WINPEPI PROGRAMS

WHATIS

MANUAL

(Version 4.61)

May 24, 2016

WHATIS is a WINPEPI program (Abramson 2004), part of the PEPI suite of computer programs for epidemiologists. (“PEPI” is an acronym for “Programs for EPIdemiologists”.)

WHATIS is a "ready reckoner" utility program, providing an expression evaluator and calculators for p-values (and their inverse), confidence intervals, and time spans. It has four modules.

FINDING WHAT YOU WANT

FINDER.PDF (provided with this program) is an alphabetical index that identifies the modules (in all WinPepi programs) that deal with a specific procedure or kind of study. It is called up by pressing F9 or clicking on “*Finder*” in any WinPepi program, or on the FINDER icon, and can be printed for easy reference.

WORDS OF CAUTION

This program offers more options than most users will need, and may display more results than are needed. Ignore the options and results you don't require.

HOW TO USE WHATIS

WHATIS can be run in any version of Windows except Windows 3.

Choose a module, by clicking on its name in the top menu, and follow the on-screen instructions**.**

Recalling results:

Click on *"View"* in the top menu to display the current session's previous results

Pasting results:

Results shown on the screen are automatically copied to the Windows clipboard**,** from which they can be pasted into a Microsoft Word or other text file at the site of the cursor (usually by pressing *Shift-Insert* or *Ctrl-V*. To ensure proper alignment of tabulated results, a Courier or similar font should be used in the text file. If the current session's previous results are recalled (by *clicking on* "View"*),* text can be marked (drag the mouse over it with button pressed) and copied to the clipboard (by pressing *Ctrl-Insert* or *Ctrl-C)* for pasting elsewhere.

Adding comments:

Click on "*Not***e**" in the top menu if you wish to add explanatory comments to be placed in the clipboard, saved, or printed with the results.

Saving results:

By default, all results of Pepi-for-Windows programs are saved in PEPI.TXT in the Winpepi folder, with a warning if it exceeds 500K. Results also go to PEPI.TMP (for display in the 'View' option); this file may be overwritten unless it is renamed on quitting WHATIS. Click on "Saving" (in the top menu) to see the default procedure or to alter it, or to find a button that opens PEPI.TXT (which can also be accessed by clicking on "*Results*" in the Winpepi portal).
[Results saved in earlier installations may be found in C:\PEPI.TXT].
TXT files can be combined by using **JOINTEXT**, supplied with the Winpepi package.

Printing:

Click on "Print". If this fails, a simple solution is to paste the currently-shown results (which have automatically been copied to the Windows clipboard) into a Microsoft Word or other text program, and print from there. To ensure proper alignment of tabulated results, a Courier or similar font should be used in the text file. Results can also be printed from one of the files in which they are automatically saved, e.g. PEPI.TXT.

A DO-IT-YOURSELF THREESOME

1. The WinPepi suite of computer programs for epidemiologists, with their manuals. Can be downloaded free at www.brixtonhealth.com

2. "Research Methods in Community Medicine: Surveys, Epidemiological Research, Programme Evaluation, Clinical Trials" (J.H. Abramson and Z.H. Abramson), sixth edition. John Wiley & Sons, 2008.

3. "Making Sense of Data: A Self-Instruction Manual on the Interpretation of Epidemiological Data" (J.H. Abramson and Z.H.Abramson), third edition. Oxford: Oxford University Press 2001.

HOW TO OBTAIN PEPI PROGRAMS

All WINPEPI (PEPI-for-Windows) and other PEPI programs can be downloaded free. The latest versions of WINPEPI programs – currently COMPARE2, DESCRIBE, ETCETERA, LOGISTIC, PAIRSetc, POISSON, and WHATIS – can be downloaded from www.brixtonhealth.com. The latest release of Version 4 of PEPI, which contains over 40 DOS-based programs (which can be used in Windows) can be downloaded from www.sagebrushpress.com/pepibook.html

COMPARE2, DESCRIBE, ETCETERA, LOGISTIC, PAIRSetc, POISSON, and WHATIS are distributed with manuals (as computer files). A printed manual is available for the earlier DOS-based programs (Abramson and Gahlinger 2001)

WINPEPI programs are provided with no liability to users and without any warranties, whether expressed or implied. They are copyrighted, but may be freely copied and distributed for personal use; they may not be exploited commercially without permission.

CALCULATOR

This is an *expression evaluator*, whose special feature is that it can save up to 24 values and 24 formulae, storing them in a disk file that it creates for this purpose, enabling stored results, constants, and formulae to be recalled and used (in the same or a subsequent session) by entering labels (a, b, etc.) that represent them. This avoids repeated entry of the same numbers (e.g. population denominators), facilitates the performance of calculations in stages (by saving intermediate results for use in the next stage), and permits recomputation of stored formulae, using new data. An optional *BMI* (body mass index) calculator is provided.

Expressions

Enter the expression to be solved, e.g.:

1367+6755

Spaces are not permissible within expressions. Use of capitals is optional. Numbers can optionally be entered in scientific notation, e.g. as 1.3E6 instead of 1300000, or 1.3E-4 instead of 0.00013.

Two or more expressions may be entered, separated by spaces, e.g.:

sq(234.2) sq(638)

Optionally, a **label** may be attached to any expression, to store the value or formula. If

a=sq(234.2) b=sq(638) c=sqrt(a+b)

is entered, the three results are stored as a, b, and c, respectively, and the formula **sqrt(a + b)** is saved as **@c.** A list of stored values and formulae can be brought up by clicking on “Memory”.

The number of decimal places displayed can be changed.

Any number of nested parentheses may be used, for example.

sqrt(7x(43-(62/(7.4-sqr(1.44-0.5)))))

The values in the innermost parentheses are computed first, and multiplication and division are performed before addition and subtraction: 4+5*6-3*2^3/(8-1) is treated as 4+(5*6)-(3*(2^3)/7).

Symbols and functions

The following symbols may be used:

+	addition
-	subtraction, or a negative value
*** , x**	multiplication
/	division
^	exponentiation; 22^3 is 22 to the power of 3, or 22-cubed; 22^(1/3) is the cube root of 22; 22^(-1) is the reciprocal of 22, i.e. 1/22.
!	factorial; this must follow the value to which it refers

The following functions may be used; they must be followed by a value or expression in parentheses, e.g. sq(12.1) or sqrt(45-22).

sq, sqr	square
sqrt	square root
ln, log	natural log
exp	antilog of natural log
lg	log to base 10
alg	antilog (exponential of log to base 10)
abs	absolute value
arctan	arctangent
cos	cosine
sin	sine

Arctan, cos, and sin refer to radians (1^{o} = *pi*/180 radians). If other trigonometric functions are needed, the formulae are:

tan	= sin/cos
cot	= 1/tan
arcsin(x)	= arctan(x/(sqrt(1-sqr(x)))
sec	= 1/cos
cosec	= 1/sin
arccos(x)	= arctan(sqrt(1-sqr(x))/x)

Constants

Two constants may be used (in addition to those set by labelling values):

e	2.71828...; the base of natural logs
pi	3.14159...

Labels (for storage of values or formulae)

To store a value or formula in memory, all that need be done is to prefix a label. The label may be any letter except '**e**' or '**x**', and must be prefixed to the value or formula, with '=' and no spaces, e.g. (for a value) **b**=1.334 or (for a formula) **t**=ln(4.3)x4. The entered or computed value will then be "remembered" until it is erased or the label is re-allocated. The label can be used to represent the value in subsequent expressions entered on the same line or when WHATIS is used again, whether in the same computer session or a later one.

Constants, such as population denominators for the calculation of rates, can be labelled and stored for later use.

The use of labels avoids repeated entry of the same value. For example, suppose you wish to compute 95% and 99% confidence limits for a value of 3.468, with a standard error of 1.213, using the formulae

Lower 95% limit = 3.468 - (1.213 x 1.96)
Upper 95% limit = 3.468 + (1.213 x 1.96)
Lower 99% limit = 3.468 - (1.213 x 2.576)
Upper 99% limit = 3.468 + (1.213 x 2.576)

One way is to enter the four expressions in exactly the above format:

3.468-1.213x1.96 3.468+1.213x1.96 3.468-1.213x2.576 3.468+1.213x2.576

(Note that "3.468-1.213x1.96" is equivalent to "3.468-(1.213x1.96)", because multiplication is performed before addition.)

To avoid repeated entry of 3.468 and 1.213, labels (e.g. **a** and **b**) can be attached to these two values, for use in subsequent formulae:

a=3.468 b=1.213 a-(b*1.96) a+(b*1.96) a-(b*2.576) a+(b*2.576)

Labels can also be attached to formulae. The formula can then be recalled and recomputed on a subsequent occasion by putting @ in front of the label. In the above example, the four formulae might be labelled f, g, h, and i:

a=3.468 b=1.213 f=a-(b*1.96) g=a+(b*1.96) h=a-(b*2.576) i=a+(b*2.576)

The computed values are then stored as f, g, h and i. The formulae are stored as @f, @g, @h and @i. To compute the confidence intervals of 5.555 (standard error, 2.222) it is enough to enter these two new values of a and b, and invoke formulae f, g, h, and i:

a=5.555 b=2.222 @f @g @h @i

The use of labels also permits complicated computations to be done in stages - intermediate results can be labelled for use in subsequent stages of the calculation, performed on the same line or by running the calculator again. As a simple example of a 'chain' of formulae:

p=4 q=*p3 r=q/2 s=sq(r)/2**

The stored values will then be: **p**=4; **q**=12; **r**=6; **s**=18. In addition, **p***3 is saved as **@q**, **q**/2 as **@r**, and sq(**r**)/2 as **@s.**

When a formula is recomputed, so are all the variables specifically mentioned in it. After entry of the above chain, subsequent entry of

p=2 **@r**

will change **p** to 2 and recompute both **q** and **r** (**q** is recomputed because it is mentioned in formula @r), yielding **q**=6 and **r**=3. But entry of

p=6 @s

will not recompute **s** correctly, because although **r** is specifically mentioned in formula **@s**, **q** is not (and the previous value of **q** (i.e., 12) will therefore be used. For safety, it is important to specify all the formulae in a chain - i.e.

p=6 @q @r @s

This will change the value of **p** and correctly recompute **q**, **r**, and **s**.

A label prefixed by @ cannot be used in an expression.

***P*=8 @q/2** is wrong

***P*=8 @q q/2** is correct**.**

A **list of the stored values** and formulae can be brought up by clicking on "Memory".

Entry of **f= g=** will erase the specific values f and g.

Factorial function

The factorial *n*! is the number of possible arrangements of *n* items; e.g. if there are three items (**a**, **b** and **c**), 3! = 6 , i.e. a**bc, acb, bac, bca, cab** and **cba**.

The program uses Brenner's algorithm (Ball 1978: 215) to compute factorials for numbers up to 275 and Stirling's approximation (Rothman and Boice 1982: 26) for larger numbers. (We are grateful to Ray Simons for bringing Brenner's procedure to our notice).

The program can compute factorials for positive integers up to 1,754. Factorials are also displayed for fractional numbers; these may be helpful if *gamma* functions are required, since the factorial of any positive number x may be taken as the *gamma* function at point $(x + 1)$ (Hoel 1984: 88; Abramowitz and Stegun 1970: 255).

Permutations and combinations

The program can compute permutations and combinations if the total number of items in the set is up to 1,754.

The number of possible subsets of r items (ignoring their arrangement) drawn from a set of n items is comb(n,r); for example, if n = 3 and r = 2, there are comb(3,2) = 3 possible subsets (**a** and **b**, **a** and **c**, **b** and **c**); comb(n,r) is the binomial coefficient ('n over r' or 'n binomial r'). The number of possible arrangements of a sub-set of r items drawn from a set of n items is perm(n,r); for example, if n = 3 and r = 2, there are perm(3,2) = 6 possible arrangements **(ab, ac, ba, bc, ca** and **cb**).

The formulae are:

$$\mathbf{perm}(n,r) = n! / (n-r)!$$

$$\mathbf{comb}(n,r) = \mathbf{perm}(n,r) / r!$$

Body mass index (BMI)

Calculation of the BMI requires entry of the individual’s weight and height (in kilograms and centimetres, or in pounds and inches).

The formulae are

$$\textbf{BMI = Weight in kg / (height in cm / 100)}^2$$

$$\textbf{BMI = kg / (cm / 100)}^2 \textbf{ or lbs / inches}^2 \textbf{ x 703.06957964}$$

Commonly used BMI ranges are:

underweight: under 18.5,
normal weight: 18.5 to 24.9
overweight: 25 to 29.9
obese: 30 or more

The BMI may overestimate body fat in athletes and others who have a muscular build, and may underestimate body fat in older persons and others who have lost muscle.

P-VALUE

This module displays the *probability* (*P*, *p-value*) corresponding to a given value of *z* (the standard normal deviate), *t*, *chi*-square or *F*.

It provides the one- and two-tailed *P* corresponding to absolute values of *z* and *t*, and one-tailed *P* for values of *chi*-square and *F*. For the *z* and *t* distributions, the program provides three p-values: one-tailed (the computed value of *P*), two-tailed (obtained by doubling *P*, to a maximum of 1.0), and one-tailed (1 - *P*). The last value is of interest if the association shown by the data is opposite in direction to that specified in the alternative to the null hypothesis.

The program also computes *inverse probabilities*, i.e. the z, *t*, chi-square, or *F* value corresponding to a given p-value. To obtain the *Z* or *t* value for a one-tailed *P*, the p-value should be multiplied by two before it is entered.

The program also provides the *standard normal cumulative function* corresponding to a value of z.

The program can also calibrate a P value to compute the Bayesian *minimum posterior probability* of the null hypothesis.

Minimum posterior probability of the null hypothesis (Bayes)

A P value is "the probability, under the assumption of no effect (the null hypothesis), of obtaining a result equal to or more extreme than what was actually observed" (Held 2010); i.e., it is the probability of obtaining a difference (or a trend, or a departure from zero, or a departure from homogeneity, etc., depending on what effect was tested) that is equal to or more extreme than what was actually observed. Bayesian statisticians claim that a P value may be misleading because of its misinterpretation as the probability of the null hypothesis being true (Hubbard and Bayarri 2003), and its consequent use as the basis for rejection or non-rejection of the null hypothesis.

Instead, they suggest use of the *minimum posterior probability of the null hypothesis*, derived from the P value by use of a Bayes factor (Goodman 2001). This "calibration" of the P value requires entry of the probability (prior to the test) of the null hypothesis. The effect of varying the prior probability can be examined by repeating the procedure; this repetition is recommended in order to see the effect of the choice of a prior probability and to determine robustness or sensitivity to the choice of priors (Berger and Sellke 1987). The higher the prior probability of the null hypothesis (i.e., the lower the prior plausibility of an association), the higher will be the minimum posterior probability of the null hypothesis, and the less convincing will be the evidence for the association.

Reliance on the P value usually exaggerates the evidence against the null hypothesis (Berger and Sellke 1987, Hubbard and Lindsay 2008). Reliance on the minimum posterior probability of the null hypothesis rather than the P value provided by (for example) a statistical test comparing two proportions or rates may therefore be particularly helpful when (as is often the case) the minimum posterior probability exceeds the P value. If the minimum posterior probability of the null hypothesis is large, the null hypothesis will not be rejected. However, a small minimum posterior probability does not necessarily mean that the actual posterior probability of the null hypothesis is small (Berger and Sellke 1987).

The program computes the minimum posterior probability by the Sellke-Bayarri-Berger procedure (Sellke *et al.* 2001), which Held (2010), who provides a nomogram based on the procedure, calls "perhaps the simplest and most intuitive calibration". Its use is especially recommended if there is no explicit alternative to the null hypothesis (Sellke *et al.* 2001).

The program displays the minimal Bayes factor – i.e. the minimal ratio (under certain conditions) of the posterior (data-based) odds of the null hypothesis to the prior odds. The lower its value, the stronger is the evidence against the null hypothesis. The following guidelines (Jeffreys 1961) are often used:

<0.01: decisive support
0.032–0.010: very strong support
0.10–0.032: strong support
0.32–0.10: substantial support
1.00–0.32: not worth more than a bare mention
>1.00: less credible after than before the study

METHODS

The *normal and t distribution* functions respectively are derived from FORTRAN routines by Hill (1973) and Cooper (1968). The p-values coincide with standard table values closely, to within 0.00001 in general.

The ***c****hi-square distribution* function is based on formula 26.4.8 of Abramowitz and Stegum (1970). If there is a single degree of freedom, the normal distribution function is employed. If the degrees of freedom are greater than 60, the Wilson-Hilferty approximation is used (Abramowitz and Stegum 1970: formula 26.4.17). The p-values coincide with standard table values to within 0.0001 in general.

The *F distribution* function is derived from FORTRAN routines by Cran *et al*. (1977) and Majumder and Bhattacharjee (1973), and employs a function ('Algama' - the logarithm of the *gamma* function) derived from a FORTRAN routine by Pike and Hill (1966). The p-values coincide with standard table values to within 0.001 in general.

The ***inverse F distribution*** formulae are derived from 26.5.22 and 26.6.15 of Abramowitz and Stegum (1970). Since this inverse *F* distribution function is less accurate than the *F* distribution function, its accuracy is enhanced by adapting its results to those of the latter function. After initial estimation of *F* from *P*, the corresponding p-value is back-estimated from *F*, and the *F* value is increased or decreased until its corresponding P coincides with the entered p-value. The F values coincide closely with tabulated values. If the numerator degrees of freedom = 1, an accurate F value is calculated from the t distribution by the formula

$$F = (t[P/2,DF2])^2$$

where $DF2$ = denominator degrees of freedom (Diem 1970: 167).

The *inverse normal distribution* function is derived from a FORTRAN routine by Odeh and Evans (1974). The Z values approximate standard table values very closely, to within 0.00001.

The *inverse t distribution* formula is given in section 26.7.5 of Abramowitz and Stegum (1970). The t values coincide with standard tables to two decimal places for degrees of freedom greater than 2 if P = .05 or more, and for degrees of freedom greater than 7 if P = 0.0001. The precision is decreased with smaller p-values and increased with higher degrees of freedom. If there are 7 or fewer degrees of freedom, the procedure displays exact t values if P is 0.2, 0.1, 0.05, 0.02, 0.01, 0.002, or 0.00 based on a table obtained from http://www.zweigmedia.com/RealWorld/finitetopic1/t_table.html; (for intermediate P values, n approximate interpolation is used.

The **i***nverse chi-square distribution* formula is derived from a FORTRAN routine by Best and Roberts (1975) and employs a procedure to calculate the incomplete *gamma* integral as described by Bhattacharjee (1970). The chi-square values generally coincide with the table values to three decimal places.

The standard normal cumulative distribution functions use code published by Graeme West (2004)

Minimum posterior probability of the null hypothesis

The program uses the procedure suggested by Sellke *et al.* (2001), as summarized by Held (2010):

$$BF = -2.718.P.\ln(P) \text{ if } P < 1 / 2.718\text{; otherwise, } BF = 1.$$
$$MPP = 1 / \{1 + 1 / [(BF - Q) / (1 – Q)]\}$$

where P = P value
Q = prior probability of the null hypothesis
BF = minimum Bayes factor
MPP = minimum posterior probability of the null hypothesis

C.I. (CONFIDENCE INTERVALS)

This module estimates *confidence intervals* for a variety of statistics:

1) a proportion
2) a risk, or a measure with a number-of-individuals denominator
3) a rate with a person-time denominator.
4) a risk ratio (ratio of measures with number-of-individuals denominators)
5) a rate ratio (ratio of measures with person-time denominators)
6) a difference between proportions (independent data)
7) a difference between proportions (paired data)
8) a difference between rates (with person-time denominators)
9) an odds ratio (independent samples)
10) an odds ratio (paired samples)
11) a mean, standard deviation, or variance
12) a Poisson variate
13) a ratio of two Poisson variates
14) a statistic whose C.I. can be estimated directly from its S.E.
15) a statistic whose C.I. can be estimated from the S.E. of its log

It can also compute an approximate *confidence level* for values in the ranges at or above, or at or below, any chosen point; this point might be (say) the lowest rate ratio that a study was designed to detect (hence providing a substitute for power calculations after a study's completion), or the highest rate ratio regarded as trivial.

Confidence intervals

In many instances exact 90%, 95%, and 99% Fisher's and exact mid-P intervals are provided. Many statisticians recommend the use of exact mid-P intervals (Berry and Armitage 1995).

1. Proportion

In most instances exact Fisher's and mid-P confidence intervals are provided. If the denominator is over 30,000 or (if the numerator is zero) over 15,000, exact Fisher's and approximate mid-P intervals are computed. In some instances only Fisher's intervals are computed.

If a numerator of 1 is entered, a second set of confidence intervals is computed, appropriate for use if this is the first success (e.g. detection of a case) after a series of

2. A risk, or a measure with a number-of individuals denominator

This option is appropriate for a risk, prevalence, cumulative incidence, or any other measure with a number-of-individuals (not person-time) denominator.

In most instances exact Fisher's and mid-P confidence intervals are provided. If the denominator is over 30,000 or (if the numerator is zero) over 15,000, exact Fisher's and approximate mid-P intervals are computed. In some instances only Fisher's intervals are computed.

3. A rate with a person-time denominator

Exact confidence intervals are displayed if there are 20 or fewer events, and approximate intervals of there are 20 or more events. Cohen and Yang (1994) point out that, unlike the conservative Fisher's intervals, the narrower mid-P intervals do not guarantee the nominal confidence interval in all instances, but these authors suggest that the discrepancies are of little practical importance.

4. A risk ratio (ratio of measures with number-of-individuals denominators)

This option is appropriate for comparisons of risks, prevalences, cumulative incidences, or any other measures with number-of-individuals (not person-time) denominators.

5. A rate ratio (ratio of measures with person-time denominators)

Exact Fisher's and mid-P confidence intervals are provided.

6. A difference between proportions (independent data)

This option is appropriate for comparisons of risks, per valences, cumulative incidences, or any other measures with number-of-individuals (not person-time) denominators derived from independent (unpaired) samples.

Three sets of confidence intervals are computed. The first uses the traditional method, and is appropriate only if the samples are large. The other two (using the Wilson score

7. A difference between proportions (paired data)

This option is appropriate for comparisons of correlated proportions, i.e. comparisons based on paired data.

Confidence intervals are computed by two methods – the traditional large-sample procedure, and a procedure (based on Wilson's score intervals) that is appropriate for small samples also, and is recommended by Newcombe and Altman (2000).

To compare the proportions of “Yes” in two matched samples, the pairs must be tabulated, and the numbers of “Yes-Yes”, “Yes-No”, “No-Yes”, and “No-No” mpairs must be entered.

8. A difference between rates (with person-time denominators)

The standard error of the difference and the confidence intervals are reported.

9. An odds ratio (independent samples)

Exact Fisher’s and mid-P confidence intervals are computed.

Enter the numbers of “Yes” and “No” observations in each sample: for a comparison of cases and controls, enter the numbers who are exposed and unexposed to the factor under study; in a study in which exposed and unexposed samples are compared, enter the numbers with and without the outcome condition.

10. An odds ratio based on paired samples

Appropriate for an odds ratio based on (for example) a matched case-control study.

Exact Fisher’s and mid-P confidence intervals are computed.

Enter the numbers of pairs with discrepant findings, e.g. (in a case-control study) the numbers of “case exposed, control not exposed” and “case not exposed, control exposed” pairs.

11. A mean, standard deviation, or variance

Confidence intervals are computed for the mean, standard deviation, and variance of a distribution if the mean and sample size are entered, together with either the standard deviation or the standard error of the mean.

To obtain confidence intervals for a mean only, it is sufficient to enter it with its standard error.

To obtain confidence intervals for a variance only, only the variance and sample size are required.

12. A Poisson variate

Appropriate for occurrences assumed to follow a Poisson distribution, e.g. the number of new cases of a rare disease in a population.

Enter the number of “randomly occurring” events.

13. A ratio of two Poisson variates

Appropriate for a ratio of two numbers of occurrences that are assumed to follow a Poisson distribution.

14. A statistic whose C.I. can be estimated directly from its S.E.

Appropriate if an approximately normal distribution can be assumed.

15. A statistic whose C.I. can be estimated directly from the S.E. of its log

Appropriate if an approximately lognormal distribution can be assumed.

Confidence levels for values above/below a specific point

The program computes an approximate confidence level for measurements at or below, or at or above, a chosen specific point. The selected point is included in both ranges – at or above" and "at or below" – since the distribution is assumed to be continuous (Zar 1998: 74).The results are not probabilities Goodman 1994).

If the selected point is the measure (e.g. the odds, rate or risk ratio or difference) that a study was designed to detect (often referred to as "delta"), this procedure may be a useful substitute for power calculations after a study has been performed (Smith and Bates 1992, 1994), since retrospective power calculations have a questionable validity (Zumbo and Hubley 1998).

The selected point might also be the highest value of an odds, rate or risk ratio or difference, or other measure, that is regarded as trivial – a value referred to as "zeta" by Feinstein (1998) or any other value.

The computation assumes a normal distribution (e.g. for a difference between rates) or a lognormal distribution (e.g. for an odds, rate or risk ratio). It is prudent to regard the results as approximations, both because of this assumption and because the computation assumes that the confidence interval entered is symmetric around the point estimate and that its width is a know multiple of the standard error.

Enter a 90%, 95% or 99% confidence interval for the measure, or a point estimate and standard error. These values may be based on a single sample or on a set of strata or studies. Crude values or adjusted ones (controlling for suspected confounders) can be used.

Click on "Normal" if the statistic is a risk difference or rate difference, and "Lognormal" if it an odds, risk or rate ratio. The base of a rate (100, 100, etc, need not be entered.

METHODS

Confidence intervals

1,2. Proportions and rates with number-of-individual denominators

Exact Fisher and mid-P binomial intervals are computed by a procedure from XLIM (version SP2.5) by A. Ray Simons.

If the denominator is over 30,000 or (if the numerator is zero) over 15,000, Fisher's intervals are estimated by a method based on a relationship between the *F* and binomial distributions (Brownlee 1965). This provides estimates that are close enough to be regarded as exact. Zar's formulae 24.28 and 24.29 are used (Zar 1998: 526).

If the denominator is over 30,000 or (if the numerator is zero) over 15,000, approximate mid-P intervals are computed by Vollset's formulae (Vollset 1993). The formulae for proportion *x/N* are:

Lower limit for $x = (LF[x] + LF[x + 1]) / 2$

Upper limit for $x = (UF[x] + UF[x+ 1]) / 2$

where *LF* and *UF* are the lower and upper Fisher's limits.
Vollset found that this method, "proposed to provide an easily computed alternative to the mid-P interval, has a level of conservativeness in between the mid-P and uncorrected score method." For large denominators the intervals are almost equal to the true mid-P values.

For Vollset's procedure, $LF[x]$ and $UF[x]$ are computed by Zar's formulae 24.28 and 24.29 (Zar 1998: 526), and $LF[x + 1]$ and $UF[x + 1]$ either by Zar's formulae or by Pratt's approximation to the exact method (Blyth 1986). The Pratt method is suggested by Vollset, who refers to these approximate mid-P intervals as "mean Pratt" intervals. The program uses Zar's formulae for proportions with a numerator less than 50, rates with a base of 10 or 100 and a numerator less than 50, rates with a base of 1,000 and a numerator less than 100, rates with a base of 10,000 and a numerator less than 500, and rates with a base of 10,000 or more and a numerator less than 700. Pratt's faster method is used in other instances, when it provides identical results to Zar's method, at the level of precision with which the program displays results.

The confidence intervals that are appropriate for the first success after a series of failures are based on a geometric distribution, using formula 1 of George and Elston (1993).

3. Rates with person-time denominators

Confidence intervals are computed for the numerator, assuming that it has a Poisson distribution.
Exact Fisher's and mid-P confidence intervals are displayed if there are 40 or fewer (for Fisher's) or 20 or fewer (for mid-P) events, using tabulated values from Pearson and Hartley (1966) and Cohen and Yang (1994), for Fisher's and mid-P intervals respectively. In other instances. Approximate Fisher's and mid-P confidence intervals are computed, using formulae provided by Rothman and Boice (1982, o. 29; formulae 17 and 18).

4. Risk ratios (ratios of measures with number-of-individuals denominators)

Estimated confidence intervals are computed by the method described by Morris and Gardner (2000: p. 58), using a log transformation.

5. Rate ratios (ratios of measures with person-time denominators)

Exact confidence intervals are calculated by treating the ratio of the numerator of one rate to the sum of the numerators of both rates as a binomial parameter *P*, and determining its confidence intervals by the methods described above for proportions. To estimate a confidence interval for the ratio of rate R1 to rate

R2, the upper and lower confidence limits of P are then substituted for P in the formula

$$(P.D2) / [(1 - P).D1]$$

6. Differences between proportions (independent data)

Fleiss's large-sample procedure uses formula 2.14 of Fleiss (1981). It is based on the normal distribution. The Wilson score procedures (Wilson 1927) are described by Newcombe (1998a) as methods 10 (without continuity correction) and 11 (with continuity correction). Formulae provided by Newcombe and Altman (2000: pp 49-50) are used for method 10. For method 11, the program computes the upper and lower confidence limits for the two proportions by formulae 1.26 and 1.27 of Fleiss (1981: p. 14), and substitutes them for $L1$, $L2$, $U1$, and $U2$ in Newcombe's formulae for L and U (Newcombe 1998a). The computation of Fleiss's intervals is based on z (the standard normal deviate).

7. Differences between proportions (paired data)

The large-sample method is described by Fleiss (1981, formulae 8.14 and 8.15, p. 117). The method based on Wilson's score intervals is described by Newcombe and Altman (2000, pp 52-54). The two methods are described by Newcombe (1998b) as Methods 2 and 10.

8. Differences between rates (with person-time denominators)

Approximate confidence intervals are calculated by the formulae provided by Rothman and Greenland (1998. p. 239).

9. Odds ratios (independent samples)

Exact Fisher's and mid-P intervals are computed by an efficient algorithm for calculating the coefficients of the conditional distribution (Martin and Austin 1991), using code from David O. Martin's public-domain EXACTBB program.
The logit method is described by Morris and Gardner (2000, pp. 60-62).

10. Odds ratios based on paired samples

The numbers of pairs with discrepant findings, a and b, are treated as Poisson variates. Intervals for their ratio are estimated by regarding $a / (a + b)$ as a binomial parameter (Ederer and Mantel 1974; Armitage and Berry 2002: p. 157) and computing its confidence intervals. If $a + b$ does not exceed 50, exact mid-P and Fisher's intervals are computed for this proportion; otherwise Zar's formulae 24.28 and 24.29 are used (Zar 1998: p. 526), substituting a for X (the numerator of the proportion) and $a + b$ for N (the denominator). The required confidence limits are $L1 / (1 - L1)$ and $L2 / (1 - L2)$, where $L1$ and $L2$ are the confidence limits of $a / (a + b)$.

11. Means, standard deviations, variances

The confidence interval for a mean is computed by adding or subtracting t.SE, where t is the upper (*alpha*/2)th quantile of the t distribution with N-1 degrees of freedom. If the sample size (N) is not entered but the standard error of the mean is, the normal distribution is used.

The estimation of confidence intervals for the standard deviation SD and the variance SD^2 of a distribution is described by Zar (1998, pp 110-111).
The confidence limits for the variance are $A/X1$ and $A/X2$, and those for the standard deviation are the square roots of $A/X1$ and $A/X2$, where

$A = \sqrt{}SD(N - 1)$

$X1$, $X2$ = the chi-square values, at $N - 1$ degrees of freedom, corresponding to respective probabilities of 0.95 and 0.05 for the 90% interval, 0.975 and 0.025 for the 95% interval, and 0.995 and 0.005 for the 99% interval.

12. Poisson variates

Exact Fisher's and mid-P confidence intervals are displayed if there are 20 or fewer events, using tabulated values from Pearson and Hartley (1966) and Cohen and Yang (1994), for Fisher's and mid-P intervals respectively.
Approximate Fisher's and mid-P intervals are computed if there are more than 20 events, using formulae provided by Rothman and Boice (1982, p. 29, formulae 17 and 18).

13. Ratios of Poisson variates

Intervals for the ratio of Poisson variates, *a: b*, are estimated by regarding $a / (a + b)$ as a binomial parameter (Ederer and Mantel 1974, Armitage and Berry 2002, p. 157) and computing its confidence intervals. If $a + b$ does not exceed 50, exact Fisher's and mid-P intervals are computed for this proportion; otherwise Zar's formulae 24.28 and 24.29 are used (Zar 1998, p. 526), substituting a for X (the numerator of the proportion) and $a + b$ for N (the denominator). The required confidence limits are $L1 /(1 – L1)$ and $L2/(1 – L2)$. Where $L1$ and $L2$ are the confidence limits of $a/(a + b)$.

14. Statistic whose C.I. can be calculated from its S.E.

Confidence intervals are estimated by adding or subtracting z.SE to the statistic, where z is the standard normal deviate representing the upper (*alpha/*2)th quantile of the normal distribution. For a 95% interval (*alpha* = 0.05), for example, $z = 1.96$.

15. Statistic whose C.I. can be calculated from the S.E. of its log.

Confidence intervals are estimated by adding or subtracting z.SE to the log of the statistic, where z is the standard normal deviate representing the upper (*alpha/*2)th quantile of the normal distribution and SE is the standard error of the log of the statistic; and then taking antilogs. For a 95% interval (*alpha* = 0.05), for example, $z = 1.96$.

Confidence level for values above/below a specific point

The procedure is described by Smith and Bates (1992). The confidence level is the one-tailed P value associated with the standard normal deviate (z) calculated by subtracting the point estimate from the selected point, and dividing this difference by the standard error. If a confidence interval is entered, its mid-point (on a normal or lognormal scale, as appropriate) is used as the point estimate for this purpose; the standard error is derived from the confidence interval – the width of a 95% interval, for example, is taken to be (2 x 1.96) times the standard error.

TIMESPAN

This module calculates the elapsed time between two calendar dates. It may be used for computing ages, exposure periods, follow-up periods, survival periods, gestational ages, etc. Leap years and the variable number of days per month are taken into account..

"From" and "To" dates must be entered. "Today" can replace either date. If only the year is entered, the date allocated by the program is July 1st. If only the year and month are entered, the 15th day of the month is allocated. Use a "minus" sign for a B.C. date.

The module can also calculate the dates at a specified number of months, weeks, or days before and after a given date(entered as the "From" date).

Julian day numbers are displayed. These are serial numbers given to daya, starting with Jan 1st, 4713 B.C. Their computation takes account of changes made in the calendar system.

The module may also be used to determine the day of the week on a given date, by entering a "From" date and "0" in the "Compute period of" box.

METHOD

The program determines the difference between the Julian Day numbers (Press et al. 1989, p. 11) of the two dates.

Intervals expressed in years and months are rounded off downwards to two decimal places and one decimal place respectively; the mean lengths of a year and month are taken as 365.25 and 30.4375 days respectively. The interval in weeks is the interval in days, divided by 7 and rounded off downwards to one decimal place.

REFERENCES

Abramowitz M, Stegum E (eds.) (1970) *Handbook of Mathematical Functions*. Washington, DC: National Bureau of Statistics, Applied Mathematics Series 55.

Armitage P, Berry G (1994) *Statistical Methods in Medical Research*, 3rd edn. Oxford: Blackwell Scientific Publications.

Armitage P, Berry G, Matthews JNS (2002) *Statistical Methods in Medical Research*, 4th edn. Blackwell Science..

Berry G, Armitage P (1995) *Mid-P confidence intervals: A Brief Review.* The Statistician, 44: 417-423

Ball JA (1978) *Algorithms for RPN Calculators*. New York: John Wiley & Sons.

Berger JO, Sellke T (1987) *Testing a point null hypothesis: the irreconcilability of P values and evidence.* Journal of the American Statistical Association 82: 112-122.

Best DJ, Roberts DE (1975) Algorithm AS 91. *The percentage points of the chi-squared distribution.* Applied Statistics 24: 385-388.

Bhattacharjee GP (1970) Algorithm AS 32. *The incomplete gamma integral*. Applied Statistics 19: 285-287.

Blyth CR (1986) *Approximate binomial confidence limits*. Journal of the American Statistical Association 81:843-855.

Brownlee KA (1965) *Statistical Theory and Methodology in Science and Engineering*, 2nd edn. New York: John Wiley & Sons.

Cohen G, Yang S-Y (1994) *Mid-P confidence intervals for the Poisson expectation.* Statistics in Medicine 3: 2189-2203.

Cooper BE (1968) *Algorithm AS 3. The integral of Student's t- distribution*. Applied Statistics 17: 19-20.

Cran GW, Martin KJ, Thomas GE (1977) *Remark AS R19 and Algorithm AS 109. A remark on algorithms AS63: The incomplete beta integral, and AS 64: Inverse of the incomplete beta function ratio.* Applied Statistics 26:111-114.

Diem *K* (1970) *Documenta Geigy: Scientific Tables*, 7th edn. Basle: J.R.Geigy S.A.

Ederer F, Mantel M (1974) *Confidence limits on the ratio of two Poisson variables*. American Journal of Epidemiology 100:165-167.

Fleiss JL (1981) *Statistical methods for rates and proportions,* 2nd edn. New York: John Wiley & Sons.

George VT, Elston RC (1993) *Confidence limits based on the first occurrence of an event* Statistics in Medicine12: 85-690.

Goodman SN (1994) *Confidence limits vs power calculations*. Epidemiology 5: 266-268.

Goodman SN (2001) *Of P-values and Bayes: a modest proposal.* Epidemiology 12:295-297.

Haug E.G. (2006) The complete guide to option pricing formulas. McGraw-Hill.
Held L (2010) *A nomogram for P values*. BMC Medical Research Methodology 10: 21.

Hill ID (1973) Algorithm AS 66. *The normal integral*. Applied Statistics 22: 424-427.

Hoel PG (1984) *Introduction to Mathematical Statistics*, 5th edn. New York: John Wiley & Sons.

Hubbard R, Bayarri MJ (2003) *P values are not error probabilities*. Available on the Internet at http://www.uv.es/sestio/TechRep/tr14-03.pdf

Hubbard R, Lindsay RM (2008) *Why P values are not a useful measure of evidence in statistical significance testing*. Theory & Psychology 18: 69-88.

Jeffreys H (1961) Theory of probability, 3rd edition. New York: Oxford University Press.

Majumder ML, Bhattacharjee GP (1973) Algorithm AS 63. *The incomplete beta integral*. Applied Statistics 22: 409-411.

Martin D, Austin H (1991) *An efficient program for computing conditional maximum likelihood estimates and exact confidence limits for a common odds ratio*. Epidemiology 2:359-362.

Morris JA, Gardner MJ (2000) *Epidemiological studies*. In: Altman DG, Machin D, Bryant TN, Gardner MJ. Statistics with confidence, 2nd edn. BMJ Books.

Newcombe RG (1998a) *Interval estimation for the difference between independent proportions: comparison of eleven methods*. Statistics in Medicine, 17:873-890.

Newcombe RG (1998b) *Improved confidence intervals for the difference between binomial proportions based on paired data.* Statistics in Medicine 17: 2635-2650.

Newcombe RG. Altman DG (2000) Proportions and their differences. In: Altman DG, Machin D, Bryant TN, Gardner MJ. Statistics with confidence, 2nd edn. BMJ Books.

Odeh RE, Evans JO (1974) *Algorithm AS 70. The percentage points of the normal distribution*. Applied Statistics 23: 96-97.

Pearson ES, Hartley HO (1966) *Biometrika Tables for Statisticians*, Vol. 1, 3rd edn. Cambridge: Cambridge University Press.

Rothman KJ, Boice JD Jr (1982) *Epidemiological Analysis with a Programmable Calculator*. Boston: Epidemiology Resources Inc.

Rothman KJ, Greenland S (1998) *Modern epidemiology*, 2nd edn. Lippincott-Raven.

Sellke T, Bayarri MJ, Berger JO (2001) *Calibration of p values for testing precise null hypotheses.* The American Statistician 55: 62-71.

Smith AH, Bates MN (1992) *Confidence limit analyses should replace power calculations in the interpretation of epidemiological studies*. Epidemiology 3: 449-451.

Smith AH, Bates MN (1994) *Confidence limits vs power calculations*. Epidemiology 5:268-269.

Vollset SE (1993) *Confidence intervals for a binomial proportion*, Statistics in Medicine 12:809-824.

West G (2004) Better approximations to cumulative normal functions. Available on the Internet;

Wilson EB (1927). *Probable inference, the law of succession, and statistical inference*. Journal of the American Statistical Association 22: 209-212.

Zar JH (1998) *Biostatistical Analysis, 4*th edn. Prentice-Hall.

Zumbo BC, Hubley AM (1998) *A note on misconceptions concerning prospective and retrospective power*. The Statistician 47 (part 2); 385-388.