DOES MICROSOFT FIX ERRORS IN EXCEL?

B. D. McCullough
Department of Decision Sciences, Drexel University, Philadelphia PA 19104

KEY WORDS: accuracy, reliability, statistical software

1 Is There a Demand for Inaccurate Software?

McCullough and Vinod (1999, p. 637) have argued that if users of statistical software do not demand accuracy, there is no reason for developers to supply it. Wilkinson (1994) has argued that there is a large class of users of statistical software who place a very low priority on accuracy and reliability. Taken together, the above observations suggest there may be a market for unreliable and inaccurate software. Perhaps more precisely, it can be asserted that there is a large class of users of statistical software for whom accuracy is not a priority. Such users place a higher premium on ease-of-use, ubiquity, compatibility with other products, low cost, a glitzy GUI – these are all things at which Excel excels.

Firms that seek to satisfy such users would quite wisely invest more resources in making the software easy-to-use than in making the software accurate. From an economic perspective, it may be a profit-maximizing strategy for the developer of such a package to invest few or no resources in fixing errors. Microsoft has a proven track record of not fixing errors in the statistical procedures of Excel. It is unknown whether Microsoft adopted this strategy intentionally or unintentionally.

2 Errors in Excel

That the statistical procedures in Excel are not numerically robust has long been known to the statistical computing community. The records of various internet newsgroups are replete with accounts of such errors. More formally, this lack of robustness has been known at least since Sawitzki (1994a, 1994b), which circulated in manuscript form for some time before its publication. In particular, using Wilkinson's (1985) entry-level tests of reliability, Sawitzki found many errors and documented that Excel 4.0 used an unreliable algorithm for calculation of the sample variance.

It is a very simple matter to reprogram the calculation of the sample variance – that Microsoft has not fixed this elementary error is extremely troubling. But the fact is that Microsoft has had many years and many opportunities to fix this problem and the others noted by Sawitzki – Excel 5.0, Excel 97, Excel 2000, and Excel 2002 – and for some reason did not.

The importance of this phenomenon – that Microsoft does not fix published errors in the statistical procedures of Excel – cannot be understated since Excel is a very popular package for performing statistical calculations. Searching Amazon.com for the keywords "statistics" and "Excel" turns up over 50 titles, and numerous introductory texts use Excel without having that word in the title. It is quite conceivable that more statistical calculations are performed in Excel than in all other statistical packages combined. However, Excel's growing reputation for unreliable statistical procedures is beginning to catch up with Microsoft, as evidenced by this panel discussion. As Professor Cryer pointed out, even introductory texts that use Excel are warning students against using Excel for solving "real-world" problems (Keller, 2001, p. 14; Sincich, Levine and Stephan, 2002, pp. 50-51).

McCullough (1998) proposed an intermediate-level battery of tests assessing reliability in three areas: estimation, statistical distributions, and random number generation. This methodology was applied to Microsoft Excel 97 (McCullough and Wilson, 1999), which was found wanting in all three areas. To be sure, the numerical methods for remediating most all of these problems are well-known, readily available, and practically costless. McCullough and Wilson (2001) examined Excel 2000 and Excel 2002 and found that practically nothing had changed. The only changes they noted are remarkable because Microsoft did *not* fix problems – but only made the problems less noticeable, and therefore more pernicious. It is unknown why Microsoft chose to replace bad algorithms with other bad algorithms rather than with reliable, accurate algorithms.

3 Example I: The Random Number Generator

Standard normal random deviates should rarely exceed plus or minus 4.0. Yet Excel 97 regularly produces standard normal deviates of plus of minus 5,000,000. To see this, use the pull-down menu for "Tools", "Data Analysis", and "Random Number Generator"; set "Number of Variables" to 10, set "Number of Random Numbers" to 2000, and set "Random Seed" to 123. Next, inspect cells H:253, G:121 and D:1245 to see -5,000,000 and F:1845 to see 5,000,000. These values at least have the virtue that they are so large that the user might notice that something is wrong with his results. Users of Excel 2000 and Excel 2002 will instead find the values -9.53674 and 5.364418. In the same way we expect to find values as large as 1.96 every 20 calls, we expect to see 5.364418 every 12 billion calls, and 9.53674 every 678,000,000,000,000,000,000 calls. These excessively large "standard normal deviates" not only will corrupt results, but they are more likely to produce erroneous results that go undetected by the user. Indeed, one cannot help but wonder why Microsoft implemented this particular "fix"; what was once a serious but possibly noticeable error has become one of what Kahan (2000) calls "errors of the worst kind": too small to be obvious but too big to be tolerable, and too rare to be discovered by the customary desultory testing.

It is worth noting the the Excel random number generator, even aside from the errors described in this section, should be avoided: see McCullough and Wilson (1999), Rotz, Falk, Wood and Mulrow (2001), and L'Ecuyer (2001). Additional questions about Excel's basic numerical capabilities have been raised by the United Kingdom's National Physical Laboratory (Cook, et al, 2000).

4 Example II: Inverse Normal Distribution

Knüsel (1998) observed that Excel 97 had many deficiencies in its statistical distributions, so many that he concluded, "So one has to warn statisticians against using Excel functions for scientific purposes." Knüsel's standard, to which many statistical packages adhere, is that if the package displays six digits, all six digits should be correct. (At the very least, the package should say how many digits can be trusted!) All the errors catalogued by Knüsel remained uncorrected in Excel 2000, and the only observed change in Excel 2002 was to the inverse normal distribution. Table 1 shows some exact values for the inverse normal distribution, along with values from Excel.

x	exact	Excel 97/2000	Excel 2002
0.001	-3.09023	-3.09024	-3.09025
0.0001	-3.71902	-3.71947	-3.71909
1E-5	-4.26489	-4.26546	-4.26504
1E-6	-4.75342	-4.76837	-4.75367
3E-7	-4.99122	-7.15256	-4.99152
2E-7	-5.06896	-5000000	-5.06928

Table 1: Inverse Normal Distribution.

Observe that Excel 97/2000 returns -5,000,000 as a value of the standard normal. Even after being "fixed", Excel 2002 cannot return more than three accurate digits in the tail of the distribution – an unacceptable level of accuracy. It is not difficult to compute the inverse normal correctly to six digits when $x = 0.0000002$ or even smaller. The requisite algorithms are well-known, freely available, and easy to find. Many packages can easily reproduce the exact results given in Table 1. It is unknown why Microsoft chose an unreliable algorithm inverse normal algorithm, when implementing a reliable algorithm would have been just as much work.

5 Conclusions

Perhaps it was Microsoft's strategy to satisfy the demand for inaccurate software. If so, maybe Microsoft has recently changed its strategy, because Microsoft has hired Dr. Bell and charged him with fixing the errors in the statistical procedures of Excel.

If Microsoft intends for Excel to offer reliable statistical functionality, it should act in the same fashion as the purveyors of reliable statistical software. In particular, when serving a clientele that cares about reliability and accuracy, it is desirable that a developer respond to bugs in the following way:

- Upon becoming aware of a bug, the developer makes this knowledge available to all users, by including the bug on its published "bug-list". (Users should beware any package that doesn't have a bug-list at its website. The developer is hiding something.)
- The developer makes known its plans for correcting the bug – will it be fixed in the next minor upgrade or the next major upgrade?
- When the bug is fixed, this fact is made known in the release notes that accompany the upgrade in which the bug has been fixed, and the bug-list is updated to reflect the version in which the bug was fixed.

Especially for Microsoft, the following two caveats should be added: don't replace a bad algorithm with another bad algorithm, but instead use a good algorithm; and don't try to make errors less obvious, instead fix them.

References

Cook, H. R., M. G. Cox, M. P. Dainton and P. M. Harris (2000), "Testing the Intrinsic Functions of Excel," NPL Report CISE 27/00, www.npl.co.uk/ssfm/download

Kahan, W. (2000), "Marketing versus Mathematics," manuscript, EE&CS Dept., U. of CA at Berkeley

Keller, G. (2001), *Applied Statistics with Microsoft Excel*, G. Keller, Pacific Grove, CA: Duxbury

Knüsel, L. (1998), "On the Accuracy of Statistical Distributions in Microsoft Excel 97," *Computational Statistics and Data Analysis*, 26, 375-377

L'Ecuyer, P. (2001), "Software for Uniform Random Number Generation: Distinguishing the Good and the Bad," *Proceedings of the 2001 Winter Simulation Conference*

McCullough, B. D. (1998), "Assessing the Reliability of Statistical Software: Part I," *American Statistician* **52**, 358-366

McCullough, B. D. and Berry Wilson (1999), "On the Accuracy of Statistical Procedures in Microsoft Excel 97," *Computational Statistics and Data Analysis* **31**, 27-37

McCullough, B. D. and Berry Wilson (2001), "On the Accuracy of Statistical Procedures in Excel 2000 and Excel 2002," manuscript, Decision Sciences Dept., Drexel University

McCullough, B. D. and H. D. Vinod (1999), "The Numerical Reliability of Econometric Software," *Journal of Economic Literature* **37**, 633-665

Rotz, W., E. Falk, D. Wood and J. Mulrow (2001), "A Comparison of Random Number Generators Used in Business," mimeo, Ernst and Young, 1225 Conecicut Ave NW, Washington DC, 20036

Sincich, T., D. Levine and D. Stephan (2002), *Practical Statistics by Example Using Microsoft Excel*, Upper Saddle River, NJ: Prentice Hall

Sawitzki, G. (1994a), "Testing the Numerical Reliability of Data Analysis Systems," *Computational Statistics and Data Anaylsis* **18**, 269-286

Sawitzki, G. (1994b), "Report on the Numerical Reliabililty of Data Analysis Systems," *Computational Statistics and Data Analysis (Statistical Software Newsletter)* **18**, 289-301

Wilkinson, L. (1985), *Statistics Quiz: Problems which reveal deficiencies in statistical programs*, mimeo, SYSTAT

Wilkinson, L. (1994), "Practical Guidelines for Testing Statistical Software," in P. Dirschedl and R. Ostermann, eds., *Computational Statistics*, Berlin: Physica-Verlag, 111-124