effort to ensure the reliability of the WinBUGS PACKAGE neither the Licensor nor its employees or agents may be held

The WinBUGS user manual and Examples Vol I and Vol II provide a reference for beginners and novices.

A good way to approach a problem with WinBUGS is to scan through the examples to find a problem like yours. You can then modify the code to fit your problem.

When you click on examples, you will open them in what is called a *compound document*. This organizes programs, graphs, and explanations into a single file. A good example is the surgical: institutional ranking.

WinBUGS14

File Tools Edit Attributes Info Model Inference Options Doodle Map Text Window Help

WinBUGS User Manual

**BUGS**

## WinBUGS User Manual

Version 1.4, January 2003

David Spiegelhalter<sup>1</sup>    Andrew Thomas<sup>2</sup>    Nicky Best<sup>2</sup>    Dave Lunn<sup>2</sup>

<sup>1</sup> MRC Biostatistics Unit,  
Institute of Public Health,  
Robinson Way,  
Cambridge CB2 2SR, UK

<sup>2</sup> Department of Epidemiology & Public Health,  
Imperial College School of Medicine,  
Norfolk Place,  
London W2 1PG, UK

e-mail:    bugs@mrc-bsu.cam.ac.uk    [general]  
          andrew.thomas@ic.ac.uk    [technical]

internet: http://www.mrc-bsu.cam.ac.uk/bugs

### Permission and Disclaimer

[please click here to read the legal bit](#)

More informally, potential users are reminded to be extremely careful if using this program for serious statistical analysis. We have tested the program on quite a wide set of examples, but be particularly careful with types of model that are currently not featured. If there is a problem, *WinBUGS* might just crash, which is not very good, but it might well carry on and produce answers that are wrong, which is even worse. Please let us know of any successes or failures.

© 1998-2003 MRC Biostatistics Unit, Cambridge

Example #1: Inference is required about the proportion of people in the population who are unemployed. Let us call this value “pi”.


Think: The true value of pi is in the range between 0 and 1, where 0 means no one is unemployed and 1 means everyone is.


Realistically, we might have some prior information on the value of pi. For instance, newspaper reports, economic theory, previous surveys, etc. Your prior can be as informative as you like.

To get things started, let's assume that you choose a prior as a random value from a [beta distribution](#) restricted between the values of 0.1 and 0.6.

## Specifying your prior

Create a directed graph. In WinBUGS, select Doodle > New...

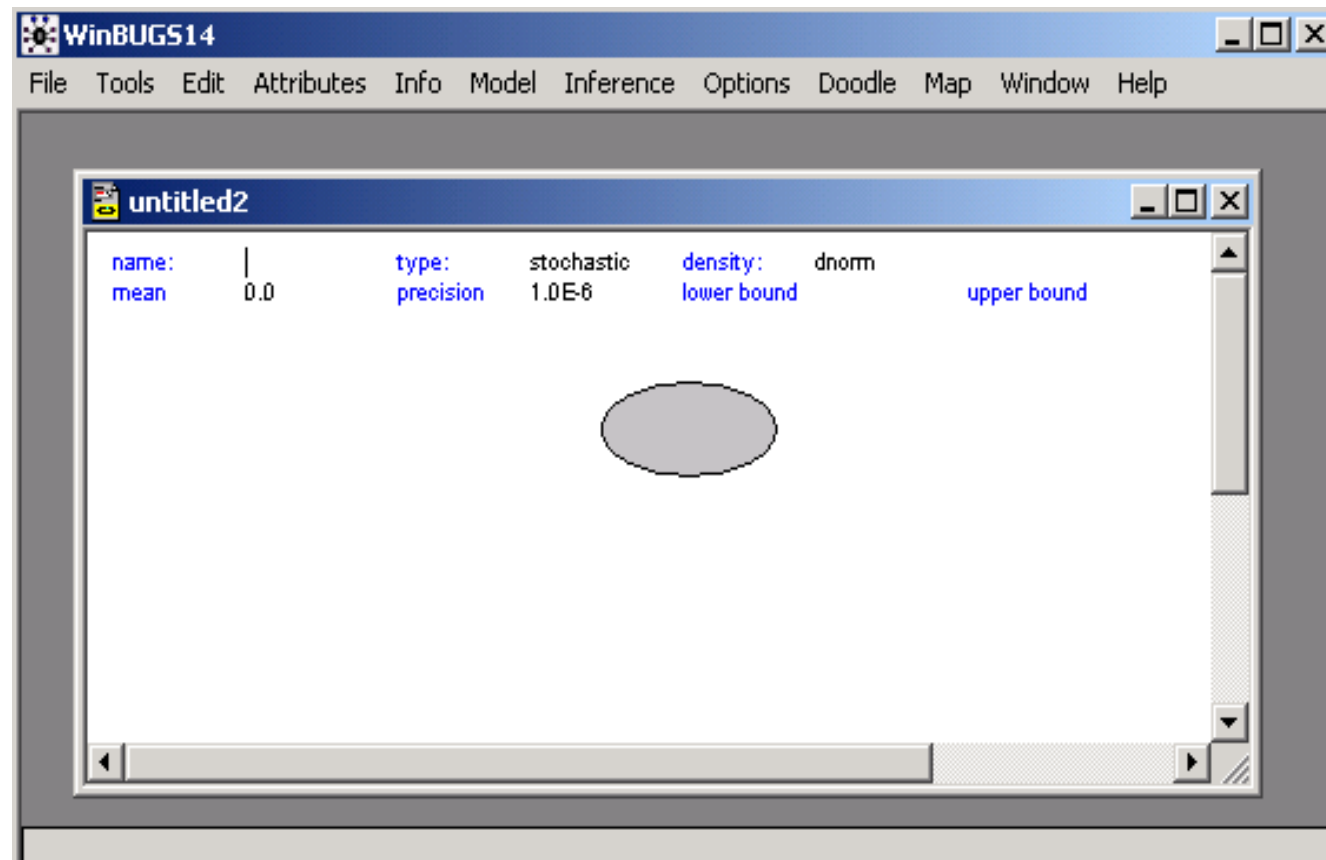
In the New Doodle dialog, click  on the default options. You will see a blank worksheet called “untitled”.



Left click anywhere in the middle of the sheet to create a node.

A node has a name and type along with other characteristics and parameters depending on its type.

Notes: To delete a node, highlight it, then Ctrl-Delete.  
To highlight a node, click on it.  
Use the Help > Doodle help to learn more.



8

Click on **name**, then type “pi”. Note the graphical node is labeled simultaneously.

Leave **type** as “stochastic”.

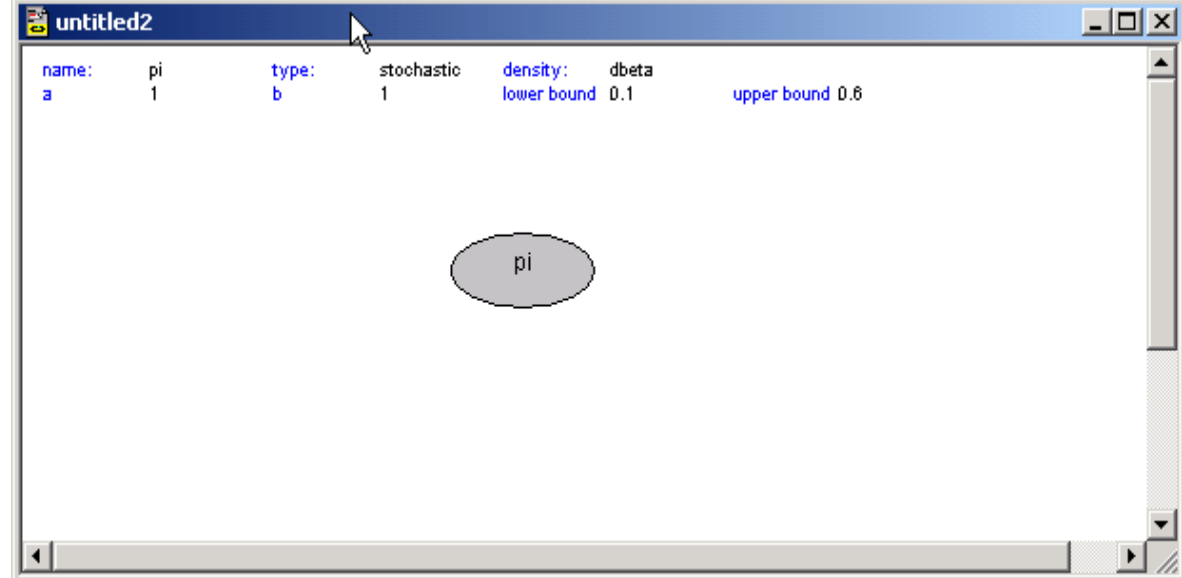
Click on **density** and change to “dbeta”. This means winBUGS will choose a random value from the beta distribution.

Click on **a** and type “1” then on **b** and type “1”. These are the parameters of the beta distribution.

Click on **lower bound** at type “0.1” then on **upper bound** and type “0.6”. These are the bounds we set on the true value of our prior.

We will begin by using winBUGS to look at samples from our prior. To do this, we need to write code using this doodle. WinBUGS executes from the code. The doodle helps us keep track of our model, which at this stage consists of a single node called pi.

On the main menu, select Doodle > Write Code.





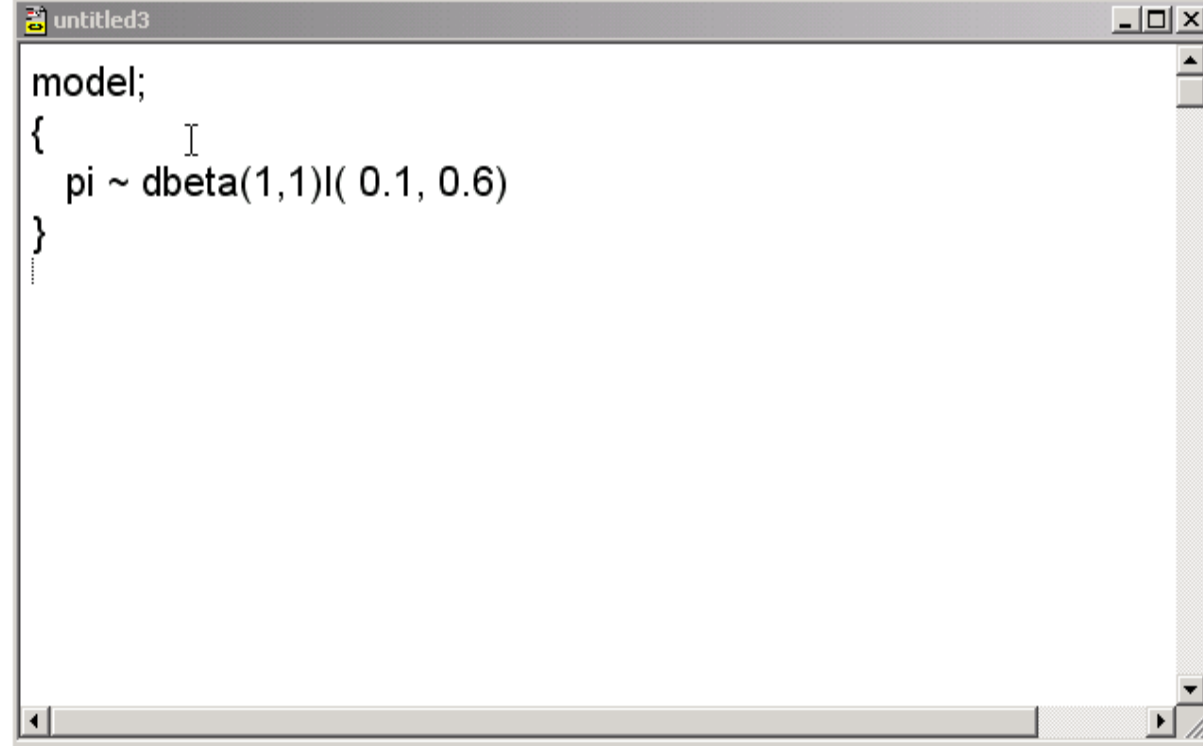
9

A new window opens that displays the model code.

Left click over the text to highlight it.  
Click on Attributes > 16 point.

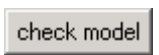
The model selects a random number from a beta distribution with parameters  $a=b=1$ , keeping only those values that lie between 0.1 and 0.6.

Note  $\sim$  is read “distributed as”.



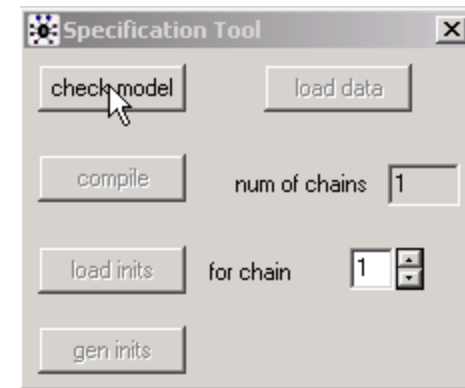
```
model;
{
  pi ~ dbeta(1,1)|( 0.1, 0.6)
}
```

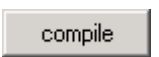
To run the model, select Model > Specification... to bring up the Specification Tool dialog box.

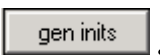
Click . In the lower left corner of the main dialog box you should see the words “model is syntactically correct”. The compile button in the Specification Tool becomes active.




model is syntactically correct



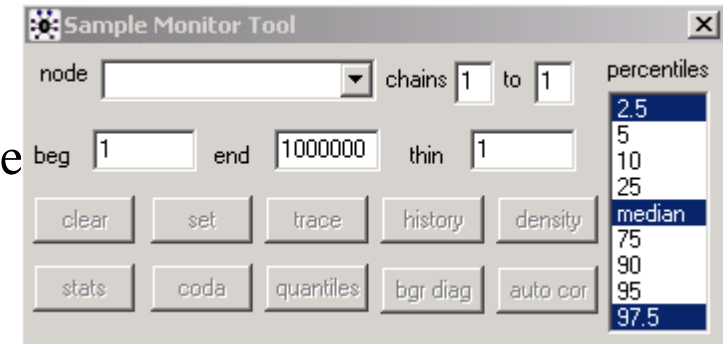
Click . In the lower left corner you should see the words “model compiled”.

Click . This creates a starting value for the model. You should see the words “initial values generated, model initialized”.


Before you produce results, select Options > Output options  
 Select the log radio button. Note here you can also select the output precision.



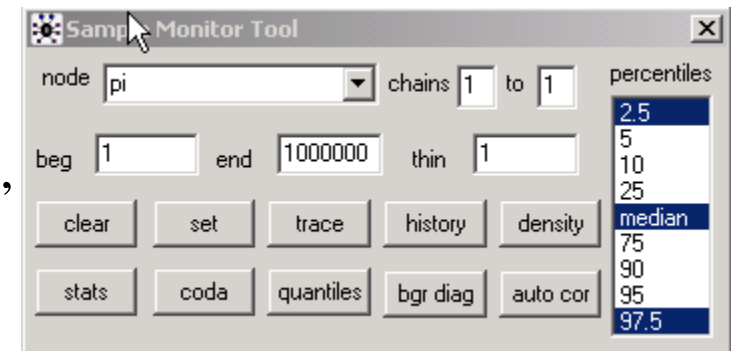
Select Inference > Samples... This brings up the Sample Monitor Tool. In the node window type “pi”, then select . Note here you can also select the percentiles of interest.



To run the model, select Model > Update... This opens the Update Tool. Change updates to 5000 then press . Watch as the iterations are counted by 100 to 5000.




Return to Inference > Samples... to open the Sample Monitor Tool. Scroll to pi in the node window, then press .





The MC error is the Monte Carlo error, it decreases as the number of samples increases. It helps in deciding when enough samples have been taken. Since we know the density is flat on top, we can reduce the wiggles by increasing the number of samples.

Try running 50,000 samples.



Node statistics								
node	mean	sd	MC error	2.5%	median	97.5%	start	sample
pi	0.3509	0.1442	5.985E-4	0.1126	0.3511	0.5874	1	50000

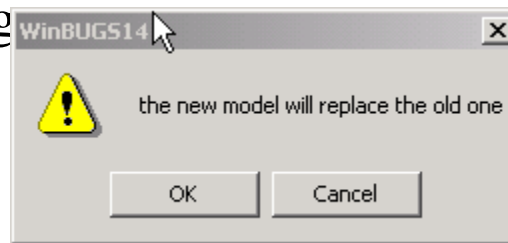
Note the smoother density and the reduction in the MC error from 0.002 with 5000 samples to 0.0006 for 50000 samples.

For comparison and practice, let's rerun the analysis with a slightly different model. Let's restrict our prior to the range between 0.2 and 0.45. In words, we are more precise about what we know concerning the value of pi.

Go back to the Doodle window and change the upper and lower bounds accordingly. Then select Write Code. Check to see if the code is consistent with the Doodle.

With the Model window highlighted (window containing the code), select Model > Specification... to open the Specification Tool. First press . You will receive the following warning

Click  .



Then click  and .

Select Inference > Samples... to bring up the Sample Monitor Tool.

Type in "pi" and the press .

Select Model > Update... to open the Update Tool. Change updates to 5000 then press .

Select Inference > Samples... Scroll to pi in the node window and press  and .

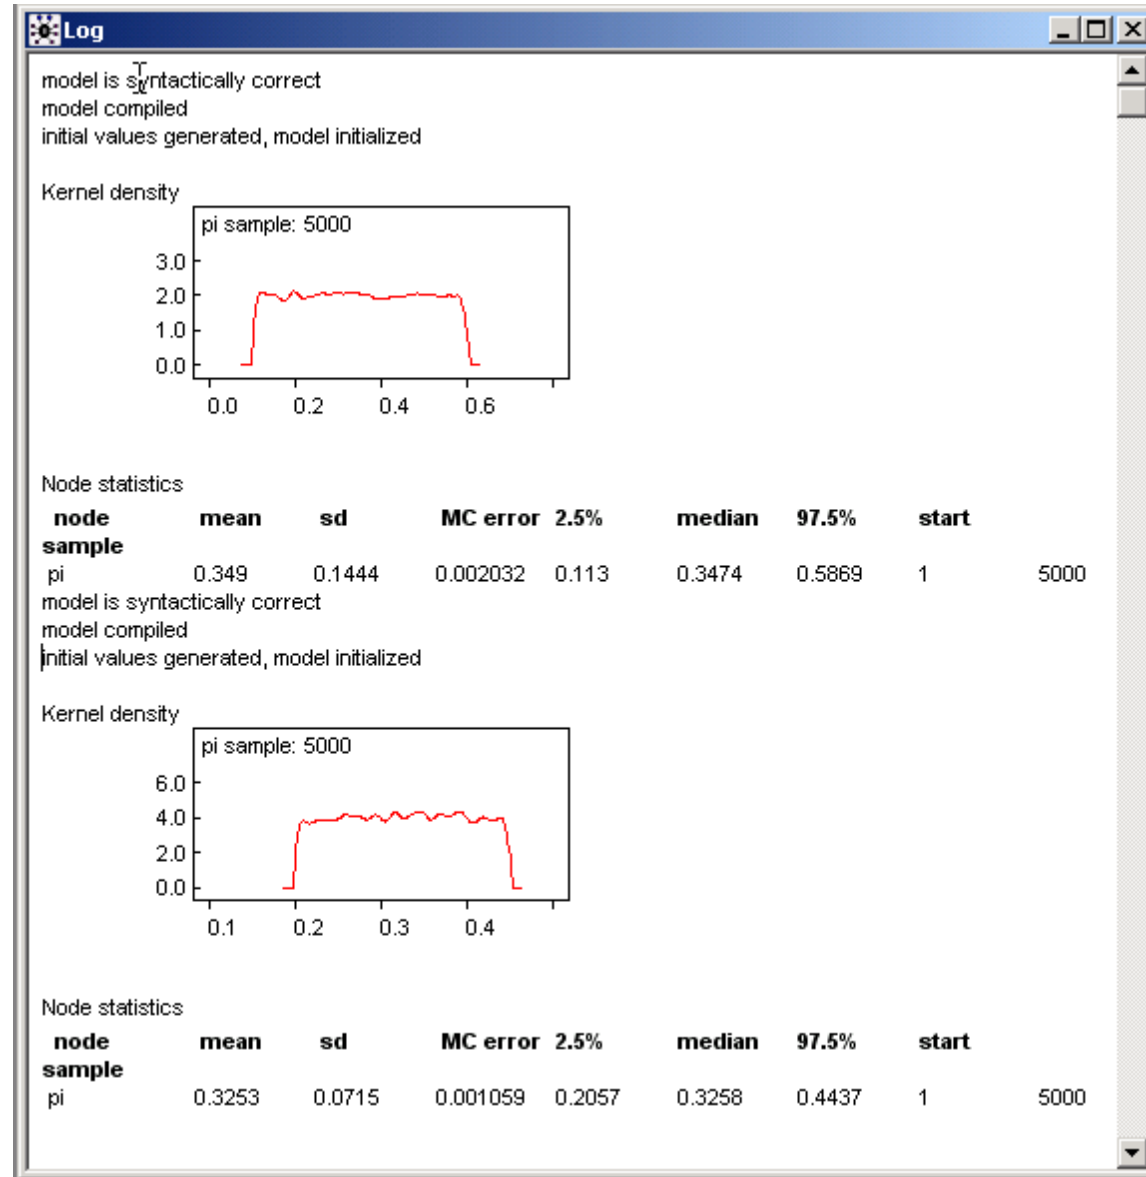
The new results are added to the Log window.

Note that the x- and y-axis scales are different in the two graphs.

We see that the range of possible values is constrained and that the mean value shifts to the left (is smaller).

The graph indicates that we believe the true value for pi is bounded, but that within the bounds any value is equally likely.

This simple example demonstrates how WinBUGS works. It shows how to start with a doodle and end up with a set of random numbers that encapsulate our belief about the unknown population parameter.



## Inference by combining your prior with data

The real power of WinBUGS comes from the ability to combine your prior beliefs with data you have in hand, thus allowing you to make inferences.

Continuing with our unemployment example, suppose you have results from a small sample of 14 people. A total of  $n = 14$  people were surveyed and  $r = 4$  of them were unemployed.

These data will have a **binomial distribution** with proportion  $\pi$  and denominator  $N$ .

Go back to your directed graph and add two additional nodes.

Add a constant node  $N$ .

Add a stochastic node  $r$  with binomial density, proportion equal to  $\pi$  and order  $N$ .

Left click in your Doodle window, change **type** to constant and **name** to  $N$ .

Left click again to add a stochastic node with **name**  $r$ , **density**  $\text{dbin}$ , **proportion**  $\pi$  and **order**  $N$ .

To add links between the nodes, click on node  $r$  to highlight it. With the Ctrl key held, click on node  $\pi$  and node  $N$ . Arrows will be added to the nodes. The arrows will point to the highlighted node. The solid arrows indicates a statistical relationship (is distributed as).

16

The directed graph describes the new model.

The number of unemployed  $r$  is estimated from our prior and data as a random variable having a binomial distribution with proportion  $\pi$  and order  $N$ .

$N$  is a constant, and the prior for  $\pi$  is a beta distribution with two parameters ( $a=b=1$ ) and restricted between 0.2 and 0.45.

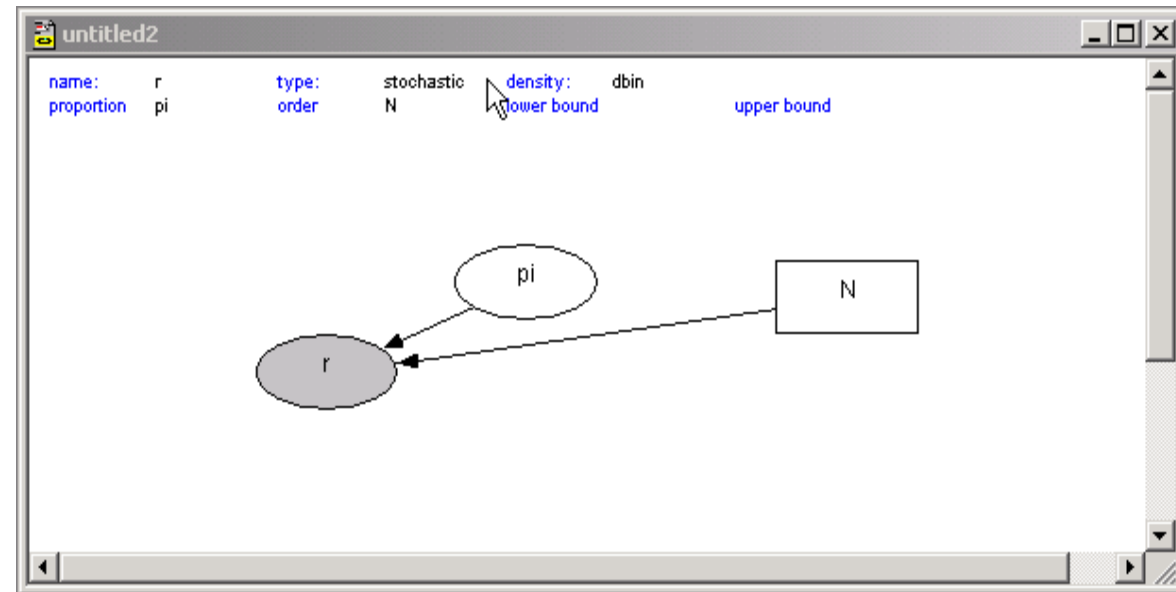
Select Doodle > Write Code

To enter the data, type the following in the model code window.


`list(N=14,r=4)`

```
model;
{
  pi ~ dbeta(1,1)|( 0.2,0.45)
  r ~ dbin(pi,N)
}

list(N=14,r=4)
```

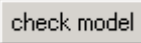


```
model;
{
  pi ~ dbeta(1,1)|( 0.2,0.45)
  r ~ dbin(pi,N)
} I
```

 When you make you add arrows, make sure the parameters of the stochastic node do not change.



Select Model > Specification...

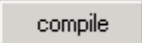

Click .

Highlight the word “list” and

press .

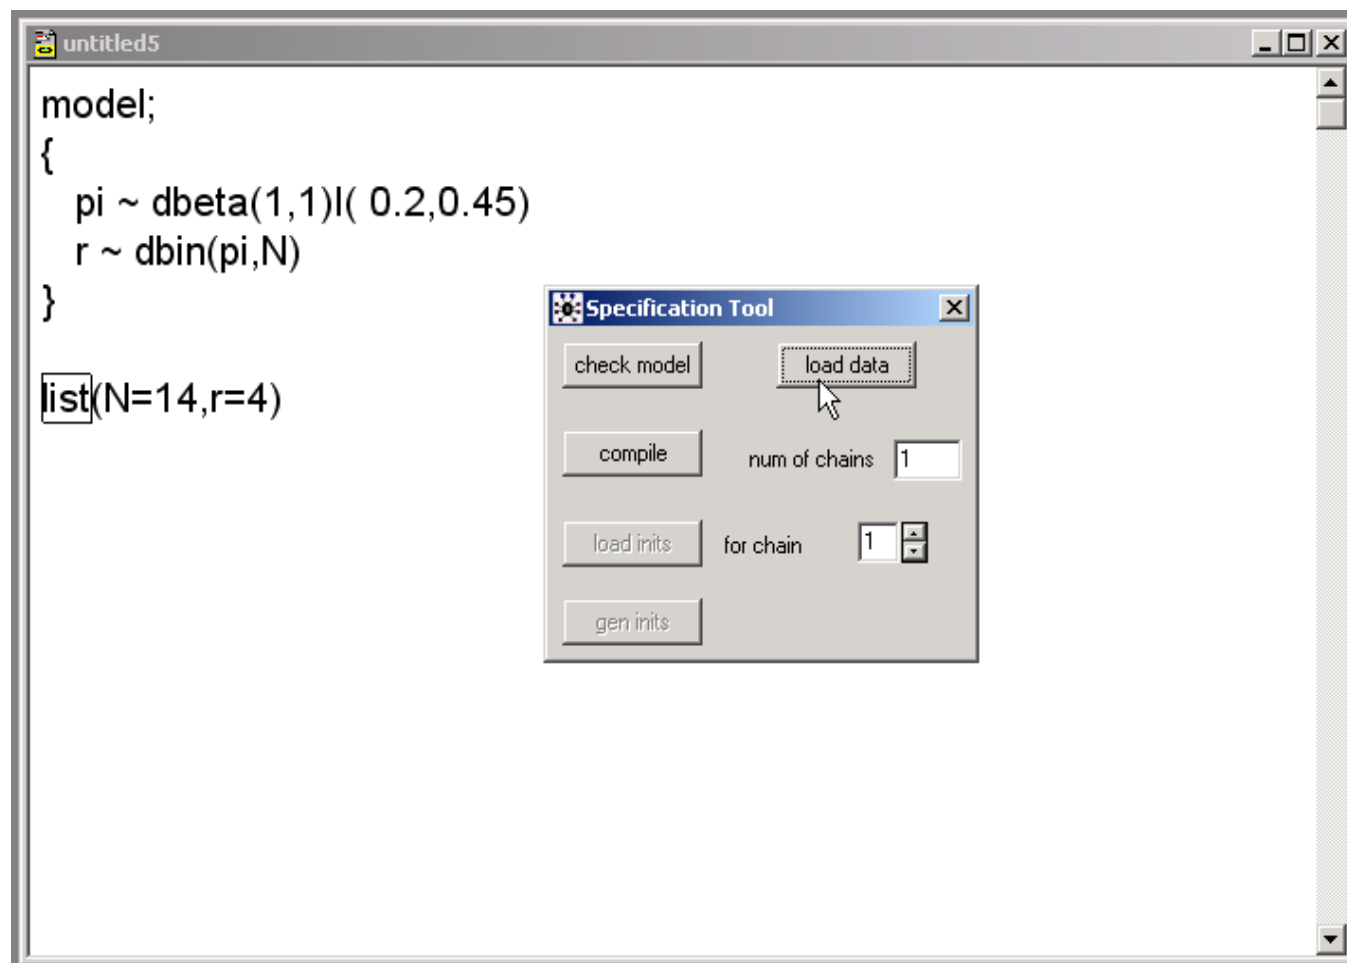
everything

is well you will see the message “data loaded”.

Press  then .

Select Inference > Samples...

Type “pi” and press .



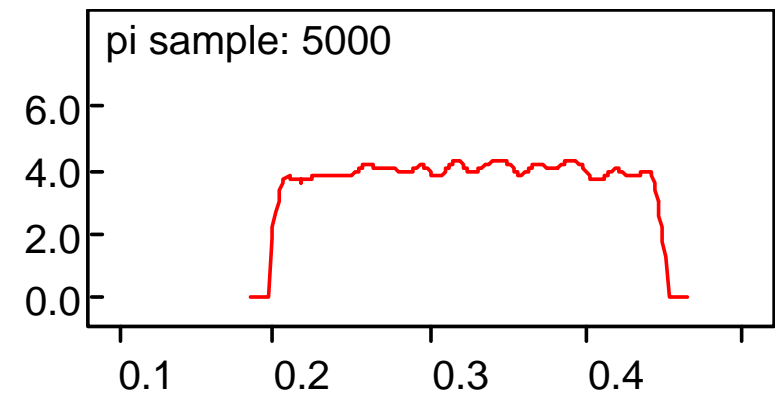
WinBUGS has data for a node with a distribution, so it will calculate the appropriate likelihood function and prior for pi, and combine them into a posterior distribution. It knows about conjugate pair of distributions, so the calculation is straightforward.

Select Model > Update... Change updates to 5000, then press .

WinBUGS generates updated samples of pi (updated from the initial) by combining the prior information on pi and the new information on pi given by the data r and N.


## Prior

node	mean	sd	MC error	2.5%	median	97.5%	start
sample							
pi	0.3253	0.0715	0.001059	0.2057	0.3258	0.4437	1



## Posterior

node	mean	sd	MC error	2.5%	median	97.5%	start
sample							
pi	0.3144	0.06675	9.218E-4	0.2072	0.3101	0.4393	1




Note how the data changes our view of the unemployment rate. For one thing, the data gives us reason to think that the unemployment rate  $\pi$  is less than 0.4. There is still plenty of uncertainty, but it is less than before we took the sample. The standard deviation (sd) of  $\pi$  from the prior is 0.072 but is 0.067 from the posterior. The 95% credible interval shrinks accordingly.

For more practice, and to see the effect of a larger sample on the posterior, rerun the model with  $n=140$ , and  $r=40$ .

## Posterior

node	mean	sd	MC error	2.5%	median	97.5%
sample						
pi	0.2876	0.03749	5.096E-4	0.2194	0.2864	0.3642



Thus as we increase our information about  $\pi$  through more data, the influence of the prior on the posterior decreases. This is seen by the fact that the posterior density looks less like the original prior distribution.

Example #2: Suppose that we take another survey, perhaps at some time later. This time 12 different people were asked and 5 said they were unemployed. What is the evidence that the underlying rates for the two surveys is really different?


Note: From the first sample,  $4/14$  or 29% of the people were not employed. From the second sample,  $5/12$  or 42% of the people are unemployed. Looking only at the percentages, the difference appears to be large.

Using WinBUGS we can determine how much we know about the differences in the rates from these two small surveys by calculating a posterior for the difference. Alternatively, we can calculate a posterior for the ratio.

## A model for the ratio and difference of two binomial samples

Use the model from example #1 and add a second set of nodes for the new survey.

Call the prior  $\pi_2$ , the number of unemployed  $r_2$ , and the number of surveyed  $N_2$ .



### Notes:

The nodes can be rearranged using click and drag.

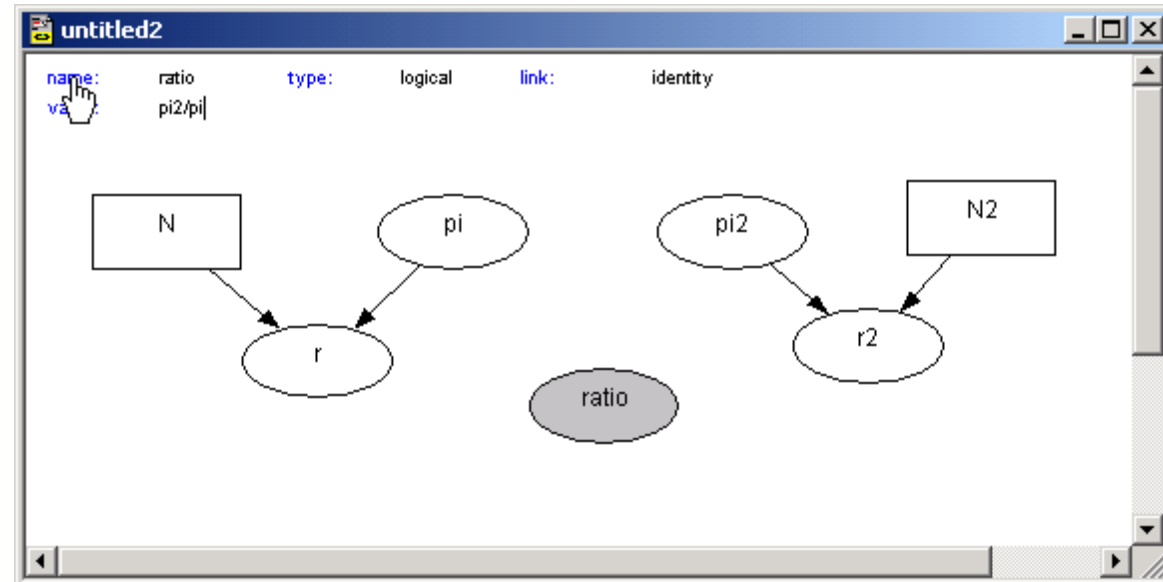
$\pi_2$  is set up the same as  $\pi_1$  using a beta density ( $a=b=1$ ) and bounded.

$N_2$  is a constant like  $N$ .

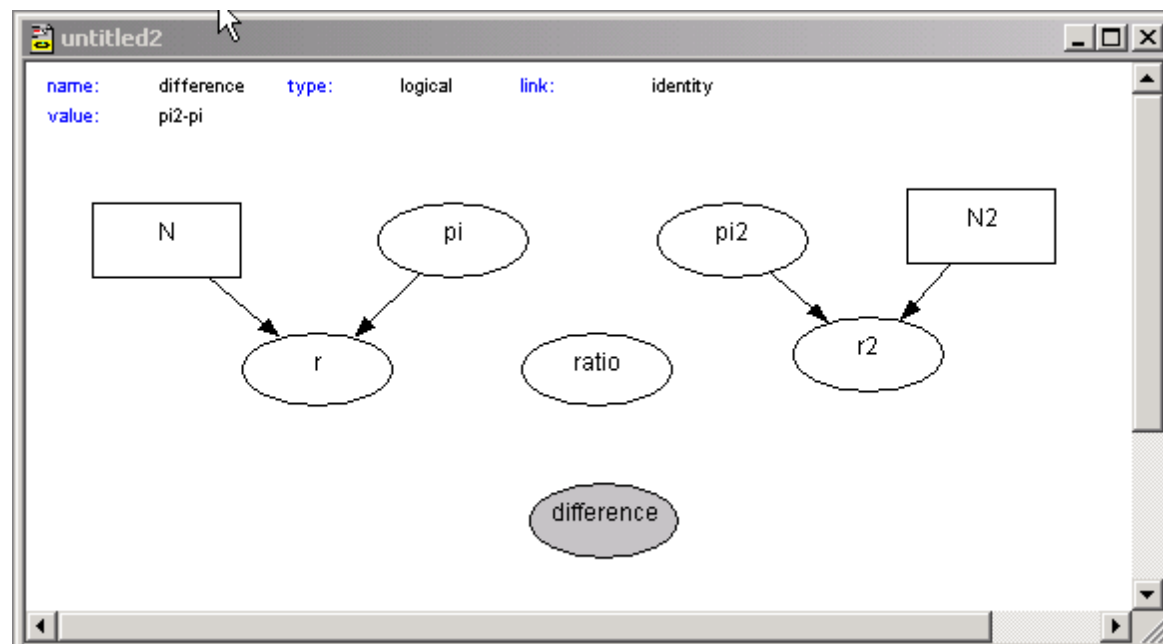
$r_2$  is set up the same as  $r$  using a binomial density with proportion  $\pi_2$  and order  $N_2$ .

To get the differences and ratios we create two logical nodes.

Left click to create a node. Change **type** to logical, **name** to ratio, leave **link** as identity, and set **value** to  $\pi^2/\pi$ .

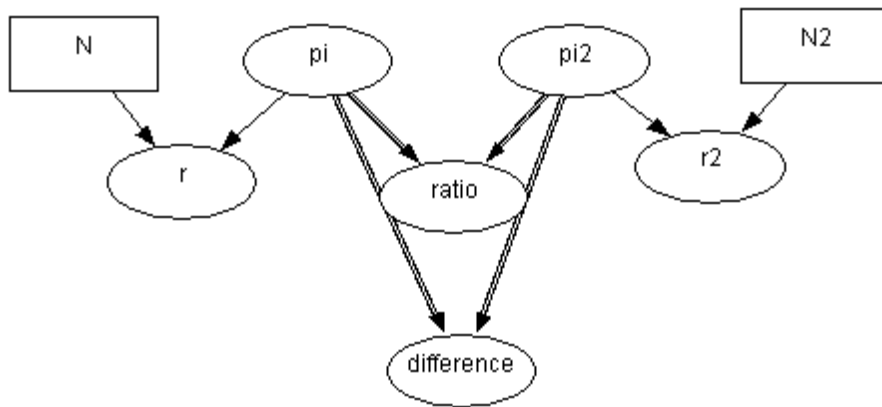


Left click to create a node. Change **type** to logical, **name** to difference, leave **link** as identity, and set **value** to  $\pi^2 - \pi$ .



Add the links to finish the model. Note that the links to the logical nodes are made with a hollow arrow which indicates a deterministic relationship as opposed to the solid arrow which indicates a stochastic relationship. The final doodle and corresponding code for the two proportion model are:

### Directed Graph



```
model;
{
  pi ~ dbeta(1,1)|( 0.2,0.45)
  r ~ dbin(pi,N)
  pi2 ~ dbeta(1,1)|( 0.2,0.45)
  r2 ~ dbin(pi2,N2)
  ratio <- pi2 / pi
  difference <- pi2 - pi
}
```

A deterministic relationship as indicated by a hollow arrow gets coded in the model using the `<-` symbol as is used in R or Splus. A statistical relationship gets coded in the model with a `~` symbol.


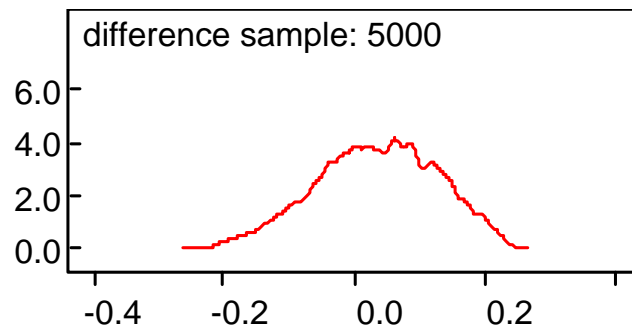
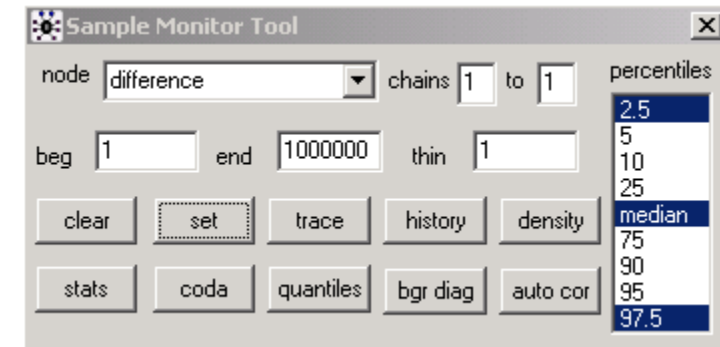

Add the data.

In the model window type: `list(N=14, r=4, N2=12, r2=5)`.

Use the Specification Tool to check the model, load the data, compile, and generate the initial values.

Use the Sample Monitor Tool to set ratio and difference.

Use the Update Tool to generate 5000 values of the ratio and difference.



node	mean	sd	MC error	2.5%	median	97.5%	node	mean	sd	MC error	2.5%	median	97.5%
sample							sample						
difference	0.03275	0.0927	0.001252	-0.155	0.03572	0.1996	ratio	1.156	0.3338	0.004347	0.6136	1.115	1.892



Does it appear as if the sample rates of unemployment are significantly different?

Hint: The area under the difference curve to the right of zero is not much larger than the area to the left of zero. Also the area to the right of one under the ratio curve is not much larger than the area to the left of one.

Also, the mean difference is close to zero and the mean ratio is close to one. The 95% credible interval for the difference in unemployment rates is  $(-0.15, 0.20)$ , which is interpreted to mean that there is a 95% probability that the difference lies somewhere in this interval. This interval includes 0. Similarly the 95% credible interval for the ratio of unemployment rates is  $(0.61, 1.90)$ , which includes the value of 1.

In Bayesian analysis, the 95% **credible interval** is the analogue of the 95% **confidence interval** in conventional statistics. However, with a Bayesian analysis we state that there is a 95% probability that the parameter is between the interval values. Whereas in a conventional analysis we state that 95% of all such intervals will contain the true, but unknown value for the parameter assuming the null hypothesis is correct.

## Summary

This tutorial describes the basics of using WinBUGS.

It explains how to make directed graphs using Doodle. A directed graph provides visual representation of a statistical model. The graphs simplify complex models, communicate the structure of the problem, and provide the basis for computation.

It explains how to compile, load data and run WinBUGS code. The code can be written automatically from the directed graph.

It explains how to view output using a log file. The density and statistics give information about the posterior distribution that can be used for drawing inferences.

The examples were easy as they relied on conjugate distributions (beta, binomial).

Also, there was only one unknown parameter in example #1 and only two unknown parameters in example #2.

WinBUGS works well for these types of problems. For more complicated problems with lots of parameters, we cannot be so sure. WinBUGS provides a set of diagnostic tools that allow you to check whether things are working properly.



Things can go wrong and the help manual contains a warning. Using examples from the manual is a good path to follow, but always look critically at your results to see if they make sense. Use the diagnostic tools.