Interfaces
Vol. 40, No. 6, November–December 2010, pp. 454–464
ISSN 0092-2102 | EISSN 1526-551X | 10 | 4006 | 0454

informs

DOI 10.1287/inte.1100.0527

Assessing What Distinguishes Highly Cited from Less-Cited Papers Published in *Interfaces*

Thomas A. Hamrick, Ronald D. Fricker Jr., Gerald G. Brown
Naval Postgraduate School, Monterey, California 93943
{tahamric@nps.edu, rdfricker@nps.edu, gbrown@nps.edu}

We evaluate what distinguishes a highly cited *Interfaces* paper from other *Interfaces* papers that are cited less often. Citations are used to acknowledge prior relevant research, to document sources of information, and to substantiate claims. As such, citations play a key role in the evolution of knowledge. More recently, citations are also being used to quantify the impact of papers and journals, a practice not without controversy, but one that motivates our work here. We find that Edelman competition papers, longer papers, tutorials, papers with larger numbers of references to prior literature, and papers with a larger number of "callouts" (a feature no longer used by *Interfaces*) tend to have a higher number of citations.

Key words: citation; bibliometrics, impact; impact factor.

Not all citations are created equal, yet they are counted equally by various *impact metrics* (syn. factors). These impact metrics are used for an increasing number of purposes, including tenure evaluation by many universities, promotions and awards throughout the scientific community, assessment and ranking of journals such as *Interfaces*, and even comparison of performance of entire schools and laboratories.

BusinessWeek (2010) calculates an "intellectual capital score" for business schools based on the number of articles published by faculty in 20 journals over the last five years and book reviews published in *The New York Times*, *The Wall Street Journal*, and *BusinessWeek*. Four of the 12 INFORMS journals are on this list.

Wally Hopp, former editor of *Management Science*, says article submissions doubled following *BusinessWeek*'s recognition, although he is mute about submission quality (Hopp 2010). There is some question of causal precedence, because *BusinessWeek* polls the schools to see which journals should be counted toward its own biyearly ranking.

Our purpose in writing this paper is to determine what properties of *Interfaces* papers have led to more citations, and thus to higher impact scores for this journal.

Edelman papers epitomize *Interfaces'* mission and are subsequently cited more often.

Background

Modern, formal use of citations in scientific literature dates back only to the nineteenth century—with earlier inconsistent use—as scholars and scientists learned to give continuity to their body of ideas (e.g., Nicolaisen 2007). Under the rubric of *bibliometrics*, citation counts have been incorporated into metrics intended to measure the impact of researchers, papers, and journals, a practice not without controversy. As David Kelton remarks, "A paper that is incorrect can lead to a great impact factor" (Kelton 2010).

Bornmann and Daniel (2006) state that Gross and Gross (1927) were the first to use citation counts to evaluate the importance of scientific work. In 1961, Eugene Garfield created the science citation index, both to provide researchers with "quick, powerful

access to the bibliographic and citation information they need to find relevant, comprehensive research data" (Thomson Reuters 2010a), and to more objectively quantify the impact of journals. Of course, "objective" is in the eye of the beholder, but the notion of standardizing impact metrics for consistency across journals and papers has merit. As Garfield (2006, p. 92) states, "Some editors would calculate impact solely on the basis of their most-cited papers so as to diminish their otherwise low impact factors."

Citations are certainly not the perfect measure by which to judge a paper because the quantity of citations is but a crude proxy for quality. This follows because papers are cited for a variety of reasons, not all of which stem from a paper's research quality or contribution.

Researchers interested in measuring scientific impact are divided. On one side are those who believe bibliometric analyses based on citation counts are useful for assessing impact because "a substantial body of literature has shown that the number of citations to scientists' publications are [sic] correlated with other assessments of scientists' impact or influence" (Bornmann and Daniels 2006, p. 46). On the other side are those who say a paper or scientist is cited for a variety of factors, not all of which are related to the scientific merit of a piece of research. These two positions are not mutually exclusive; we ignore any controversy between them because it does not affect our analyses.

Bornmann and Daniel (2006) review the literature on citing behavior and propose the following classifications of citation use: affirmational, assumptive, conceptual, contrastive, methodological, negational, perfunctory, or persuasive. Although a discussion of the details of such classifications is beyond the scope of this paper, the relevant point is that papers are cited for a variety of reasons and not all are related to quality, relevance, or impact. For additional discussion of citation behavior and related issues, see Camacho-Miñano and Múñez-Nickel (2009), Nicolaisen (2007), Bornmann and Daniel (2006), Ahmed et al. (2004), Case and Higgins (2000), Cronin (1998), Leydesdorff (1998), White and Wang (1997), and Oppenheim and Renn (1978).

Impact factors and researcher indices are functions of the number of citations that researchers, papers, or journals have received over time. They are often used as measures of the relative importance of a researcher, paper, or journal within a particular field; those with higher impact factors are deemed to have more influence or to be of more relevance than those with lower ones. At the journal level, Thomson Reuters calculates the *journal impact factor* (JIF) as the average number of citations per article over the previous two years (Thomson Reuters 2010b, Garfield 1994). Other indexes assess the median age of citations used in papers or journals or the number of citations over a certain period after publication. At the researcher level, the *h-index* (Hirsch 2005) is based on a researcher's most frequently cited papers and is intended to simultaneously account for quality of scientific publication and a broad publication record. The *g-index* (Egghe 2006) is an alternate method that embellishes the *h-index* for ranking researchers based on their global publication record. Other indices include Zhang's *e-index*, the *age-weighted citation rate*, and numerous other variations of the *h-index*.

The use of citations to measure impact is not restricted to citation counts in journals. O'Leary (2009) used citations of journals in patents as a way to measure journal impact and the value and relevance of particular business school disciplines to business. Numerous lists of the most highly cited papers in various disciplines have been (and continue to be) published, including Frogel (2010) for astronomy, Patsopoulos et al. (2006) for medicine, Ryan and Woodall (2005) for statistics, and Shapiro (1985) for law. Ryan and Woodall (2005) suggest some characteristics of the most frequently cited papers in a family of journals, rather than one as we do here. Fernandez-Alles and Ramos-Rodriquez (2009) consider a similar question to our own in an analysis of the journal, *Human Resource Management*, and find that the articles cited most often tend to have a macroscopic approach centered on an analysis of an organization's performance.

Impact factors and indices are subject to all the criticisms and weaknesses of any metric or statistic, particularly when they are used to rank or are used improperly, over interpreted, or over applied. Adler et al. (2009) offer an excellent exposition of these issues. We agree with and echo their concern that excessive reliance on the use of impact factors, and

bibliometrics in general, to rank research programs, journals, and scientists oversimplifies such assessments and can be erosive.

In this paper, we consider citation counts as our sole metric and evaluate what distinguishes a highly cited *Interfaces* paper from other *Interfaces* papers that are cited less often. We avoid the controversy about bibliometrics in that we do not suggest that the highly cited papers are better than those that are not highly cited. Rather, our interest is in exploring whether there is an association between a paper's characteristics and whether or not it is highly cited. Of course, some characteristics may be directly related to a paper's scientific merit or impact. However, others may simply be due to other qualities, such as a paper's readability, accessibility to a wider audience, type of subject matter, or even length.

Interfaces is published bimonthly. It was started in 1970 by The Institute of Management Sciences (TIMS) and is now published by INFORMS. Norden (1970) describes the genesis for the journal as stemming from a worthwhile body of applied operations research that might not meet the requirements of classical refereeing, but is still worthy of publication for managers and practitioners. Norden (p. 1) states that *Interfaces* "will publish digests and highlights of current, interesting, and useful work." Forty years later, we feel *Interfaces* still gives OR practitioners a place to publish interesting and helpful articles to better the policies and decisions of practitioners. Although the style of the journal has changed slightly over the years, its focus and scope have prevailed.

The Data

We assembled a database of 20 years of *Interfaces* papers from 1989–2008 using the electronic tables of contents on the INFORMS website (http://pubsonline.informs.org/). The data include the title of each paper, key words, lead author and affiliation, total number of authors, year of publication, volume and issue number, starting and ending page numbers, type of paper, and whether or not the paper was a finalist in the Edelman competition. The database focuses on *Interfaces* papers, columns, and editorials, and does not include practice abstracts or book reviews. Our database contains 1,133 papers: 1,005 articles, 99 columns, and 29 editorials (see Figure 1).

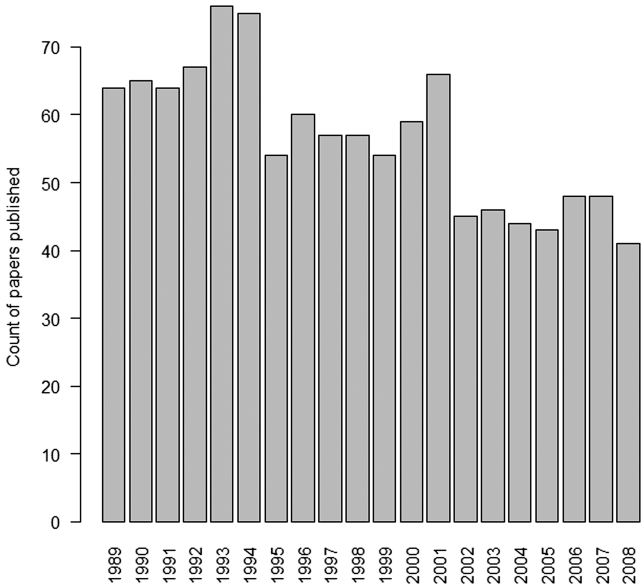

Figure 1: The graph shows the distribution of the number of papers included in our study by year of publication from 1989–2008.

Onto this list, we merge the citation counts that we downloaded from the ISI Web of Knowledge (WoK) website (http://apps.isiknowledge.com/) on August 15, 2009. By using WoK, we restrict our focus to citations from traditional and largely scholarly sources. WoK has its flaws; for example, it does not index the INFORMS journal *Decision Analysis* (a discovery we conveyed to WoK). However, it has the advantage of ensuring that the counts are based on unique citations. Other electronic tools that use Web searches, such as Google Scholar (http://scholar.google.com) or Publish or Perish (http//www.harzing.com/pop.htm), frequently contain duplicate citations resulting from a source being posted (or, at least, accessed) on the Web more than once. Furthermore, we find some of the Web sources that these tools count are less than either scholarly or authoritative.

As more material moves to the Web, including the newer online journals, one might argue that the use of WoK misses relevant citations. However, our goal is to gain insight into what separates highly cited *Interfaces* papers from less frequently cited ones; therefore, the precise quantification of citations is of less relevance than an authoritative and reliable method for classifying the papers into those with large numbers

of citations from those with fewer numbers. In addition, given that we are looking back over 20 years to a period before the widespread use of the Web, WoK provides a more consistent citation-count procedure for comparing the older and newer papers.

Downloading the citation counts is relatively simple using WoK; all one needs to do is enter the journal title and choose some dates. Merging the data, on the other hand, is not as simple because of the large variations in paper titles and author names between the INFORMS tables of contents and the WoK data. Ultimately, we chose to merge the two data sets based on a unique key for each paper; the key consists of volume number, issue number, beginning page number, and ending page number. This key was sufficient to allow us to successfully merge 90 percent of our data. Because of various errors and mismatches, such as page-number errors or papers split into multiple parts, we had to manually resolve the remainder prior to doing the merge. When we finished, the two databases matched, with the exception of eight papers that are not listed in WoK. For these we used Google Scholar to retrieve citation counts where, as of September 2009, six papers had zero citations and two had one citation each.

There is a lesson here for current authors and editors: getting citations right is vital to maintaining the continuity of our tree of knowledge.

For the 1,005 articles, 99 columns, and 29 editorials from 1989–2008, the mean number of citations is 6.6. Just under 25 percent of all the articles, columns, and editorials have zero citations, another 14 percent have one citation; and 11 percent are cited twice. Thus, slightly less than 50 percent of the papers are cited two or fewer times. Across all the papers, the median citation count is three and the 75th percentile is seven. By contrast, the six most often cited papers each have 100 or more citations, with the top one in our count garnering 238 citations. We list these papers in order of the number of citations.

1. Glover, F. 1990. Tabu Search: A Tutorial (238 citations).

2. Arntzen, B., G. Brown, T. Harrison, L. Trafton. 1995. Global Supply Chain Management at Digital Equipment Corporation (192 citations).

3. Saaty, T. 1994. How to Make a Decision: The Analytic Hierarchy Process (125 citations).

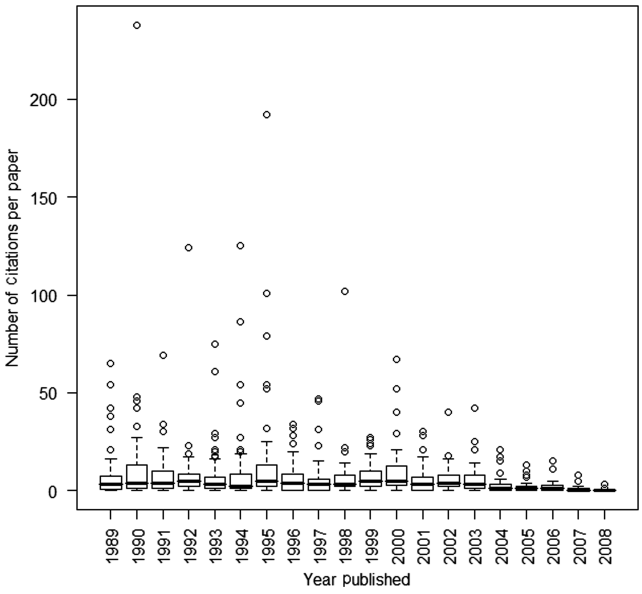

Figure 2: The graph shows the distribution of the number of citations per paper by year. The papers with large numbers of citations are clearly visible as outliers.

4. Smith, B., J. Leimkuhler, R. Darrow. 1992. Yield Management at American Airlines (124 citations).

5. Fylstra, D., L. Lasdon, J. Watson, A. Waren. 1998. Design and Use of the Microsoft Excel Solver (102 citations).

6. Lee, H., C. Billington. 1995. The Evolution of Supply-Chain-Management Models and Practice at Hewlett-Packard (101 citations).

In Figures 2 and 3 we see, not surprisingly, that the number of citations decreases in more recent years. This is most likely just an artifact of the lag in publishing; the papers within the past five years, and perhaps even the past decade, have lower counts because many of the papers that will cite them have yet to be published.

Analyses and Results

We present two analyses. The first evaluates the number of citations by various paper characteristics derived from the data in the tables of contents for all 1,133 papers. The second evaluates the number of citations by characteristics derived from the papers themselves, where we compare the 59 most frequently cited papers (see Appendix A) against a sample of

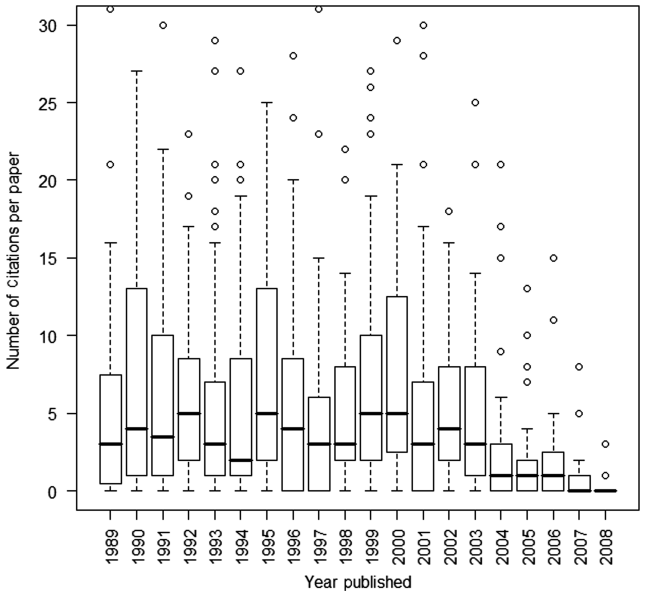

Figure 3: The graph shows the distribution of the number of citations per paper by year with the *y*-axis truncated at 30 to make the structure of the boxplots more visible.

papers with no citations. These 59 papers are defined as being at or above the 95th percentile for the citation count for the year in which they were published. The papers with no citations were drawn by year to correspond to the year distribution of the papers cited most frequently.

In our analysis of all the papers, we evaluate the following characteristics.

- Paper length (number of pages)
- Edelman competition paper or not
- Number of authors
- Type of paper (article, column, tutorial, or editorial)
- Lead-author affiliation (academic or not)

The most statistically significant factor ($p < 0.0001$) associated with the number of citations is whether the paper is a tutorial. We identified nine papers that can be classified as tutorials in our data. These are a small fraction of the papers we examine; however, they are the only ones to have the word "tutorial" either in their title or content descriptor. The tutorials have a mean citation count of 44.0. Glover (1990) again has the most citations; however, even if we exclude Glover's paper, the difference in mean citation counts is still significant ($p < 0.004$), with the remaining tutorials having a mean citation count of 19.8. In Appendix B, we list the tutorials by year, journal number, and then lead author, and we include the number of citations for each.

The next most statistically significant factor associated with the number of citations is paper length (see Figure 4). The figure shows an obvious positive correlation ($r = 0.45$), although it also shows that papers with zero citations span almost the entire range of paper lengths. On average, across all the papers, each additional page of an article is associated with an increase of one additional citation ($p < 0.001$). However, that average is heavily influenced by the highly cited papers. When we focus on those papers at or below the 75th percentile of the number of citations, we see that every seven-page increase in a paper's length is associated with one additional citation ($p < 0.001$). However, regardless of the citation rate per page, longer papers are associated with more citations.

Edelman papers are also positively associated with an increased number of citations. On average, across all the papers, an Edelman competition paper has

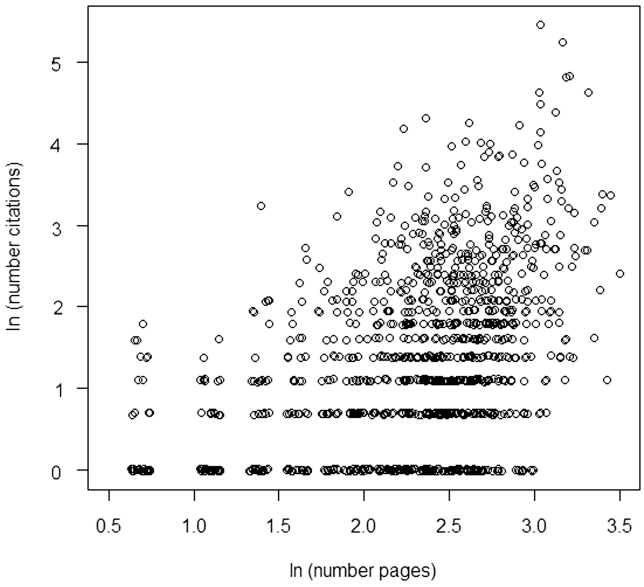

Figure 4: We plotted the number of citations vs. number of pages for each paper (on a started log-log scale and jittered to mitigate symbol overplotting).

an average of 3.3 citations more than a non-Edelman paper ($p < 0.02$). As with paper length, this average is heavily influenced by the highly cited papers. When we focus on those papers at or below the 75th percentile, we see that the mean number of citations for papers from the Edelman competition is statistically indistinguishable from the mean for non-Edelman papers. Edelman competition papers comprise a disproportionate fraction of the upper 25th percentile (21 percent of these are Edelman competition papers, although these papers comprise only 11 percent of all papers). Figure 5 shows that the entire distribution of the number of citations is clearly shifted upwards for the Edelman competition papers when compared to the non-Edelman papers.

Interestingly, the number of authors is sometimes mildly negatively associated with the citation count. With a simple linear model fit to all the papers, on average, each additional author is associated with a *decrease* of 0.73 citations ($p < 0.01$) after accounting for the effects of paper length and Edelman competition status. This relationship holds when we look at those papers at or below the 75th percentile of citation counts, although the effect is less with each additional author and is associated with a decrease of 0.11 citations ($p < 0.03$). However, when fitting a model that includes year indicators to account for temporal differences, such as in editorial policy, the number of authors becomes insignificant.

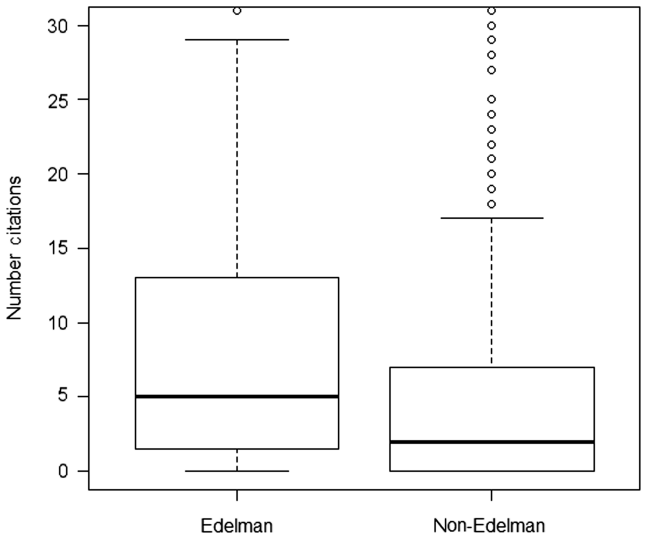

Figure 5: The graph shows side-by-side boxplots of the number of citations for Edelman competition papers vs. non-Edelman papers.

Finally, across all the papers, columns are associated with an increase in citation counts, although this association disappears when we look at the papers at or below the 75th citation-count percentile. Little to no statistical difference exists between articles and editorials; nor is there any difference based on lead-author affiliation.

To assess the impact of these characteristics on the number of citations, we fit a linear model with the started log of the number of citations as the dependent variable and paper length, Edelman competition status, type of paper, number of authors, lead-author affiliation, and year indicators as independent variables. The coefficient of determination (R^2) varies depending on the parameterization of the model; however, the largest R^2 is 0.35 across all the models. This indicates that the characteristics we evaluate account for at best about one-third of the variation in the number of citations.

We also evaluate a subset of the papers based on characteristics derived from the papers themselves. Because this was a manual process of literally counting various paper characteristics, we did not attempt to evaluate all 1,133 papers. Rather, we compared the most highly cited papers by year to a subset of papers that had no citations, but were from the same year and were comparable in length. The characteristics we assessed are the number of:

- Figures
- Tables
- Numbered equations
- "Callouts"
- References

A "callout" (i.e., "pull quote" or "lift-out quote") is a phrase or sentence from the paper, perhaps paraphrased, that is displayed prominently in 14-point font, as shown in callout instances, somewhere in the paper. Generally, these callouts are inserted by the editorial staff to call attention to potentially interesting aspects of a paper. Callouts were present in the first 15 years of papers in our *Interfaces* database; the practice was evidently discontinued in 2004.

Comparing the highly cited papers and the noncited papers, the mean number of figures, tables, and numbered equations is statistically indistinguishable between the two groups. However, the mean number of callouts is statistically different ($p < 0.004$). Focusing on the 15 years of papers with callouts, the highly cited papers had an average of 2.7 callouts; the noncited papers had an average of 1.6 callouts.

Callouts were present in the first 15 years of papers in our *Interfaces* database; the practice was discontinued in 2004.

The number of references given by the authors in their work also shows a significant association ($p < 0.006$) with the number of citations their work subsequently receives. The highly cited papers have an average of 26.3 references; the noncited papers have an average of 15.4.

In addition to the data we developed, we asked INFORMS for comparative data for all 12 of its journals. Table 1 shows WoK impact metrics ranked in decreasing order of INFORMS subscription counts.

Interfaces is ranked third in subscriptions, but last in impact factors. Evidently, *Interfaces* attracts many more paying subscribers than citations counts can explain. *Interfaces'* mission is to publish a worthwhile body of applied operations research that is dedicated "to improving the practical application of OR/MS to decisions and policies in today's organizations and industries" (INFORMS online 2010) Evidently, subscribers are attracted by this.

Journal	2009 impact	5 yr imp	Art inf	Cites
Mgmt Sci*	2.23	4.13	2.45	20,103
Opns Res*	1.58	2.67	1.68	9,661
Interfaces	0.84	1.27	0.53	1,552
Mktg Sci*	2.19	4.25	2.17	3,996
Org Sci	3.13	5.78	2.77	8,404
Math of OR	1.05	1.61	1.40	2,895
Inf Sys Res*	1.79	4.89	1.88	3,037
Trans Sci	1.48	3.60	1.53	2,884
MSOM	2.15			950
J on Comp	1.32	1.50	0.76	1,139

Table 1: The data show INFORMS journals indexed by the *ISI Web of Knowledge* (accessed July 1, 2010) and listed in order of decreasing subscription counts, with impact scores and total citations.
Note. An asterisk denotes those distinguished by *BusinessWeek* in its evaluation of business schools.

Discussion

In this paper, we have assessed some of the characteristics that differentiate highly cited *Interfaces* papers from papers that are cited less often. Our findings inevitably lead to additional questions. For example, paper length is positively associated with the number of citations. Do longer papers garner more citations simply because they communicate more information that is citable? Could it be that important or significant topics, which would naturally tend to be cited, also take more pages to develop and communicate, or is there some other aspect that is driving both length and number of citations? For example, are editors and reviewers perhaps more tolerant of longer good papers?

Similar questions arise when thinking about why Edelman competition papers are cited more than non-Edelman papers. An attractive hypothesis is that Edelman papers epitomize *Interfaces'* mission and, as such, the journal's readership cites these papers more often in subsequent research. On the other hand, Edelman papers surely benefit from prescreening and additional preparation that is part of the Edelman competition, Edelman competition finalist status engenders interest, and the INFORMS promotion of polished video presentations surely attracts a wider audience to these papers.

Our sample of only nine tutorial papers may not justify any conclusion about the impact of tutorials, but these papers do seem to attract many more citations. The tutorial style of the papers is not necessarily solely what attracted the increased number of citations. As we understand it, the authors of these particular tutorials were recruited, likely for their expertise in their particular subject area, but perhaps also for their reputation, writing ability, and possibly other factors such as new research contributions that are in fashion but not necessarily accessible to the average *Interfaces* reader.

The observed increase in citation rates could be due to any of these factors. Similarly, we note that four of the nine tutorials come from a single issue of *Interfaces*

(1990, Volume 20, Issue 4). This suggests another factor that may have had some effect on the increased citations. Regardless, the disparity between mean citations is interesting, and the citation rates of the tutorials demonstrate that readers found value in these types of papers. This suggests that the *Interfaces* editorial staff may want to commission more such papers.

The association between callouts and number of citations is unexpected. Do callouts induce citations because the highlighted phrases attract the attention of readers who would not otherwise have read the papers? Or, is it possible that callouts reflect inherently interesting aspects of a paper that would have been cited regardless of being highlighted in a callout? *Interfaces* should consider reinstituting callouts.

Given that callouts were discontinued in 2004, something of a natural experiment is currently in progress. Figure 3 shows a substantial drop in the number of citations from 2004 on. Some of this is undoubtedly because of the recent publication of these papers, but the dramatic decrease from 2003 to 2004 is suggestive of a possible callout effect. Given the confounding in the data between paper age and the discontinuation of callouts, as well as possible confounding with other temporal effects, it is premature to conclude that callouts promote increased citations. However, in another five years or so, when these newer papers have had a chance to season, and particularly if *Interfaces* reinstitutes callouts, it would be worth taking another look at the data to see if a difference in the number of citations for these papers remains.

Finding that papers that make more references to prior literature tend to be cited more often is not surprising. Is a paper with more references better researched? We think so. The lack of literature review supporting some papers surprises us.

Interfaces may want to commission more tutorials.

In this paper, we limit the analysis to objectively quantifiable characteristics such as paper length and author affiliation, although when we planned our research we also considered evaluating less easily and sometimes more subjectively quantifiable characteristics. For example, we considered characteristics such as the primary research area and the reputations of the authors as well as the clarity of the writing and the importance of the research topic.

We ultimately chose not to address some characteristics for the purely practical reason that the level of effort it would take to collect and compile the data was beyond our capacity to execute. Indeed, the biggest frustration that arose during our research was our inability to convert electronic copies of papers into an easily editable form with existing technology so that we could use off-the-shelf tools to assess the quality and readability of the writing. We hope to complete such analysis when this becomes possible in the future, as we trust it will.

We also chose not to address some more qualitative characteristics of the papers because we were not comfortable quantifying them with simple metrics. This is exactly what bibliometrics seeks to do; however, we are wary of assessments based on oversimplified metrics.

Conclusions

Our research for this article has given us reason to reread and think about a lot of good papers, and revise the list of papers we expect our students to know. That has been pure pleasure.

Interfaces should consider reinstituting callouts (as we have done in this paper).

As for the *BusinessWeek* selections and their potential influence, we believe it is significant that the journals used in the rankings are chosen by business school professors, not by MBA students or other *BusinessWeek* patrons, who would potentially prefer *Interfaces* over the selected INFORMS journals. To us, *Interfaces* is the *Harvard Business Review* of analysis, full of real-world, contemporary, well-edited case studies. We wonder what *BusinessWeek* would discover if it put its journal selection in the hands of business school graduates, readers, and their employers, rather than business school professors.

We advise *BusinessWeek* to reconsider its journal selection criteria to consider not just the opinion of—well, theoretically, purely academically competitive

professors—but also the opinion of business school students and their prospective employers. We suspect such a change would highlight *Interfaces* for its many practical allures, including the useful and successful application of quantitative methods and the distinguishing *Interfaces* requirement that reported results be certified by a beneficiary client. After all, intellectual capital is most useful if it works in the real world. Perhaps Shewhart (1931, p. 18) said it best: "... the fact that the criterion we happen to use has a fine ancestry of highbrow statistical theorems does not justify its use. Sucn justification must come from empirical evidence that it works."

Appendix A. Top 59 Most Highly Cited *Interfaces* Papers from 1989–2008

The following list shows the 59 most highly cited *Interfaces* papers from 1989–2008. We developed this list from those papers with citation counts in the top 5 percent by year of publication. The papers are ordered by year and by issue number. Note that 18 of these papers are Edelman papers from the first issue of the year.

1989

Taylor, P., S. Huxley. A break from tradition for the San Francisco police: Patrol officer scheduling using an optimization-based decision support system. **19**(1) 4–24 (42 citations).

Wind, Y., P. Green, D. Shifflet, M. Scarbrough. Courtyard by Marriott: Designing a hotel facility with consumer-based marketing models. **19**(1) 25–47 (38 citations).

Abara, J. Applying integer linear programming to the fleet assignment problem. **19**(4) 20–28 (65 citations).

Gershkoff, I. Optimizing flight crew schedules. **19**(4) 29–43 (54 citations).

1990

Cohen, M., P. Kamesam, P. Kleindorfer, H. Lee, A. Tekerian. Optimizer: IBM's multi-echelon inventory system for managing service logistics. **20**(1) 65–82 (42 citations).

Eom, H., S. Lee. A survey of decision support system applications (1971–April 1988). **20**(3) 65–79 (48 citations).

Glover, F. Tabu search: A tutorial. **20**(4) 74–94 (238 citations).

Bertsekas, D. The auction algorithm for assignment and other network flow problems: A tutorial. **20**(4) 133–149 (46 citations).

1991

Fetter, R. Diagnosis related groups: Understanding hospital performance. **21**(1) 6–26 (34 citations).

Anbil, R., E. Gelman, B. Patty, R. Tanga. Recent advances in crew-pairing optimization at American Airlines. **21**(1) 62–74 (69 citations).

Zahedi, F. An introduction to neural networks and a comparison with artificial intelligence and expert systems. **21**(2) 25–38 (30 citations).

1992

Smith, B., J. Leimkuhler, R. Darrow. Yield management at American Airlines. **22**(1) 8–31 (124 citations).

Liberatore, M., R. Nydick., P. Sanchez. The evaluation of research papers (or how to get an academic committee to agree on something). **22**(2) 92–100 (19 citations).

Zangwill, W. The limits of Japanese production theory. **22**(5) 14–25 (17 citations).

Keeney, R., T. McDaniels. Value-focused thinking about strategic decisions at BC Hydro. **22**(6) 94–109 (23 citations).

1993

Kaplan, E., E. O'Keefe. Let the needles do the talking! Evaluating the New Haven needle exchange. **23**(1) 7–26 (61 citations).

Andrews, B., H. Parsons. Establishing telephone-agent staffing levels through economic optimization. **23**(2) 14–20 (29 citations).

Martin, C., D. Dent, J. Eckhart. Integrated production, distribution, and inventory planning at Libbey-Owens-Ford. **23**(3) 68–78 (27 citations).

Lee, H., C. Billington, B. Carter. Hewlett-Packard gains control of inventory and service through design for localization. **23**(4) 1–11 (75 citations).

1994

Cariño, D. L., T. Kent, D. H. Myers, C. Stacy, M. Sylvanus, A. L. Turner, K. Watanabe, W. T. Ziemba. The Russell-Yasuda Kasai model: An asset/liability model for a Japanese insurance company using multistage stochastic programming. **24**(1) 29–49 (86 citations).

Subramanian, R., R. Scheff, J. Quillinan, D. Wiper, R. Marsten. Coldstart: Fleet assignment at Delta Air Lines. **24**(1) 104–120 (45 citations).

Sharda, R. Neural networks for the MS/OR analyst: An application bibliography. **24**(2) 116–130 (54 citations).

Saaty, T. How to make a decision: The analytic hierarchy process. **24**(6) 19–43 (125 citations).

1995

Arntzen, B., G. Brown, T. Harrison, L. Trafton. Global supply chain management at Digital Equipment Corporation. **25**(1) 69–93 (192 citations).

Geoffrion, A., R. Powers. Twenty years of strategic distribution system design: An evolutionary perspective. **25**(5) 105–127 (79 citations).

Lee, H., C. Billington. The evolution of supply-chain-management models and practice at Hewlett-Packard. **25**(5) 42–63 (101 citations).

1996

Leachman, R., R. Benson, C. Liu, D. Raar. IMPReSS: An automated production-planning and delivery-quotation system at Harris Corporation—semiconductor sector. **26**(1) 6–37 (28 citations).

Mulvey, J. Generating scenarios for the Towers Perrin investment system. **26**(2) 1–15 (32 citations).

Rosenhead, J. What's the problem? An introduction to problem structuring methods. **26**(6) 117–131 (34 citations).

1997

Geraghty, M., E. Johnson. Revenue management saves National Car Rental. **27**(1) 107–152 (31 citations).

Camm, J., T. Chorman, F. Dill, J. Evans, D. Sweeney, G. Wegryn. Blending OR/MS, judgment, and GIS: Restructuring P&G's supply chain. **27**(1) 128–142 (46 citations).

Ackerman, F., C. Eden, T. Williams. Modeling for litigation: Mixing qualitative and quantitative approaches. **27**(2) 48–65 (47 citations).

1998

Porco, T., S. Blower. Designing HIV vaccination policies: Subtypes and cross-immunity. **28**(3) 167–190 (22 citations).

Fylstra, D., L. Lasdon, J. Watson, A. Waren. Design and use of the Microsoft Excel Solver. **28**(5) 29–55 (102 citations).

1999

Epstein, R., R. Morales, J. Serón, A. Weintraub. Use of OR systems in the Chilean forest industries. **29**(1) 7–29 (26 citations).

Bermon, S., S. Hood. Capacity optimization planning system (CAPS). **29**(5) 31–50 (27 citations).

2000

Dowlatshahi, S. Developing a theory of reverse logistics. **30**(3) 143–155 (52 citations).

Guide, V., J. Jayaraman, R. Srivastava, W. Benton. Supply-chain management for recoverable manufacturing systems. **30**(3) 125–142 (67 citations).

Klausner, M., C. Hendrickson. Reverse-logistics strategy for product take-back. **30**(3) 156–165 (40 citations).

2001

Geoffrion, A., R. Krishnan. Prospects for operations research in the e-business era. **31**(2) 6–36 (28 citations).

Sodhi, M. Applications and opportunities for operations research in Internet-enabled supply chains and electronic marketplaces. **31**(2) 56–69 (21 citations).

Green, P., A. Krieger, Y. Wind. Thirty years of conjoint analysis: Reflections and prospects. **31**(3) S56–S73 (30 citations).

Lustig, I., J.-F. Puget. Program does not equal program: Constraint programming and its relationship to mathematical programming. **31**(6) 29–53 (30 citations).

2002

Blake, J., J. Donald. Mount Sinai Hospital uses integer programming to allocate operating room time. **32**(2) 63–73 (16 citations).

Ledyard, J., M. Olson, D. Porter, J. Swanson, D. Torma. The first use of a combined-value auction for transportation services. **32**(5) 4–12 (40 citations).

Benjaafar, S., S. Heragu, S. Irani. Next generation factory layouts: Research challenges and recent progress. **32**(6) 58–76 (18 citations).

2003

Hohner, G., J. Rich, E. Ng, G. Reid, A. Davenport, J. Kalagnanam, H. Lee, C. An. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. **33**(1) 23–35 (42 citations).

Guide, V., T. Harrison, L. Van Wassenhove. The challenge of closed-loop supply chains. **33**(6) 3–6 (25 citations).

2004

Kapuscinski, R., R. Zhang, P. Cabonneau, R. Moore, B. Reeves. Inventory decisions in Dell's supply chain. **34**(3) 191–205 (21 citations).

Sheffi, Y. Combinatorial auctions in the procurement of transportation services. **34**(4) 245–252 (15 citations).

Pil, F., M. Holweg. Linking product variety to order-fulfillment strategies. **34**(5) 394–403 (17 citations).

2005

Metty, T., R. Harlan, Q. Samelson, T. Moore, T. Morris, R. Sorensen, A. Schneur, et al. Reinventing the supplier negotiation process at Motorola. **35**(1) 7–23 (13 citations).

Sahoo, S., S. Kim, B. Kim, B. Kraas, A. Popov Jr. Routing optimization for Waste Management. **35**(1) 24–36 (10 citations).

van den Briel, M., J. Villalobos, G. Hogg, T. Lindemann, A. Mulé. America West Airlines develops efficient boarding strategies. **35**(3) 191–201 (8 citations).

2006

Sandholm, T., D. Levine, M. Concordia, P. Martyn, R. Hughes, J. Jacobs, D. Begg. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. **36**(1) 55–68 (11 citations).

Brown, G., M. Carlyle, J. Salmerón, K. Wood. Defending critical infrastructure. **36**(6) 530–544 (15 citations).

2007

Ahuja, R., K. Jha, J. Liu. Solving real-life railroad blocking problems. **37**(5) 404–419 (5 citations).

Fildes, R., P. Goodwin. Against your better judgment? How organizations can improve their use of management judgment in forecasting. **37**(6) 570–576 (8 citations).

2008

Wright, M., J. Armstrong. The ombudsman: Verification of citations: Fawlty towers of knowledge? **38**(2) 125–139 (3 citations).

Appendix B. Tutorials from 1989–2008

Bertsekas, D. 1990. The auction algorithm for assignment and other network flow problems: A tutorial. **20**(4) 133–149 (46 citations).

Glover, F. 1990. Tabu search: A tutorial. **20**(4) 74–94 (238 citations).

Lustig, I. J., J-F. Puget. 2001. Program does not equal program: Constraint programming and its relationship to mathematical programming. **31**(6) 29–53 (30 citations).

Marsten, R., R. Subramanian, M. Saltzman, I. Lustig, D. Shanno. 1990. Interior point methods for linear programming: Just call Newton, Lagrange, and Fiacco and McCormick! **20**(4) 105–116 (33 citations).

Rubin, D., H. Wagner. 1990. Shadow prices: Tips and traps for managers and instructors. **20**(4) 150–157 (12 citations).

Seal, K. C., Z. H. Przasnyski. 2003. Using technology to support pedagogy in an OR/MS course. **33**(4) 27–40 (2 citations).

Sen, S., J. Higle. 1999. An introductory tutorial on stochastic linear programming models. **29**(2) 33–61 (24 citations).

Sobol, M. G. 1991. Validation strategies for multiple regression analysis: Using the coefficient of determination. **21**(6) 106–120 (1 citation).

Tovey, C. A. 2002. Tutorial on computational complexity. **32**(3) 30–61 (10 citations).

Acknowledgments

The authors dedicate this paper to Rick Rosenthal (1950–2008) and Mike Rothkopf (1939–2008), two scholars who cared deeply about this journal and its impact.

References

Adler, R., J. Ewing, P. Taylor. 2009. Citation statistics (with discussion). *Statist. Sci.* **24**(1) 1–14.

Ahmed, T., B. Johnson, C. Oppenheim, C. Peck. 2004. Highly cited old papers and the reasons why they continue to be cited. Part II, The 1953 Watson and Crick article on the structure of DNA. *Scientometrics* **61**(2) 147–156.

Bornmann, L., H.-D. Daniel. 2006. What do citation counts measure? A review of studies on citing behavior. *J. Documentation* **64**(1) 45–80.

BusinessWeek. 2010. Full-time MBA rankings. Accessed July 3, 2010, http://www.businessweek.com/bschools/content/nov2008/bs20081113_320726.htm.

Camacho-Miñano, M., M. Núñez-Nickel. 2009. The multilayered nature of reference selection. *J. Amer. Soc. Inform. Sci. Tech.* **60**(4) 754–777.

Case, D. O., G. M. Higgins. 2000. How can we investigate citation behavior? A study in reasons for citing literature in communication. *J. Amer. Soc. Inform. Sci.* **51**(7) 635–645.

Cronin, B. 1998. Metatheorizing citation. *Scientometrics* **43**(1) 45–55.

Egghe, L. 2006. Theory and practise of the *g*-index. *Scientometrics* **69**(1) 131–152.

Fernandez-Alles, M., A. Ramos-Rodriquez. 2009. Intellectual structure of human resources management research: A bibliometric analysis of the journal *Human Resource Management*, 1985–2005. *J. Amer. Soc. Inform. Sci. Tech.* **60**(1) 161–175.

Frogel, J. A. 2010. Astronomy's greatest hits: The 100 most cited papers in each year of the first decade of the 21st century (2000–2009). Accessed June 6, 2010, http://arxiv.org/abs/1005.5377v1 on.

Garfield, E. 1994. The Thomson Reuters impact factor. Accessed February 24, 2010, http://thomsonreuters.com/products_services/science/free/essays/impact_factor/.

Garfield, E. 2006. The history and meaning of the journal impact factor. *J. Amer. Medical Assoc.* **295**(1) 90–93.

Gross, P. L. K., E. M. Gross. 1927. College libraries and chemical education. *Science* **66** 385–389.

Hirsch, J. E. 2005. An index to quantify an individual's scientific research output. *Proc. National Acad. Sci.* **102**(46) 16569–16572.

Hopp, W. 2010. Personal communication with author.

INFORMS Online. 2010. *Interfaces*. Accessed June 29, 2010, http://www.informs.org/Journal/Interfaces.

Kelton, W. 2010. Personal communication with author.

Leydesdorff, L. 1998. Theories of citation? *Scientometrics* **43**(1) 5–25.

Nicolaisen, L. 2007. Citation analysis. *Ann. Rev. Inform. Sci. Tech.* **41** 609–641.

Norden, P. V. 1970. A message from the president. *Interfaces* **1**(1) 1.

O'Leary, D. E. 2009. Business school research: Measuring value contribution through citations of journals in patents. *Interfaces* **39**(6) 516–526.

Oppenheim, C., S. P. Renn. 1978. Highly cited old papers and the reasons why they continue to be cited. *J. Amer. Soc. Inform. Sci.* **29**(5) 225–231.

Patsopoulos, N. A., J. P. Ioannidis, A. A. Analatos. 2006. Origin and funding of the most frequently cited papers in medicine: Database analysis. *British Medical J.* **332** 1061–1064

Ryan, T. P., W. H. Woodall. 2005. The most cited statistical papers. *J. Appl. Statist.* **32**(5) 461–474.

Shapiro, F. R. 1985. The most cited law review articles. *California Law Rev.* **73**(5) 1540–1554.

Shewhart, W. A. 1931. *Economic Control of Quality of Manufactured Product*. D. Van Nostrand Company, New York.

Thomson Reuters. 2010a. Science citation index. Accessed February 21, 2010, http://thomsonreuters.com/products_services/science/science_products/a-z/science_citation_index.

Thomson Reuters. 2010b. Introducing the impact factor. Accessed March 2, 2010, http://thomsonreuters.com/products_services/science/academic/impact_factor.

White, M. D., P. Wang. 1997. A qualitative study of citing behavior: Contributions, criteria, and metalevel documentation concerns. *Library Quart.* **67**(2) 122–154.