XVI-1. RESULTS OF NEW TESTS ON STATISTICAL DISTRIBUTIONS, DISCRETES

Section	Outline	Topic/Function	Page Start	Page End	Tables	Figures
16-1	1	**Discrete**	2	2		
16-1	1.1	**Component Functions**	2	2		
16-1	1.1.1	COMBIN	2	3		1
16-1	1.1.2	FACT	3	4		
16-1	1.1.3	PERMUT	4	6		2
16-1	1.2	**Density**	6	6		
16-1	1.2.1	BINOMDIST	6	8		3, 4
16-1	1.2.2	HYPGEOMDIST	9	12	1,2	
16-1	1.2.3	NEGBINOMDIST	12	14	3	
16-1	1.2.4	POISSON	15	18		5, 6, 7
16-1	1.3	**Cumulative**	18	18		
16-1	1.3.1	BINOMDIST	18	30	4, 5, 6	8, 9, 10, 11, 12, 13, 14, 15
16-1	1.3.2	POISSON	30	35	7, 8	16, 17
16-1	1.4	**Inverse**	35	35		
16-1	1.4.1	CRITBINOM	35	38	9	

1.DISCRETE FUNCTIONS

1.1.COMPONENT FUNCTIONS

1.1.1 COMBINATIONS: COMBIN

The function is =COMBIN(number, number_chosen)

This function calculates the number of combinations that can be formed by a set group size (number-chosen) from a given population (number). It also calculates binomial coefficients. An integer number of combinations are returned.

There were no changes made to this function.

EXCEL 2000, 2003 AND 2007

The function calculates exact values within an allowable range. Figure 16-1 shows the allowable range. The figure just shows the lower boundary. The upper boundary is a mirror image of the lower, since COMBIN is symmetric. The starting point on the left is 1029, and 514.

Figure 16-1: COMBIN, Allowable Range of Input Parameter Values

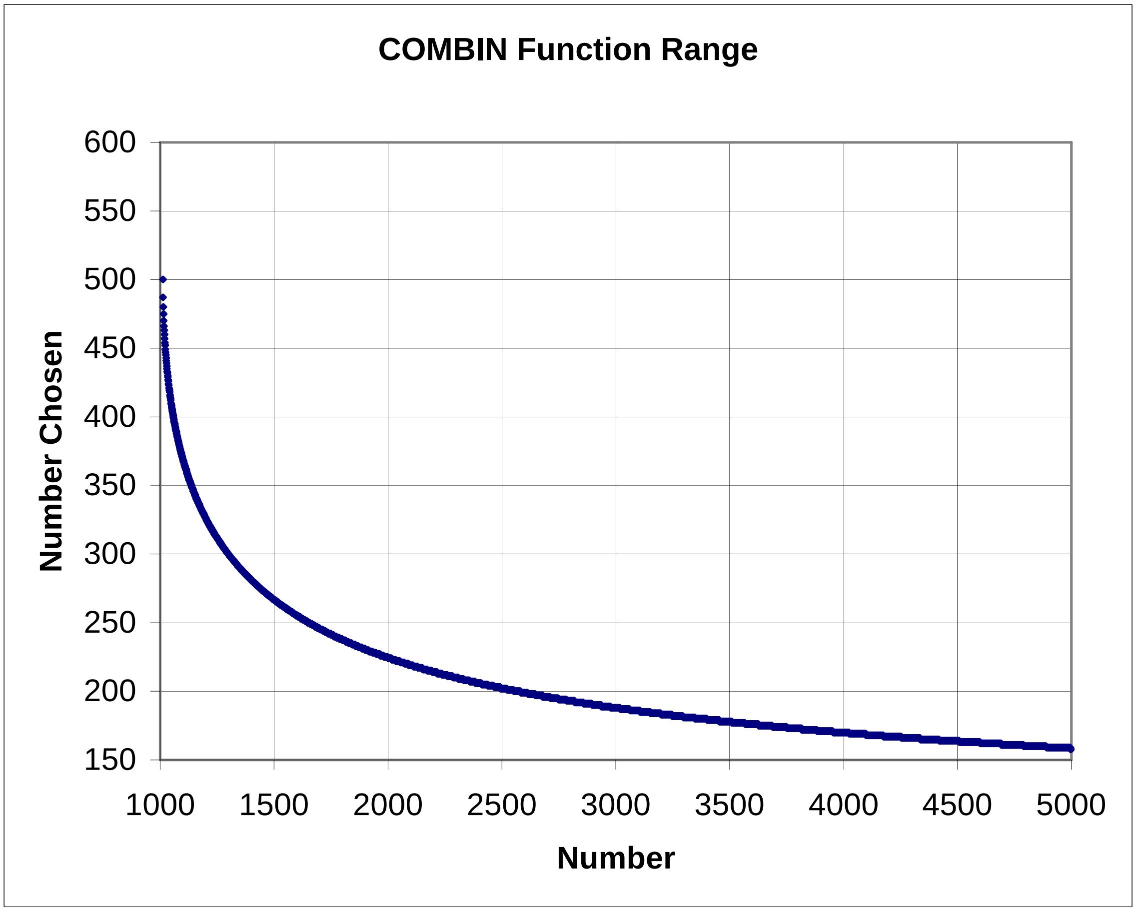

The region above the lower curve and the region below the upper curve (the mirror image of the lower curve) represents a region where the true value is larger than the capability of IEEE-754 double precision numbers. (See section 3 on Excel limitations) The region then represents true values larger than this. The allowable range of number_chosen values lie below and to the left of the curve. Input parameter pairs representing points above the curve will result in #NUM!-returns or other error codes. The point of symmetry, where all 'number-chosen' values result in numerical returns is at 'number' equal to 1029. At 'number' equal to 1030, any 'number_chosen' value from 500 to 531 results in #NUM!-returns.

COMBIN will return numbers up to a limit on the size of the input parameter 'number'. This limit on the curve of figure 16-1 is 'number' equal to 2.1475E+09, and 'number_chosen' equal to 37. The function will return #NUM! for any value of 'number_chosen' (even 1), when 'number' exceeds 2.1475E+09.

Returned combination numbers are fully accurate to 13 digits, floating point up to the limits.

RELIABILITY ASSESSMENT OF COMBIN IN EXCEL 2000, 2003 AND 2007

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

RECOMMENDED COMBIN USAGE IN EXCEL 2000, 2003 AND 2007

Range of number Values	Range of number_chosen	Restrictions	Round Level	Basis
1 to 1029	1 to number	None	14	Floating Point
1030 to 2.1475E+09.	Below the limit of 14-1, from 514 to 1	Size of number_chosen	>12	Floating Point.

1.1.2 FACTORIAL: FACT

The function is =FACT(number)

This function calculates the factorial product of a number. There were no changes made to this function for Excel 2000 or 2003.

EXCEL 2000, 2003 AND 2007

The function is a straightforward multiplication of terms, and is fully accurate to 13 digits. The function will not return any numerical results if the input parameter is larger than 170. For an input of 171, the value exceeds MAXFP.

RELIABILITY ASSESSMENT OF FACT IN EXCEL 2000, 2003 AND 2007

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

RECOMMENDED FACT USAGE IN EXCEL 2000, 2003 AND 2007

Range of number Values	Restrictions	Round Level	Basis
1 to 170	None	13	Floating Point

1.1.3 PERMUTATIONS: PERMUT

The function is =PERMUT(number, number_chosen)

This function calculates the number of permutations that can be made within a group of size “number_chosen” given a population size, “number”.

There were no changes made to the function for Excel 2000, 2003 or 2007.

EXCEL 2000, 2003 AND 2007

The function calculates exact values within an allowable range. Figure 16-2 shows the allowable range. The curve starts on the left at the point 170, 170, and decreases as shown. The boundary line from 0,0 to 170,170 is not shown.

Figure 16-2: PERMUT, Allowable Range of Input Parameter Values

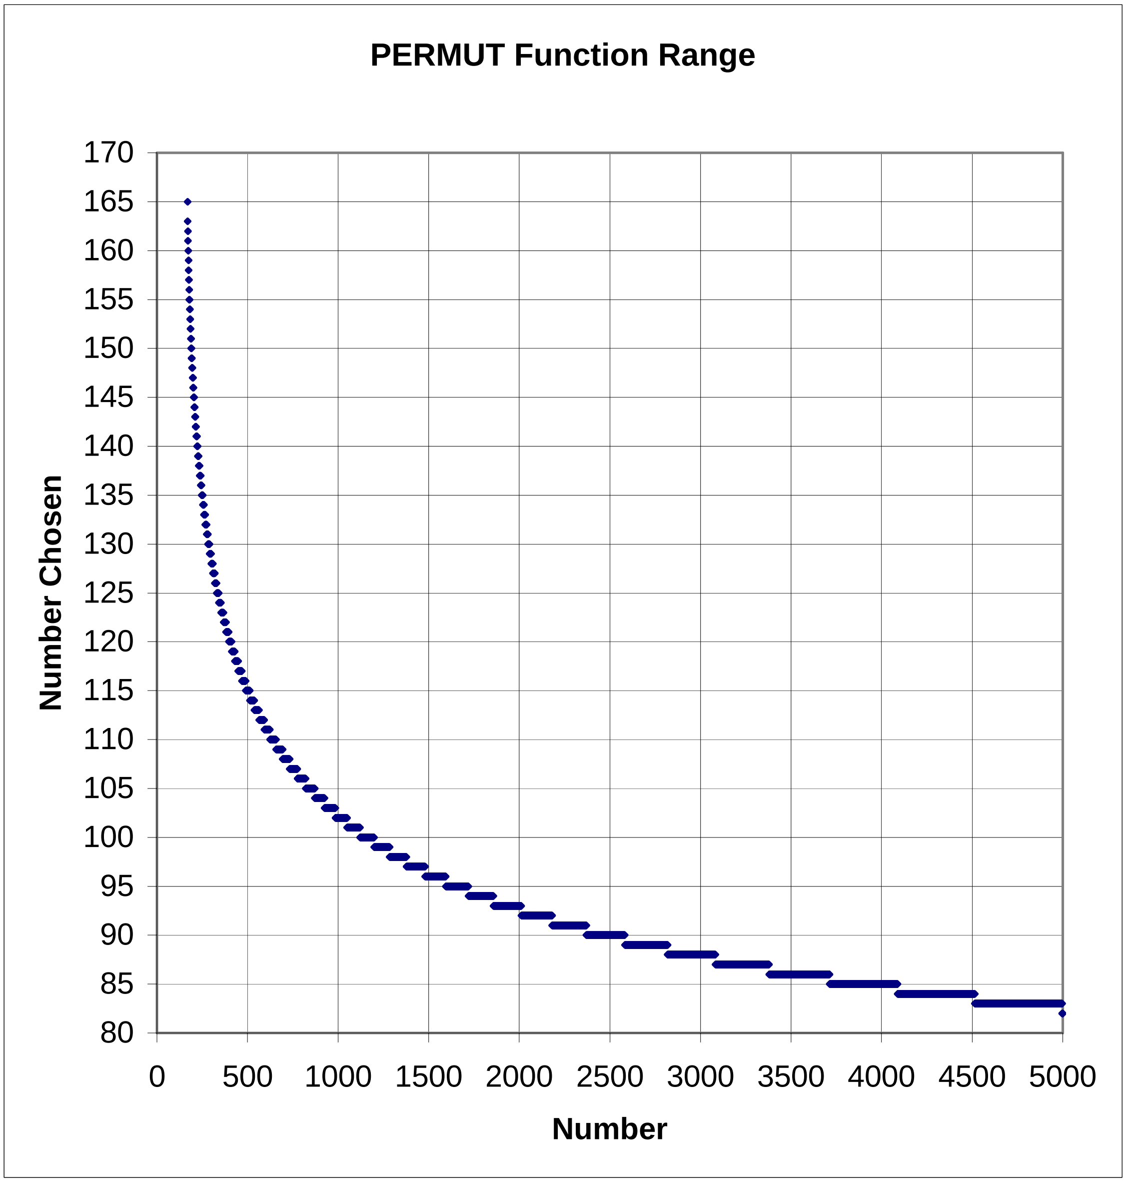

The region above the lower curve represents a region where the true value is larger than the capability of IEEE-754 double precision numbers (See section 3 on limits). Beyond this #NUM! is returned. The region above the curve then represents true values larger than rmax. If PERMUT is used in some intermediate equation where succeeding terms reduce the size to be less than rmin, then an alternate equation in terms of logs has to be used. An alternate equation using the GAMMALN and LN functions is recommended.

PERMUT will return numbers up to a ‘number’ limit of 2.1476E+09. The end point on the figure 16-2 curve is at “number’ equal to 2.1475E+09 and ‘number_chosen’ equal to 33. Any values above, below and to the right of this point will return #NUM!.

The PERMUT function is a direct multiplication and division operation and is fully accurate to 13 digits.

RELIABILITY ASSESSMENT OF PERMUT IN EXCEL 2000, 2003 AND 2007

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

RECOMMENDED PERMUT USAGE IN EXCEL 2000, 2003 AND 2007

Range of Number Values	Range of Number_chosen	Restrictions	Round Level	Basis
1 to170	1 to 170	None	14	Floating Point
171 to 2.1476E+09	33	Below the line of figure 16-2	Unknown	Floating Point

1.2 DENSITY FUNCTIONS

1.2.1 BINOMIAL DISTRIBUTION, DENSITY: BINOMDIST

The function is =BINOMDIST(number_s, trials, probability_s, FALSE)

The density function is shown in Help, as the terms that make up the cumulative sum.

EXCEL 2000

In KBA 827459, Microsoft said, “Knüsel (1998) documented instances where BINOMDIST does not return a numeric answer and yields #NUM! instead because of a numeric overflow. When BINOMDIST returns numeric answers, they are correct. BINOMDIST returns #NUM! only when the number of trials is greater than or equal to 1030. There are no computational problems if $n < 1030$. In practice, such high values of *n* are unlikely. With such a high number of independent trials, a user may want to approximate the Binomial distribution by a normal distribution (if $n*p$ and $n*(1-p)$ are sufficiently high, for example, each is greater than 30) or by a Poisson distribution otherwise.

For the non-cumulative case, BINOMDIST(*x*, *n*, *p*, false) uses the following formula

```
COMBIN(n,x)*(p^x)*((1-p)^(n-x))
```

In a set of 1000 calls to BINOMDIST with the density option, 746 were numbers the rest were #NUM. Out of the 746 numbers, 206 were zero and 58 were false zeros, that were true small tails. The remaining 540 numbers had LRE values of around 13. There was one low LRE value of 7.84.

EXCEL 2003 AND 2007

In KBA 827459, Microsoft said, “Because Microsoft has diagnosed when an overflow causes BINOMDIST to return #NUM! and knows that BINOMDIST is well behaved when overflow does not occur; Microsoft has implemented a conditional algorithm in Excel 2003. The algorithm uses BINOMDIST code from earlier versions of Excel (the computational formula mentioned earlier in this article) when $n < 1030$. When $n >= 1030$, Excel 2003 uses the alternative algorithm that is described later in this article. Typically, COMBIN overflows because it is astronomical, but p^x and (1-p)^(n-x) are each infinitesimal. If it were possible to multiply them together, the product would be a realistic probability between 0 and 1. But because existing finite arithmetic cannot multiply them, an alternative algorithm avoids the evaluation of COMBIN."

This code is discussed under the cumulative version of BINOMDIST.

Some tests were made on the accuracy of the probability of individual events. This is the BINOMDIST FALSE option. The return is a density value. Figures 16-3 and 16-4 show the results of some density test values. The low range again is from 0 to 1029 trials, and the high range is from 1030 to 10000 trials. The 1030 demarcation was set by Microsoft as the point where COMBIN starts returning #NUM!.

Figure 16-3: BINOMDIST-Density, Excel 2003 and 2007, Accuracy, Low Range

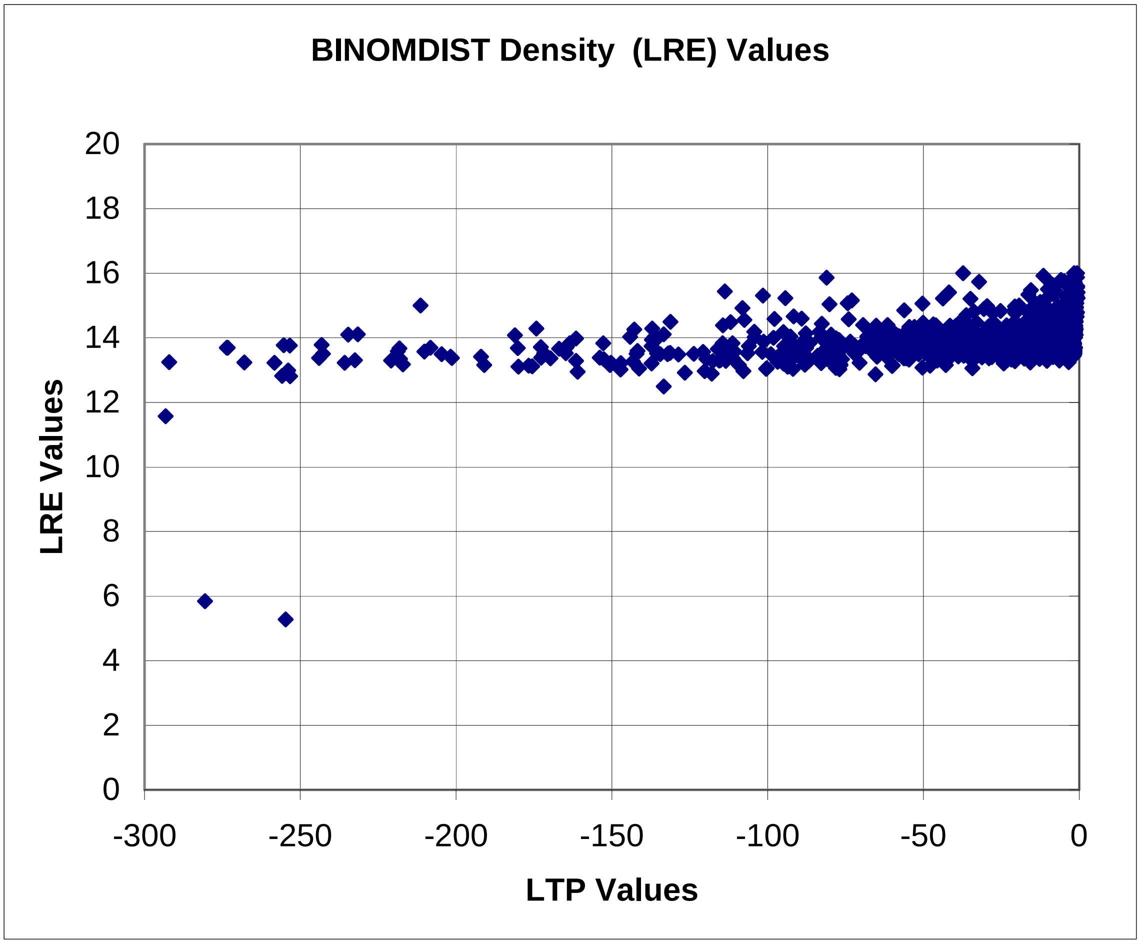

Figure 16-4: BINOMDIST-Density, Excel 2003 and 2007, Accuracy, High Range

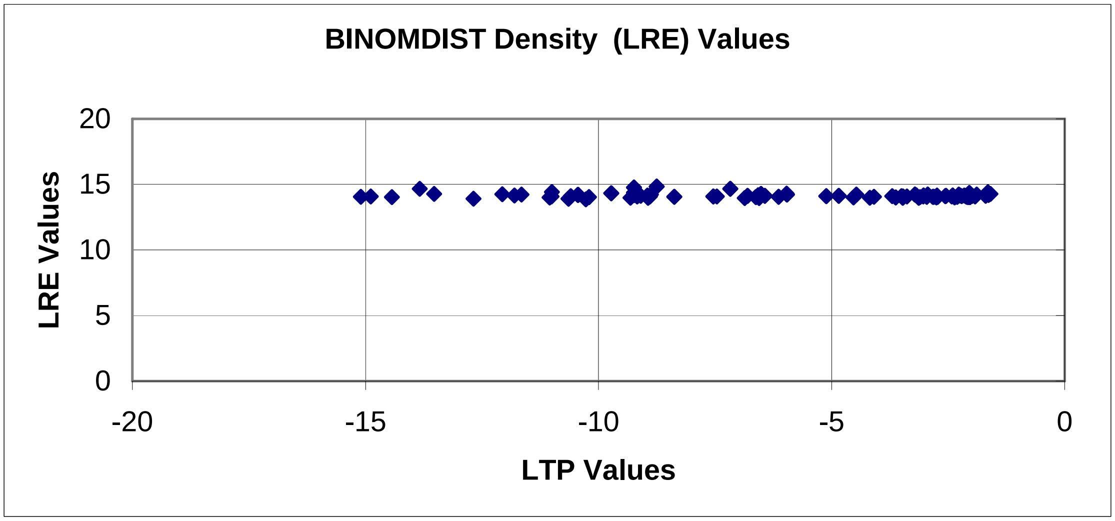

RELIABILITY ASSESSMENT OF BINOMDIST (DENSITY) IN EXCEL 2000

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

RECOMMENDED EXCEL 2000 BINOMDIST(DENSITY) USAGE

Range of x Values	Range of Number of Trials	Range of Number of Successes	Restrictions	Round Level	Basis
0 to 1	0 to 1030	0 to #Trials	No zero output values	13	Floating Point

RELIABILITY ASSESSMENT OF BINOMDIST DENSITY 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 BINOMDIST DENSITY USAGE FOR EXCEL 2003 AND 2007

Range of Probability of Single Event Success	Range of Number of Trials	Range of Number of Successes	Restrictions	Round Level	Basis	Percent of Output Values Below Round Accuracy
0 to 1	0 to 1030	0 to #Trials		13	Floating Point	0.1
0 to 1	1030 to 10000	0 to #Trials		13	Fixed Point	0

1.2.2 HYPERGEOMETRIC DISTRIBUTION, DENSITY: HYPGEOMDIST

The function is =HYPGEOMDIST(sample_s, number_sample, population_s, number_population)

This function is a discrete probability density function and is described in Help. In KBA 828515, Microsoft stated, "The HYPGEOMDIST Help file provides a formula to compute HYPGEOMDIST in Excel 2003 (where COMBIN(n, k) returns the number of combinations of size k in a population of size n):

```
COMBIN(population_s, sample_s) * COMBIN(number_population -
population_s, number_sample - sample_s) / COMBIN(number_population,
number_sample)
```

There is no cumulative version of HYPGEOMDIST."

As such it is very sensitive to floating point overload and underload. #NUM occurs frequently even for moderate problems. This was Knüsel's (1999) complaint. It is a reliability problem.

EXCEL 2000

Population sizes above 1030 will give false zeros. HYPGEOMDIST has the same problem that BINOMDIST had/has.

If the function were coded to use log gamma functions, then there would be no overload for any size problem. The accuracy of the log gamma method for typical problems gives LRE values of about 13-14, when compared to HYPGEOMDIST outputs. The loss of accuracy is primarily due to the effects of adding a*LOG(a) terms in the log gamma functions. Knüsel (1989) however points out that the logarithms introduce errors due to the fact that the result is the difference between large values of about the same magnitude, and actually may only end up with 6 accurate digits.

The reliability problem can be fixed by making the following changes.

Column A: Value for population size

Column B: Value for number of population successes

Column C: Sample size.

Column E: Number of successes in the sample.

Column E: =HYPGEOMDIST(Dn,Cn,Bn,An) "where n is the row number"

Column F: =IF(ISNUMBER(En),En,0) "if HYPGEOMDIST returns #NUM!, then a zero is put into the cell"

The probability value that is needed is in column F. The problem is underflow, and the correct answer is zero.

Tests with the reference function pdf_hypergeometric, all gave LRE values greater than 13.1 when the above fix was inserted and population size was limited.

Population sizes above 1030 will give many false zeros. HYPGEOMDIST will not return the correct small tail areas. The reliability fix has nothing to do with actual function returns of zero.

EXCEL 2003 AND 2007

KBA 828515 describes the changes made to HYPGEOMDIST. Microsoft said, “Because an overflow causes HYPGEOMDIST to return #NUM! and HYPGEOMDIST is well behaved when overflow does not occur; Microsoft has implemented a conditional algorithm in Excel 2003. The conditional algorithm uses HYPGEOMDIST code from earlier versions of Excel (the computational formula that involves COMBIN that is mentioned previously) when number_population is less than 1,030. In cases where number_population is greater than or equal to 1,030, Microsoft has implemented an alternative plan. Pseudo code is provided in the "Appendix" section in this article. The alternative plan for HYPGEOMDIST is in the same spirit as the alternative plan for BINOMDIST, CRITBINOM, and NEGBINOMDIST. The plan is used to avoid the evaluation of COMBIN with a first argument that is greater than or equal to 1,030.”

All the faults of BINOMDIST as a result of this change are inherited by HYPGEOMDIST. Table 16-1 is a summary of results from testing.

Table 16-1: Results of Random Number Input Tests, HYPGEOMDIST, Excel 2003 and 2007

SUMMARY			
Lower limit of Random Populations	1030	1	1
Upper Limit of Random Populations	10000	1030	2060
True Zero as Zero	1035	0	183
True Zero as #NUM!	77	796	577
Sum True Zeros	1112	796	760
#NUM!-returns, True Value >0	1996	3270	3093
False Zeros As Zero Below Threshold	819	0	261
False Zeros As Zero Above Threshold	563	0	20
Sum False Zeros	3378	3270	3374
Sum Total Zero Returns	2417	0	464
#NUM! Return Above Threshold	1919	2474	2516
#NUM Return Below Threshold	77	796	577
Sum #NUM!-returns	1996	3270	3093
NZ Values Below Threshold	82	533	446
NZ Values Above Threshold	505	1159	982
Sum NZ Returns	587	1692	1428
Matched Ones	0	38	15
False Ones (True NZ)	0	0	0
Sum One Returns	0	38	15
Total Returns	5000	5000	5000
Minimum LRE	12.99	12.73	12.99
Average LRE	14.99	14.99	14.97
Standard Deviation of LRE Values	0.54	0.67	0.66

	FLOATING POINT	FLOATING POINT	FLOATING POINT
Number of LRE Values Used	587	1692	1428
NZ Values Not Round Matched			
Round Set Number, N	12	12	12
Number NZ values Not Matched at N-1	0	0	0
Number NZ Values Not Matched at N	1	4	2
Number NZ Values Not Matched at N+1	3	24	22

The problem with the excessive #NUM!-returns still exists in Excel 2003 and 2007. 56% of all returns are #NUM!-returns. There are a lot of #NUM!-returns that are not true zeros. There are an excessive number of false zeros. 57% of all zero returns are false zeros. The only redeeming quality in HYPGEOMDIST is that when a NZ number appears, it is accurate to at least 11 digits, and the appearance of a one, is a true one. Both zero returns and #NUM!-returns cannot be assumed to be true zeros.

These proportions will change if you take a fixed-point view where everything below a p value of 1E-11 is zero and all NZ numbers are accurate. In this case, all zeros are true zeros, and all ones are true ones. #NUM!s are still indeterminate as to a definable value.

Table 16-2 shows the results of applying the Excel 2003 function to the Excel 2000 fault table from Knuesel (1997)

Table 16-2: Results of HYPGEOMDIST, Excel 2003 and 2007 on Knüsel's Data Set

Row Number	Reports	# White Balls In Sample	#White Balls in Population	# Black Balls in Population	Population Size	Sample Size
Symbol		k	N1, M	N2	N	n
1	EXCEL	225	515	515	1030	500
2	EXCEL	250	515	515	1030	500
3	EXCEL	275	515	515	1030	500

Table 16-2 Continued:

	Reported ELV	HYPGEOMDIST (EXCEL 2003 AND 2007)	pdf_hyp(Smith), Reference	LRE(on ELV))	LRE(on HYPGEOMDIST)
1	3.86527000E-04	3.86527295E-04	3.86527295E-04	6.12	14.42
2	4.97072000E-02	4.97071823E-02	4.97071823E-02	6.45	14.39
3	3.86527000E-04	3.86527295E-04	3.86527295E-04	6.12	14.42

Excel 2003 and 2007 corrects the failures previously reported.

RELIABILITY ASSESSMENT OF HYPGEOMDIST IN EXCEL 2000

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
High for populations over 1030	Frequent, all underflows. Can be corrected for.	None	None	None

RECOMMENDED HYPGEOMDIST USAGE IN EXCEL 2000

Range of Populati on Size	Range of Number of population successes	Range of sample size	Range of the number of successes in the sample	Restrictions	ROUND level	Basis
0 to 1030	1 to population size	0 to population size	0 to sample size	Use of the above reliability fix	13	Floating Point

RELIABILITY ASSESSMENT OF HYPGEOMDIST IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
High, up to 57%	Frequent, up to 56%.	None	None	None

RECOMMENDED HYPGEOMDIST USAGE IN EXCEL 2003 AND 2007

Range of Population Size	Range of Number of population successes	Range of sample size	Range of the number of successes in the sample	Restrictions	ROUND level	Basis
0 to 1E+300	1 to population size	0 to population size	0 to sample size	Only valid for NZ returns	11	Floating Point

1.2.3 NEGATIVE BINOMIAL DISTRIBUTION, DENSITY: NEGBINOMDIST

The function is =NEGBINOMDIST(number_f, number_s, probability_s)

This is only a density function. It returns the single probability of a specified sample result, given the true probability of drawing a single item success out of the population. Microsoft describes it as, “The NEGBINOMDIST(number_f, number_s, p) function returns the probability of exactly number_f failures before the number of successes reaches number_s in independent Bernoulli trials, each of which has an associated probability p of success and probability 1-p of failure."

This distribution is a discrete distribution similar to the binomial distribution. Because of the factorials, there will be overloads even for moderately sized problems.

EXCEL 2000

In accuracy tests based on pdf_negbonomial, the reference function, NEGBINOMDIST consistently gave LRE values above 12.6. The function will return #NUM for large numbers of failures (over several hundred). Since it has three input parameters, there was no obvious boundary where #NUM was consistently outputted. Input numbers of failures up to 9000 sometimes returned numerical values. Whenever there was a numerical output, the accuracy was in excess of 11.

EXCEL 2003 AND 2007

KBA 828361 describes the changes made to NEGBINOMDIST for Excel 2003. "Microsoft has implemented a conditional algorithm in Excel 2003 because of the overflow issue and because NEGBINOMDIST is well-behaved when the overflow does not occur. The conditional algorithm uses NEGBINOMDIST code from earlier versions of Excel (the computational formula involving COMBIN) when number_f + number_s – 1 < 1030. When number_f + number_s – 1 >= 1030, Microsoft implemented an alternative plan to use the formula that calls BINOMDIST."

Accuracy tests were made on this function based on random number inputs. Values of the probability of success varied from 0 to 1, the number of successes varied from 0 to 500 for the low range and 5000 for the high range, and the number of failures varied from 0 to 500 for the low range and to 5000 for the high range. The low range was set so that the sample size was always less than 1030. The worksheet contained 100 rows. The reference function was pdf_negbinomial.

During the testing (>150 worksheet regenerations) there were no #NUM returns, indicating that the overload problem was fixed. The underload problem was not fixed. When a number was returned, its accuracy as an LRE value was consistently above 12.6.

The split between the sample sizes below 1030 and above 1030 occurs also in NEGBINOMDIST. False zeros represent small tail values that are actually greater than rmin.

Table 16-3: NEGBINOMDIST- Density, Excel 2003 and 2007 Outputs, Random Number Inputs

Function Returns	Low Range	High Range
True zeros as zero	477	2941
True zeros as #NUM!	33	0
False zeros as zero, below cutoff	605	1601
False zeros as #NUM!, below cutoff	0	0
False zeros as zero above cutoff	5	14
False zeros as #NUM!, above cutoff	0	0
Correct 0-1 values below cutoff	2220	0
Correct 0-1 values above cutoff	1660	444
True ones	0	0
Total	5000	5000
Minimum LRE	12.5	12.8

False zeros below cutoff again include all p values calculated as non-zero values that are changed internally to zero on function output

RECOMMENDED NEGBINOMDIST USAGE IN EXCEL 2000

Range of Number of Failures	Range of Number of Successes	Range of the probability of population successes	Restrictions	ROUND level	Basis
0 – 9000	0 - 9000	0 to 1	Non-zero, numerical outputs only	11	Floating Point

RELIABILITY ASSESSMENT OF NEGBINOMDIST IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
Yes, high up to 32%	Small, 0.7%	No	No	No

RECOMMENDED NEGBINOMDIST USAGE IN EXCEL 2003 AND 2007

Range of Number of Failures	Range of Number of Successes	Range of the probability of population successes	Restrictions	ROUND level	Basis
0-10000	0-10000	0 to 1	No 0 p values	12	Floating Point

1.2.4 POISSON DISTRIBUTION, DENSITY: POISSON

The function is =POISSON(x, mean, FALSE)

Microsoft said, "When cumulative = TRUE, the function POISSON(*x*, *mu*, cumulative) returns the probability that a POISSON random variable with mean *mu* takes on a value less than or equal to *x*. When cumulative = FALSE, POISSON returns the probability that such a random variable takes on a value exactly equal to *x*. The POISSON distribution is frequently used to model the number of occurrences of certain events such as the number of customers who arrive in a queuing facility or the number of proofreading errors in an article. Because the POISSON distribution is used to count in this manner, *x* must be a non-negative integer. "

EXCEL 2000

In KBA 828130, Microsoft states, "For the non-cumulative case POISSON(*x*, *mu*, false) uses the following formula:

```
EXP(-x)*(mu^x)/FACT(x)
```

Overflow occurs when *mu*^*x* is too large. This does not occur if *mu*^*x* < 10^290 (or equivalently *x**LOG10(*mu*) < 290). FACT(*x*) also must not overflow. *x* <= 170 guarantees this. However, earlier versions of Excel do not look for these conditions."

The maximum X value in FACT is 170. Given x=170, the maximum allowable value for mu is 65. Values beyond these result in error code returns.

EXCEL 2003 AND 2007

In KBA 828130, Microsoft stated," Because Microsoft has diagnosed when overflow causes POISSON to return #NUM! and knows that POISSON is well behaved when overflow does not occur, Microsoft has added a conditional algorithm in Excel 2003. The algorithm uses POISSON code from earlier versions of Excel (the computational formula mentioned earlier in this article) when *x**LOG10(*mu*) < 290 and *x* <= 170. When *x**LOG10(*mu*) >= 290 or *x* > 170, Excel implements an alternative plan described later in this article. The alternative plan calculates an unscaled sum of probabilities of each possible observed value. This unscaled sum of probabilities is used later for scaling purposes. The algorithm also calculates an unscaled value of the probability that you want POISSON to return. Finally, it uses the scaling factor to return a correct POISSON value. The algorithm takes advantage of the fact that the ratio of successive terms of the form EXP(-*x*)*(*mu*^*x*)/*x*! has a simple form. The algorithm works as detailed in the pseudo code that is in the following steps. This approach is similar to the method used for the BINOMDIST, CRITBINOM, HYPGEOMDIST, and NEGBINOMDIST functions."

Figure 16-5 shows some LRE values for the new POISSON density function. The data is from random x and mean values, with means from 0 to 100, and x values from 1 to 1000.

Figure 16-5: POISSON, Excel 2003 and 2007 Density Accuracy

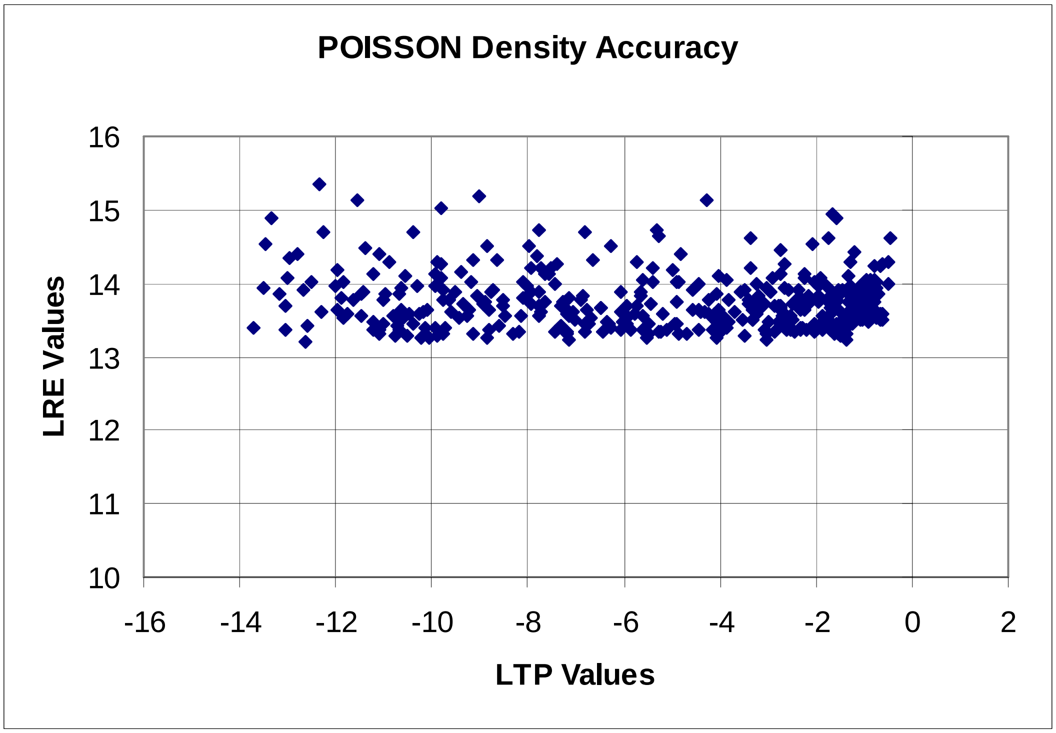

The points show the truncation that is built into the function. All density values less then about 1E-15 are returned as zeros. There were no small tail density values returned. Overall accuracy is good.

When the parameter range of the mean and of x is increased to 10,000, the results as shown in figure 16-6, are entirely different.

Figure 16-6: POISSON, Excel 2003 and 2007 Density Accuracy, Expanded Parameter Range

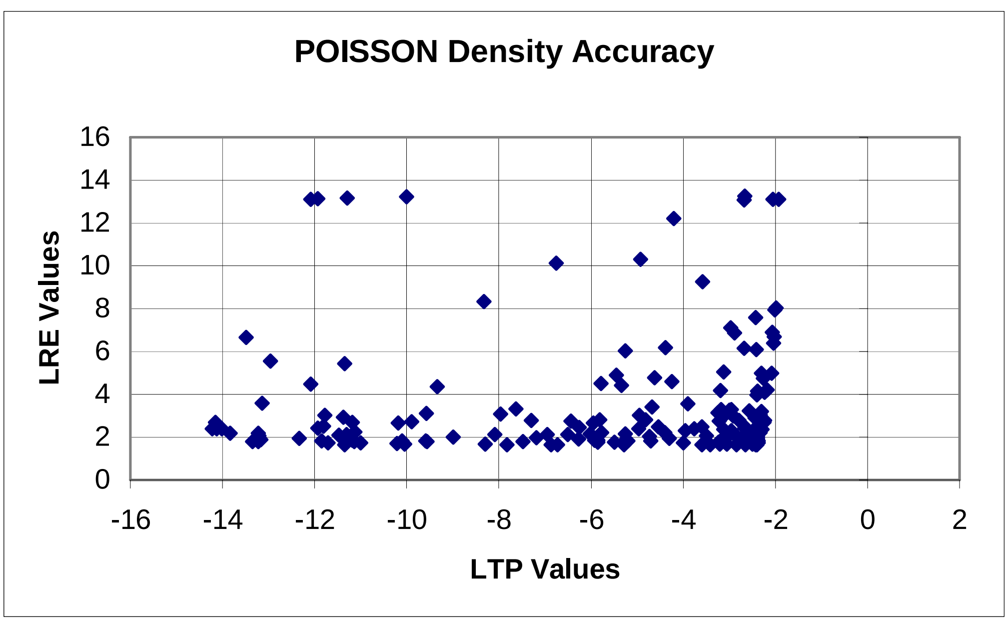

If the error measure is changed from the relative LRE to the absolute LAE, then the errors have a different appearance. Figure 16-7 shows some LAE values for the same extended range parameter inputs.

Figure 16-7: POISSON, Excel 2003 and 2007 Density Accuracy, Expanded Parameter Range

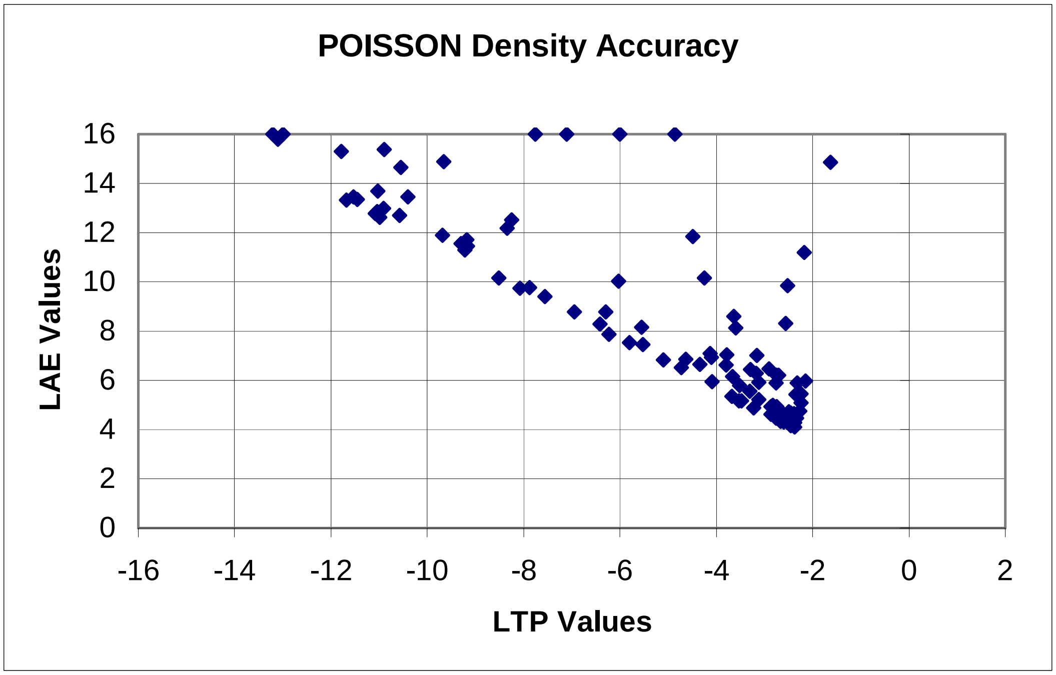

Figure 16-7 shows that for a density value in fixed point, the two digits after the leading zeros represent accurate values. Digits beyond these two are inaccurate. Both figures 13-6 and 13-7 show that the accuracy here is relative. The LRE measure indicates about 1.5 digits are accurate, while the LAE measure indicates that about 2.5 digits are accurate

Poisson function densities should not be obtained from POISSON if the mean is greater than 1000 or the x value is greater than 1000. The errors are just too high.

RECOMMENDED POISSON (DENSITY) USAGE IN EXCEL 2000

Range of x Values	Range of test parameter	Restrictions	ROUND level	Basis
0 to 170	X * LOG10(mean) < 290	Range of Input Parameters	13	Floating Point

RELIABILITY ASSESSMENT OF POISSON (DENSITY) IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
Treat as values <1E-15		No	No	No

RECOMMENDED POISSON (DENSITY) USAGE IN EXCEL 2003 AND 2007

Range of x Values	Range of mean values	Restrictions	ROUND level	Basis
1 to 1000	0 to 100	Values >1E-15	13	Floating Point
1 to 10,000	1 to 10,000	Values >1E-15	2	Floating Point

1.3 DISCRETE CUMULATIVE FUNCTIONS

1.3.1 BINOMIAL DISTRIBUTION, CUMULATIVE: BINOMDIST

This function is =BINOMDIST(number_s, trials, probability_s, TRUE)

This function is calculated as described under help. For a limited number of N values, the function is directly calculated from direct multiplications and divisions. When this limit is reached, the Excel 2000 version returns #NUM! The Excel 2003 and 2007 function added a different algorithm to obtain p values for inputs greater than the limit.

EXCEL 2000

Depending on the actual values being entered, BINOMDIST may overload (the response is #NUM). This is directly the result of COMBIN input values being above the curve shown in figure 16-1 (or below its mirror image). Microsoft stated in KBA 827459 that overload would occur when the number of trials is over 1030. Knüsel (1999) reported an overload with 1030 trials and p=0.5. The value 1030 lies at the extreme left of the allowable region (figure 16-1) at 515 successes. BINOMDIST will output values within the range of double precision numbers way beyond 5000 trials, providing the number of success are less than the figure 16-1 boundary.

The cumulative BINOMDIST was tested by comparison to Smith's function using random number generations of sample sizes (0 to 1030), number of successes and binomial distribution p values. In general, LRE values were 13 and greater, but occasionally an LRE value of 11 occurred. There is a region of low tail probabilities, where BINOMDIST will output zero, when the true value is greater than zero.

Another test was done using random values of "Probability of Success" from 0 to 1, random "Number of trials" from 1 to 2,000, and random "Number of Successes" from the # of trials to 1. This is looking at extreme values. The cumulative reference function was cdf_binomial. There were 1000 lines (calls to BINOMDIST cumulative and density) generated. Figure 16-8 represents the values. True zeros, true ones and #NUM returns were not included.

Figure 16-8: Accuracy of BINOMDIST-Cumulative, Excel 2000

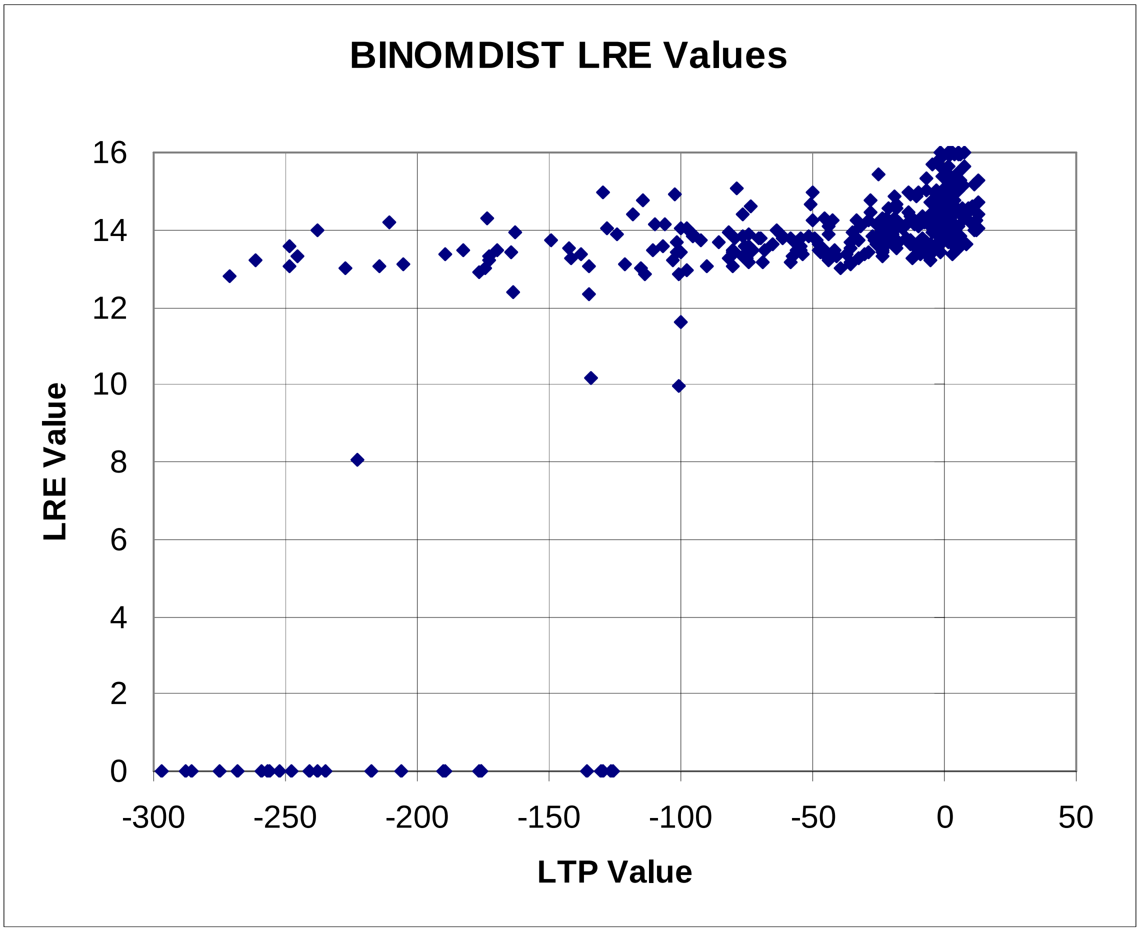

Most of the returns were accurate. 26 out of the 639 that were numerical returns were false zero's, giving a false zero proportion of about 4%. Additional observations were:

Out of the 362 #NUM returns, the lowest number of trials was 1031.

Most of the number of trials greater than 1030 were in the #NUM category. Of the correct, non-zero numbers, 45 (7%) were with a number of trials greater than 1030. The maximum number of trials that gave correct results was 1978.

EXCEL 2003 AND 2007:

The Excel 2003 version is considerably different from the earlier versions. In the 2003 version the function as shown in help, was the algorithm used where the number of trials is less than or equal to 1030. Above this value, a unique approach was taken to only sum up the p^n * (1-p)^(n-s) terms below the maximum (S1) and sum up the terms above the maximum (S2) and return the ratio of S1/(S1+S2) as the cumulative p value. This approach allowed the function to exceed the prior version limits. KBA 827459 describes this and other changes made for Excel 2003. The new algorithm for larger numbers of trials is an interesting approach. However in its actualization, there are some shortcomings.

The key to understanding the 2003 BINOMDIST is to recognize that two functions were smashed together, and carefully covered to look like one function. To make any sense of test results, the two parts have to be tested separately. As it turns out, both have entirely different error structures, and have different accuracies.

One part is the older pre-2003 function for all number of trials less than 1030. The second part is the new added in algorithm to provide p values for cases where the number of trials exceeds 1030. Although KBA 827459 does not give Knüsel (1998) the credit for developing the method, the algorithm is a version of the method developed in section 6.4, Computation of Upper Tail Probabilities of Knüsel (1998). The example is on the Poisson distribution, and Microsoft adapted it for the binomial distribution. The algorithm given in KBA 827459 was transferred to a VBA function (FIDDLEB[1]), and was tested separately from BINOMDIST. The results are shown below. The BINOMDIST code that causes the output of the algorithm to be degraded as shown in figure 16-8, comes directly from this new algorithm.

Testing the function and the information in KBA 827459, indicates that there is a threshold value somewhere between 1E-15 and 1E-16 where the output p value abruptly goes to zero. The threshold is variable and is inherent in the new algorithm. The threshold depends on whatever value Microsoft set an internal variable to (called “EssentiallyZero”). For purposes of testing, the threshold value or cutoff was set to 1E-16 as a reference function output value (2E-16 is the smallest fixed point number). This is the lower boundary on the fixed-point number space.

For numbers of trials below 1030, the function will report small tail areas below 1E-16. For numbers of trials above 1030, the function will return all p values less than this threshold as zero. This is a built in discontinuity.

A test was made based on 5000 random points in the following ranges: population size from 3 to 1030 (low range) or 1030 to 10000 (high range), number of successes in population from 0 to population size, sample size from 0 to population size, and number of successes in the sample from 0 to sample size. The random numbers were all uniformly distributed within these ranges. One run is one combination of the values, together with both BINOMDIST and reference function outputs.

Table 16-4: BINOMDIST-Cumulative Outputs, Excel 2003 and 2007, Random Number Inputs

Function Returns	Low Range	High Range
True zeros as zero	246	2940
True zeros as #NUM!	0	0
False zeros as zero, below cutoff	205	1562
False zeros as #NUM!, below cutoff	0	0
False zeros as zero above cutoff	0	26
False zeros as #NUM!, above cutoff	0	0

[1] Knüsel’s development is a basic theme. Microsoft did variations on it. In classical music, variations on a theme from another composer, played on a solo instrument is common. Hence the “term” fiddle” seems appropriate.

Function Returns	Low Range	High Range
Correct 0-1 values below cutoff	1152	0
Correct 0-1 values above cutoff	2741	886
True ones	926	4586
Total	5000	10000
Minimum LRE	12.9	13.7

The false zeros above cutoff were in all cases, due to the BINOMDIST output value being below cutoff, when the true value was greater than cutoff. At this point the BINOMDIST output is off by a factor near 10 from the true value. Figure 16-9 is a plot of LRE measurements on random data inputs from the low range, where the number-of-trials do not exceed 1030. The true zeros, the false zeros and the unity points were left out.

Figure 16-9: Accuracy of BINOMDIST-Cumulative, Excel 2003 and 2007, for Numbers of Trials less than 1030

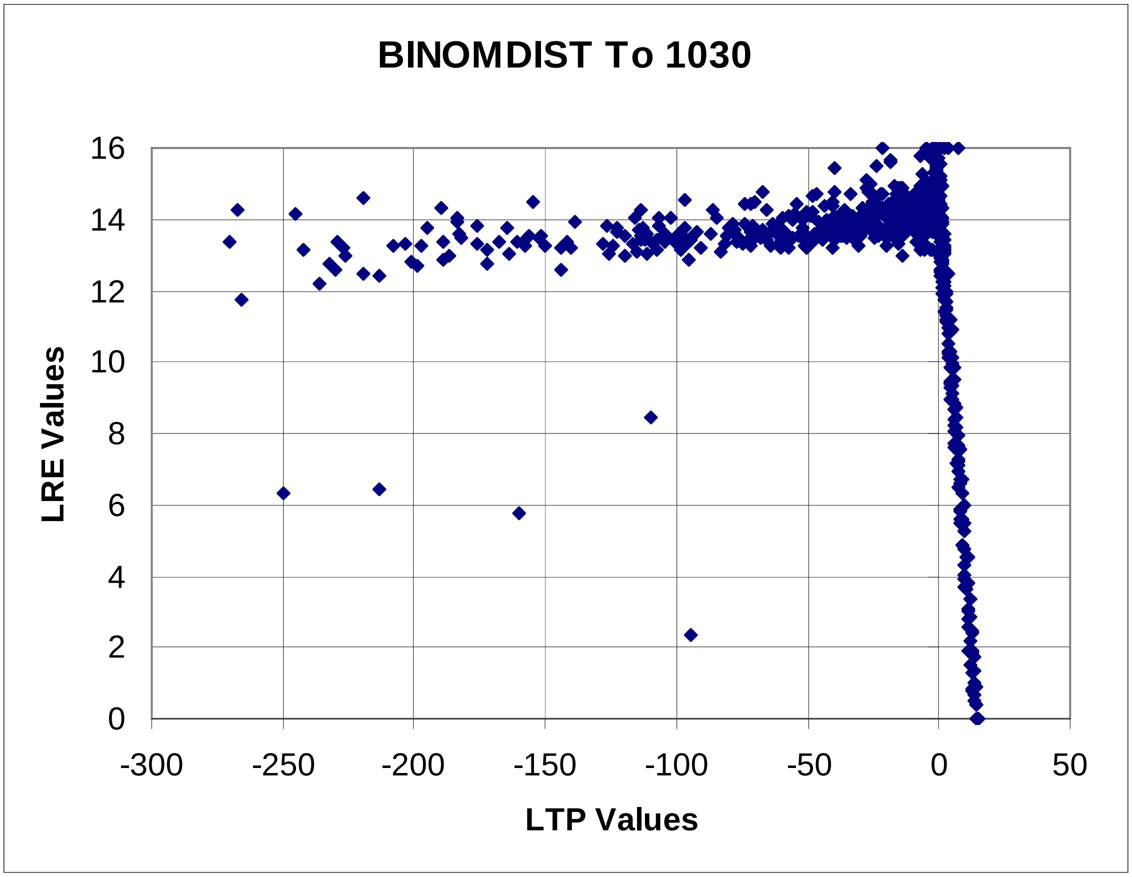

The false zeros (not shown) and the seven low LRE points in figure 16-9 come from basic defects in the algorithm. In the original algorithm, the series terms are summed up until a term below the threshold occur or a limit on the number of terms is reached, then the summation stops. The sum is undervalued, and results in low accuracy of the output. The seven points in figure 16-9 come from this defect.

Figure 16-10 is a close up of the central region. It very clearly shows the loss in accuracy when the p values exceed 0.5. The values used for plotting come from test method 3. False zeros are excluded.

Figure 16-10: Accuracy of BINOMDIST-Cumulative, Excel 2003 and 2007, for Numbers of Trials less than 1030, Central Region

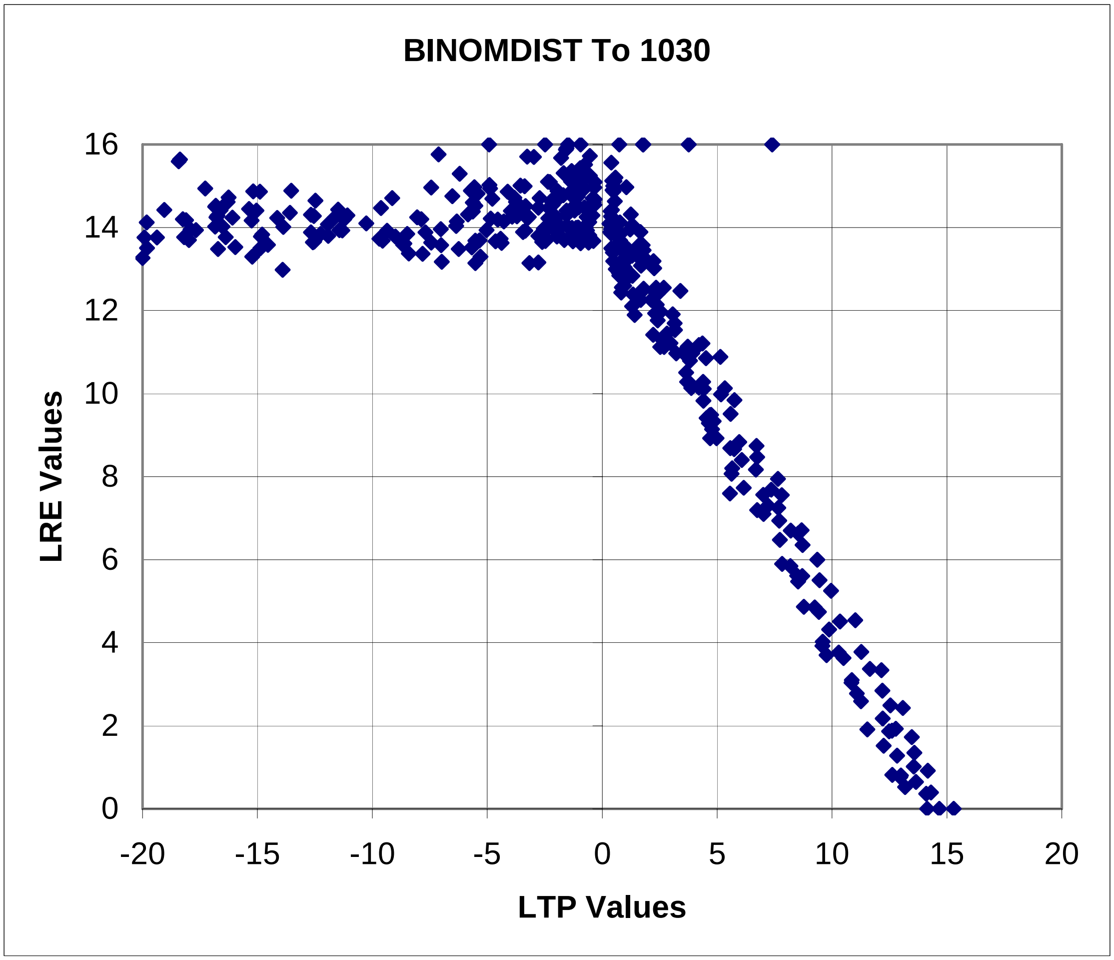

This older function will give small tail p values down to 1E-305. When a zero (0) is output, there is a chance of 45% that the zero is wrong (i.e. the tail is greater than zero). If your standard is to consider these small tail areas below threshold as zero, then the false zero occurrences is not a problem.

Figures 16-9 and 16-10 are different from figure 16-8, which was a test on Excel 2000, and test method 2 was used. The test on Excel 2000 did not have any bounds on the maximum number of trials. It is evident that in the older code, numbers would come out for numbers of trials greater than 1030.

Figure 16-11 is a plot of LRE measurements on random data inputs where the number of trials exceeds 1030. This is the region where the new code described in KBA 827459 works.

<u>Figure 16-11: Accuracy of BINOMDIST-Cumulative, Excel 2003 and 2007, for Numbers of Trials Greater Than 1030, Central Region</u>

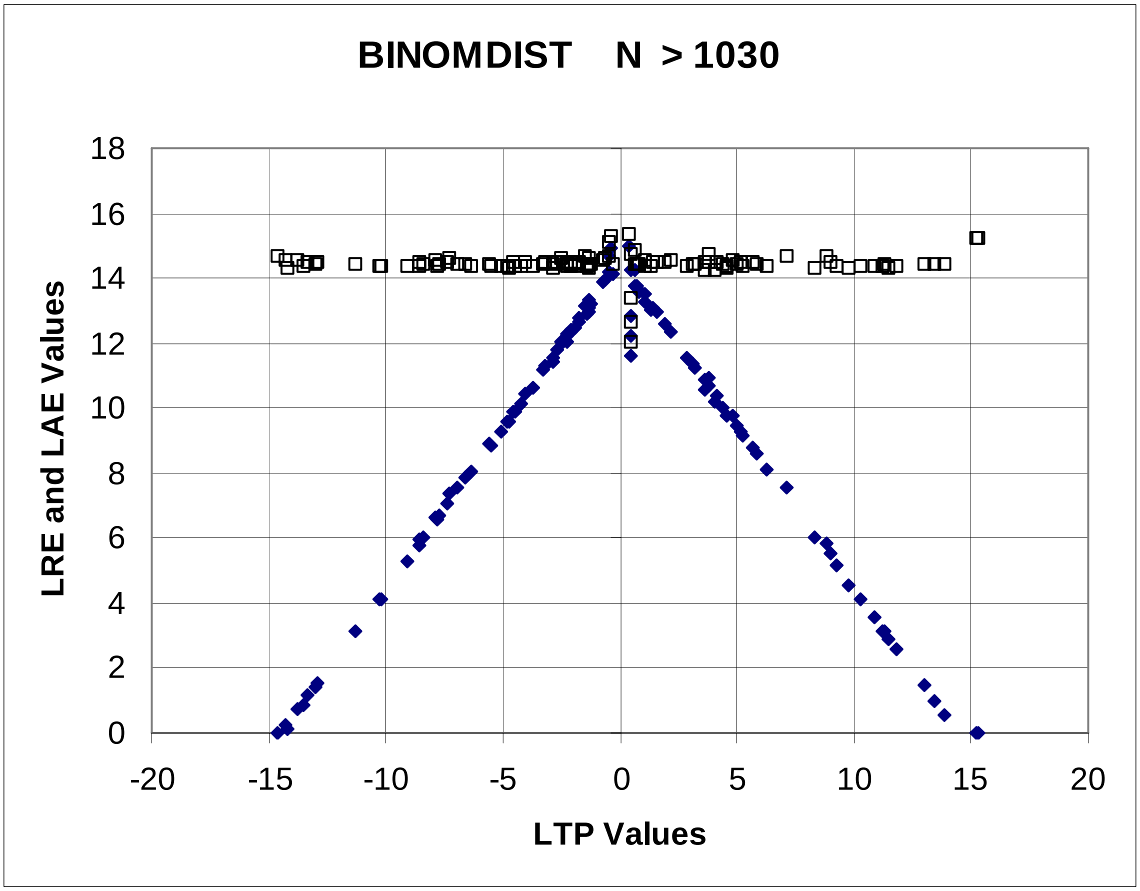

There are two sets of points here. The open boxes are LAE measures (Method 1) and the solid diamonds are LRE measures (Method 3). The function returns zero for any true probability value less than about 1E-14, and 1.0 for any true probability greater than 1-1E-14. For N > 1030, BINOMDIST will not return proper left tail values in terms of a floating point reference.

The points vividly show the declining relative accuracy in p values, as the tails get smaller. This is the relative measure. Note that the 15 digit fixed point accuracy claimed by Microsoft is not true, it is more like 14 digits outside of the –2 LTP to +2 LTP range, and about 11 inside this range.

The new algorithm was tested to examine its accuracy, separate from its BINOMDIST shell. For reference, this new function is called "FIDDLEB". Figure 16-12 is the output of FIDDLEB, using the value of the loop exit variable "*EssentiallyZero*" as 1E-12 as suggested in KBA 827459. The comparison is to the reference functions cdf_binomial and comp_cdf_binomial.

Figure 16-12: Accuracy of FIDDLEB-Cumulative, for Numbers of Trials Greater Than 1030, Central Region

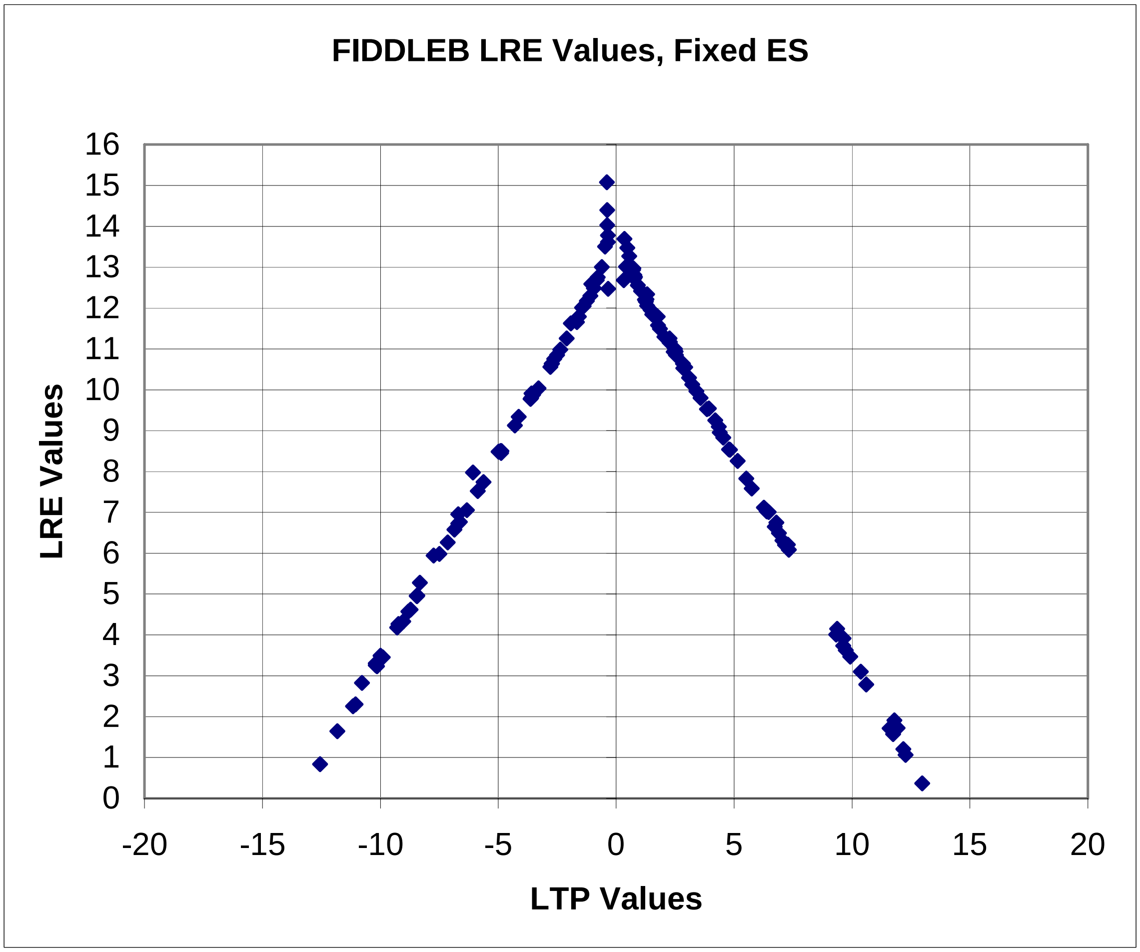

The inverted v shape of the points is almost identical to that in figure 16-12. The differences in vertical location suggest that Microsoft used a value smaller than 1E-12 as *EssentiallyZero*. It is evident that the poor accuracy of BINOMDIST comes from this algorithm.

Figure 16-13 is another view of the algorithm in which the loop exit is not a fixed value, but a relative value (*EsentiallyZero* times *CurrentValue*).

Figure 16-13: Accuracy of FIDDLEB-Cumulative, for Numbers of Trials Greater Than 1030,

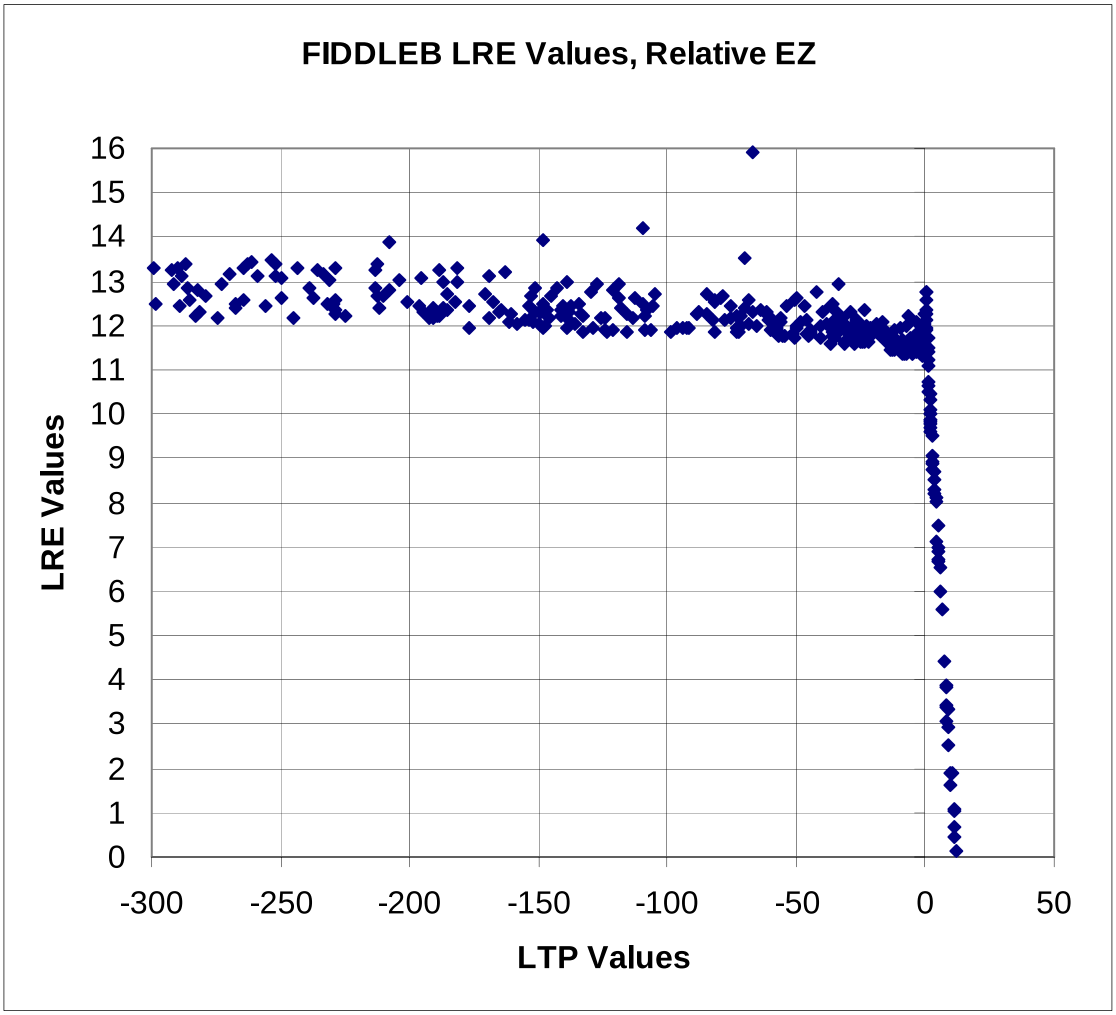

Making this change drastically improves the accuracy of FIDDLEB. Accurate values of the tails are now available, where before, tail areas were not available. Figure 16-14 is an expanded view of the points in the central region. For this data *EssentiallyZero* was changed to a relative 1E-16. This figure should be compared to figure 16-13.

Figure 16-14: Accuracy of FIDDLEB-Cumulative, for Numbers of Trials Greater Than 1030, Central Region.

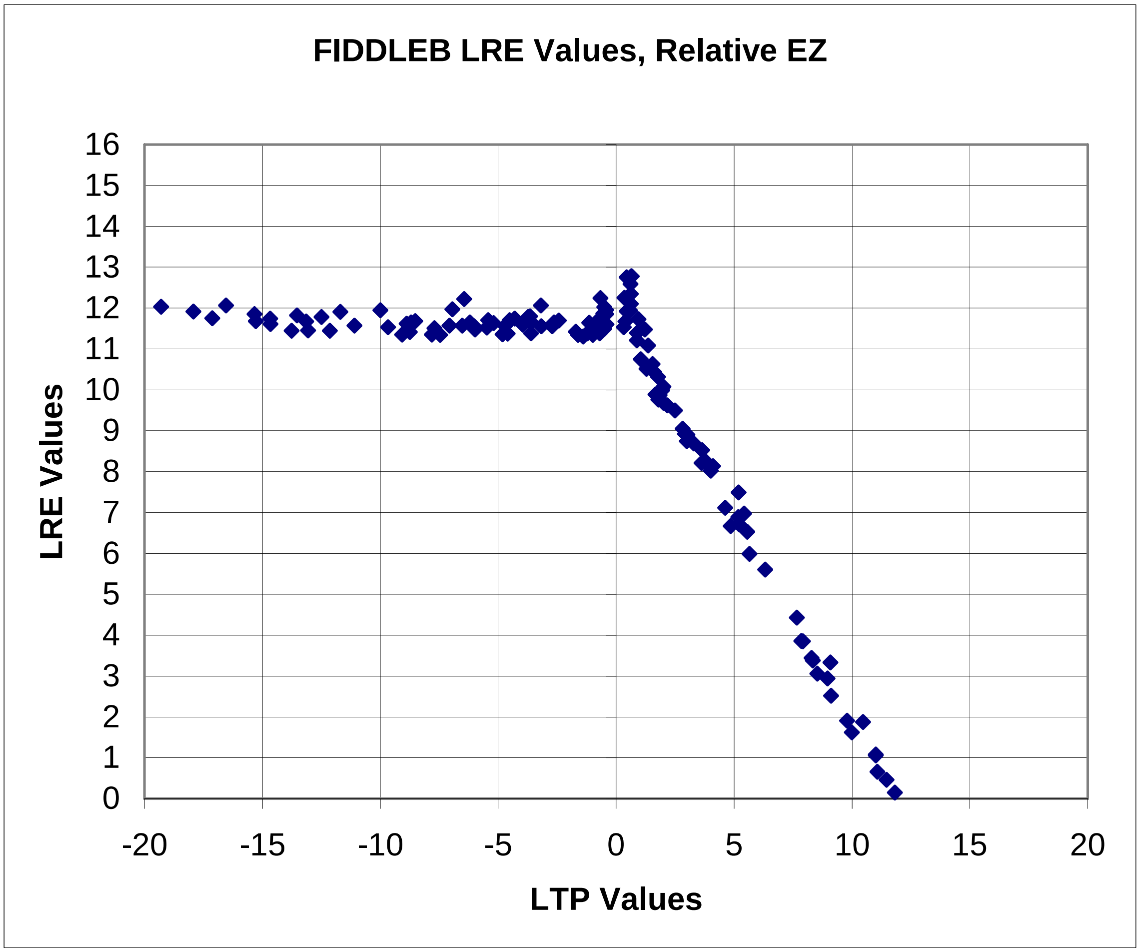

The decline in accuracy for LTP values greater than zero is an inherent characteristic of upper tail values where preceding nines reduce the accuracy of the tail area by complementation. Only by use of a separate q function can this be eliminated. Actually, by interchanging variables as described in table 9-3, accurate values of q can be obtained. If this method is used, then the points to the left of the zero axis can be mirrored on the right side, giving accuracies of at least 11 digits for all possibilities. Excel is penalized here, because a separate function as stated by Knüsel (1999) is required for q values.

Table 16-5 is a summary of the output details from both versions.

Table 16-5: FIDDLEB-Cumulative Outputs, Random Number Inputs

EssentiallyZero	1E-12, Fixed	1E-12, Relative
True zeros as zero	1445	1499
True zeros as #NUM!	0	0
False zeros as zero, below cutoff	773	0
False zeros as #NUM!, below cutoff	0	0
False zeros as zero above cutoff	22	0
False zeros as #NUM!, above cutoff	0	0
Correct 0-1 values below cutoff	0	437
Correct 0-1 values above cutoff	418	761
True ones	2346	2303
Total	5000	5000
Average LRE	8.4	11.5

FIDDLEB in the relative mode provides accurate p values for N=1 to N=10,000, without any false zeros or #NUM!-returns.

Other investigators such as Knüsel using test method 5 are likely to test BINOMDIST with a different set of input parameters. Table 9-9 gives some error results using these input parameter values.

Table 16-6 is an overall tally of results of 10,000 runs of BINOMDIST with sample size varying randomly between 1 and 5500, the number of success as random varying from sample size to 0 and p values random between 0 and 1. Values were uniformly distributed within these ranges. The rounding (N) was to 12 decimal digits to the right of the true decimal point.

Table 16-6: BINOMDIST-Cumulative, Excel 2003 and 2007 Output Analysis

SUMMARY	Total Set
True Zero as Zero	2044
True Zero as #NUM!	0
Sum True Zeros	2044
#NUM!-returns, True Value >0	0
False Zeros As Zero Below Threshold	1687
False Zeros As Zero Above Threshold	24
Sum False Zeros	1711
Sum Total Zero Returns	3755
#NUM! Return Above Threshold	0
#NUM Return Below Threshold	0
Sum #NUM!-returns	0
NZ Values Below Threshold	452
NZ Values Above Threshold	1878
Sum NZ Returns	2330
Matched Ones	3890
False Ones (True NZ)	25

SUMMARY	Total Set
Sum One Returns	3915
Total Returns	10000
Minimum LRE	0.12
Average LRE	13.28
Standard Deviation of LRE Values	2.76
Number of LRE Values Used	2330
NZ Values Not Round Matched	
Round Basis	FIXED POINT
Round Set Number, N	12
Number NZ values Not Matched at N-1	1
Number NZ Values Not Matched at N	11
Number NZ Values Not Matched at N+1	59

If the number-of-NZ-values-not-matched are converted to probabilities, based on the total number of NZ values, then the probabilities in terms of standard normal z values can be plotted versus the rounding digits. Figure 16-15 is the result

Figure 16-15: Relationship Between the Proportion of NZ Values Rounded to N Digits That Are Inaccurate

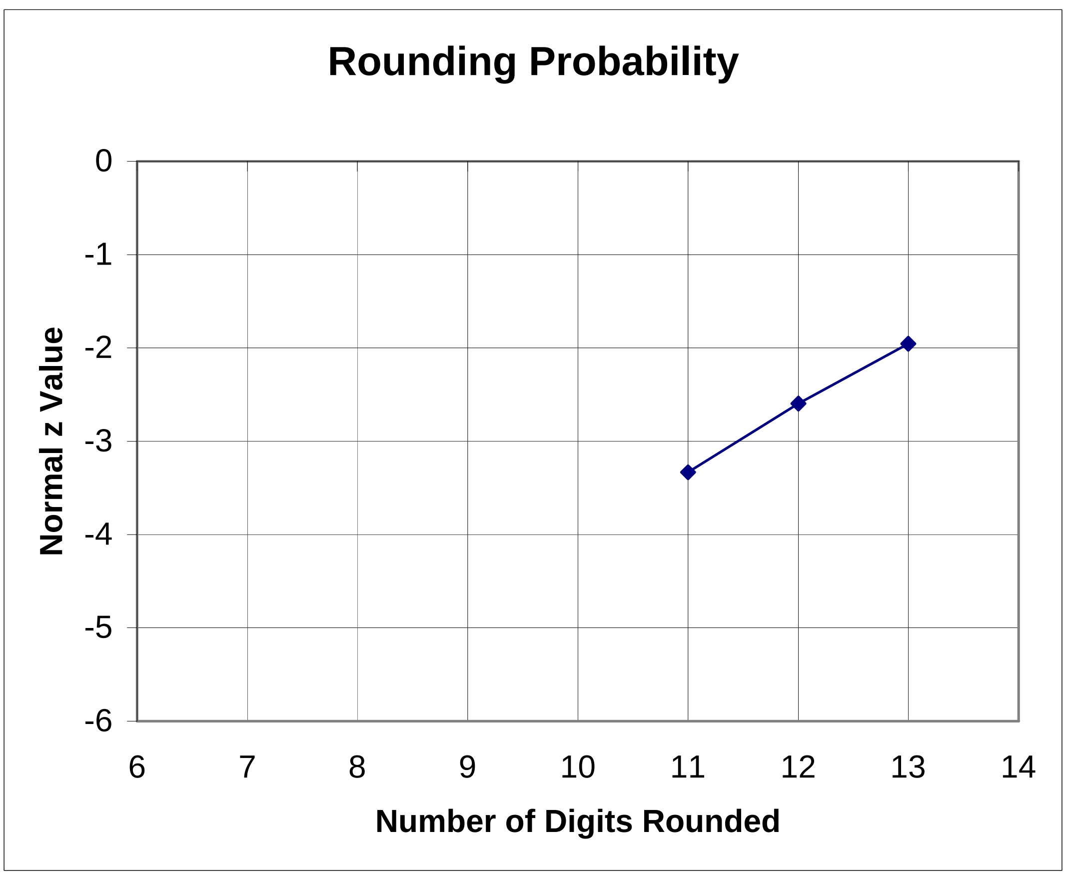

This suggests that the approach of “Accurate to 12 decimal digits“ is not a correct view, but should be restated in terms of probability. In this case, the statement “Accurate to12 decimal digits 99.53% of the time’ fits the test results.

Table 16-7 shows several things. One is that an analysis based on likely user input parameter values results in values entirely different from method 5 tests as shown in table 12-6. It also shows that if only NZ numbers are evaluated for accuracy, then if the number-of-NZ-digits are rounded to 11, then this will be accurate, 99.95 percent of the time. It also shows that the number of rounded digits representing accurate digits is above that shown in table 16-6. It also strongly suggests that the concept of n digits being accurate is a probability issue, not an absolute issue. It also suggests that method 5 gives biased results. Figure 16-15 shows that even with a fixed point rounding of 11, there are still output values that occur that are not accurate to 11 digits, but at a very small percentage of NZ values. It also shows that BINOMDIST will return a zero when the true p value is greater than 1E-16 about 0.6% of the time. Given the 11 digit rounding level, then the false zero and zeros above threshold issues are moot.

FIDDLEB with a small relative *EssentiallyZero* is about the best accuracy that can be obtained with Microsoft’s algorithms. It fixes the problem of false zeros. The problem of the failure to provide p values for very large n and s values still has to be solved. It is unfortunate that the accuracy of the actual product is that shown in figure 16-11 because Microsoft made a wrong decision.

From this we can conclude several things:

1. Smith’s reference function and ELV agree within the limit of his displayed figures (from 5 to 7 digits).
2. The new BINOMDIST will never pass any method 5 test, because there will always be the finding that some input parameter set will give an inaccurate 11th digit.
3. Microsoft needs to do some work on their algorithm to handle large values of input parameters, especially for the benchmark set, rows 4 thru 12 of table 16-6. The failure is in Microsoft’s algorithm as published in KBA 827459.
4. Microsoft needs to change the BINOMDIST closure from fixed to relative, to pass rows 1-3, 18-22, 27, and 38-41 of table 16-6. The issue the fact that BINOMDIST is one function that requires a fixed point rounding method, and the others can be floating point rounded. This is a non-consistent approach to the issue of displayed digits, as long as Excel allows floating point displays, and continues to show p values as floating point numbers in the add-in Tool Pak outputs.
5. A right tail (q) function (which is not really needed, since it can be calculated by interchanging parameters as described in table 9-3) should be part of the suite, to satisfy Knüsel (1999).

RELIABILITY ASSESSMENT OF BINOMDIST IN EXCEL 2000

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

RECOMMENDED EXCEL 2000 BINOMDIST USAGE

Range of x Values	Range of Number-of-Trials	Range of Number-of-Successes	Restrictions	Round Level	Basis	Percent of Output Values Below Round Accuracy
0 to 1	0 to 2000	0 to #Trials	No zero output values	12	Floating Point	5
0 to 1	0 to 1000	0 to #Trials	No zero output values	13	Floating Point	1

RELIABILITY ASSESSMENT OF CUMULATIVE BINOMDIST IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
Very Frequent. About 17% of the time	None	None	None	None

RECOMMENDED EXCEL 2003 AND 2007 CUMULATIVE BINOMDIST USAGE

Range of Probability of Single Event Success	Range of Number-of-Trials	Range of Number-of-Successes	Restrictions	Round Level	Basis
0 to 1	0 to 6000	0 to #Trials	P values <1E-16 are all considered as zero	10	Truncated Floating Point

1.3.2 POISSON DISTRIBUTION, CUMULATIVE: POISSON

The function is =POISSON(x, mean, TRUE)

The cumulative Poisson distribution is computed in POISSON as a straightforward sum of Poisson density terms, or equivalents.

There are two ways to obtain cumulative values, one the simple sum of terms as described above, and two from the relationship to other functions. Microsoft uses the first method and as such is limited by the overflow/underflow characteristics of floating point numbers.

EXCEL 2000

The POISSON function was not well coded for robust performance and is highly susceptible to floating point number limits. This is why POISSON will suddenly output #NUM in a range of computations. It has nothing to do with "central probabilities" as stated by Knüsel.(1999). The #NUM will show up in tails just as well.

There is a region in the POISSON function where the function outputs #NUM. I call this "rejection". If a number comes out, I call it "acceptance". If the number is zero and the true p value small, this is then a LRE problem. Table 16-7 gives boundary values for the rejection region.

Table 16-7: POISSON-Cumulative, Excel 2000 Function Reject Region

Poisson Mean value	Accept	Reject	LRE at accept
1	170	171	15.65
20	170	171	14.48
100	154	155	13.38
200	133	134	13.20
300	124	125	13.48
400	118	114	14.35
500	114	115	13.73
600	110	111	13.32
700	108	109	13.19
800	106	107	0
900	104	105	0
1000	102	103	0

Straight lines between these points represents good approximate limits.

The second method to obtain cumulative POISSON values is not sensitive to the overflow problem and can be used instead of POISSON to obtain acceptable values. Although not exact, the CHIDIST function or the GAMMADIST function will give LRE values of 6 or more. For example, if the number of events are in cell A2 and the mean in cell B2 then if you enter =CHIDIST(2*B2,2*(A2+1)) in the cell for Poisson cumulative probability, you will get the correct probability with an accuracy of LRE = 6 or more.

When POISSON works, the p value is accurate to LRE values greater than 13, except when the output is either 1 or zero. The high LRE values hold up to about p = 1E-160. Below this the LRE values drop quickly to zero.

EXCEL 2003 AND 2007

KBA 828130 describes changes made to earlier versions for the Excel 2003 version. The changes were to incorporate the algorithm given by Knüsel (1998) in section 6.4, on

“Computation of Upper Tail Probabilities”. The change was intended to correct the problem of overflow in the original algorithm.

Microsoft stated that the overflow problem that gave #NUM!-returns occurred when the mean to the x power was too large. They stated the problem was that when the term (mean^x) exceeded 1E+291 or when x was greater than 170, overflow would occur. What their actual algorithm actually does is unknown, but of the three elements in each term (FACT(x), (mean^x) and EXP(-x)), the first is limited to a maximum of 170 (returns 7.2574E+306), the second (given x=170) is limited to a mean of no more than 65.0 (returns 1.56772769E+308), and there is no limit on x for the third term. The LOG10 function on 1.56772769E+308 is 308.195. The 290 value does not seem to correspond to these limits.

The new algorithm is very similar to FIDDLEB, and is shown in KBA 828130. Note that there is a typo error in the algorithm, since if coded exactly as shown; it will give some really weird outputs. The version used in the following tests follows Knüsel’s 6.4 development and is called FIDDLEP. It will produce accurate results.

As discussed under BINOMDIST, the two algorithms are different and have entirely different error structures. Although Microsoft states the dividing line is the 170/290 line, actual parameter values above these limits return values and error patterns (low tail probabilities and false zeros) characteristic of the old algorithm. Therefore there is no clean line as there is in BINOMDIST, and only a single view, that of testing with x from 0 to N and with the mean from 0 to V, where X and V are large, is possible.

Figure 16-16 is the results of method 3 on the Cumulative POISSON. Both x and mean values were varied randomly. It is very similar to BINOMDIST figure 16-11.

Figure 16-16: Accuracy of POISSON-Cumulative, Excel 2003 and 2007

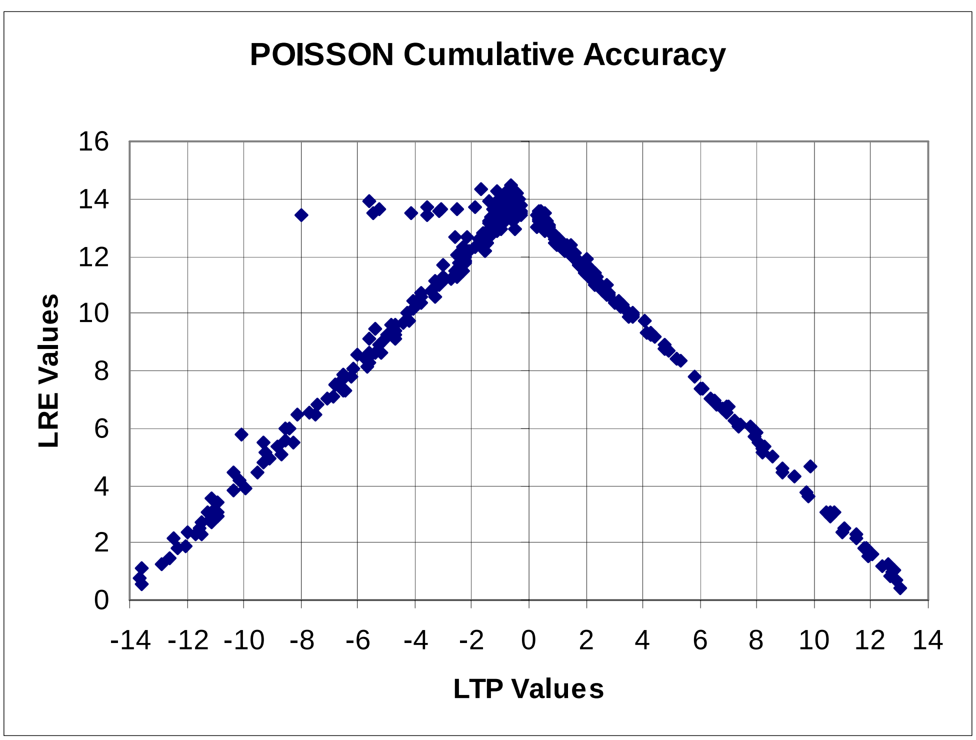

The horizontal point group above the 13-line come from the original algorithm. The other inverted V points come from the new algorithm.

Evaluation of POISSON in Excel 2003 should be based on the LAE column (last). Figure 16-16 shows that accuracy should be referenced to a fixed-point basis. As constructed, POISSON can only return zero tail p values when p is below a threshold of about 1E-15.

Figure 16-17 shows the accuracy of the FIDDLEP algorithm in the relative mode. In the absolute mode, the points are an inverted V similar to FIDDLEB in the absolute mode, as shown in figure 16-16. As discussed above, FIDDLEP differs only slightly from FIDDLEB, since they are the same theme, differing only in the key being played.

Figure 16-17: Accuracy of FIDDLEP-Cumulative, Relative Mode.

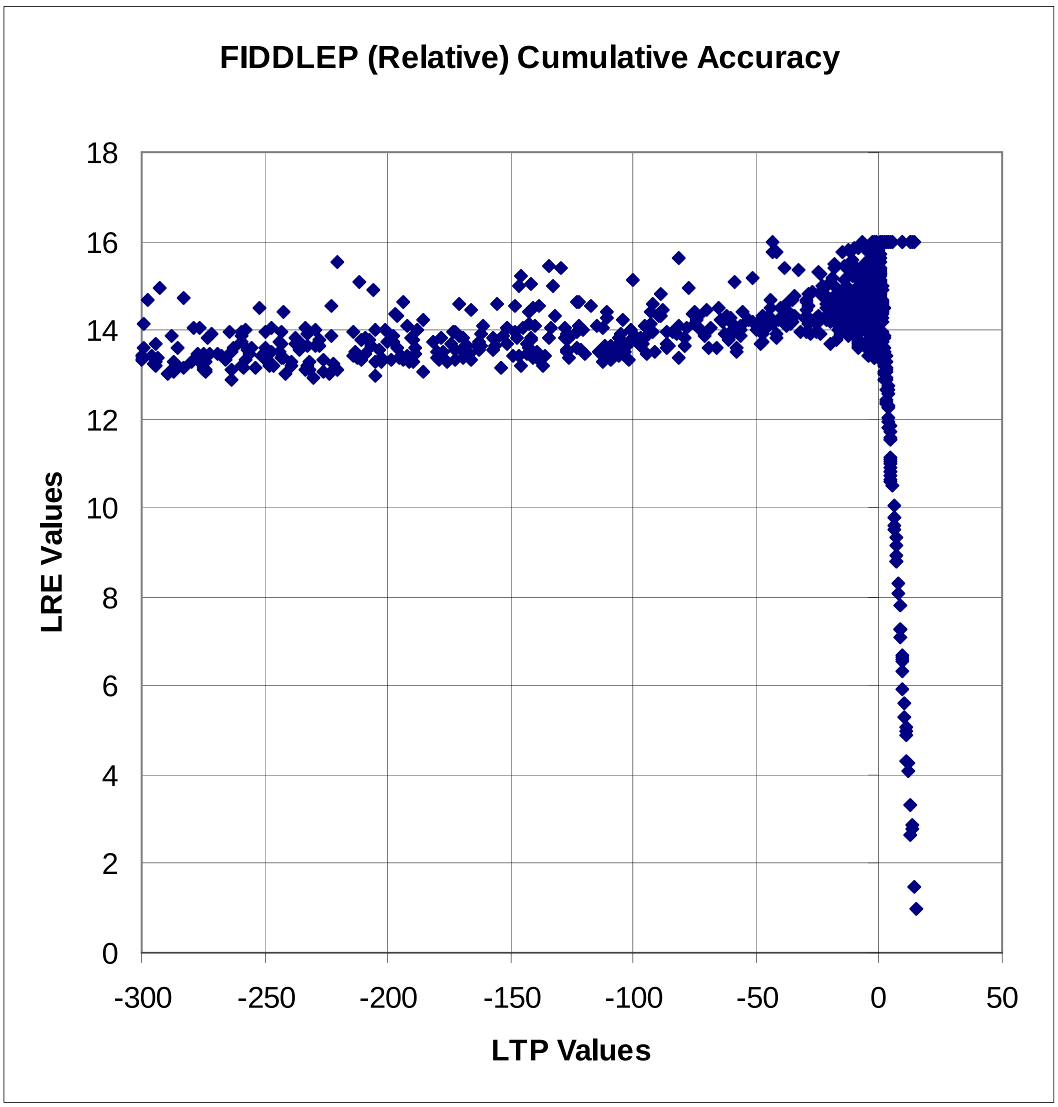

The improvement in accuracy is noticeable. The ranges for x were from 0 to 1000, and the mean from 0 to 100. Mean values above 100 resulted in many more zeros and ones, giving few plotable points. There were no #NUM! or false zero returns

Table 16-8: Summary of POISSION and FIDDLEP, Excel 2003 and 2007 Test Results

SUMMARY	POISSION	FIDDLEP
True Zero as Zero	0	425
True Zero as #NUM!	0	0
Sum True Zeros	0	425
#NUM!-returns, True Value >0	0	0
False Zeros As Zero Below Threshold	1731	0
False Zeros As Zero Above Threshold	201	0
Sum False Zeros	1932	0
Sum Total Zero Returns	1932	425
#NUM! Return Above Threshold	0	0
#NUM Return Below Threshold	0	0
Sum #NUM!-returns	0	0
NZ Values Below Threshold	0	3748
NZ Values Above Threshold	2623	609
Sum NZ Returns	2623	4490
Matched Ones	393	85
False Ones (True NZ)	52	0
Sum One Returns	445	85
Total Returns	5000	5000
Minimum LRE	0.30	0.80
Average LRE	8.76	13.90
Standard Deviation of LRE Values	4.12	0.97
Number of LRE Values Used	2623	4357
NZ Values Not Round Matched		
Round Basis	Fixed Point	Floating Point
Round Set Number, N	10	11
Number NZ values Not Matched at N-1	0	1
Number NZ Values Not Matched at N	1	5
Number NZ Values Not Matched at N+1	6	38

The results show that using the new algorithm <u>for all values</u> and converting it to a relative closure is an improvement in the performance of POISSON.

Microsoft should change the POISSON function to use the new algorithm in the relative mode. A change to deal with large input values must be added. There is no reason why POISON can't be changed to properly return dependable values for the benchmark points.

RECOMMENDED EXCEL 2000 POISSON (CUMULATIVE) USAGE

Range of x Values	Range of test parameter	Restrictions	ROUND level	Basis
0 to 170	X * LOG10(mean) < 290	Range of Input Parameters	8	Floating Point

RELIABILITY ASSESSMENT OF POISON IN EXCEL 2003 AND 2007

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

RECOMMENDED EXCEL 2003 AND 2007 POISSON USAGE

Range of x Values	Range of Mean Values	Restrictions	ROUND level	Basis
0 to 1000	0 to 1000	P values <1E-16 are all considered as zero	11	Truncated Floating Point

1.4 DISCRETE INVERSE FUNCTIONS

1.4.1 BINOMIAL DISTRIBUTION, INVERSE: CRITBINOM

The function is =CRITBINOM(trials, probability_s, alpha)

CRITBINOM is the inverse of BINOMDIST. The inputs are sample size, the population single event probability of success, and an alpha (probability) value. It is a cumulative function.

EXCEL 2000

This function has limitations. If the alpha values are within the range of 0.00001 to 0.99999, then the function will return the correct k values. Outside of this range, the k values are likely to be off by one or more units. Since it iterates BINOMDIST to obtain a solution, the overload aspects of BINOMDIST for a large number of trials may occur.

A test of this function in comparison to Smith's corresponding function was made, using random values of the population probability, alpha values and number of trials. The outputs of both functions are integer values. The results showed no differences in outputs in 10,000 calls. Again the range was limited to less than 1030 trials.

EXCEL 2003 AND 2007

The function was changed with an improved closure algorithm. The change is described in KBA 828117.

Microsoft said, “The approach to improvements in Excel 2003 is exactly the same as with BINOMDIST: use existing pre-Excel 2003 code if n < 1030 and switch to an alternative plan if n >= 1030. The remainder of the discussion in this section deals with only the case where n >= 1030. The alternative plan is built in the same way as for BINOMDIST: find the modal value m (m = approximately n*p), assign an unscaled probability of 1 to m, find unscaled probabilities of m+1, m+2, m+3, ... stopping when such probabilities become infinitesimal, find unscaled probabilities of m-1, m-2, m-3, ... stopping when such probabilities become infinitesimal. Finally, scale the appropriate probabilities.

“CRITBINOM is a kind of inverse function for a discrete distribution that is similar to NORMSINV for the continuous standard normal distribution. NORMSINV is computed through a search process that frequently calls NORMSDIST as it homes in on the result. Because BINOMDIST is computationally expensive, you want to avoid a similar process for CRITBINOM that would repeatedly call BINOMDIST. The procedure establishes an initial guess, executes the code below (similar to a single call to BINOMDIST), and then "tweaks" (adjusts) the guess to arrive at a final answer.

“First, find a guess by using a normal approximation to the Binomial distribution. You can assume a normal distribution with the same mean and standard deviation as the Binomial, namely mean n*p and standard deviation SQRT(n*p*(1-p)). This approximation should be reasonably accurate as long as n*p*(1-p) > 30, say; and this will be the case with n >= 1030 unless p is very close to 0 or very close to 1. Because you require only an approximate value, use a quick approximation to NORMSINV instead of calling NORMSINV itself. The approximation comes from 26.2.23 in Abramowitz, M. and I. A. Stegun, *Handbook of Mathematical Functions*, Dover, 1972, p. 933.

“CRITBINOM has been thoroughly tested for accuracy. However, only casual anecdotal testing has been done to investigate how close the initial Guess is to the correct answer and how many times the Guess has to be increased or decreased. The normal approximation generally provides an excellent value of Guess; in our limited casual tests, we never had to increase or decrease the initial Guess by more than 2.”

“In pass 1, the function was run over 10,000 times. The number of trial was random between 1 and 10000, the probability of success randomly varied from 0 to 1 and alpha randomly varied from 0 to 1.The function returned identical values to the reference function. Since the returned values are integers, all numerical returns had an LRE of 16. There were a few #NUM returns, representing a #NUM rate of less than 0.1%. There were no #NUM returns from the crt_binom reference function. Table 16-11 gives parameter values that resulted in a #NUM! return.

Table 16-9: Summary of Inputs to CRITBINOM, Excel 2003 and 2007, For #NUM Returns

Number of Trials	Probability of Success	Alpha value
6003	0.998204	0.99951846
7306	0.836596224	0.002180192
5015	0.874876222	0.995134767
4679	0.906213314	0.994669647
9831	0.916451856	0.993101166
Number of Trials	Probability of Success	Alpha value

5130	0.73842427	0.000530713
5851	0.824195244	0.001895629

In KBA 828117, Microsoft said, “Inaccuracies in earlier versions of Excel occur only when the number of trials is greater than or equal to 1030. In such cases, CRITBINOM returns #NUM! in earlier versions of Excel. This issue occurs because one term in a sequence of terms to be multiplied together when you evaluate BINOMDIST overflows. This issue has been corrected in Excel 2003 and 2007 by implementing an alternative procedure that is described earlier when such an overflow would otherwise occur.”

Microsoft didn’t completely fix the #NUM! problem, as they said.

In pass 2, the number of trials (n) was increased substantially, up to the 1E+10 range and probability of one event success (ps) varied over several ranges from 0 to 1, from 0 to 1E-03, from 0 to 1E-09 and from 0 to 1E-09. This was a one-row test to explore the upper limits. Most of the time the function would return a value. However there were frequent #NUM!-returns. When #NUM! occurred, the relationship between n, ps and alpha values that caused it, was not clear.

One combination, an n value of about 8.23E+08, an alpha of 0.99, and ps equal to 0.76 caused a lock-up. The lock-up occurred with other value combinations. One possibility is that the n * ps product in the 6E+08 range together with a high alpha results in Excel lockup. There is some error in the function code. In all cases it was with high n values. Clearing the lock-up required clicking the Excel X button in the upper corner, closing out Excel and loosing the current workbook.

Excel very carefully blocked the ESC key from stopping the calculations when the calculations only involve Excel internal subroutines and functions. The ESC key works to stop calculations (end, debug, continue MSG box) for error tracing in external VBA add-ins. When there is a fault in one of the Excel built-in functions, and a loop occurs, the only recourse is to close out Excel and lose everything.

CRITBINOM will accept ps values less than 1E-15. There does not appear to be an input test on the size of n. There are the standard input tests on ps and alpha.

RELIABILITY ASSESSMENT OF CRITBINOM IN EXCEL 2000

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

RECOMMENDED EXCEL 2000 CRITBINOM USAGE

Range of Number of Trials	Range of Single trial Success p Values	Restrictions	Accuracy	Probability of Values Being Inaccurate
1 to 1030	0 to 1	Disregard #NUMs	The whole number is fully accurate	0

RELIABILITY ASSESSMENT OF CRITBINOM IN EXCEL 2003 AND 2007

False Zeros	Non-Numeric Returns	Gross Errors	Logic Traps and Loops	System Halts and Crashes
None	About 0.1% of the time	None	Yes	Yes

RECOMMENDED EXCEL 2003 AND 2007 CRITBINOM USAGE

Range of Number of Trials	Range of Single trial Success p Values	Restrictions	Accuracy	Probability of Round Values Being Inaccurate
1 to 10000	0.0001 to 1	Disregard #NUMs	The whole number (integer) is fully accurate	0