Pooled Prevalence Calculator
User Guide

Contents

Contents 1
Introduction 1
Overview 2
Bayesian vs frequentist methods 3
Important Assumptions 4
Pooled prevalence estimates are biased! 5
Pooled prevalence for fixed pool size and perfect tests 6
  Method 1 6
  Method 2 7
Pooled prevalence for fixed pool size and tests with known sensitivity and specificity 8
  Method 3 8
  Method 4 9
Pooled prevalence for fixed pool size and tests with uncertain sensitivity and specificity . 10
  Method 5 10
Pooled prevalence for variable pool size and perfect tests 12
Pooled prevalence using a Gibbs sampler 12
Estimated true prevalence using one test with a Gibbs sampler 14
Estimated true prevalence using two tests with a Gibbs sampler 16
Estimation of alpha and beta parameters for prior Beta distributions and summarisation of Beta distributions for specified alpha and beta parameters 18

Introduction

Pooled or group testing is a testing strategy where samples from a number of individuals are aggregated into a single sample (or pool), which is then tested for the disease or agent of interest. Pooled testing strategies have been proposed for identification of infected individuals (classification), for demonstration of herd or area infection status and for estimation of prevalence of infection.

The utilities provided on this web-site are designed to assist with the estimation of prevalence from the testing pof pooled samples. When used for this purpose, individual samples are aggregated into pools for testing. Individual-level prevalence is then estimated based on the number of individuals represented in each pool and the test result for that pool. Estimation methods are provided for both fixed and variable pool size and to adjust estimates for imperfect sensitivity and specificity of the test used.

Pooled testing has some significant advantages over individual testing. In particular, pooled testing provides:

- significant cost savings, particularly where the cost of collecting additional samples is small relative to the cost of the test or if prevalence is low;
- increased precision of estimates compared to individual testing where the same number of tests is undertaken, particularly if prevalence is less than about 30%;

- increased precision of estimates compared to individual testing where test sensitivity and/or specificity are less than 1; and
- reduced bias in estimates when assumed values for sensitivity and specificity are not equal to the true values.

Pooled testing strategies also have a number of disadvantages, including:

- logistical requirements and cost for pooling of samples, either in the laboratory or in the field;
- the requirement of some strategies for re-testing of samples from positive pools either as smaller pools or individually, increasing requirements for sample volumes, storage and handling;
- the possible effect of dilution on test performance, particularly test sensitivity cannot be ignored, and any appropriate for imperfect sensitivity should include consideration of any effect of dilution on test sensitivity;
- estimates may be biased to a variable degree, depending on prevalence, number of pools and pool size; and
- The methods depend on a number of major assumptions, and may result in biased estimates if the assumptions are violated. These assumptions include that:
    - the outcome is assumed to follow a binomial distribution - clustering or overdispersion of the positive outcome can result in substantial bias in the resulting estimate;
    - sample size is small compared to the population;
    - the health status of each individual is independent of the status of others, both within and between pools;
    - test sensitivity and specificity are 100% (for some methods);
    - samples are assumed to be mixed homogeneously in the pools and any sub-samples taken for testing are equally representative of all of the individuals contributing to each pool; and
    - all pools represent the same number of individuals (except for the variable pool-size method).

Additional utilities are also provided for the estimation of true prevalence based on testing of individual samples (unpooled) using either one or two tests of inperfect and uncertain sensitivity and specificity. These utilities use Bayesian methods to account for uncertainty about the true values of sensitivity and specificity for the test(s) used.

Overview

The programs available on this site provide a suite of utilities for estimating prevalence and to assist in the design of sampling and pooling strategies for the estimation of disease prevalence from the testing of pooled (or individual) samples. The various programs have been

implemented in the statistical software environment "R", with your web browser used to pass input values to the program and to display the output of the resulting analysis.

To use the programs, enter the desired input values in each text box on the input screen and click on submit. Example values are already displayed in the input boxes, but can be over-written with your own values. Alternatively, you can use the default values to experiment and see how the program works. All input values are checked before processing, to ensure that they are valid and within the ranges specified in the accompanying description. An informative error message is displayed if invalid input values are entered, and the program will not run until these values are corrected. Input values for parameters that can be represented as percentages, proportions or probabilities (prevalence, sensitivity, specificity, confidence limits) must be entered as proportions (decimal numbers between 0 and 1). Similarly, output results for these parameters wil also be expressed as proportions.

Error checking is limited to the numerical validity of input values, for example by checking valid ranges or that counts are input as integers. It is therefore possible to enter inappropriate or unlikely values which could result in non-sensical output. It is the user's responsibility to ensure that input values are appropriate and that results are meaningful.

Some of the utilities included in this web-site use simulation to estimate parameter values or to evaluate proposed testing strategies. These simulations require multiple iterations (runs) of the model to produce the required result. For the Bayesian analysis, many iterations are required to allow the model to converge on the true parameter values, and additional iterations are then required for inference about the value. In most cases a minimum of 10,000 iterations is recommended (in some cases 20,000 - 50,000 may be better), with 2,000 - 5,000 iterations discarded to allow for convergence of the model. For other programs, 5,000 - 10,000 iterations is usually sufficient. Because of the large number of iterations required, some of these simulations make take several minutes (or longer) to complete. For long simulations that are likely to result in a time-out on your browser you will be offered an option to have the results emailed to you.

Output from each program is returned to your web-browser in a standard format. This starts with a brief description of the analysis/method used, followed by a summary of input values and finally a summary table of results. For most programs, graphical representations and text files of detailed results are available for most analyses by clicking on the appropriate icon in the results table. Text files of results can be either opened directly in MS Excel or saved on your PC in a tab-delimited format.

A summary description and brief help is provided on the input page for each program, with a more detailed description provided for al the programs in this user guide.

Bayesian vs frequentist methods

The analytical methods provided on this site all fall into one of two broad categories of statistical mehtods: frequentist or Bayesian.

Frequentist methods use conventional statistical techniques to calculate maximum-likelihood estimates of true prevalence and confidence limits, in a similar manner to standard techniques used to analyse conventional survey data. Frequentist methods are conceptually simpler to understand and are usually computationally easier to implement (and take less computer time to run). They do not take account of any existing knowledge of the likely prevalence,

although some methods do allow for adjustment of estimates for imperfect sensitivity and specificity of the tests used.

On the other hand, Bayesian analysis uses simulation (using a Gibbs sampler) to derive a posterior probability distribution(s) for the parameter(s) of interest - usually true prevalence but distributions for sensitivity, specificity and other parameters are also generated.

Bayesian methods also have the advantages that:

- pre-existing estimates of prevalence can be incorporated in the analysis to increase confidence in the results;
- imperfect sensitivity and specificity of tests and uncertainty about their true values are incorporated explicitly in the procedure; and
- the lower probability limit of the estimated prevalence can never be negative.

Briefly, a Bayesian approach allows the combination of any prior information available on test sensitivity and specificity and estimated prevalence of disease with the results of testing, to produce a posterior probability distribution of the estimated true prevalence (and other measures such as test sensitivity and specificity) that best fits the combination of prior distributions and observed testing results. Bayesian methods were initially developed for estimating prevalence from individual testing and were subsequently extended for use with pooled testing strategies.

However, Bayesian estimates can be seriously affected by the use of inappropriate prior distributions (inaccurate estimates and/or overconfidence in the values) for prevalence, sensitivity or specificity and therefore must be used with care. Wherever possible prior estimates should be based on real data and should be appropriately weighted (wide probability limits) to ensure that any errors do not dominate the data, causing inaccurate results. See the Glossary for more details on Bayesian methods and the Beta distribution and its parameters.

Bayesian also rely on simulation rather than analytical methods, and therefore can take some time to run, depending on the number of iterations used. It is also that sufficient iterations are run to allow convergence of the Bayesian model and to support inference from the results.

Outputs from the Gibbs sampler are revised estimates of prevalence, test sensitivity, test specificity and other parameters as posterior probability distributions.

Important Assumptions

As with conventional methods for estimating prevalence, pooled testing methods and bayesian methods for estimating prevalence depend on a number of important assumptions. Violation of these assumptions will result in biased estimates. Key assumptions are that:

- the outcome is assumed to follow a binomial distribution - clustering or overdispersion of the positive outcome can cause substantial bias in the resulting estimate,
- the health status of each individual is independent of the status of others, both within and between pools,
- sampling of the population is by simple random sampling,

- individual samples are allocated into groups (pools) by random selection,
- sample size is small relative to the population being sampled,
- assumed values for sensitivity and specificity are appropriate,
- dilution of individual samples by pooling has no effect on sensitivity or specificity estimates (or the estimates used take any dilution effects into account),
- samples are assumed to be mixed homogeneously in the pools and any sub-samples taken for testing are equally representative of all of the individuals contributing to each pool,
- all pools represent the same number of individuals (except for the variable pool-size method),
- assumed prior distributions for prevalence, sensitivity and specificity for Bayesian methods are appropriate and
- assumed true prevalence for sample size calculations and simulations is appropriate.

Pooled prevalence estimates are biased!

For all of the frequentist methods for estimating prevalence from pooled samples, p is an upwardly biased estimator of the true prevalence. The magnitude of the bias is decreased with lower prevalence (p), increased numbers of pools (m), smaller pool size (k) and increased total sample size (n). Bias also increases as the probability that all pools will test positive increases.

In general, bias is negligible for $m > 30$. However, even for quite small values for m and k the bias may also be quite small, particularly if p is low.

The estimated prevalence (p) is also sensitive to (and may be biased by) errors in the assumptions of perfect test sensitivity or specificity. p is particularly sensitive to errors in sensitivity as p increases and if k is too large. Clustering or overdispersion of positive individuals in the sampled population can also result in substantial bias in prevalence estimates.

The actual bias in any estimate depends on the true prevalence, pool size and the number of pools and can be estimated for any particular pooling strategy using simulation methods. Simulation utilities are provided for both fixed and variable pool-size strategies to assist in evaluating the potential bias in proposed pooling strategies.

Bias can be minimised by ensuring an adequate total sample size, by testing a larger number of pools of smaller size, rather than vice versa or by testing several individual samples in addition to the pooled samples (using the variable pool size method).

Pooled prevalence for fixed pool size and perfect tests

These methods all use frequentist approaches to estimate prevalence and confidence limits, assuming a fixed pool size and perfect (100%) test sensitivity and specificity, as described below.

Method 1

This method (Method 2 from Cowling et al. (1999) or Sacks et al. (1989)) assumes 100% test sensitivity and specificity and fixed pool size. Confidence limits are based on a normal approximation and may be <0 for low prevalence values.

Prevalence is estimated as:

$$p = 1 - (1 - x/m)^{1/k}$$

and the standard error (SE(p))is estimated as the square root of the variance, given by:

$$Var(p) = ((x/m) \times (1 - x/m)^{2/k - 1})/(mk^2)$$

where:

- p = estimated prevalence,
- k = pool size,
- m = the number of pools tested and
- x = the number of positive pools.

Asymptotic confidence limits are calculated using the normal approximation:

$$CL = p \pm Z_{\alpha/2} \times SE(p)$$

where $Z_{\alpha/2}$ is the standardised normal variate corresponding to the desired confidence limit.

Required inputs for this method are:

- pool size,
- number of pools tested,
- number of pools positive and
- desired upper and lower confidence limits for the estimate.

Pool size, number of pools and number of pools positive must be positive integers and the number of positive pools must be less than the number of pools tested. Upper and lower confidence limits must be >0 and <1.

Outputs include:

- a point estimate of animal-level prevalence,
- upper and lower asymptotic confidence limits,
- the standard error for the estimate and

- graphic and text files of estimates and confidence limits for all possible results (download by clicking on the appropriate icon).

The algorithm used to estimate prevalence and confidence limits fails if either all or none of the pools are positive. In these cases the point estimates are 100% and 0% respectively.

Method 2

This method (Method 3 from Cowling et al. (1999)) assumes 100% test sensitivity and specificity and fixed pool size. Exact confidence limits are calculated based on binomial theory, so that confidence limits are never <0 or >1.

Prevalence and variance are estimated as for Method 1:

$$p = 1 - (1 - x/m)^{1/k}$$

and:

$$Var(p) = ((x/m) \times (1 - x/m)^{2/k - 1})/(mk^2)$$

where:

- p = estimated prevalence,
- k = pool size,
- m = the number of pools tested and
- x = the number of positive pools.

Exact confidence limits are estimated by calculating the corresponding binomial confidence limits for the proportion of positive pools and then transforming these back to individual-level prevalence values using the equation for estimating prevalence from Method 1.

Required inputs for this method are:

- pool size,
- number of pools tested,
- number of pools positive and
- desired upper and lower confidence limits for the estimate.

Pool size, number of pools and number of pools positive must be positive integers and the number of positive pools must be less than the number of pools tested. Upper and lower confidence limits must be >0 and <1.

Outputs include:

- a point estimate of animal-level prevalence,
- upper and lower exact (binomial) confidence limits,
- the standard error for the estimate and
- graphic and text files of estimates and confidence limits for all possible results (download by clicking on the appropriate icon).

The algorithm used to estimate prevalence and confidence limits fails if either all or none of the pools are positive. In these cases the point estimates are 100% and 0% respectively.

Pooled prevalence for fixed pool size and tests with known sensitivity and specificity

These methods use frequentist approaches to estimate prevalence and confidence limits, assuming a fixed pool size and a test with known values for sensitivity and specificity, as described below.

Method 3

This method (Method 4 from Cowling et al. (1999)) assumes a fixed pool size and that test sensitivity and specificity are known exactly (no uncertainty about their values). Confidence limits are based on a normal approximation and may be <0 for low prevalence values.

Prevalence is estimated as:

$$p = 1 - ((Se - x/m)/(Se + Sp - 1))^{1/k}$$

and the standard error (SE(p)) is estimated as the square root of the variance, given by:

$$Var(p) = ((1 - p)^{2/k - 2}/mk^2) \times ((x/m \times (1 - x/m))/(Se + Sp - 1)^2)$$

where:

- p = estimated prevalence,
- k = pool size,
- m = the number of pools tested,
- x = the number of positive pools,
- Se = the sensitivity of the test and
- Sp = the specificity of the test.

Asymptotic confidence limits are calculated using the normal approximation:

$$CL = p \pm Z_{\alpha/2} \times SE(p)$$

where $Z_{\alpha/2}$ is the standardised normal variate corresponding to the desired confidence limit.

Required inputs for this method are:

- pool size,
- number of pools tested,
- number of pools positive,
- assumed sensitivity of the test,
- assumed specificity of the test and
- desired upper and lower confidence limits for the estimate.

Pool size, number of pools and number of pools positive must be positive integers and the number of positive pools must be less than the number of pools tested. Sensitivity and specificity must be >0 and <=1 and upper and lower confidence limits must be >0 and <1.

Outputs include:

- a point estimate of animal-level prevalence,
- upper and lower asymptotic confidence limits,
- the standard error for the estimate and
- graphic and text files of estimates and confidence limits for all possible results (download by clicking on the appropriate icon).

Estimates are only valid if the proportion of positive pools is greater than the false positive rate (1 - Specificity) and less than or equal to the true positive rate (Sensitivity). Invalid results are indicated by NA in the results table.

Method 4

This method (Method 5 from Cowling et al. (1999)) assumes a fixed pool size and that test sensitivity and specificity are known exactly (no uncertainty about their values). Exact confidence limits are calculated based on binomial theory, so that confidence limits are never <0 or >1.

Prevalence and variance are estimated as for Method 3:

$$p = 1 - ((Se - x/m)/(Se + Sp - 1))^{1/k}$$

and:

$$Var(p) = ((1 - p)^{2/k - 2}/mk^2) \times ((x/m \times (1 - x/m))/(Se + Sp - 1)^2)$$

where:

- p = estimated prevalence,
- k = pool size,
- m = the number of pools tested,
- x = the number of positive pools,
- Se = the sensitivity of the test and
- Sp = the specificity of the test.

Exact confidence limits are estimated by calculating the corresponding binomial confidence limits for the proportion of positive pools and then transforming these back to individual-level prevalence values using the equation for estimating prevalence from Method 3.

Required inputs for this method are:

- pool size,
- number of pools tested,

- number of pools positive,
- assumed sensitivity of the test,
- assumed specificity of the test and
- desired upper and lower confidence limits for the estimate.

Pool size, number of pools and number of pools positive must be positive integers and the number of positive pools must be less than the number of pools tested. Sensitivity and specificity must be >0 and <=1 and upper and lower confidence limits must be >0 and <1.

Outputs include:

- a point estimate of animal-level prevalence,
- upper and lower exact (binomial) confidence limits,
- the standard error for the estimate and
- graphic and text files of estimates and confidence limits for all possible results (download by clicking on the appropriate icon).

Estimates are only valid if the proportion of positive pools is greater than the false positive rate (1 - Specificity) and less than or equal to the true positive rate (Sensitivity). Invalid results are indicated by NA in the results table.

Pooled prevalence for fixed pool size and tests with uncertain sensitivity and specificity

This method uses frequentist approaches to estimate prevalence and confidence limits, assuming a fixed pool size and a test with unknown (uncertain) values for sensitivity and specificity, as described below.

Method 5

This method (Method 6 from Cowling et al. (1999)) assumes fixed pool size but unknown test sensitivity and specificity. Uncertainty associated with the point estimates of test sensitivity and specificity is incorporated through the inclusion of additional variance associated with the sample size used to determine the values used for these parameters. The smaller the sample size, the greater the uncertainty about the true values for sensitivity and/or specificity and hence the greater the uncertainty about the resulting prevalence estimate. Confidence limits are based on a normal approximation and may be <0 for low prevalence values.

Prevalence is estimated as for Method 3:

$$p = 1 - ((Se - x/m)/(Se + Sp - 1))^{1/k}$$

and the standard error (SE(p)) is estimated as the square root of the variance, given by:

$$\begin{aligned}\mathrm{Var}(p) = &(((Se - x/m)/(Se + Sp - 1))^{2/k - 2})/(k^2 \times (Se + Sp - 1)^2) \\ &\times (x/m \times (1 - x/m)/m + (Se \times (1 - Se)/n1) \\ &\times (1 - (Se - x/m)/(Se + Sp - 1))^2 \\ &+ (Sp \times (1 - Sp)/n2) \times ((Se - x/m)/(Se + Sp - 1))^2)\end{aligned}$$

where:

- p = estimated prevalence,
- k = pool size,
- m = the number of pools tested,
- x = the number of positive pools,
- Se = the sensitivity of the test,
- Sp = the specificity of the test,
- n1 = the sample size for estimating the sensitivity of the test and
- n2 = the sample size for estimating the specificity of the test.

Asymptotic confidence limits are calculated using the normal approximation:
$CL = p \pm Z_{\alpha/2} \times SE(p)$

where $Z_{\alpha/2}$ is the standardised normal variate corresponding to the desired confidence limit.
Required inputs for this method are:

- pool size,
- number of pools tested,
- number of pools positive,
- assumed sensitivity and specificity of the test,
- sample sizes for estimating the sensitivity and specificity of the test and
- desired upper and lower confidence limits for the estimate.

Pool size, number of pools, number of pools positive and sample sizes for estimating sensitivity and specificity must be positive integers and the number of positive pools must be less than the number of pools tested. Sensitivity and specificity must be >0 and <=1 and upper and lower confidence limits must be >0 and <1.

Outputs include:

- a point estimate of animal-level prevalence,
- upper and lower asymptotic confidence limits,
- the standard error for the estimate and
- graphic and text files of estimates and confidence limits for all possible results (download by clicking on the appropriate icon).

Estimates are only valid if the proportion of positive pools is greater than the false positive rate (1 - Specificity) and less than or equal to the true positive rate (Sensitivity). Invalid results are indicated by NA in the results table.

Pooled prevalence for variable pool size and perfect tests

This method uses generalised linear modelling to calculate maximum-likelihood estimates of prevalence and confidence limits where multiple different pool sizes are used. The method assumes 100% test sensitivity and specificity. See Williams & Moffitt (2001) for more details. Prevalence estimates can be calculated for sampling strategies with up to 10 different pool sizes used.

Required inputs for this method are:

- pool size for each pool size used,
- number of pools tested for each pool size used,
- number of pools positive for each pool size used and
- desired level of confidence in the estimate.

The desired confidence level must be a decimal number >0 and <1 (for example, 0.99 = 99% or 0.95 = 95%). Pool sizes and numbers of pools must be positive integers (>0) and the numbers of positive pools must be non-negative integers (>=0). The number of positive pools must be less than or equal to the corresponding number of pools tested.

Outputs include:

- a point estimate of animal-level prevalence and
- upper and lower asymptotic confidence limits for the estimate for the specified level of confidence.

Pooled prevalence using a Gibbs sampler

This method uses a Bayesian approach and a Gibbs sampler iterative model to estimate the posterior distribution of the true animal-level prevalence of infection for a given pool size, number of pools tested, number of pools positive and estimated test sensitivity and specificity. See Cowling et al. (1999) (Method 7) or {Mendoza-Blanco et al. (1996) for more details.

Required inputs for this method are:

- pool size,
- number of pools tested,
- number of pools positive,
- alpha and beta parameters for prior Beta distributions for:
    - assumed true prevalence,
    - estimated test sensitivity and

    - estimated test specificity.

- the number of iterations to be simulated in the Gibbs sampler,
- the number of iterations to be discarded to allow convergence of the model and
- lower and upper probability (confidence) limits for summarising the output distributions.

Pool size, number of pools and number of pools positive must be positive integers and the number of positive pools must be less than the number of pools tested. Alpha and beta parameters for prevalence, sensitivity and specificity must be >0 and upper and lower confidence limits must be >0 and <1. The number of iterations and the number discarded must both be positive integers (>0) and the number discarded must be less than the number of iterations.

The Gibbs sampler is used to estimate the posterior probability distributions of true prevalence, sensitivity and specificity that best fit the data and the prior distributions provided.

Prior estimates of the true prevalence and test sensitivity and specificity may be based on expert knowledge or on previous data. These estimates are specified as Beta probability distributions, with parameters alpha and beta. Beta probability distributions are commonly used to express uncertainty about a proportion based on a random sample of individuals. In this situation, if x individuals are positive for a characteristic out of n examined, then the alpha and beta parameters can be calculated as alpha = x + 1 and beta = n - x + 1. Alternatively, alpha and beta can be calculated using the *Beta distribution utility*, provided estimates of the mode and 5% or 95% confidence limits are available from expert opinion.

Outputs from the Gibbs sampler are posterior probability distributions for animal-level prevalence, test sensitivity and test specificity, described by their:

- minimum,
- lower probability limit,
- median,
- upper probability limit,
- maximum,
- mean and
- standard deviation.
- histogram and density chart (download by clicking on the appropriate icon).
- text files of parameter estimates for all iterations of the Gibbs sampler (download by clicking on the appropriate icon).

Because the Gibbs sampler estimates prevalence iteratively, based on the data and the prior distributions, it may take a number of iterations for the model to converge on the true value. Therefore, a specified number of initial iterations must be discarded (not used for estimation) to allow the model to converge on the true values. This number must be sufficient to allow

convergence, and should be at least 2000 - 5000. It is also important to carry out an adequate number of iterations to support inference from the results. Suggested minimum values for the total number of iterations and the number to be discarded are provided, but can be varied if desired.

This analysis may take several minutes to complete, depending on the number of iterations required.

Estimated true prevalence using one test with a Gibbs sampler

This analysis uses a Bayesian approach and Gibbs sampler to estimate the true animal-level prevalence of infection based on testing of individual (not pooled) samples using a test with imperfect sensitivity and/or specificity. The analysis requires prior estimates of true prevalence, test sensitivity and test specificity as Beta probability distributions, and outputs posterior distributions for prevalence, sensitivity and specificity. See Joseph et al. (1995) for more details.

Required inputs for this analysis are:

- the number of samples tested,
- the number of samples positive,
- alpha and beta parameters for prior Beta distributions for:
    - assumed true prevalence,
    - estimated test sensitivity and
    - estimated test specificity.
- the number of iterations to be simulated in the Gibbs sampler,
- the number of iterations to be discarded to allow convergence of the model,
- lower and upper probability (confidence) limits for summarising the output distributions and
- starting values for the number of true-positive and false-negative test results (the numbers of truly infected individuals among the test-positive and test-negative groups, respectively).

The number of samples tested must be a positive integer and the number of positive samples must be an integer >=0 and <= the number of samples tested. Alpha and beta parameters for prevalence, sensitivity and specificity must be >0 and upper and lower confidence limits must be >0 and <1. Starting values for the numbers of true positives and false negatives must be integers >= zero and <= the number of positive samples and the number of negative samples, respectively. The number of iterations and the number discarded must both be positive integers (>0) and the number discarded must be less than the number of iterations.

The Gibbs sampler is used to estimate the posterior probability distributions of true prevalence, sensitivity and specificity that best fit the data and the prior distributions provided.

Prior estimates of the true prevalence and test sensitivity and specificity may be based on expert knowledge or on previous data. These estimates are specified as Beta probability distributions, with parameters alpha and beta. Beta probability distributions are commonly used to express uncertainty about a proportion based on a random sample of individuals. In this situation, if x individuals are positive for a characteristic out of n examined, then the alpha and beta parameters can be calculated as alpha = x + 1 and beta = n - x + 1. Alternatively, alpha and beta can be calculated using the *Beta distribution utility*, provided estimates of the mode and 5% or 95% confidence limits are available from expert opinion.

Outputs from the Gibbs sampler are posterior probability distributions for:

- animal-level prevalence,
- test sensitivity,
- test specificity,
- positive and negative predictive values for the test,
- positive and negative likelihood ratios for the test and
- the numbers of true-positive and false-negative test results.

These distributions are described by their:

- minimum,
- lower probability limit,
- median,
- upper probability limit,
- maximum,
- mean and
- standard deviation.
- histogram and density chart (download by clicking on the appropriate icon).
- text files of parameter estimates for all iterations of the Gibbs sampler (download by clicking on the appropriate icon).

Because the Gibbs sampler estimates prevalence iteratively, based on the data and the prior distributions, it may take a number of iterations for the model to converge on the true value. Therefore, a specified number of initial iterations must be discarded (not used for estimation) to allow the model to converge on the true values. This number must be sufficient to allow convergence, and should be at least 2000 - 5000. It is also important to carry out an adequate number of iterations to support inference from the results. Suggested minimum values for the total number of iterations and the number to be discarded are provided, but can be varied if desired.

This analysis may take several minutes to complete, depending on the number of iterations required.

Estimated true prevalence using two tests with a Gibbs sampler

This analysis uses a Bayesian approach and Gibbs sampler to estimate the true animal-level prevalence of infection based on testing of individual (not pooled) samples using two (2) tests with imperfect sensitivities and/or specificities. The analysis requires prior estimates of true prevalence and sensitivity and specificity for both tests as Beta probability distributions. Outputs are posterior probability distributions for prevalence, sensitivity and specificity. The analysis assumes that the two tests are independent, conditional on disease status. See Joseph et al. (1995) for more details.

Required inputs for this analysis are:

- the number of samples tested,
- the number of samples in each cell of the 2x2 table of comparative test results,
- alpha and beta parameters for prior Beta distributions for:
    - assumed true prevalence,
    - estimated sensitivity for both tests and
    - estimated specificity for both tests.
- the number of iterations to be simulated in the Gibbs sampler,
- the number of iterations to be discarded to allow convergence of the model,
- lower and upper probability (confidence) limits for summarising the output distributions and
- starting values for the number of truly infected individuals in each cell of the 2x2 table of comparative test results.

The number of samples tested must be a positive integer and the number of positive samples must be an integer >=0 and <= the number of samples tested. Alpha and beta parameters for prevalence, sensitivities and specificities must be >0 and upper and lower confidence limits must be >0 and <1. Starting values for the numbers of truly infected individuals in each cell of the 2x2 table must be integers >= zero and <= the number of results in that cell. The number of iterations and the number discarded must both be positive integers (>0) and the number discarded must be less than the number of iterations.

For this analysis, the observed results of testing with two tests concurently can be described in a 2x2 table as follows:

	Test 2:	
Test 1:	+ve	-ve
+ve:	*a*	*b*
-ve:	*c*	*d*

where a, b, c & d are the observed number of sample results in each cell. A proportion of these samples in each cell will be from truly infected animals, depending on true prevalence and test sensitivities and specificities. The Gibbs sampler is used to estimate the true number of infected animals represented in each of the cells (Y1, Y2, Y3 & Y4) and hence to generate

posterior probability distributions for true prevalence, and test sensitivities and specificities that best fit the data and the prior distributions provided.

Prior estimates of the true prevalence and test sensitivity and specificity may be based on expert knowledge or on previous data. These estimates are specified as Beta probability distributions, with parameters alpha and beta. Beta probability distributions are commonly used to express uncertainty about a proportion based on a random sample of individuals. In this situation, if x individuals are positive for a characteristic out of n examined, then the alpha and beta parameters can be calculated as alpha = x + 1 and beta = n - x + 1. Alternatively, alpha and beta can be calculated using the *Beta distribution utility*, provided estimates of the mode and 5% or 95% confidence limits are available from expert opinion.

Outputs from the Gibbs sampler are posterior probability distributions for:

- animal-level prevalence,
- test sensitivity for both tests,
- test specificity for both tests,
- positive and negative predictive values for both tests,
- the numbers of truly infected individuals (Y1, Y2, Y3 & Y4) in each cell of the 2x2 table describing the comparative test results.

These distributions are described by their:

- minimum,
- lower probability limit,
- median,
- upper probability limit,
- maximum,
- mean and
- standard deviation.
- histogram and density chart (download by clicking on the appropriate icon).
- text files of parameter estimates for all iterations of the Gibbs sampler (download by clicking on the appropriate icon).

Because the Gibbs sampler estimates prevalence iteratively, based on the data and the prior distributions, it may take a number of iterations for the model to converge on the true value. Therefore, a specified number of initial iterations must be discarded (not used for estimation) to allow the model to converge on the true values. This number must be sufficient to allow convergence, and should be at least 2000 - 5000. It is also important to carry out an adequate number of iterations to support inference from the results. Suggested minimum values for the total number of iterations and the number to be discarded are provided, but can be varied if desired.

This analysis may take several minutes to complete, depending on the number of iterations required.

Estimation of alpha and beta parameters for prior Beta distributions and summarisation of Beta distributions for specified alpha and beta parameters

This utility calculates the alpha and beta parameters for prior Beta distributions, based on the values specified for the mode and 5th or 95th percentile of the distribution. See Suess et al. (2002) for more details. An accompanying utility allows the calculation of summary values for Beta distributions for specified alpha and beta parameters.

Beta distributions are a type of probability distribution that is commonly used to describe uncertainty about the true value of a proportion, such as sensitivity, specificity or prevalence. They are appropriate distributions to express uncertainty about the prior values for prevalence, sensitivity or specificity in the Gibbs sampler (Joseph et al., 1995; Vose, 2000). When used for this purpose, the Beta distribution can be defined by the two parameters, alpha and beta (written as Beta(alpha, beta)), with alpha = x + 1 and beta = n - x + 1, where x is the number of positive events out of n trials. As n increases, the degree of uncertainty (the width of the distribution) about the estimated proportion (x/n) decreases. The mean value of the distribution can be calculated as alpha/(alpha+beta) and the mode is (alpha-1)/(alpha+beta-2).

Required inputs for the parameterisation utility are:

- mode and 5/95th percentile for at least one distribution and
- mode and 5/95th percentile for each distribution required, up to a maximum of six distributions.

Required inputs for the summarisation utility are:

- alpha and beta parameters for at least one distribution and
- alpha and beta parameters for each distribution required, up to a maximum of six distributions.

If any input cell is blank or contains an invalid value results will not be calculated for that or any subsequent distributions.

For parameterisation of Beta distributions, all values entered must be >0 and <1. It is suggested that, where the mode is less than 0.5, you enter the 95th percentile, and where the mode is greater than 0.5 enter the 5th percentile. For example, if the most likely value for the prior estimate of prevalence is 0.1, and you are 95% confident that the true prevalence is less than 0.2, enter a mode of 0.1 and a 95th percentile value of 0.2. Conversely, if the most likely value for test sensitivity is 0.95 and you are 95% confident that the true value is >0.75, use the values 0.95 and 0.75. For summarisation of Beta distributions, all parameter values must be positive real (decimal or integer) numbers.

Output from the parameterisation utility is a table of input values (mode and percentile) and corresponding alpha and beta values for each distribution requested. Output from the summarisation utility is a table of alpha and beta values and corresponding percentile values,

mean, mode and standard deviation for each distribution requested. Density curves for each distribution are also generated and can be viewed or downloaded by clicking on the icon.