

XVI-4. RESULTS OF NEW TESTS ON STATISTICAL DISTRIBUTIONS, CONTINUOUS INVERSE

Section	Outline	Topic/Function	Page Start	Page End	Tables	Figures
16-4	2.4	Inverse	2	2		
16-4	2.4.1	BETAINV	2	7		42, 43, 44, 45, 46, 47
16-4	2.4.2	CHIINV	7	11		48, 49, 50, 51
16-4	2.4.3	FINV	11	16		52, 53, 54, 55, 56
16-4	2.4.4	GAMMAINV	16	20		57, 58, 59, 60
16-4	2.4.5	LOGINV	20	20		
16-4	2.4.6	NORMINV	20	20		
16-4	2.4.7	NORMSINV	21	25		61, 62, 63, 64
16-4	2.4.8	TINV	25	28		65, 66, 67, 68

2.4 CONTINUOUS INVERSE FUNCTIONS

In the tests on the inverse functions, LTP values refer to the probability values used as inputs to the functions.

2.4.1 BETA DISTRIBUTION, INVERSE: BETAINV

The function is =BETAINV(probability, alpha, beta, A, B)

In KBA 828299, Microsoft said, “The BETAINV(p, alpha, beta, A, B) function is the inverse function for BETADIST(x, alpha, beta, A, B). For any particular x, BETADIST(x, alpha, beta, A, B) returns the probability that a Beta-distributed random variable (with the parameters alpha, beta, A, and B) is less than or equal to x. In other words, BETADIST returns the cumulative probability that is associated with x. If A and B are removed, they are assumed to be 0 and 1, respectively.”


“The BETAINV(p, alpha, beta, A, B) function returns the value of x where BETADIST(x, alpha, beta, A, B) returns p. Therefore, BETAINV is evaluated by a search process that returns the appropriate value of x by evaluating BETADIST for various candidate values of x until it finds a value of x where BETADIST(x, alpha, beta, A, B) is "acceptably close" to p.”

EXCEL 2000

The CISE (CISE 27/99) report shows BETAINV only at one a and one b value (0.5 and 0.5), with x between 0.01 and 0.99. The performance outside of this range was not explored. It shows the characteristic decrease in accuracy as p approaches zero. Knüsel (1999) did not comment on this distribution.

Tests on BETAINV for Excel 2000 are shown in figure 16-42.

Figure 16-42: BETAINV, Excel 2000, Accuracy



The black triangles represent LRE points that come from large values of a and b, in the range of 1000 to 10000. The open squares all represent points where both alpha and beta did not exceed 10. All the points lie above the theoretical curve.


EXCEL 2003 and 2007

In KBA 828299, Microsoft says:

“No changes were made in Excel 2003 to BETADIST. The only change that affects BETAINV is that "acceptably close" is redefined in the search process to be much closer. The search now continues until the closest possible value of x is found (within the limits of finite precision arithmetic in Excel). The resulting x should have a BETADIST(x, alpha, beta, A, B) value that differs from p by about 10^{-15} . “

The workbook for the tests on BETAINV was saved and rerun under 2003. Figure 16-43 is the result. The open squares are the same open squares as shown in figure 16-42. The dark triangles in figure 16-43 are the new points calculated under Excel 2003. All the points are from random numbers with both alpha and beta values within the range from 0 to 10, and p in the range from 0 to 1.

Figure 16-43: BETAINV, Excel 2003 and 2007, Accuracy



It is obvious that there was no change, in spite of Microsoft's claim in KBA 828299 to have improved the inverse routine and applied it to the beta distribution. If there was an improvement in closure, then the points in figure 16-43 should be greater than the points in 16-42. It continues to show the basic 8 LRE accuracy of BETADIST.

A general bound on LRE values can be made, based on approximate minimum LRE values. Figures 16-44 and 16-45 show a general p accuracy bound. It shows the effect of the arbitrary truncation. The LTP values are from the function input p values. The LRE values are approximate minimum values from a set of 200 random alpha and beta values within the range from 0 to 1000. If the set size is increased from 200 to about 15,5000, then figure 16-45 represents the data. The figure 16-44 curve is somewhat below that shown in figure 16-43, roughly about 1 to 2 LRE units lower in the skirts.

Figure 16-44: BETAINV, Excel 2003 and 2007, General Minimum LRE Bounds out of a Random Set of 200



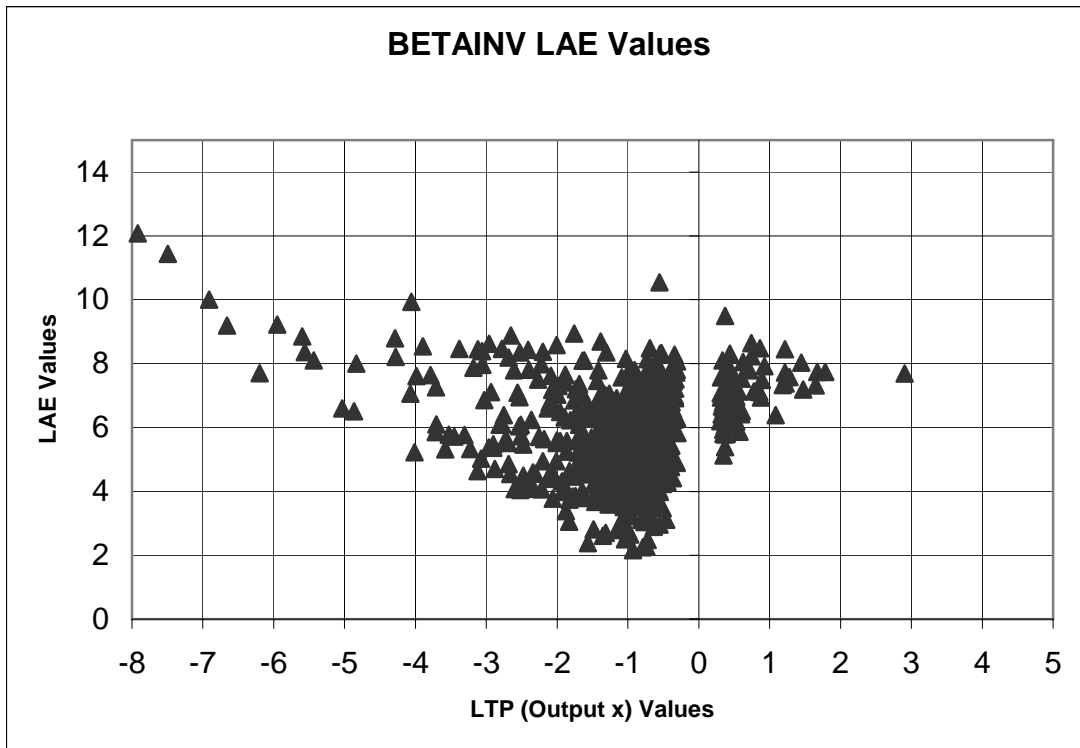


Figure 16-45: BETAINV, Excel 2003 and 2007, General Minimum LRE Bounds out of a Random 1Set of 15,500



The basic test criteria in all of the testing in this article is LRE values. If Microsoft were basing closure on an absolute deviation, then the chart would look different. Figure 16-46 is data generated from BETAINV with the LTP value derived from the function output, and the LAE (absolute) value again from the function output. The comparison again is with the reference function and with figure 16-43.


Figure 16-36: BETAINV, Excel 2003 and 2007, Accuracy, Different Plot Variables



This chart shows that the inverse function output values have increasing LAE values, as the size of the output p value gets smaller. Compare figure 16-46 with figures 9-3 and 9-4. The inverted bowl shape of the LAE characteristic (from figure 9-3) is evident in figure 16-46. This supports Microsoft's statement, "The resulting x should have a BETADIST(x, alpha, beta, A, B) value that differs from p by about 10^{-15} ". The important aspect in figure 16-46 is the large cluster of very low LAE values in the central region about 0.5.

The claim on the accuracy of BETAINV by Microsoft has to do with "a dog-chasing-his-tail" test. The claim is about the gap between the head and tail end, and is stated to be very small. The next chart is doing exactly what the claim says, generating BETAINV output values and putting it back into BETADIST, then comparing the two p values. The first p value is the input to BETAINV and the second is the output from BETADIST. This is all of course with the same alpha and beta values. Figure 16-47 shows the results, again based on LAE values.

Figure 16-47: BETAINV-BETADIST Loop, Excel 2003 and 2007, Accuracy



The chart shows four clusters, each with a different range of p values. The cluster on the left is with p values in the range of 0 to 0.000001. The p ranges have to be resized to get

points across the chart. The gap is just the fact that values from the 0.0001 and 0.00001 clusters were not generated.

The fact that the points all lie on a horizontal band, and they are far from 15, indicates that the BETADIST-BETAINV set does not even come close to Microsoft's claim. The central closure tolerance remains the same as it was under Excel 2000.

RELIABILITY ASSESSMENT OF BETAINV IN EXCEL 2000

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
None	None	None	None	None

RECOMMENDED BETAINV USAGE, EXCEL 2000

Range of p Values	Range of a Values	Range of b Values	Restrictions	ROUND level	Basis
0.00001 to 0.99999	0 to 10000	0 to 10000	Output >0	5	Fixed Point

RELIABILITY ASSESSMENT OF BETAINV IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
None	None	None	None	None

RECOMMENDED BETAINV USAGE, EXCEL 2003 AND 2007

Range of p Values	Range of Alpha Values	Range of Beta Values	Restrictions	ROUND level	Basis
0.001 to 0.999	0 to 10000	0 to 10000	Output >0	4	Fixed Point
0.00001 to 0.99999	0 to 10000	0 to 10000	Output >0	2	Fixed Point

2.4.2 CHI-SQUARE DISTRIBUTION, INVERSE: CHINV

The function is =CHIINV(probability, degrees_freedom)

In KBA 828313, Microsoft said, "CHIINV(p, df) is the inverse function for CHIDIST(x, df). For any particular x, CHIDIST(x, df) returns the probability that a Chi-Square-distributed random variable with df degrees of freedom is greater than or equal to x."

"The CHIINV(p, df) function returns the value x where CHIDIST(x, df) returns p. Therefore, CHIINV is evaluated by a search process that returns the appropriate value of x by evaluating CHIDIST for various candidate values of x until it finds a value of x where CHIDIST(x, df) is "acceptably close" to p. "

CHIINV is computed from GAMMAINV, from the p and q relationship these have to each other.

$$\text{CHIDIST}(x,v) + \text{GAMMADIST}(x, v/2, 2, \text{TRUE}) = 1$$


A complement of the input p value is used as the input to GAMMAINV, and the proper v value is returned. The complement limits the range of small p values that can be input into CHIINV. Figure 16-48 shows this limit.

EXCEL 2000

CHIINV is limited in that it will not return a chi square value if the q value is less than 1E-06 or the degrees of freedom are 10,000 or more. The help statement “If *degrees_freedom* <1 or *degrees_freedom* >= 10¹⁰, CHIINV returns the #NUM! Error value.” is misleading, since #NUM! is returned for calls to this function with df>1182 and p=0.1 or less. 1182 represents the limit in df values.

Knüsel (1999) has a general criticism that a displayed number of 9-10 digits here may only have two-digit accuracy. In his example, the LRE for df=10 and p=1E-06 is 2.68. The following figure shows the error for some other q values. The reference function used was comp_inv_chi_sq. Random p and v values were used to generate the points

Figure 16-48: CHINV, Excel 2000 Accuracy



The gap around zero is the normal LTP gap described under the definition of this value. CHIINV did output numbers for many combinations of large df values (1000 to 8451) and small p values ($\sim 1E-07$). The outline of the lower edge of the groups shows the truncating effect of the inverse algorithm, discussed above under the theoretical curve.


In general CHIINV has limited accuracy and this needs to be considered in using the function to determine the validity of hypotheses.

EXCEL 2003 and 2007

In KBA 828313, Microsoft said, “No changes were made in Excel 2003 to CHIDIST. The only change affecting CHIINV was to redefine "acceptably close" in the search process to be much closer. The search now continues until the closest possible value of x is found (in the limits of Excel's finite precision arithmetic). The resulting x should have a CHIDIST(x, df) value that differs from p by about $10^{(-15)}$.”

Figure 16-49 shows the improvement in LRE values around the central region.


Figure 16-49: CHIINV, Excel 2003 and 2007 Accuracy Within the Central region



In Excel 2003, there is a lower limit on the p value that can be input. The limit is $3E-07$. P values below this value, result in #NUM returns for all df values. When a p value of one is input, CHIINV will return values. Anything above one, returns #NUM.

Figure 16-50 shows minimum LRE values for random df values from 1 to 1000. LRE values for dfs greater than 100, generally show slightly higher LRE values. The figure comes from a set of 15,500 random “lines”.

Figure 16-50: CHIINV, Excel 2003 and 2007, Average LRE Values Over the Entire p Range



The downward trend on the right side, generally follows the relationship: $17 = LTP + LRE$. When $p=1$, the value returned is in error, since true values for the inverse function do not exist. The cut-off on the right is built into CHIINV. It will return a zero for any p value less than $1E-07$ on input.

Figure 16-51 shows a major reliability problem with CHIINV.

Figure 16-51: CHIINV, Excel 2003 and 2007, Reliability

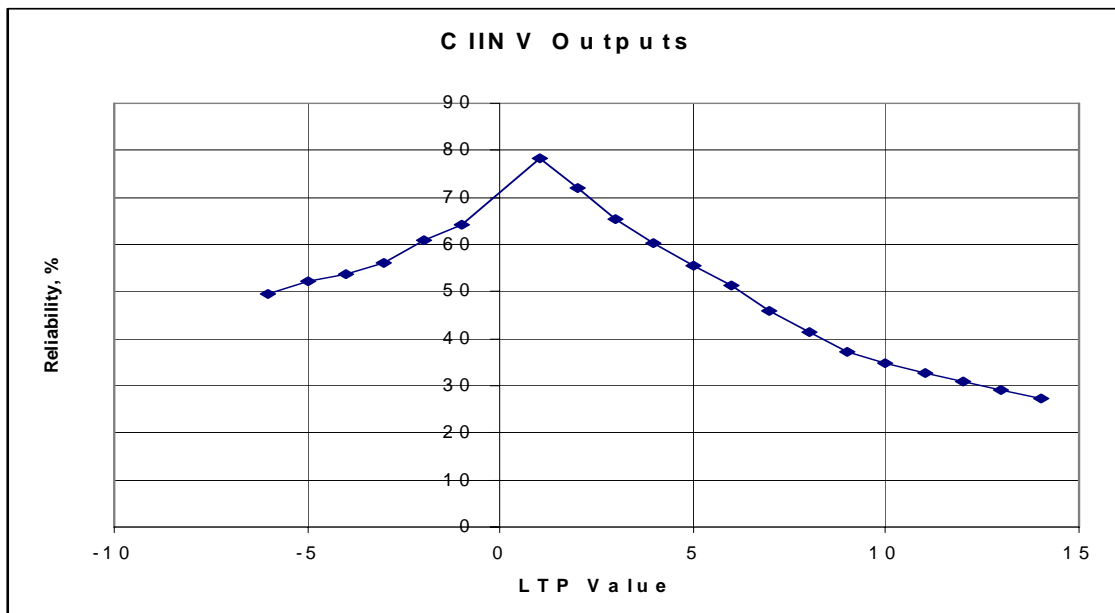


Figure 16-51 shows that the great majority of CHIINV return are #NUM. To have a chance of only 60% in getting an eta value from a CHIINV return is pretty bad. At p values approaching 1, the reliability drops to 25%. The problem is in the poor coding of the closure routine. There are regions where the number of allowable iterations are exceeded, giving #NUM returns. The region around 920 to 940 dfs, generally always returns #NUM.

The fault is that one cannot get a CHIINV eta value at all for the given input values.

RECOMMENDED CHIINV USAGE, EXCEL 2000

Range of p Values	Range of df Values	Restrictions	ROUND level	Basis
0.000001 to 0.999999	0 to 10000	Output > 0	7	Fixed Point

RELIABILITY ASSESSMENT OF CHIDIST IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
None	Excessive	None	None	None

RECOMMENDED CHIINV USAGE, EXCEL 2003 AND 2007

Range of p Values	Range of df Values	Restrictions	ROUND level	Basis
5E-07 to (1-1E-10)	0 to 800	Outputs > 0	6	Fixed Point

2.4.3 F DISTRIBUTION, INVERSE: FINV


IN KBA 828332, Microsoft said, “The FINV(p, df1, df2) function is the inverse function for FDIST(x, df1, df2). For any particular x, FDIST(x, df1, df2) returns the probability that an F-distributed random variable with df1 and df2 degrees of freedom is greater than or equal to x.

The FINV(p, df1, df2) function returns the value of x where FDIST(x, df1, df2) returns p. Therefore, FINV is evaluated by a search process that returns the appropriate value of x by evaluating FDIST for various candidate values of x until it finds a value of x where FDIST(x, df1, df2) is "acceptably close" to p. “

EXCEL 2000

The accuracy of the F distribution inverse function FINV is essentially dependent on the input p value. The following chart shows the effect.

Figure 16-52: FINV, Excel 2000 Accuracy



The vertical scale is in LRE units, and the horizontal scale is in LTP values. The best accuracy is about 6-7 at p values close to 0.5. At low p values the accuracy drops off considerably to the 1-2 range. This characteristic is dependent to some extent on the degrees of freedom values, with accuracy increasing as df values increase. The lower LRE values generally occur when at least one df is 10 or less.

In a general sense it is similar to the theoretical error chart of figure 9-1.

EXCEL 2003 and 2007

In KBA 828332, Microsoft said, "No changes were made in Excel 2003 to FDIST. However, the definition of "acceptably close" was changed in the search process so that it is much closer to the actual value. The search now continues until the closest possible value of x is found (within the limits of finite precision arithmetic in Excel). The resulting x should have an FDIST(x, df1, df2) value that differs from p by about $10^{(-15)}$."

The results of random variable tests are shown in figure 16-43. The degrees of freedom were varied from 1 to 1000. The ranges of the p values were changed to generate the groups shown in the figure.

Figure 16-53: FINV, Excel 2003 Accuracy




Figure 16-53 shows the improvement made in FINV for Excel 2003. There was no limit to the p values, and FINV resulted in #NUM free returns over the entire range of p values. There is a limit in the size of the allowable F value, set at 1E+09. Consequently when this is reached, 1E+09 is returned. This rarely occurs, and when it is observed, it should be taken as a wrong F value. In the data generated for figure 16-53, it occurred about 6 times out of 2844 points.

Figure 16-54 shows an approximate lower bound for minimum LRE values as p is varied. Again the descending line at the right basically represents the limitations of a 15 decimal number in representing p values close to 1.

Figure 16-54: FINV, Excel 2003 and 2007 Accuracy






Figure 16-55 shows that FINV can take a full range of p inputs. The degrees of freedom were randomly varied from 1 to 10,000. Values were not set in to show the region above p values of 0.5.

Figure 16-55: FINV, Excel 2003 and 2007 Full Range Input Accuracy



The scattered points below the 13 line are primarily from large df values. Figure 16-56 is the central region of figure 16-55.

Figure 16-56: FINV, Excel 2003 and 2007 Central Region of Figure 16-55



The lower points in figure 16-57 from 0 to -13 lie slightly above the boundary line of figure 16-55. This shows that there was consistency between the two methods to show FINV accuracy limits.

FINV will give good output values (LRE values to 13) with p input values down to 1E-300. There is a general increasing trend in the accuracy from LTP values of -15 to -300. That's pretty good for an inverse function.

RECOMMENDED FINV USAGE, EXCEL 2000

Range of p Values	Range of df1 Values	Range of df2 Values	Restrictions	ROUND level	Basis
0.0001 to 0.9999	0 to 10000	0 to 10000	P value range	6	Fixed Point

RELIABILITY ASSESSMENT OF FINV IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
None	None	None	None	None

RECOMMENDED FINV USAGE, EXCEL 2003 AND 2007

Range of p Values	Range of df1 Values	Range of df2 Values	Restrictions	ROUND level	Basis
1E-300 to 0.5	1 to 10000	1 to 10000	P value range	8	Floating Point
1E-300 to 1E-50	1 to 1000	1 to 1000	P value range	10	Floating Point

2.4.4 GAMMA DISTRIBUTION, INVERSE: GAMMAINV

In KBA 828334, Microsoft said, “The GAMMAINV(p, alpha, beta) function is the inverse function for the GAMMADIST(x, alpha, beta, TRUE) function. The last argument to GAMMADIST indicates whether the cumulative distribution function is wanted; for using GAMMADIST in the evaluation of GAMMAINV, the cumulative = TRUE case is wanted. For any particular x, GAMMADIST(x, alpha, beta, TRUE) returns the probability that a GAMMA-distributed random variable with parameters alpha and beta is less than or equal to x.”


“The GAMMAINV(p, alpha, beta) function returns the value x for which GAMMADIST(x, alpha, beta, TRUE) returns p. Therefore, GAMMAINV is evaluated by a search process that returns the appropriate value of x by evaluating GAMMADIST for various candidate values of x until it finds a value of x for which GAMMADIST(x, alpha, beta, TRUE) is "acceptably close" to p.”

“Special cases of the Gamma distribution include the Exponential distribution when alpha equals 1 and the Erlang distribution when alpha is a positive integer greater than 1. Additionally, when n is a positive integer, GAMMADIST(x, n/2, 2, TRUE) returns 1 – CHIDIST(x, n). Therefore the Chi-Square distributions (for any number of degrees of freedom), the Exponential distributions, and the Erlang distributions are all special cases of the family of Gamma distributions. “

EXCEL 2000

Figure 16-57 shows the results of tests on GAMMAINV

Figure 16-57: GAMMAINV, Excel 2000 Accuracy



The accuracy is poor.

EXCEL 2003 and 2007

In KBA 828334 Microsoft said, “No changes were made in Excel 2003 to GAMMADIST. The only change affecting GAMMAINV was to redefine "acceptably close" in the search process to be much closer. The search now continues until the closest possible value of x is found (within the limits of the finite precision arithmetic of Excel). The resulting x should have a GAMMADIST(x, alpha, beta, TRUE) value that differs from p by about 10^{-15} . “

Several tests of the new function were made using random value inputs. In figures 13-58 and 13-59, the alpha values ranged from 0 to 1000. In figure 16-59, Beta was varied from 0 to 10 to represent the three-parameter version, and in figure 16-59, beta was held at 1 to represent the two-parameter version. Each figure is based on 15,500 points

Figure 16-58: Two Parameter GAMMAINV, Excel 2003 and 2007 Accuracy, Central Region





Figure 16-59: Three Parameter GAMMAINV, Excel 2003 and 2007 Accuracy, Central Region




For the choice of parameters, there is no difference between the two figures.

Adding scatter plot points just results in a grouping above the lines.

The GAMMAINV – GAMMADIST test loop shows a small gap between the x value input and output. LRE values run from about 13 to 15.

Figure 16-60 shows the full range accuracy. The lower points that curve upward from 7 represent random alpha values from 0 to 1000. The somewhat horizontal line of points represents random alpha values from 0 to 10.

Figure 16-60: GAMMAINV, Excel 2003 and 2007 Full Range Accuracy



It is evident that the alpha value has a strong effect on the accuracy. Since in general applications, the typical alpha values used is unknown, the expected accuracy can range from 7 to 10, over the full range of p value inputs.

GAMMAINV will put out good values at low p values, down to 1E-300. At that level LRE values are about 10. This is a good robust inversion.

RECOMMENDED GAMMAINV USAGE, EXCEL 2000

Range of p Values	Range of alpha Values	Restrictions	ROUND level	Basis
0.00001 to 0.99999	0 to 1000	Output > 0.001	7	Fixed Point

RELIABILITY ASSESSMENT OF GAMMAINV IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
Rare	Rare	None	None	None

RECOMMENDED GAMMAINV USAGE. EXCEL 2003 AND 2007

Range of p Values	Range of alpha Values	Range of beta Values	Restrictions	ROUND level	Basis
1E-300 to 1E-3	1 to 1000	1 to 100	P value range	10 to 7, 7 is safe	Floating Point
0.001 to (1-1E-10)	1 to 1000	1 to 100	P value range	7	Floating Point

2.4.5 LOG NORMAL DISTRIBUTION, INVERSE: LOGINV

This function takes the input p value and uses NORMSINV to return a z value. LOGINV then computes an x value from the input mean and standard deviation values as described in Help.

The accuracy then depends only on the conversion of p values to z values, and this is discussed under the accuracy of NORMSINV above. There was no separate testing done on LOGINV. The expected number of accurate digits on the returned p value are shown at the end of the section on NORMSINV.

2.4.6 NORMAL DISTRIBUTION, INVERSE: NORMINV

This function takes the input p value and uses NORMSINV to return a z value. NORMINV then computes an x value from the input mean and standard deviation values as described in Help.

The accuracy then depends only on the conversion of p values to z values, and this is discussed under the accuracy of NORMSINV above. There was no separate testing done on NORMINV. The expected number of accurate digits on the returned p value is shown at the end of the section on NORMSINV.


2.4.7 STANDARD NORMAL DISTRIBUTION, INVERSE: NORMSINV

EXCEL 2000

NORMSINV was tested against Smith's `inv_normal` function. This reference function had tight closure and was able to give values identical to the Q values listed at the end of table 26.6 in Abramowitz (1964).


The errors in NORMSINV based on Smith's `inv_normal` function are shown in figure 16-61. The point groups were chosen to illustrate the errors in NORMSDIST, and do not show the complete region between -7 and $+7$.

Figure 16-61: NORMSINV, Excel 2000 Error in Output



The regular, accented points correspond to the theoretical limits discussed in section 9. The points below the theoretical limits line indicate the poor accuracy of NORMSDIST. The symmetry of the data indicates that a polynomial set covering sigma values from -8 to 0 is the computation, and the full range obtained by complementation. Figure 16-62 shows more detail over the 0.0001 to .9999 region.

Figure 16-62: NORMSINV, Excel 2000 Error in Output Over The 0.0001 To 0.9999 Region



The very marked drop in accuracy in the central region around 0.5 is due to the fact that the function is iterating at a wall (0.5), and consequently can't properly project the z value for the next iteration. This is a known problem for all gradient search methods when the minimum is exactly at a boundary.

On the whole, NORMSINV is a poor representation. It does not come anywhere near to a minimum requirement of LRE values to be generally greater than 6.

A major problem with NORMSINV is that it will not return credible sigma values for any p values less than 1E-06 or more than 1 - 1E-06 (returns +-5000000). Where the interest of a company is in 6 Sigma quality, it would be useful if Excel provided accurate probabilities at least up to the limit of floating point numbers. This would allow great flexibility in exploratory studies on data and the processes that produce the data.

Knüsel's (1999) criticism is on both the limiting p range and on the loss of accuracy.

If the criteria is at least an LRE value of 6 (minimum), then none of Excel's normal functions will meet the requirement. The user will have to obtain add-in functions. In other words, NORMSDIST and NORMSINV are failures.

EXCEL 2003 and 2007


In KBA 826772, Microsoft said, "The NORMSDIST function has been improved in Excel 2003. Improved refinements in the search process were introduced in Microsoft Excel 2002. An article by Knüsel (see note 2) discusses numeric deficiencies in the NORMSINV function in Microsoft Excel 97. These deficiencies persisted as documented by Knüsel until the improvements in the search process in Excel 2002 made results better, but still not in complete agreement with Knüsel's".

The accuracy of the function is improved from the Excel 2000 version, but shows the limitation of NORMSDIST on which it depends.

Figure 16-63 shows the general accuracy over the accept region on the returned z value.

The function will output a -30 value for all inputs less than 4.905E-198. This is a definite bound and error problem, since users will put in smaller values, and still get -30 returned. The function does not show any #NUM returns, except for p values of 1 or zero, or negative values. NORMSINV will return z values up to 1 but not 1. The maximum is an input of 0. (...15-9's) which returns 7.93473987.


Figure 16-63: NORMSINV, Excel 2003 and 2007 Accuracy Over The Full Range Of Inputs



There is a similarity to figures 13-37 and 13-38, the characteristics of the dependent NORMSDIST function.

Figure 16-64 shows an expanded central region around zero, of the points in figure 16-63.

Figure 16-64: NORMSINV, Excel 2003 and 2007 Accuracy Over The Central Region



The figure shows the basic scatter of NORMSDIST outputs that are shown in figures 13-37 and 13-38. It also shows the boundary problem shown in figure 16-38 and discussed below figure 16-38.

The accuracy of returned z values is complicated, and depends on the view taken and whether p values greater than 0.5 should be imputed, or their complement imputed. Based on a less than 0.5 view, and on the LRE minimum value, the returned z value is accurate to about 8 decimal digits (floating point), if the p value is not the result of a complementation. Based on a greater than 0.5 view, the number of accurate z digits (floating point) is equal to minus 0.867 times the LTP value of the complement plus 16. If the LTP value of the imputed p value is less than 2.5, then the returned z value is accurate to 15 decimal digits (fixed point).

The simplest way is to state that the z value is accurate to 8 floating point digits over the entire range. However it is clear that it is much more accurate then this when $abs(z)$ is less than 0.1, and when z is less than -7.

RECOMMENDED NORMSINV USAGE, EXCEL 2000

Range of p Values	Range of df1 Values	Range of df2 Values	Restrictions	ROUND level
1E-13 to (1-1E-09)	1 to 10000	1 to 10000	P value range	8

RELIABILITY ASSESSMENT OF NORMSINV EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
None	None	None	None	None

RECOMMENDED NORMSINV USAGE, EXCEL 2003 AND 2007

Range of p Values	Restrictions	ROUND level	Basis
1E-30 to {1-1E-10}	Only p values within this range	8	Floating Point
4.9E-198 to 1E-30	Only p values within this range	13	Floating Point

2.4.8 T DISTRIBUTION, INVERSE: TINV

In KBA 828340, Microsoft said, “TINV(p, df) is the inverse function for TDIST(x, df, 2). The last argument in TDIST is the appropriate number of tails in the student's (see note 1) t distribution; this is set to 2 throughout this article. For any particular positive value of x, TDIST(x, df, 2) returns the probability that a t-distributed random variable with df degrees of freedom is greater than or equal to x or is less than or equal to -x.”


”The TINV(p, df) function returns the value x for which TDIST(x, df, 2) returns p. Therefore, TINV is evaluated by a search process that returns the appropriate value of x by evaluating TDIST for various candidate values of x until it finds a value of x for which TDIST(x, df, 2) is "acceptably close" to p.”

TINV calls TDIST, and iterates on TDIST p values using varied t values until the difference in input and TDIST p values are less than a small value, epsilon. For Excel 2000, epsilon was 3E-07, and for Excel 2003 it is 2E-15.

EXCEL 2000

TINV has the same limitations that the other inverse functions have. P-values input less than 1E-06, or more than 1 – 1E-06, will not result in any valid outputs. For p values below 1E-06, the output is 5000000 and above 0.999999 it will be #VALUE or a multiple of 1.4E-07. Within this range, the accuracy is as shown in figure 16-65. The points may indicate unusually poor accuracy, but as shown, all points lie above the theoretical error line.

Figure 16-65: TINV, Excel 2000 Errors



The regular points are those of the theoretical curve discussed above. All test point lie either on this boundary or above.

Knüsel's (1999) comments are directed to the inaccuracies, even though up to 15 digits can be displayed (format as being independent of number of significant figures).

EXCEL 2003 and 2007

In KBA 828340, Microsoft said, "No changes were made in Excel 2003 to TDIST. The only change that affects TINV was to redefine "acceptably close" in the search process to be much closer. The search now continues until the closest possible value of x is found (within the limits of the finite precision arithmetic in Excel). The resulting x should have a TDIST(x, df, 2) value that differs from p by about 10^{-15} ."

KBA 827743 describes some problems with the new closure value in TINV returning #NUM! and VALUE! "To calculate inverse functions, TINV uses an iterative search technique. If the search does not converge after 100 iterations, the function returns the #N/A error value"

The new TINV will accept p values from $1E-306$ to 0.9999 without returning #NUM! or #VALUE!. In all the testing done on TINV, there was no #NUM! or #VALUE! returned.

The output of TINV has an upper limit of $1E+07$. All t values greater than this are returned as $1E+07$. Therefore the return of $1E+07$ should be viewed as an error, not as a value to be used.

The testing involved both the wide p value region ($p = 1E-306$ to 0.5) and the narrow region ($p = 0$ to 1). The uniform random variable distribution essentially does not allow for narrow region points with p values less than $1E-05$ from occurring. The wide region used random values of the 10s exponent from 0 to -306 to obtain small p values.

The reference function was `cdf_beta` since there was no other comparable reference function.

Figure 16-66: TINV, Excel 2003 and 2007 Accuracy, Full Range




Figure 16-67 shows the details around the zero LTP region.

Figure 16-67: TINV, Excel 2003 and 2007 Accuracy, Region Around $p = 0.5$

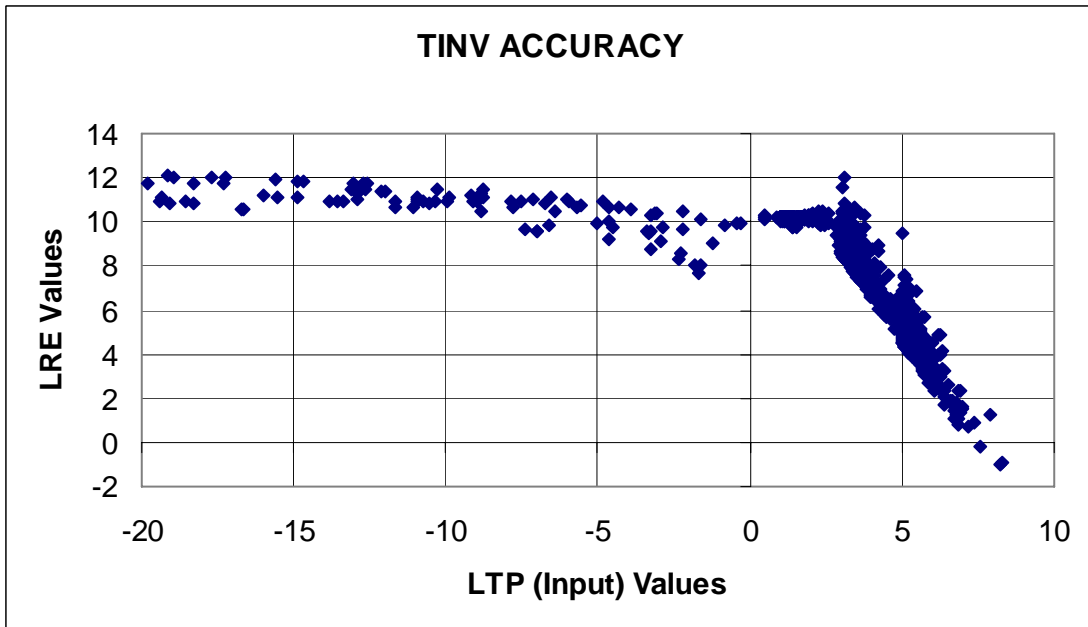



Figure 16-68 shows the same thing, except it shows a lower boundary of LRE values. Figure 16-68 is based on 15,500 points over the central region, where df values ranged from 1 to 1000.

Figure 16-68: TINV, Excel 2003 and 2007 Accuracy, Central Region,



The points show a very odd behavior to the right of the LTP zero line.

A boundary line on the sloping data points left of zero in figure 16-59 is approximately $14 = LRE + 2 * LTP$. The slope being larger than 1 indicates that the loss in accuracy is not due to complementation, but true errors in TINV.

RECOMMENDED TINV USAGE, EXCEL 2000

Range of p Values	Range of df Values	Restrictions	ROUND level	Basis
1E-06 to 0.999999	1 to 1000	P value range	6	Fixed Point

RELIABILITY ASSESSMENT OF TINV IN EXCEL 2003 AND 2007

False 1E+07 Values	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
A problem.	None	None	None	None

RECOMMENDED TINV USAGE, EXCEL 2003 AND 2007

Range of p Values	Range of t Values	Range of df Values	Restrictions	ROUND level	Basis
1E-300 to 0.999	0 to 1E+07	1 to 10000	All 1E+07 returns are in error,	6	Floating Point