

Impact of errors in operational spreadsheets[☆]

Stephen G. Powell^{*}, Kenneth R. Baker, Barry Lawson

Tuck School of Business, Dartmouth College, Hanover, NH 03755, USA

ARTICLE INFO

Article history:

Received 29 February 2008

Received in revised form 22 January 2009

Accepted 9 February 2009

Available online 21 February 2009

Keywords:

Spreadsheets

Error classification

Decision support

End-user computing

ABSTRACT

All users of spreadsheets struggle with the problem of errors. Errors are thought to be prevalent in spreadsheets, and in some instances they have cost organizations millions of dollars. In a previous study of 50 operational spreadsheets we found errors in 0.8% to 1.8% of all formula cells, depending on how errors are defined. In the current study we estimate the *quantitative impacts* of errors in 25 operational spreadsheets from five different organizations. Within these 25 spreadsheets we identified 381 potential errors, of which 117 (31%) were confirmed as errors by the developers of the spreadsheet. Among these confirmed errors, 47 (40%) had no quantitative impact on the results. Among the remaining 70 confirmed errors, the largest error was \$100 million; however, 9 of the 25 spreadsheets tested had no errors at all.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Spreadsheets are ubiquitous in organizations and some cannot function without them [3]. Given the importance of spreadsheets one might assume that these computer systems are created by professional programmers. However, in many organizations this is not the case. In fact, it is a reasonable conjecture that most spreadsheets are built by business experts, not computer experts.

Professional programmers understand the difficulties of creating error-free code and are trained to avoid errors. Most spreadsheet developers, being largely self-taught, are less aware of the dangers that errors pose. A widely cited estimate [7] is that 5% of all formulas contain errors, although recent research suggests that figure might be closer to 1% [11]. But the fundamental question is not how *common* errors are but how *significant the impacts* of errors are. After all, we make numerous errors in all facets of our lives, but rarely are those errors consequential [13].

The best way to determine the impact of errors in a spreadsheet would be to measure how the errors affect the decisions made using the spreadsheet. Although this may be the ideal measure, it would be extremely difficult to carry out in practice. In this paper we take a more limited but also more practical approach: we audit 25 operational spreadsheets, identify potential errors, confirm those errors with the spreadsheet developer, correct each error, and determine the quantitative impact on the outputs.

The paper is organized as follows. In the next section we summarize the relevant research, which relates to error taxonomies,

the frequency of errors, and their impacts. Then we describe our research: how we selected our sample of spreadsheets, how we identified errors, and how we measured the impact of errors. We then describe our results, first in general and then by organization and spreadsheet. The subsequent section offers a discussion of our results and a number of insights we gleaned. The paper concludes with a short summary and suggestions for future research.

2. Literature review

We recently summarized the research literature on spreadsheet errors [9]. That paper organized the literature into five areas:

- categories of errors
- impact of errors
- frequency of errors
- creation and prevention of errors
- detection of errors

Our current research focuses on the first three issues.

A number of different taxonomies of spreadsheet errors have been offered by previous researchers [12,14]. Most of these are not suitable for our purposes because auditors cannot distinguish errors in different categories. For example, we cannot tell whether an error in a formula is due to lack of knowledge about Excel or about the application domain. We summarized the existing literature on error taxonomies as follows:

- Classifications are offered without specifying the context or purpose for which the classification is intended.
- The existing classifications do not include sufficient examples of specific errors that satisfy each category.
- Classifications are not rigorously tested to demonstrate that multiple users can consistently classify actual errors into the proper categories.

[☆] This work was performed under the sponsorship of the U.S. Department of Commerce, National Institute of Standards and Technology. Reproduction of this article, with the customary credit to the source, is permitted.

^{*} Corresponding author.

E-mail address: sgp@dartmouth.edu (S.G. Powell).

- The boundary between quantitative errors and qualitative errors remains vague.

The categories of errors we used in this research are described in the following section.

A number of studies have attempted to measure the frequency of errors in spreadsheets. Panko [7] summarized the results of seven field audits in which operational spreadsheets were examined, typically by an outsider to the organization. His results showed that 94% of spreadsheets have errors and that the average cell error rate (the ratio of cells with errors to all cells with formulas) was 5.2%. However, in our analysis of these studies we concluded that these estimates were unreliable given that the studies were not sufficiently documented or comparable. A better estimate of the cell error rate based on these studies would have been 1.3% of cells, but only 13 spreadsheets would have been included in this average. Our own efforts to measure the incidence of errors [11] concluded that 0.8% to 1.8% of all formulas contained errors, depending on the definition of errors used. (Note that corrected cell and spreadsheet error rates now appear on Panko's website [8]).

Ironically, the least studied aspect of spreadsheet use is the quantitative impact that errors have on the results. What little evidence we have on the impact of errors in audited spreadsheets is largely anecdotal. For example, Hicks (cited in [7]) reported that the errors found in the audit of a single large spreadsheet caused the results to be off by 1.2%. Clermont [2] found errors in an average of 3.03% of cells in three large spreadsheets but reported that “we did not find any tremendous erroneous result values that might have had severe negative effects on the company.” Lukasic (personal communication cited in [7]) found a 16% error in the results of one of two spreadsheets audited. Panko's two interviewees [7] suggested that 5% of spreadsheets had “extremely serious” errors.

Caulkins et al. [1] interviewed 45 managers about spreadsheet usage in their organizations and 57% agreed with the statement “Spreadsheet errors are a significant threat to decisions.” This suggests that many managers are concerned about the effects of errors, although it does not help us estimate the risks quantitatively.

Another source of information on the impact of errors is the set of stories that have appeared in the press documenting losses due to mistakes involving spreadsheets. The European Spreadsheet Risks Interest Group (EUSPRIG) maintains a web page (<http://www.eusprig.org/stories.htm>) that documents dozens of these cases. Many of these incidents involve errors in formulas, but others involve errors in the use of spreadsheets, such as in sorting, or even errors in interpreting results. These types of errors, while potentially serious, are beyond the scope of the current research.

We can summarize the literature on the impacts of errors in spreadsheets as follows:

- No systematic studies have been reported that measure the quantitative impacts of errors in spreadsheets.
- Studies of the frequency of errors suggests around 1% of formulas are erroneous.
- Errors can be classified in many ways and no one classification has become accepted.

3. Research design

Our interest in this study was on the impact of errors in completed, operational spreadsheets, not errors made in a laboratory setting or errors made during the development of a spreadsheet. In a previous study [11] we detected errors in fifty spreadsheets using a formal auditing procedure, and we used an approach that did not require access to the developers. This approach allowed us to audit a large number of spreadsheets, but it also meant that we could not check our understanding of a model with the developer. In the current study, we worked with 25 spreadsheet developers within five different

organizations. The developers each completed a survey (see Appendix) describing their spreadsheet's purpose and design, as well as any unusual or special formulas or assumptions. Two researchers then independently audited each spreadsheet and pooled their results. Finally, we debriefed the developer on each issue our audit raised, and categorized each issue as (1) an error, (2) a poor practice, or (3) not an error. Confirmed errors were then corrected, and the change in the output cell was recorded as the quantitative measure of the impact of the error.

3.1. Sample spreadsheets

We first identified five organizations that were willing to provide volunteers and to have their spreadsheets audited. This included two consulting companies, a large financial services firm, a manufacturing company, and an educational institution. These organizations identified five volunteers, each of whom provided one spreadsheet for auditing. We provided the following specifications to the volunteers for help in choosing a spreadsheet to audit:

- contains 3–10 worksheets
- contains 250–10,000 formulas
- occupies 200–1000 kb of memory
- has been in use for no more than a year
- contains no complex Visual Basic code
- is well understood by the developer
- has no broken external links

Not all the files we audited conform to these specifications. In fact, the average number of worksheets in our sample was 15.2 (the range was from 2 to 135 sheets) and the average size in kilobytes was 1463 (the range was 45 to 7450 kb). Many of the spreadsheets in our sample were larger on one or more of the dimensions than specified above.

While our sample of spreadsheets is not strictly *random*, it is certainly *representative* of the general population of spreadsheets (with the caveats cited above). The sample includes spreadsheets from different types of organizations, spreadsheets created by developers with different degrees of expertise, and spreadsheets that span a broad range from small and simple to large and complex. As our results show, the quality of these spreadsheets ranges from very high to very low.

3.2. Error taxonomy

One of the challenges of spreadsheet error research is how to categorize errors. As we pointed out earlier, many different error classifications have been offered. Most of these suffer from the same flaw: errors that arise from different causes cannot be distinguished by an auditor. For example, when we encounter an error in a formula we can rarely determine whether the error was due to sloppy typing, lack of domain knowledge, lack of Excel knowledge, a change made by a subsequent user, or an unknown cause. We can, however, easily detect some formulas that give the wrong result. We can also identify many practices that are likely to cause errors as the spreadsheet is used or that will simply make it harder than necessary to use the spreadsheet productively.

After considerable testing we settled on the following six error types that our experience with auditing suggested were well-defined in theory and could be identified with high reliability in practice (more information on this error taxonomy is available in [10]):

- logic error — a formula is used incorrectly, leading to an incorrect result
- reference error — a formula contains one or more incorrect references to other cells
- hard-coding numbers in a formula — one or more numbers appear in formulas and the practice is sufficiently dangerous

- copy/paste error — a formula is wrong due to inaccurate use of copy/paste
- data input error — an incorrect data input is used
- omission error — a formula is wrong because one or more of its input cells is blank.

We should point out that hard coding was identified in our previous study as the most common poor practice. In the present study we ignored hard coding in most instances, on the grounds that it was unlikely to represent an outright quantitative error. However, there were a few instances in which contradictory inputs were hard-coded, and we did cite those as potential errors.

3.3. Auditing protocol

The auditing protocol we developed for our previous study is a highly-detailed document that specifies the steps to take in auditing a spreadsheet for size, complexity, several types of qualitative features, and errors. (More information on the protocol is available in [10]. A complete description of the protocol itself is available at <http://mba.tuck.dartmouth.edu/spreadsheet/index.html>.)

This protocol evolved over several months of testing. During this time we trained auditors and tested the protocol ourselves and through our assistants on dozens of operational spreadsheets. Our current study used a very similar approach but focused less on gathering data about the spreadsheets and more on finding potential errors. The typical approach was to review the survey provided by the developer, especially the portion describing the various worksheets and their interrelationships. The next step was to run the auditing software *Spreadsheet Professional* (<http://www.spreadsheetinnovations.com>), which provides summary maps for each sheet and reports on the location of potential errors. We then examined each sheet in turn, first looking at the map for suspicious cells or ranges and scanning the reports on potential errors. Then we inspected the sheet itself, first determining the location of data and formulas and subsequently auditing every unique formula and most copied formulas.

In this study we recorded data only on cells that were potentially erroneous. For each problematic cell or range we recorded the following information:

- location: cell address or range
- type of error
- how it was discovered (whether by map analysis, error tests, code inspection, or sensitivity testing)
- description of the possible error

After we had discussed the potential error with the developer, we then recorded how the issue was resolved (Error, Poor Practice, No Error). For Errors, we also recorded the cell used to measure the impact and the absolute and percentage changes in that cell when the error was corrected.

3.4. Measuring impacts

Determining the quantitative impact of an error in a spreadsheet is not as straightforward as it might appear. First, some errors occupy a single cell while others occupy many cells. Do we consider each cell as a separate error and measure the impact of correcting it alone, or do we correct all the cells with a similar error and measure the overall impact? Second, some error cells influence many other cells while others impact no other cells. When a cell impacts many other cells it is not always obvious which of the impacted cells to use to measure the effect. Third, it is not always clear how to correct an error. For example, if erroneous inputs were used do we replace them with average inputs or extreme inputs? Finally, it is necessary to decide whether to measure errors in absolute or relative terms, and whether to combine

all the errors in a given workbook into one overall error or to treat them separately.

In this study we chose to measure the effect of each error separately and then to report the *maximum* error in a workbook. In most cases we corrected all the cells with a given type of error, treating this as one error with a single (overall) impact. When such an error impacted only the erroneous cells themselves, we computed the *maximum* change from the base case and took that as our error estimate. When such an error impacted a single dependent cell, we measured the impact of correcting all the error cells on that one cell. When an error cell had many dependent cells, we identified the most important dependent cell and measured the impact of correcting the error on that cell. Obviously, the magnitude of the error impacts we identified are somewhat dependent on the measurement system used, although we believe our choices are justifiable and our results fairly reflect the quantitative accuracy of the workbooks tested.

4. Results

Table 1 summarizes our results. In column 1 we have used a two-digit code to label each spreadsheet. For example, spreadsheet 3.4 is the fourth spreadsheet from organization 3. The table gives the following information for each spreadsheet:

- number of issues raised in our audits
- number of errors confirmed in interviews
- number of errors with non-zero quantitative impact
- maximum percentage impact
- maximum absolute impact

Within this sample of 25 spreadsheets we identified a total of 381 issues. After we discussed these issues with the developers we found that nine spreadsheets had no errors; among the remaining 16 spreadsheets we found a total of 117 errors. Of these 117 errors, 47 had zero quantitative impact, leaving 70 errors with non-zero impact.

Many different kinds of issues turned out not to be errors after we discussed them with the spreadsheet developers. For example, if we found one odd formula in a column of identical formulas we flagged it as an issue. Typically the developers would then tell us that this case

Table 1
Errors and impacts in 25 spreadsheets.

Organization-workbook	Size (kb)	#Issues	#Errors	Errors with no impact	Maximum percentage impact	Maximum absolute impact
1.1	842	7	3	3	0.0%	\$0
1.2	3291	50	6	1	28.8%	\$32,105,400
1.3	3556	18	7	4	137.5%	\$110,543,305
1.4	1376	4	1	1	0.0%	\$0
1.5	1678	0	0	0	NA	NA
2.1	3180	19	6	1	3.6%	\$13,909,000
2.2	520	27	11	4	16.0%	\$74,000,000
2.3	394	6	0	0	NA	NA
2.4	156	30	4	1	416.5%	\$10,650,000
2.5	1734	40	2	0	NA	\$0
3.1	389	19	2	0	5.3%	\$238,720
3.2	177	1	1	1	0.0%	\$0
3.3	47	11	2	0	15.6%	\$4,930,000
3.4	990	6	1	1	0.0%	\$0
3.5	286	23	1	1	0.0%	\$0
4.1	606	27	22	10	116.7%	\$13,355,445
4.2	576	8	4	2	141.8%	\$272,000
4.3	1571	0	0	0	NA	NA
4.4	2116	1	0	0	NA	NA
4.5	1932	79	44	17	39.1%	\$216,806
5.1	7450	2	0	0	NA	NA
5.2	415	2	0	0	NA	NA
5.3	1900	0	0	0	NA	NA
5.4	309	0	0	0	NA	NA
5.5	427	1	0	0	NA	NA
Totals		381	117	47		

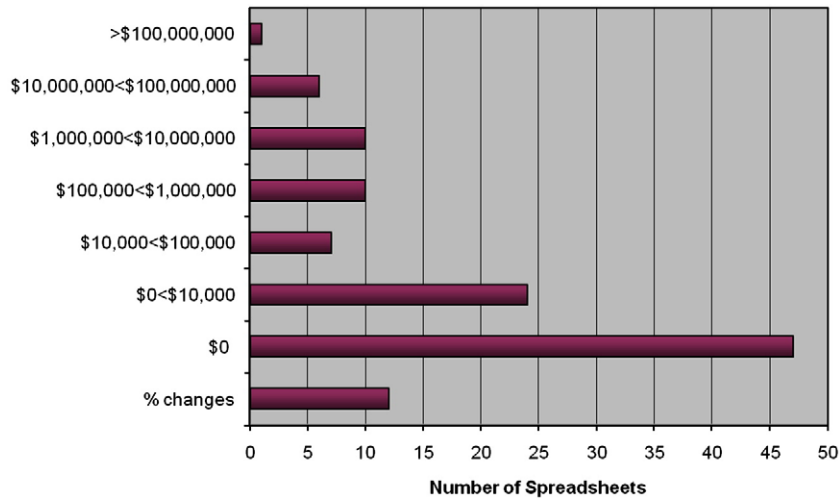


Fig. 1. Absolute impact of errors.

was calculated differently from all the rest and that it was correct, although there was no such indication in the spreadsheet itself. Most of these cases, in fact, involved inconsistent formulas in a column or row.

Errors with zero quantitative impact also arose in a variety of ways. Sometimes a formula referred to the wrong cell for an input, but the numerical value in that cell was the same as in the correct input cell. We categorized this case as an error, but one with zero impact. We recognize that this is a dangerous practice, which could lead to a nonzero error during the use of the spreadsheet. However, for the purposes of this study we categorized it as an error with zero impact.

As we pointed out above, there are two ways to measure the impact of errors: absolute and relative. Absolute impacts are important because they tell us how large errors are in the units of the spreadsheet. However, they cannot be compared easily across workbooks, since a million dollar error may be trivial in one spreadsheet and catastrophic in another. Relative (or percentage) errors more accurately reflect the significance of an error, but they have their shortcomings as well. One problem with relative errors is that percentage changes cannot be determined when the initial value is zero; another is that percentage changes in percentages are not generally as meaningful as percentage changes in dollar amounts. Accordingly, we present our results here in both absolute and relative terms.

Fig. 1 shows the distribution of absolute error impacts. The most salient point to draw from this figure is that 47 of the errors we found had zero impact on the spreadsheet. This often came about when a formula had an erroneous reference, but both the erroneous and the correct input cells had the same value. Thus when the error was fixed the results did not change.

Returning to Fig. 1, we see that 12 of the errors involved percentages; among these the average absolute change was 22%. Twenty-four of the remaining 58 errors involved absolute errors less than \$10,000. However, some errors were huge: the largest single absolute error we found was over \$100 million!

Fig. 2 shows the distribution of percentage error impacts. (The N/A category includes four errors in cells with an initial value of zero, for which a percentage change is not defined.) Of course, 47 of the 117 errors had zero relative impact. Fifty of the 66 remaining errors were less than 10% of the initial value. As with absolute errors, there are some very large errors: four, in fact, were over 100%.

Our evidence suggests that spreadsheet practice is very different among the five organizations we studied. Fig. 3 shows the distribution of issues, confirmed errors, and errors with non-zero impact in the five organizations we studied. In the five spreadsheets from Organization 5 we could identify only five issues to discuss with the developers and no errors were identified among those five issues. Organization 5 is a small consulting company with highly educated employees and a

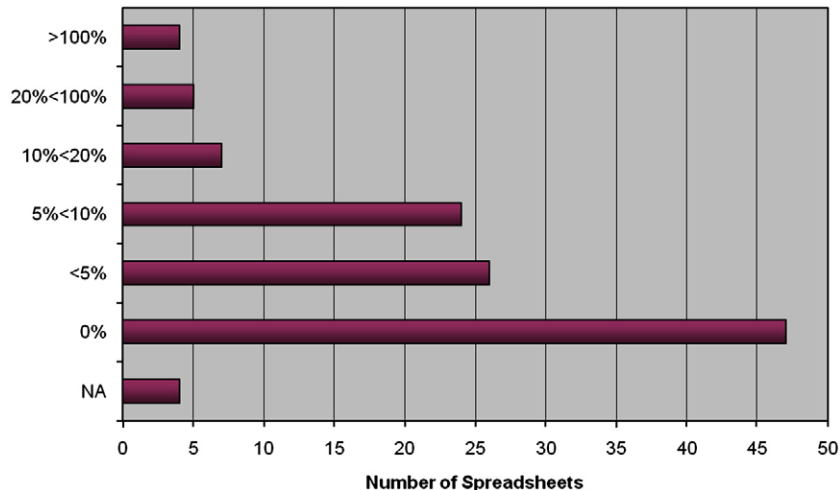


Fig. 2. Relative impact of errors.




Fig. 3. Issues, errors, and errors with impact by organization.

culture that demands excellence. The spreadsheets we audited from this firm were works of art: thoughtfully designed, well documented, easy to understand, and error free.

Organization 4 had two spreadsheets with no errors and two with 22 and 44 errors, respectively. The quality of the spreadsheet practice in this organization clearly depends on just where one looks. In this case we found both the best of practice and the worst of practice in offices just a few miles apart.

In Organization 3, which is another consulting company, all the spreadsheets we audited had errors but in three cases no error had a measurable impact on the results. Even in the remaining two spreadsheets the errors were few in number and fairly small in terms of impact.

Organizations 1 and 2 are both very large. One is a financial firm and the other is a manufacturing firm. Some of the spreadsheets we audited from these companies were astonishingly large and complex. Perhaps for this reason, only two of the ten we audited were error-free (although four had no errors with impact). The quality of spreadsheet practice in both of these companies was inconsistent, with inadequate attention paid to spreadsheet design, simplicity, ease of use, documentation, and consistency.

We can summarize our findings as follows:

- Some organizations have spreadsheets that are essentially error-free.
- Within a single organization, spreadsheet practice can range from excellent to poor.
- Some organizations use spreadsheets that are rife with errors and some of these errors are of substantial magnitude.
- Many errors have zero impact, or impact unimportant calculations.
- There is little apparent correlation between the importance of the application or the risk involved and the quality of the spreadsheet.
- Few spreadsheets contain errors that, in the eyes of their developers, would destroy their usefulness.

5. Discussion

Many writers have observed that the spreadsheet, for all its attractiveness to end-users, is in some ways a dangerous programming platform. Not only are the logic of the model and the numbers commingled, but the physical layout permits unstructured design. It is not surprising that amateur programmers, who lack structured design methods, make errors when using such free-form software [4,6].

Our research has shown that errors come in more varieties than perhaps even the most extensive taxonomy can encompass. Because

the spreadsheet platform is so unstructured, and because end-users generally follow unique designs, errors and poor practices can arise in thousands of different guises. Error researchers must inevitably use their judgment in deciding what is an error and what is not. Thus we should be skeptical of claims about the frequency and impact of errors based on rough averages and casual research.

Another general observation that our research supports is that spreadsheet auditing is more difficult and limited than might have been anticipated. We knew in advance that we would not be able to identify errors in problem formulation or in the use of the model by auditing the spreadsheet itself, although these types of errors may be the most consequential. But even within the narrower domain of our audits we encountered limitations. First, the data used in most spreadsheets is undocumented and there is no practical way to check it. Even the original developer would have difficulty checking the data. Second, many spreadsheets are so chaotically designed that auditing (especially of a few formulas) is extremely difficult or impossible. Finally, we have found that many spreadsheets do not have just a few key outputs but are used to calculate hundreds or thousands of results. This makes it difficult to measure the impact of a particular error unambiguously.

One important generalization our work supports is that many errors are benign: they either have no impact on the results, the quantitative impact is very small, or the effect is on a vestigial portion of the spreadsheet that is no longer of interest. One can conjecture that this is the result of a sensible attitude toward errors on the part of spreadsheet developers. Perhaps developers look out for errors that impact the key outputs, and are generally good at correcting them. However, they pay less attention to inconsequential errors and therefore more of these survive to be observed. And, as we know from our interviews, many developers do not clean up their spreadsheets before they move on to other tasks.

One factor that might explain the substantial differences within and among companies in the quality of their spreadsheets is the degree of risk involved. We might hypothesize that companies devote their best resources to high-risk spreadsheets, and fewer resources to low-risk ones, although in our study we did not observe this. (There is some evidence in a user survey we conducted that would support this conjecture [5].) We did not measure this feature of the spreadsheets we audited (which would certainly be difficult to do), but our impression is that if such a correlation exists at all it is weak. For example, one of the best designed spreadsheets we audited was used to help with daily staffing of nurses and doctors in a medical practice. The spreadsheet was elegantly engineered, error-free, and easy to use.

But little was at risk: an error in this spreadsheet at worst would assign the same person to adjacent shifts or to too many consecutive days. Errors of this type would almost certainly be caught, and their impact would be negligible. Nevertheless, the spreadsheet was nearly perfect. By contrast, we also audited spreadsheets in use at a major financial firm for calculating tax liabilities (measured in the billions) to various state and national entities. These spreadsheets were astonishingly complex, difficult to understand, difficult to work with, and error-prone. So factors other than risk appear to explain spreadsheet quality.

Another observation that helps to understand our results is that many of the developers we worked with were not especially surprised or devastated when we pointed out potential errors. Sometimes the reaction was that they knew the formula “wasn’t quite right,” but they saw that it gave the right *answer* and thus was acceptable. Sometimes the reaction was that the result was “close enough,” or that the result in question was no longer used, or not important. So developers seem to have a sense of what level of accuracy is appropriate for a given spreadsheet. (It is another question entirely as to whether their perceptions are correct, and whether the spreadsheets are actually as accurate as they need to be.)

An experienced auditor can rather quickly detect a spreadsheet that is likely to have errors. The major symptoms we observed of poor spreadsheet practice are the following:

- chaotic design
- embedded numbers
- special cases
- non-repeating structures
- complex formulas

Chaotic design refers to a poorly structured physical layout of the formulas and data. Numbers embedded in formulas, while not necessarily direct causes of errors, are strongly correlated with the presence of other problems. Special cases refers to designs in which similar results are calculated in slightly different ways, which requires great care in building and checking formulas. Non-repeating structures includes designs in which the formulas in a row or column change structure repeatedly, precluding the use of copying and pasting. In the hands of experts, complex formulas can be used to great effect. But in the hands of novices the same formulas can be error-prone.

Why is spreadsheet practice sometimes so poor? We cannot know for sure, but we did gather some anecdotal evidence during our interviews. When asked what kept them from building better spreadsheets, our developers typically cited one or more of the following reasons:

- time pressure
- unstructured design
- changing specifications
- lack of testing
- lack of relevant knowledge and skills

Time pressure was the most often cited reason for poor spreadsheet practice. Many spreadsheets are built under great time pressure, which precludes use of some of the most effective methods for avoiding errors. Managers of spreadsheet developers should be aware of the effects of putting their employees under excessive time pressure. Another commonly cited factor was unstructured design: either the spreadsheet design was inherited from a poorly designed predecessor spreadsheet, or it grew organically during the project without ever consciously being designed. Another complaint was changing requirements: if the designer had only known from the start what the spreadsheet was going to be used for, he or she could have designed it more appropriately. We also observed that very few of our developers used any formal approach to testing their spreadsheets; in fact, most of them did no testing at all. Finally, in some cases we could

observe directly that the cause of an error was lack of relevant knowledge, either of the problem domain or of spreadsheet tools. However, it was remarkable how rare this cause appeared to be. Most developers could see quickly that a particular formula was an error, once we pointed it out to them. Only very rarely did we have to explain to them why it was an error, or how to fix it.

Finally, we offer a comment on the importance of good spreadsheet design both to auditing, and to good spreadsheet usage in general. Our work makes us extremely conscious that a well-designed spreadsheet is *simple*, *consistent*, and *general*. Simplicity means a logical use of worksheets and a logical and intuitive layout of each individual sheet. Simplicity makes building and auditing easier. Consistency means, for example, that a single formula can be copied down an entire row or across an entire column. Such a formula can easily be checked. Rows or columns in which the formulas change structure constantly often hide errors. Generality means that the spreadsheet is built to handle all of the likely combinations of inputs that users will want to use. The opposite is a workbook in which individual cases are calculated separately, which makes it difficult to keep inputs consistent across cases.

6. Summary and future research

We audited 25 operational spreadsheets from five different organizations. We identified 381 potential errors, of which 117 were confirmed as errors by the spreadsheet developer and 70 had a non-zero impact on the output. Among these remaining errors, seven exceeded \$10 million.

Our research suggests, but does not prove, a number of generalizations about the quality of spreadsheets in organizations. First, it appears that some organizations build spreadsheets that are essentially error-free. In others, however, spreadsheet practice can range from excellent to poor. Clearly, some organizations use spreadsheets that are rife with errors and some of these errors are of substantial magnitude. On the other hand, many errors have zero impact, or impact unimportant calculations. There appears to be little correlation between the importance of the application or the risk involved and the quality of the spreadsheet. Finally, few spreadsheets contain errors that, in the eyes of their developers, would destroy their usefulness.

This research can be extended in several directions. Additional studies like this one that cut across several organizations would help to highlight differences in practice. Of particular interest would be studies that attempted to correlate individual or corporate practices around spreadsheets with the quality of the results. For example, do organizations that train their spreadsheet developers in auditing have fewer and less consequential errors? Studies that go into more depth in a single organization would also be helpful. Studies of the overall decision process, of which spreadsheets are a part, would be especially powerful in revealing how spreadsheets, and particularly spreadsheet errors, influence decision quality.

Appendix. Spreadsheet Audit Information Form

The purpose of this survey is to record essential information on a workbook to be audited as part of the Tuck Spreadsheet Engineering Research Project auditing research.

1. Please describe the developer of the spreadsheet.
 - a. name
 - b. company name
 - c. email address
 - d. physical address
2. Please describe the spreadsheet.
 - a. filename
 - b. number of worksheets
 - c. number of kb

3. Please describe the business problem that is addressed with this spreadsheet.
4. Please describe the purpose of the spreadsheet; e.g. it forms the basis for an investment decision, it is used for cash flow analysis, etc.
5. Please list the key outputs of the spreadsheet.
6. Please list the key inputs of the spreadsheet.
7. Who will use the spreadsheet? Please check all that apply.
 - Developer only
 - Limited number of users who receive instructions from the developer
 - Limited number of users who receive the spreadsheet without instructions from the developer
 - Unlimited number of users
 - Other (please specify)
8. How will the spreadsheet be used? Please check all that apply
 - For one decision only
 - For repeated decisions
 - For continuous updating
 - Will be expanded later
 - Sensitivity analysis of key inputs is planned
 - Other (please specify).
9. Over what period of time will the spreadsheet be used? Please provide your best estimate.
10. Please check all that apply.
 - Spreadsheet uses macros or user-defined functions
 - Spreadsheet has external links to other workbooks
 - Spreadsheet has protected cell or worksheets
 - Spreadsheet uses Data Validation or Scenario Manager
 - Spreadsheet contains hidden worksheets or cells (please list names of hidden worksheets)
11. Please list each worksheet, its purpose, and the sheets it takes data from.
12. Please describe any complex calculations or calculations that use unusual business logic.

References

- [1] J. Caulkins, E. Morrison, T. Weidemann, Spreadsheet errors and decision making: evidence from field interviews, *Journal of End User Computing* 19 (3) (2006) 1–23.
- [2] M. Clermont, A Spreadsheet Auditing Tool evaluated in an industrial context, *Proceedings of the European Spreadsheet Risks Interest Group Conference Cardiff Wales, 2002*, pp. 35–46.
- [3] G. Croll, The importance and criticality of spreadsheets in the city of London, Paper delivered at European Spreadsheet Risks Interest Group Conference, Greenwich, UK, 2005.
- [4] J. Kreie, T. Cronin, J. Pendley, J. Renwick, Applications development by end-users: can quality be improved? *Decision Support Systems* 29 (2000) 143–152.
- [5] B. Lawson, K. Baker, S. Powell, and L. Foster-Johnson, A Comparison of Spreadsheet Users with Different Levels of Expertise, *Omega* 37 (2009) 579–590.
- [6] T. McGill, J. Klobas, The role of spreadsheet knowledge in user-developed application success, *Decision Support Systems* 39 (2005) 355–369.
- [7] R. Panko, What We Know About Spreadsheet Errors, 2006 <http://panko.cba.hawaii.edu/ssr/My Papers/whatknow.htm>, accessed September 2.
- [8] R. Panko, What We Know About Spreadsheet Errors, 2008 <http://panko.shidler.hawaii.edu/SSR/My Papers/whatknow.htm>, accessed August 11.
- [9] S. Powell, K. Baker, and B. Lawson, A Critical Review of the Literature on Spreadsheet Errors, *Decision Support Systems* 46 (2008) 128–138.
- [10] S. Powell, K. Baker, B. Lawson, An Auditing Protocol for Spreadsheet Models, *Information & Management* 45 (5) (2008) 312–320.
- [11] S. Powell, K. Baker, and B. Lawson, Errors in Operational Spreadsheets, *Journal of End User and Organizational Computing* (forthcoming).
- [12] M. Purser, D. Chadwick, Does an awareness of differing types of spreadsheet errors aid end-users in identifying spreadsheet errors? *Proceedings of the European Spreadsheet Risk Interest Group Annual Conference, Cambridge, UK, 2006*, pp. 185–204.
- [13] J. Reason, *Human Error*, Cambridge University Press, Cambridge, UK, 1990.
- [14] K. Rajalingham, D. Chadwick, B. Knight, Classification of spreadsheet errors, *Proceedings of the European Spreadsheet Risks Interest Group Annual Conference, Greenwich, England, 2000*, pp. 23–34.

Steve Powell is a Professor at the Tuck School of Business at Dartmouth. His primary research interest lies in modeling production and services processes, but he has also been active in research in energy economics, marketing, and operations. At Tuck, he has developed a variety of courses in management science, including the core Decision Science course and electives in the Art of Modeling, Business Process Redesign, and Applications of Simulation. He originated the Teacher's Forum column in *Interfaces*, and has written a number of articles on teaching modeling to practitioners. He is the academic director of the INFORMS Annual Teaching of Management Science Workshop. In 2001 he was awarded the INFORMS Prize for the Teaching of Operations Research/Management Science Practice. He is the co-author with Kenneth Baker of *The Art of Modeling with Spreadsheets* (Wiley, 2009).

Ken Baker is a faculty member at Dartmouth College. He is currently Nathaniel Leverone Professor of Management at the Tuck School of Business and also adjunct professor at the Thayer School of Engineering. At Dartmouth, he has taught courses relating to decision science, manufacturing management, and environmental management. Over the years, much of his teaching and research has dealt with production planning and control, and he is widely known for his textbook *Elements of Sequencing and Scheduling*, in addition to a variety of technical articles. He has served as the Tuck School's associate dean and directed the Tuck School's management development programs in the manufacturing area. In 2001 he was named a Fellow of INFORMS's Manufacturing and Service Operations Management (MSOM) Society, and in 2004 a Fellow of INFORMS. He is the co-author with Stephen Powell of *The Art of Modeling with Spreadsheets* (Wiley, 2009).

Barry Lawson is a research associate at the Tuck School of Business at Dartmouth and is also a visiting scholar in the geography department of the college. He founded and has served as president of Barry Lawson Associates, a consulting firm, since 1978. As visiting scholar, he coordinates the development of an atlas of the upper Connecticut River Watershed in New Hampshire and Vermont. As research associate at Tuck he serves as the program manager for the Tuck Spreadsheet Engineering Research Project. Lawson has taught in graduate programs at Boston University and Wayne State University as well as in short courses at Bentley College. He has moderated a host of public hearings for local, state and federal governments on controversial environmental and energy- and waste-related projects, and has considerable experience in group facilitation, conflict resolution and simulation design.