Statistical Science
1994, Vol. 9, No. 1, 94–108

Citation Patterns in the Journals of Statistics and Probability

Stephen M. Stigler

Abstract. This is a study of the use of citation data to investigate the role statistics journals play in communication within that field and between statistics and other fields. The study looks at citations as import-export statistics reflecting intellectual influence. The principal findings include: there is little variability in both the number and diversity of imports, but great variability in both the number and diversity of exports and hence in the balance of trade; there is a tendency for influence to flow from theory to applications to a much greater extent than in the reverse direction; there is little communication between statistics and probability journals. The export scores model is introduced and employed to map a set of journals' bilateral intellectual influences onto a one-dimensional scale, and the Cox effect is identified as a phenomenon that can occur when a disciplinary paper attracts a large degree of attention from outside its discipline.

Key words and phrases: Citations, bibliometrics, Gini index, Herfindahl index, Simpson's Index, quasi-symmetry, Bradley-Terry model, journals.

1. INTRODUCTION

Each article in any scientific journal contains a list of references, what are termed *citations* to previous work. There can be many reasons for an author to cite a particular work: to acknowledge a source of ideas, methods, or data; to point to related work or a review of earlier work; to declare alliance with another author's approach; or to criticize, perhaps even totally refute, the referenced work. All of these represent some type of intellectual influence, and counts based upon lists of citations provide quantitative expression—measures—of the flow of intellectual influence in the scientific literature. This investigation is an attempt to explore the use of citation counts to provide insight into the flow of intellectual influence (and a better understanding of the nature of the intellectual influence captured by citation counts) in the journals of statistics and probability.

There are many problems with the use of citation counts as a measure of influence, some obvious and well known, others more subtle. For example, in some instances the propensity for self-citation may become so extreme as to skew the counts (although it could be countered that the authors in question are honestly acknowledging their most important intellectual influence), and some citations may exist for no purpose other than to flatter (but even this is a form of declaration of alliance and hence of influence). There are also technical measurement problems: citation indexes usually key on only the first of a set of multiple authors, and slight variations in spelling or a different choice of initials may lead to the creation of a new individual "author." There are major sources of variation in the life history of a paper's citations. Some papers attract many early citations but are subsequently ignored, while others are noticed only slowly and achieve their peak of acknowledged influence 20 or more years after publication. Among the subtler problems encountered are those identified by Robert K. Merton as Obliteration By Incorporation (the tendency for some particularly influential works to be so absorbed by the literature that they are rarely cited directly at all; Merton, 1968, pp. 27–28, 35–37) and the Matthew Effect (a form of contagion—"For unto every one that hath shall be given, and he shall have abundance: but from him that hath not shall be taken

Stephen M. Stigler is the Ernest DeWitt Burton Distinguished Service Professor, Department of Statistics, University of Chicago, 5734 University Avenue, Chicago, Illinois 60637. He is a former editor of the Journal of the American Statistical Association and is current (1993–94) president of the Institute of Mathematical Statistics.

away even that which he hath," from the Gospel According to St. Matthew as quoted in Merton, 1973, page 445).

These potential problems are all real, but they are greatest at the level of the individual, when the object of study is a small set of papers or authors. Even there, the difficulties are rarely as severe as feared, and the resulting studies can be revealing. (See MacRoberts and MacRoberts, 1986, for a catalogue of potential problems and Stigler, 1987, and Zuckerman, 1987, for rejoinders and further discussion. For one use in an historical study of citations by an individual, Pierre Simon Laplace, see Stigler, 1978.)

In this study, the difficulties in individual level data are largely avoided by taking the *journal* as the unit of analysis, an approach that was used successfully by Stigler and Stigler (1992) in a study of the interrelationships among journals in economics. Earlier, Garfield (1972) had suggested that studying citations to journals could be useful for science policy studies, and he discussed some of the benefits and new difficulties that occur. Pinski and Narin (1976) looked at a variety of ad hoc measures of influence for physics journals, using weights derived from the eigenvalues of the cross-citation matrix.

By treating the journal as the unit, some of the major sources of variation are diminished by aggregation, and technical problems (such as the spelling of authors' names or multiple authors) are avoided entirely. That is not to say that no problems occur, because there are difficulties that are peculiar to journals as units. Journals may (and do) change name or split, and new journals will necessarily attract fewer citations than do the old, *ceteris paribus* (no issue of a journal can be cited before the first). One additional advantage, however, is the existence of a good source of data.

2. THE DATA

Since 1963 the Institute for Scientific Information has published indexes of academic journals. Initially only journals of the physical and biological sciences were included, but the coverage was subsequently expanded to the social sciences and to the arts and humanities. Within each area, the coverage of journals has changed over time, often through the addition of journals, but also to reflect changes such as title cancellation, name change, or the amoeba-like splitting of existing journals. The present coverage cannot be said to be complete; however, all major journals are included and it is a nontrivial task to discover even an obscure journal that escapes their net. The principal publications of the Institute are its citation indexes: *Science Citation Index, Social Sciences Citation Index, The Arts and Humanities Citation Index*. They also publish *Current Contents* and offer other services keyed to their now-massive data base, which extends considerably before 1963.

Since the citation indexes started to appear they have been the subject of much statistical analysis. For examples, see Elkana et al. (1978), Garfield (1979), and most issues of the journal *Scientometrics*. A great deal of this work has been concerned with the practices of individual scholars—the patterns of the growth and decay of citation counts for separate articles, tendencies for networks of scholarly communication to develop, and so forth. Since it deals with journals rather than articles or scholars as the unit of analysis, this study will use a compilation prepared annually by the Institute called *SCI Journal Citation Reports: A Bibliometric Analysis of Science Journals in the ISI Data Base* (edited by Eugene Garfield), henceforth referred to as *JCR*. The principal tabulations presented in these volumes, which after 1988 are published only on microfiche, are two extensive tables. The first table gives for each journal cited a list of the major source journals for its citations, together with the counts of the citations its volumes received, broken down by year. The other table presents much the same information but arranged by citing journal rather than cited journal.

The citation counts of *JCR* are broken down by year back for a decade prior to publication, with the counts of citations made in the year in question to papers published in earlier years given in aggregate form. For each journal cited (among those covered by the data base), the counts of citations[1] are given for all citing journals that produce at least 6 citations during the year with the additional provisos that at least 15 citing journals are included in all cases and no more than about 100 are included. The range of citation counts thus varies. Almost complete records are given for some infrequently cited journals while only the 100 major producers of citations are given for some large and frequently cited journals, such as *Science*, *Nature* or *Psychological Bulletin*.

3. CITATIONS AS IMPORT-EXPORT STATISTICS

When one journal, say *Journal A*, prints a paper containing a citation to work previously published in another, *Journal B*, we may consider this as indicating an instance of the *export* of intellectual influ-

[1]If a journal article cites another journal article, whether in the same journal or not, a single count is recorded, no matter how many times the cited paper is referred to in the same article.

ence from *Journal B* and the *import* of intellectual influence by *Journal A* subject to those caveats already mentioned. In one sense, this use of language sounds paradoxical, but it is nonetheless appropriate. The journal that performs the act of printing the citation is the importer; the journal whose article lies passively in the library, perhaps having been printed many years before, is engaged in a current act of exporting. The exporter/producer may even have gone out of business many years before the export is actually shipped! The importer/consumer need not be a passive consumer at all, but may play a powerful role of creating and defining a new market for a long-ago printed "product." This market metaphor has obvious limitations (the "product" is potentially inexhaustible, and the "currency" of reputation that is exchanged for ideas is not the usual type of monetary instrument), but thinking of citations as a market suggests a number of ways of looking at that market quantitatively. The first group of these we shall consider refers to the characteristics of the individual journals. Subsequently, we shall examine the statistics of bilateral trade between journals.

4. UNIVARIATE DATA: SIMPLE COUNTS

The simple aggregate counts are given in Table 1 for the 33 statistical journals included in this study. The counts given are fairly straightforward with the exception of the journals of the Institute of Mathematical Statistics, *The Annals of Statistics* and *The Annals of Probability*.[2] These journals were created in 1973 as successors to *The Annals of Mathematical Statistics*. One might naively expect that the number of citations dating to before 1973 would not be large, but then one would be strikingly mistaken. The average number of citations in 1987–89 to issues of *The Annals of Mathematical Statistics* dating from before 1973 was 2240. This was sufficient to put that journal in fifth place among the journals included and ahead of its progeny, *The Annals of Statistics* and *The Annals of Probability*, and a large number of other major (and all minor) publications. One of the entries in Table 1 gives the counts for *The Annals of Mathematical Statistics* and *The Annals of Statistics* combined, effectively treating *The Annals of Statistics* as a single successor to *The Annals of Mathematical Statistics* with *The Annals of Probability* as a newly created journal. For a later study, the citations to *The Annals of Mathematical Statistics* were allocated among the two successors; see the Appendix.

As is evident from the table, *The Journal of the American Statistical Association* (*JASA*) is the single most cited statistics journal, although if *The Annals* of the IMS are combined (as they were before 1973) they would exceed *JASA*. *JASA* exhibits the largest citing audience, referred to by some 1063 sources.[3] The four most cited core statistics journals now published as single journals are, in descending order of citation, *JASA, The Annals of Statistics, Biometrika* and the *Journal of the Royal Statistical Society* (*Series B*) (*JRSS*(*B*)). The four most cited core probability journals are *The Annals of Probability, The Journal of Applied Probability, Stochastic Processes and their Applications* and *The Theory of Probability and its Applications* (the translation of the Russian journal *Theorija Verojatnostei*).[4] All other highly cited journals (such as *Econometrica, Biometrics, Technometrics, Psychometrika* and *The Journal of Econometrics*) aim at and find significant audiences beyond professional statisticians and probabilists. Indeed, the fact that statistics journals attract easily twice as many citations as probability journals is plausibly due as much to their finding a large audience beyond statisticians as to the larger size of their professional community: two characteristics that are no doubt related. The diversity of citation patterns and the interrelations among citations within the professional community will be discussed later.

5. BALANCE OF TRADE

Aggregate citation counts give only a crude indication of the role a journal plays in the professional literature. A high count may indicate that the journal is a leader in the publication of newly created knowledge and that it wields substantial intellectual influence in the field's development. But this makes no allowance for a large variety of factors. Some of these factors are arguably directly related to intellectual influence and should not be "corrected" for: whether the journal has a limited geographical distribution or a small number of subscribers, the size of the journal, the number of issues a year and the

[2]At the time these data were recorded (1988), the Institute's newest journal, *Statistical Science* (founded 1986), was not yet included. Its *Annals of Applied Probability* was not begun until 1991.

[3]The number of journals given as citing the IMS's *Annals* (in Table 1, column 3) is of necessity an estimate, since a portion of the count was given in aggregate form, and for those counts, overlap of citing journals was impossible to determine exactly. The errors in these estimates are unlikely to be as large as 100.

[4]One source of bias in these counts for probability journals is the fact the journal *Probability Theory and Related Fields* was until 1986 known as *Zeitschrift für Wahrscheinlichkeitsrechnung*, and citations to it under that title were not tabulated in the *Index*.

TABLE 1
Aggregate citation data for 33 journals in statistics and probability. The counts given are the average yearly totals for 1987–89, except as noted, and cover all citations recorded in issues of journals in the ISI data base

Journal	#cites received	#jnls citing*	#cites sent	#jnls cited*
J. Amer. Statist. Assoc.	5436	1063	2580	824
Econometrica	4689	661	1329	482
Ann. Statist. and Ann. Math. Statist.	4462	625est	2017	536
Biometrics	3916	954	1612	449
Biometrika	3564	795	1306	352
Ann. Math. Statist.	2240	494	0	0
Ann. Statist.	2222	237	2017	536
J. Roy. Statist. Soc. Ser. B	2207	564	686	278
Psychometrika	1656	568	666	281
Technometrics	1575	573	746	246
Ann. Probab.	1084	136	1561	443
Probab. Theory Related Fields	985**	147**	1636	515
J. Econometrics	939	164	1211	438
J. Roy. Statist. Soc. Ser. C	803	301	686	255
Amer. Statist.	781	305	1131	473
J. Appl. Probab.	778	150	965	371
Stochastic Process. Appl.	684	218	942	355
J. Roy. Statist. Soc. Ser. A	635	290	573	322
Theory Probab. Appl.	614	182	707	399
Statistics in Medicine	567	228	NA	NA
Comm. Statist. Theory Methods	491	110	3433	940
J. Multivariate Anal.	350	81	1100	371
Ann. Inst. Statist. Math.	300	111	642	244
Internat. Statist. Rev.	267	109	746	333
Scand. J. Statist.	248	102	328	101
Sankhyā Ser. A	238	65	251	152
J. Statist. Plann. Inference	202	53	1380	467
Statistician	148	100	335	209
Statist. Probab. Lett.	133	54	866	314
Comm. Statist. Simulation Comput.	120	43	861	305
Canad. J. Statist.	110	41	461	229
Sankhyā Ser. B	93	27	239	138
South African Statist. J.	52	24	97	39
Biometrical J.	44	20	914	384

* Based on 1987–88 data
** Based on 1987 data

age of the journal. Even some of these would be desirable to incorporate in the analysis, for example, by checking whether a journal's influence is proportional to its distribution or its size. Others are more problematical: whether it is a review journal or specializes in a relatively narrow technical area, how large a professional community it serves and whether the patterns change if one focuses exclusively on citations within the professional community narrowly defined.

If we regard citations as import-export statistics, this suggests considering measures of the balance of trade for individual journals as a way of compensating for the different size of these "markets" in scientific influence. As is true in the case of economic markets, there are many ways of measuring " balance of trade," and they correspond to different concepts of "balance." The simplest is the ratio of counts of citations received in a given year divided by counts of citations given; that is, the measure of exports (citations received) divided by imports (citations made). Table 2 gives this ratio as BOT for the journals of this study, omitting those journals (like *The Annals of Mathematical Statistics*) for which no import statistics were available. Since the count of imports (citations made) is a measure of the size of the journal, BOT may be thought of as a way of adjusting the aggregate citation count for journal size. Another measure of size is page count, and Table 2 also gives (as "Cites pp") the results of this adjustment, where the number of pages published in 1988 was taken as the size. The two adjustments are in reasonably close agreement, with the anomalous exception of the *American Statistician*, a nonresearch journal publishing a small number of large pages with long reference lists and a diverse audience.

TABLE 2

Citation data for 31 journals in statistics and probability. "BOT" is Balance of Trade, the (Citations Received 1987–89)/(Citations Sent 1987–89) ratio; that is, BOT = (Exports 1987–89)/(Imports 1987–89). "Cites pp" is the number of citations per page published (1988 basis); that is, (Citations Received 1987–89)/(Pages published 1988). "Impact" is essentially the average number of citations per article for the previous two years, computed separately for each of the three years 1987–89 and the three impact factors averaged. Journals marked by an asterisk commenced publication after 1970; only one of these began publication after 1977, Statistics and Probability Letters (begun in 1982)

Journal	BOT	Cites pp	Impact
Econometrica	3.53	3.14	2.64
J. Roy. Statist. Soc. Ser. B	3.22	4.64	1.21
Biometrika	2.73	4.43	1.10
Psychometrika	2.49	2.74	1.03
Biometrics	2.43	3.23	1.10
Ann. Statist. and Ann. Math. Statist.	2.21	2.56	1.14
J. Amer. Statist. Assoc.	2.11	4.37	1.19
Technometrics	2.11	3.30	0.92
J. Roy. Statist. Soc. Ser. C	1.17	1.63	0.62
J. Roy. Statist. Soc. Ser. A	1.11	1.10	1.49
*Ann. Statist.**	1.10	1.28	1.14
Sankhyā Ser. A	0.95	0.52	0.17
Theory Probab. Appl.	0.87	0.81	0.15
J. Appl. Probab.	0.81	0.92	0.48
*J. Econometrics**	0.78	0.79	0.79
*Scand. J. Statist.**	0.76	0.78	0.52
*Stochastic Process. Appl.**	0.73	0.68	0.46
*Ann. Probab.**	0.69	0.58	0.87
Amer. Statist.	0.69	2.61	1.20
South African Statist. J.	0.53	0.24	0.21
Ann. Inst. Statist. Math.	0.47	0.36	0.22
Statistician	0.44	0.30	0.25
Sankhyā Ser. B	0.39	0.21	0.12
Internat. Statist. Rev.	0.36	0.89	1.01
*J. Multivariate Anal.**	0.32	0.25	0.39
Probab. Theory Related Fields	0.27	0.44	0.72
*Canad. J. Statist.**	0.24	0.18	0.26
*J. Statist. Plann. Inference**	0.15	0.16	0.24
*Statist. Probab. Lett.**	0.15	0.23	0.27
*Comm. Statist. Theory Methods**	0.14	0.11	0.16
*Comm. Statist. Simulation Comput.**	0.14	0.08	0.14
Biometrical J.	0.05	0.04	0.06

An obvious objection to this measure is that it does not represent trade in currently produced information. A journal that has recently undergone substantial enlargement or upgrading will not yet have that growth/improvement reflected in the export statistics; an older journal may command a large share of the export market yet not have produced any worthy product in many years. A new journal will not have the extensive stock of older issues exporting influence that a long-established journal will have. An alternative measure of balance of trade, one which attempts to deal with this problem, is one published by the various *Citation Indexes* under the name *impact factor*; this is included in Table 2. The impact factor for *Journal A* is computed each year by taking the ratio of the citations recorded that year from any of the journals in the data base to issues of *Journal A* published during the preceding two years, divided by the number of citable articles published by *Journal A* during those same two years. It is thus relatively insensitive to changes in the size of the journal, and captures one sort of balance of recent trade, where the denominator is the number of articles rather than the number of citations (if all articles produced the same number of references, these two numbers would be proportional). The drawback, of course, is that this measure entirely misses the archival impact of the journals and gives much greater weight to those publications of a more ephemeral nature or to those publications more concerned with debates about current issues than with research. This measure is also unequipped to deal with long publication lags, so that if it were to take two years from submission to publication (a lag that is far from unknown), a journal may contain no citations to publications less than three years old; therefore its "imports" would always fall outside the range used.

TABLE 3
Diversity measures for 33 journals in statistics and probability. The counts given are the average yearly totals for 1987–88, except as noted. "Conc" is the Gini Concentration Index

Journal	Conc***	Conc old***	Conc new***	% Self Citation
Biometrics	0.9	0.6	2.8	6.3
J. Roy. Statist. Soc. Ser. A	1.0	0.7	2.0	5.7
J. Amer. Statist. Assoc.	1.2	0.9	2.8	6.2
J. Roy. Statist. Soc. Ser. B	1.2	0.9	3.1	2.8
Statistics in Medicine	1.3	0.7	2.3	NA
Econometrica	1.5	1.2	1.9	5.5
Psychometrika	1.6	1.0	4.5	7.4
Stochastic Process. Appl.	1.8	1.1	2.5	7.4
Biometrika	2.0	1.4	4.6	7.1
Statistician	2.0	1.9	2.2	9.0
J. Roy. Statist. Soc. Ser. C	2.2	1.8	2.9	9.0
Amer. Statist.	2.2	1.5	2.8	11.3
Technometrics	2.3	1.8	4.8	7.6
Internat. Statist. Rev.	2.5	2.5	14.3	2.4
Ann. Math. Statist.	2.5	2.5	NA	0.0
J. Econometrics	2.8	2.2	3.3	11.3
Scand. J. Statist.	3.7	6.2	3.5	5.7
J. Appl. Probab.	5.0	4.3	5.8	17.7
J. Multivariate Anal.	5.1	4.0	10.2	13.7
Sankhyā Ser. A	5.1	5.8	5.5	11.2
Ann. Statist.	5.3	4.9	5.7	16.2
Ann. Inst. Statist. Math.	6.0	4.7	10.5	10.8
Statist. Probab. Lett.	6.0	2.3	7.4	18.7
Ann. Probab.	6.1	4.3	7.3	16.9
Canad. J. Statist.	6.4	29.6	6.7	13.9
South African Statist. J.	8.6	12.6	11.9	17.5
Probab. Theory Related Fields	8.7	NA	8.5	21.0
Theory Probab. Appl.	8.8	6.1	17.0	26.7
Comm. Statist. Simulation Comput.	11.2	16.7	11.1	19.3
J. Statist. Planning Inference	12.8	10.7	13.2	29.9
Comm. Statist. Theory Methods	13.3	6.8	15.0	34.8
Sankhyā Ser. B	22.7	22.4	24.6	7.8
Biometrical J.	34.0	50.0	36.1	56.4

*** Based on 1988 data

6. EXTENT AND DIVERSITY OF THE MARKET

The journals in this study differ greatly in the extent and diversity of their export markets. We might expect that the highly technical, theoretical journals tend to export influence to (are cited by) a relatively small group of similarly technical journals and that the more applied journals have a more diverse export market. The data support this. But other characteristics of the extent of the markets may be less obvious.

Table 3 reports the results of the more revealing of the analyses performed. The table is ordered according to the value of the Gini concentration index and computed on the basis of 1988 data for the citations received from journals in the ISI data base. The Gini concentration index is known in economics as the Herfindahl index (Stigler, 1968, Chapter 4), in population biology as Simpson's measure of diversity, in cryptology as the repeat rate, and doubtless by many other names as well. It is given by

$$\text{Gini Concentration} = 100 \times \sum_i s_i^2,$$

where s_i is the fraction of its total citations that the journal received from source i, and the summation is over all sources (or rather, and this makes no essential difference in the results reported, over sources contributing at least five citations). This index may be thought of as the chance (expressed in per cent terms) that if two citations are selected at random with replacement from those received during the year they will come from the same journal.

Generally speaking, the more applied journals have highly diverse markets (Gini index less than 2.0), and the more theoretical probability journals are much more concentrated (Gini index 5.0 to 10.0) with the theoretical statistics and the review journals in an intermediate range. There are exceptions to these generalizations. The journal of statistical theory *JRSS(B)* has a quite diverse market—indeed, this journal is an anomaly in many of our comparisons for a reason that will be discussed later. Another class of exceptions is the group of the five most concentrated journals. These are (with the exception of the then fairly isolated East German journal *Biometrical Journal*) journals of statistical theory that are primarily importers, and in most cases their Gini concentrations are driven by a relatively high proportion of self-citations. *Communications in Statistics—Theory* is such a large importer that its imports even eclipse the self-citations of several journals. It provides 19.8% of the citations received by the *Annals of the Institute of Statistical Mathematics*, 11.2% and 45.6% for *Sankhyā A* and *B*, 25.2% for *Communications in Statistics—Simulation*, 10.5% for *Technometrics*, 9.1% for *The ISI Review*, 8.2% for *Biometrika* and 6.4% for *JASA*.

Table 3 also presents the results from separately computing the Gini indexes for citations to older issues (those at least a decade old in 1988) and newer issues (those less than a decade old). These highlight the fact that archival issues export their influence to a much broader and more diverse market than do the more recent issues, in all cases except for a few of the newer journals whose earlier issues were presumably not widely circulated.

Two other measures that could be interpreted as giving an indication of the range of diversity are the counts, given in Table 1, of the total number of journals citing the indicated journal in a year (the range of the export market), and the total number of journals cited in a year (the range of the import market). As would be expected, the range of the export market rather closely reflects the Gini index for the journal. The range of the import market, however, is remarkably small, whether the raw count is used, or the ratio of cited journals to citations sent. Exports vary widely across journals; imports are fairly constant.

This latter finding is somewhat surprising. That the number of imports is fairly constant might be expected: different journals publish comparable numbers of articles in a year, and different articles include comparably long reference lists. But that these reference lists should include reference to a comparable total number of *different* journals suggests that the producers of articles—the authors—draw their resources from an equally broad spectrum of sources in the aggregate, whether they are working in abstract theory or on applied methodology. For example, in 1987–88 *The Annals of Probability* cited an average of 443 different journals while *Biometrics* cited an average 449 different journals. During the same period, *Biometrics* exported to *seven times* as many journals as did *The Annals of Probability*. (In part, this can be attributed to the differing ages of the two journals; if only recent citations are counted the figure of 7 to 1 drops to about 4 to 1.) These two journals are of approximately equal length, and some of the variation in imports that is observed for other journals can be attributed to the number of pages published or, as here, differing ages. But even allowing for these factors, considerable diversity remains, making the remaining uniformity remarkable: not only does each article draw upon, on average, a like amount of acknowledged source material, the material is drawn from an equally diverse set of sources.

7. BILATERAL TRADE IN INTELLECTUAL INFLUENCE: THE EXPORT SCORES MODEL

The univariate statistics for separate journals tell us about their characteristics, but they do not begin to capture the patterns of interrelation among the journals. By carrying the analogy to economic trade a step further, we can construct import-export tables whose analysis gives interesting insights into some of those patterns. Because of the sparseness of some links—some journals "trade" with others so infrequently that the counts corresponding to their bilateral relationship are not readily available—this analysis will be restricted to the "high count" journals—those with large numbers of exports, imports or both—and to subsets of journals within specialized areas that do trade among themselves relatively frequently.

Table 4 describes the bilateral trade among eight statistics journals, giving the total citation count for the years 1987–89. The figures here represent only those citation "trades" where both parties were among these eight. Thus, we see by comparing Tables 1 (which gives yearly averages) and 4 (which gives three-year totals) that, for example, *JASA* received a total of $3 \times 5436 = 16308$ citations in 1987–89 from all sources in the *ISI* data base and that 3361 or 21% of them were from the eight "high count" journals of Table 4, 1072 of these (or 6.6% of the total) being self-citations.

There are $\frac{8 \cdot 7}{2} = 28$ different bilateral trade links involving these eight journals, and there is no a priori guarantee that they can be summarized in any succinct manner. We shall see, however, that a quite simple model, which we will call the export scores model, captures most of the information contained in this and several other bilateral trade tables, employing one less parameter than the number of journals included. The idea is a simple one. We will determine an export score, S_i, for each journal with the following interpretation: if a citation involves

TABLE 4
Cross-citations involving statistics journals, giving total citations for the years 1987–89. Rows correspond to citing journal, columns to cited journal, so a total of 155 citations involving Biometrics citing The Annals of Statistics were recorded in 1987–89. Fitted values for the export scores model are shown for cells with Pearson residuals above 1.7 in absolute value

	AnnSt	*Biocs*	*Bioka*	*ComSt*	*JASA*	*JRSSB*	*JRSSC*	*Tech*	Totals
AnnSt	1623	42	275 (305.7)	47	340	179	28	57	2591
Biocs	155	770	419	37	348	163	85 (103.5)	66	2043
Bioka	466	141 (169.1)	714	33	320	284	68	81	2107
ComSt	1025	273	730	425	813	276	94	418	4054
JASA	739	264	498	68	1072	325	104	117	3187
JRSSB	182	60	221	17	142	188	43 (33.0)	27 (40.8)	880
JRSSC	88	134 (115.5)	163	19	145	104	211	62	926
Tech	112	45	147	27 (40.8)	181	116	41	386	1055
Totals	4390	1729	3167	673	3361	1635	674	1214	16843

two of the journals represented in Table 4 and we were only told that it involves *Journal A* and *Journal B*, then the logodds that it is A exporting to B (i.e., B citing A) rather than B exporting to A is given by the difference of scores $S_A - S_B$;

$$\text{logodds}(A \text{ exports to } B \mid A \text{ and } B \text{ trade}) = S_A - S_B.$$

Because only differences of scores are involved, the baseline is arbitrary; and we resolve this ambiguity by assigning the score 0.0 to one of the journals: *The Annals of Statistics* in the case of Table 4. Thus, only seven scores remain to be determined.

TABLE 5
The export scores for the journals of Table 4, with the loglikelihood ratio statistics, all on 21 df. The model was fit to the data all together, and also separately to the older (a decade old) and newer citations

	Older	Recent	All
AnnSt	0.00	0.00	0.00
Biocs	−1.27	−1.07	−1.19
Bioka	−.23	−.44	−.35
ComSt	−4.89	−2.28	−3.27
JASA	−.93	−.70	−.81
JRSSB	.37	−.40	−.06
JRSSC	−1.73	−1.03	−1.30
Tech	−1.06	−.91	−.98
	$G^2 = 38.9$	$G^2 = 48.8$	$G^2 = 58.5$

If the citations of Table 4 are viewed as arising from a multinomial sampling model (a step whose appropriateness is certainly not obvious, but will be defended nonetheless), then many statisticians will recognize the export scores model as a familiar one presented under an unfamiliar name. For example, it is then equivalent to the *Bradley-Terry model* for paired comparisons (see David, 1988). It is also equivalent to the extraction of a subset of the parameters from the *model of Quasi-Symmetry* for square contingency tables (see Goodman, 1965, 1979; Caussinus, 1966). It is also equivalent to the estimation of what are called *gravity parameters* for a mobility table in demography (see McCullagh and Nelder, 1989, pp. 270–273). (Some of the interrelationships among these models are explored in Fienberg and Larntz, 1976.) Because none of these names conveys the idea we emphasize and none enjoys exclusive use, we have introduced a new name for what is arguably an old idea.

The interpretation of the fitted model is very simple. For example, refer to the scores given in the column "All" of Table 5, obtained by fitting the model to the data of Table 4. There we have

$$\text{Export score for } \textit{Biometrika} = -0.35,$$
$$\text{Export score for } \textit{JASA} = -0.81,$$

leading to the estimate that

logodds(*Biometrika* exports to *JASA*|*Biometrika* and *JASA* trade) $= -0.35 - (-0.81) = 0.46$,

or odds of $\exp(0.46) = 1.58$ to 1 in favor of the trade being an export from *Biometrika* to *JASA* rather than vice versa. The qualitative interpretation is simply this:

> The larger the export score, the greater the propensity to export intellectual influence.

It is worth noting that this model has two significant bonuses. First, by focusing upon bilateral trade, it is entirely insensitive to journal-level self-citations. Second, it focuses upon a limited subset of journals and is not at all affected by the possibility of large counts involving groups of journals outside the subset, a possibility that we shall see can have a significant effect upon univariate studies.

8. POSSIBLE SIZE BIAS IN THE EXPORT SCORES MODEL

If it can be argued that the export scores model fits the data, the model would be an appealing one—a complicated set of pairwise relationships would be summarized by a linear ordering of influence. A bivariate structure would be reduced to a simple univariate scale susceptible to easy interpretation. But what of the possibility that the scores will unduly reflect journal size and that the "propensity to export intellectual influence" that they purport to capture is confounded with the editor's indulgence and the publisher's extravagance? Fortunately, that proves not to be the case.

One way to see the absence of a size bias in the export scores model is to ask how the odds that the scores represent are affected by the construction of an artificial merger of two journals of similar characteristics. If the simple act of merger into a larger journal increases the appearance of intellectual influence, then the score would be reflecting, at least in part, the sheer size of the journal; and we would be forced to either abandon this measure or move to the question of what it tells us about the optimum design of scientific journals. But it turns out that merger does not affect scores, all else being equal.

To demonstrate this lack of size bias, let A', A'' and B be three journals. Suppose that A' and A'' are aggregated into a single journal A (where $A = A' \cup A''$); how does this affect the scores? The answer is simple. The odds on the direction of citation involving the journal B with respect to the combined journal A is a weighted average of the corresponding odds for the constituent parts A' and A''.

To prove this, let "AcB" represent the event that a citation involving A and B has A citing B. Then the odds for the constituent parts are

$$O' = \frac{P(A'cB)}{P(BcA')}$$

and

$$O'' = \frac{P(A''cB)}{P(BcA'')},$$

and the odds for the combined journal are

$$\begin{aligned} O &= \frac{P(AcB)}{P(BcA)} \\ &= \frac{P(A'cB) + P(A''cB)}{P(BcA)} \\ &= \frac{P(BcA')\frac{P(A'cB)}{P(BcA')} + P(BcA'')\frac{P(A''cB)}{P(BcA'')}}{P(BcA)} \\ &= \lambda O' + (1-\lambda)O'', \end{aligned}$$

where $\lambda = P(BcA')/P(BcA)$.

Thus, size alone cannot produce the appearance of influence. If a large journal A were randomly broken into two identical parts A' and A'', we must necessarily have

$$O' = O'' = O.$$

9. FITTING AND INTERPRETING THE EXPORT SCORES MODEL

To anyone who has studied citations, it seems simplistic to consider them as arising from independent multinomial trials, homogeneous within the categories of a modest cross-classification. Citations travel in clusters, and the life-histories of citations received by a single paper differ widely by the characteristics of the individual papers. In addition, even if citations could be well-described by a simple sampling model, there is no a priori reason to suppose that a complex of bilateral relationships can be mapped onto a linear scale as we propose to do. Journals could well exhibit what in the theory of voting are called Condorcet cycles, where journal A tends to cite B tends to cite C tends to cite A.

Our approach is this: despite a priori misgivings about the applicability of multinomial sampling, we shall employ this hypothesis to fit the export scores model, and we will even cautiously base an examination of the goodness of fit on it. We will find that despite worries about the inappropriateness of that sampling model (all of which, involving clustering of some sort, would lead in the direction of lack of fit), the fits we find either are excellent or show departure in ways that are amenable to simple explanations. One conclusion will be that, despite these misgivings about the appropriateness of a multinomial hypothesis, the treatment of citations aggregated to the level of journal effaces the inhomogeneities that we expect to find at the level of the individual paper or author. (We will discuss a partial exception to this, to be termed the *Cox effect*, below.)

We should emphasize that while the multinomial hypothesis will be employed to fit the model and evaluate the goodness-of-fit, it is not essential to the interpretation of the fitted values or the export scores. The export scores themselves can be interpreted as they were introduced above, in terms of a bet on the directionality of citation of a randomly selected citation. The multinomial hypothesis is one dynamic mechanism upon which this random selection can be conceptually based, but it is not the only one. It does link our interpretation to the social mechanism behind the counts and hence makes it relevant to inference about that larger world.

TABLE 6
Cross-citations involving statistics journals, grouped as 'theoretical', 'hybrid', and 'applied/methodological', giving total citations for the years 1987–89. Rows correspond to citing journal, columns to cited journal

	AnnStat, Bioka, JRSS(B)	JASA	Biocs, Tech, JRSS(C)	Totals
AnnStat, Bioka, JRSS(B)	4132	802	547	5481
JASA	1562	1072	485	3119
Biocs, Tech, JRSS(C)	1467	674	1800	3941
Totals	7161	2548	2832	12541

TABLE 7
The export scores for the journals of Table 6; the loglikelihood ratio statistic is $G^2 = 0.01$ on 1 df, a very good but not extraordinary fit

AnnStat, Bioka, JRSS(B)	0.00
JASA	−.66
Biocs, Tech, JRSS(C)	−.99

10. FITTING THE MODEL: STATISTICS JOURNALS

Because of the relationship between the export scores models and the model of quasi-symmetry, it can be easily fitted as a loglinear model. Bishop, Fienberg and Holland (1975, pp. 289–291) and Agresti (1990, pp. 354–355 and problem 10.27, p. 382) explain one way this can be done. Create an artificial third dimension to the two dimensional table consisting of the original table transposed about the diagonal, then fit the resulting three-dimensional table by the loglinear model with all two-factor interactions. The export scores are given as the two-factor interactions involving the original rows and the artificial third dimension, the log likelihood ratio statistic G^2 is half the nominal one, and the degrees of freedom for an $n \times n$ table is $(n-1)(n-2)/2$. The following analyses were all performed this way using the package GLIM.

The model was fit to the data of Table 4 in three ways: as the data are presented (*all* citations), and disaggregated into citations to *recent* issues (those published in the preceding decade) and citations to *older* work (papers published more than a decade before the citation was recorded).[5] Although the journals studied here are all well-established [all but the *Communications in Statistics* (1973) date from before 1960 and most from before 1950], looking at the scores based on only recent citations gives results unaffected by differences in the journals' ages. The results are presented in Table 5. The scores for all citations show *The Annals of Statistics* as the leading exporter of intellectual influence (as measured by citations) closely followed by *JRSS(B)* and *Biometrika* and then by *JASA* and *Technometrics*. *Biometrics* and *JRSS(C)* (*Applied Statistics*) occupy an intermediate position, and *Communications in Statistics—Theory and Methods* falls at an extreme even when only recent citations are used. The scores for recent citations tend to be more closely grouped, and those for older citations more dispersed, suggesting that tendencies towards a hierarchy of influence become more pronounced as the role of the papers becomes more archival. The *JRSS(B)* exhibits a slight anomaly here, and a partial explanation for this is given in Section 14. (The fact that the scores in the aggregated category *All* are intermediate between those of *Older* and *Recent* is a consequence of the result in Section 8).

The fitted values in all three cases capture the main message of the data, but in no case can the fit be said to be good when judged in terms of sta-

[5]The model was also fit separately for each single year's data, to test for stability over time; those results, which were quite satisfactory, are not shown.

TABLE 8

Cross-citations involving probability journals, giving total citations for the years 1987-88. Rows correspond to citing journal, columns to cited journal, so a total of 121 citations involving Journal of Applied Probability citing The Annals of Probability were recorded in 1987-88

	AnnPr	*PrTh/ZW*	*StochPr*	*JAP*	*ThPrApp*	*AnnSt*	Totals
AnnPr	468	255	33	46	72	74	948
PrTh/ZW	333	322	47	47	72	76	897
StochPr	208	155	93	76	40	41	613
JAP	121	31	37	283	26	35	533
ThPrApp	101	60	23	38	344	63	629
AnnSt	76	81	13	14	50	1009	1243
Totals	1307	904	246	504	604	1298	4863

TABLE 9

The export scores for the journals of Table 8, with the loglikelihood ratio statistics, all on 10 df. The model was fit to the data all together, and also separately to the older (a decade old) and newer citations

	Older	Recent	All
AnnPr	0.00	0.00	0.00
PrTh/ZW	−0.60	−0.21	−0.35
StochPr	−1.80	−1.30	−1.51
JAP	−0.79	−0.44	−0.62
ThPrApp	−0.41	−0.68	−0.46
AnnSt	−0.23	−0.17	−0.19
	$G^2 = 49.5$	$G^2 = 10.4$	$G^2 = 35.6$

tistical significance with respect to the multinomial sampling model. The reason for this is apparent by examining the largest residuals. Table 4 shows the fitted values for those cells with the largest Pearson residuals: the cells that are responsible for the lack-of-fit. Of the seven cells in question, four of them show a lower count than predicted and involve citations between journals published on opposite sides of the Atlantic Ocean! Despite the international character of academic statistics, there appears to be a residual geographic effect that is not accounted for in the model. This effect is small, however, compared to what the model does account for.

11. THEORY AND APPLICATIONS

The results of Section 8—that export scores are not biased by merging journals—can be put to constructive use. Table 6 is based on the seven main exporting journals of Table 4 but with the three principal theoretical journals merged together and the three applied journals similarly merged. The remaining journal, *JASA*, combines theory and application in one package and is placed in a class by itself. The fit here is excellent by any criterion, and the export scores testify to a striking pattern. There is a strong tendency for the theoretical journals to export intellectual influence to the applied. The estimated logodds that a citation involving a theoretical journal and an applied journal has theory being cited by applications is $0.00-(-0.99) = 0.99$, or odds of 2.7 to 1. Thus, we have striking evidence supporting a fundamental role for basic theory that runs strongly counter to the sometimes voiced claim that basic theory is not relevant to applicable methodology. Instead, we have a well-marked trail of the influence of theory on the applied journals. One might well ask then, what influences theory? The present data do not give a full answer, but they do suggest that much of the influence is internal to theory. If there is appreciable influence from outside, it is not strongly from the applied and methodological journals or not of a type that is captured by citations. Some may view this as a negative finding; however, we prefer to emphasize the positive aspect: that a strong and important role for basic theory in the scientific discourse of this important field has been decisively demonstrated.

12. PROBABILITY JOURNALS

An examination of the citations involving probability journals produces one unexpected result. Tables 8 and 9 show the bilateral trade among five major probability journals and *The Annals of Statistics*. The unexpected result is related to the fact that *The Annals of Statistics* is the only statistics journal included for the simple reason that it is the only statistics journal that generated sufficient counts in the ISI data base to permit inclusion. Even the counts for its trade with probability journals are fairly low. That is the unexpected result; not the direction of what trade there is but the lack of any significant volume of trade at all! We are accustomed to refer to "probability and statistics" in one breath; concep-

TABLE 10
Cross-citations involving econometrics journals, giving total citations for the years 1987–89. Rows correspond to citing journal, columns to cited journal, so a total of 466 citations involving Biometrika citing The Annals of Statistics were recorded in 1987–89

	AnnSt	*Bioka*	*Ecmica*	*J Ecncs*	*JASA*	*JRSS(B)*	Totals
AnnSt	1623	275	38	15	340	179	2470
Bioka	466	714	32	17	320	284	1833
Ecmica	106	39	675	105	66	19	1010
J Ecncs	147	81	493	278	195	69	1263
JASA	739	498	97	52	1072	325	2783
JRSS(B)	182	221	20	15	142	188	768
Totals	3263	1828	1355	482	2135	1064	10127

tually, the link is extremely close. But at the level of the journal article, there is remarkably little communication between them. If the counts of Table 8 are collapsed to a 2×2 table (the two categories being "probability journals" and "Annals of Statistics"), the logodds ratio is $\theta = 49.7$; if the statistics category were enlarged to include *Biometrika*, this increases to $\theta = 62.3$. Journals of probability stand for the most part as intellectual islands in the strait between statistics and mathematics, with stronger bridges to the mainland of mathematics than to statistics.

What trade there is between the two *Annals* is fairly even, with a slight tendency for probability exports to exceed statistics exports. Among the probability journals themselves, there is a hierarchy with *The Annals of Probability* as the leading exporter, then *Probability Theory and Related Fields* (formerly *Zeitschrift für Wahrschleinlichkeitsrechnung*), the Russian journal *Theory of Probability* and *The Journal of Applied Probability* and finally *Stochastic Processes* a distant fifth (the odds are 4.5 to 1 that it cites the *Annals* rather than vice versa). The lack of fit in this case is due to the fact that several of these journals are not really general journals, but they deal in narrower and only partially overlapping specialties: applied probability and the theory of stochastic processes. The export scores model can represent the main drift of the trade involved, but it is insufficiently rich to capture the detailed structure of the market.

13. ECONOMETRICS AND ECONOMICS JOURNALS

Another question of interest that can be addressed by the study of export scores is the communication between statistics and other fields. A major area of interaction is that between statistics and economics through the interface that has come to be known as econometrics. Tables 10 and 11 relate to one portion of this link, the communication of statistics and econometrics, and Tables 12 and 13 to the other, the communication of econometrics and economics. Table 12 is based on a larger study of communication in economics reported in Stigler and Stigler (1992).

TABLE 11
The export scores for the journals of Table 10; the loglikelihood ratio statistic is $G^2 = 17.37$ on 10 df, a good fit

AnnSt	0.00
Bioka	−.45
Ecmica	−.61
J Ecncs	−2.13
JASA	−.86
JRSS(B)	−.12

The trade between statistics and econometrics is much as one would expect from the earlier results. The more theoretical statistics journals are exporters to econometrics, much more so in the case of *The Journal of Econometrics* than *Econometrica*. The odds are better than eight to one that *The Journal of Econometrics* cites *The Annals of Statistics* than vice versa. *Econometrica* is actually a slight net exporter to *JASA* (odds of 5 to 4) but an importer from *The Annals* (odds of 9 to 5). Within economics, the tendency for theory to export to applications continues; here, as was found in Stigler and Stigler (1992), *Econometrica* is the leading exporter. Thus as far as bilateral trade in the commodity "statistical theory and methodology" goes, we can reconstruct a directional trade route from the more theoretical statistics journals to econometrics to economics. But here, as with the probability journals, there is relatively little trade. If Table 10 is collapsed to a 2×2 table, the categories being "statistics" and "econometrics," the logodds ratio is $\theta = 56.8$: a value of the same order of magnitude as that for probability journals.

TABLE 12

Cross-citations involving economics journals, giving total citations for the years 1987–90. Rows correspond to citing journal, columns to cited journal, so a total of 440 citations involving the Economic Journal citing the American Economic Review were recorded in 1987–90

	AER	*EJ*	*Ecmca*	*JET*	*JPE*	*QJE*	*RESt*	*RESd*	Totals
AER	1090	108	423	136	652	275	120	209	3013
EJ	440	369	458	103	375	160	55	208	2168
Ecmca	157	44	894	269	170	92	35	229	1890
JET	123	22	669	754	126	100	5	252	2051
JPE	373	61	321	108	646	155	75	155	1894
QJE	301	43	204	98	230	220	32	114	1242
RESt	416	66	416	25	310	98	574	97	2002
RESd	140	40	379	189	166	94	12	291	1311
Totals	3040	753	3764	1682	2675	1194	908	1555	15571

14. THE COX EFFECT

In these analyses we have been treating the journal as the unit of analysis. The interpretation of the results would be complicated if the counts recorded for any single journal depended primarily upon citations attracted by any single paper or small group of papers. Examination of the data shows that when they refer to a homogeneous data base—when we concentrate upon citations from source journals in statistics—the problem is not serious. But there is a special situation where the problem can arise, namely when a paper attracts significant attention in an area of application where a much higher rate of citation is the norm. In particular, this can happen when a paper treats a method that becomes widely used in the biomedical sciences. We will call this effect, where the attraction of a single paper to readers outside the field of study has a large effect upon its own journal's counts, the *Cox effect*, after the author of the paper in this study where this effect is most pronounced.

The average numbers of citations per year over 1987–89 are given in Tables 14 and 15 for six papers where the counts are given separately for all sources, for the 33 statistics journals of Table 1 and for the eight statistics journals of Table 4. Three of these papers are the three contributing the largest share to their journal's counts from all sources, and three are "high count" papers whose primary appeal is to statistics sources, included for comparison. By far the major impact is associated with D. R. Cox's (1972) *JRSS*(*B*) paper that introduced methods for regression analyses for life tables and censored data. Cox (1972), with nearly 600 citations a year, accounts for more than a quarter of the total citations received by *JRSS*(*B*) in a year. The second greatest impact is that of E.L. Kaplan and P. Meier's (1958) *JASA* paper introducing the "Kaplan-Meier" estimator for estimating survival curves from censored data. Kaplan and Meier (1958), with over 800 citations a year, actually continues to receive more citations than Cox (1972); but the impact upon a larger journal is much less.[6] The third largest impact is that of David Duncan's (1955) *Biometrics* paper on multiple range tests, a methodological discussion that is now all but ignored in statistics journals. In contrast, even highly cited papers that do not attract significant attention in the biomedical literature such as Box and Cox (1964); Dempster, Laird and Rubin (1977); and Efron (1979) do not have nearly so large an impact upon their journal's counts.

TABLE 13

The export scores for the journals of Table 12; the loglikelihood ratio statistic is $G^2 = 25.4$ on 21 df, an excellent fit

AER	0.00
EJ	−1.32
Ecmca	1.00
JET	.12
JPE	.44
QJE	.04
RESt	−1.26
RESd	.41

The three top papers do have an impact upon the statistical analyses of Tables 1 through 3. Since this impact, although concentrated, is nonetheless real, no attempt has been made to remove it. What about the analyses of Tables 4 through 13? Here the greatest impact is that of Cox (1972), which still is responsible for 4.6% of the citations *JRSS*(*B*) receives in a year from the eight journals of Table 4.

[6] By comparison, the most cited paper in the biomedical literature in 1988 received nearly 10,000 citations.

TABLE 14
The Cox effect for all sources. The average numbers of citations 1987–89 *for six high-count papers, and their percentages of their journal's yearly average count (from Table* 1)

Paper and Journal	Citations—All Sources 1987–89 Paper's Yearly Average	Journal's Yearly Average	Paper's Percent
Cox (1972) *JRSS(B)*	581	2207	26.3
Kaplan and Meier (1958) *JASA*	813	5436	15.0
Duncan (1955) *Biometrics*	453	3916	11.6
Box and Cox (1964) *JRSS(B)*	89	2207	4.0
Dempster, Laird, Rubin (1977) *JRSS(B)*	85	2207	3.8
Efron (1979) *Annals of Statistics*	59	4462	1.3

TABLE 15
The Cox effect for statistics sources. The average numbers of citations 1987-89 *from the* 33 *journals of Table* 1 *for six high-count papers, and from the* 8 *journals of Table* 4, *with the latter's percentages of the journal's yearly average count from these eight sources (from Table* 4)

	Citations–33 Stat Journals	Citations—All Sources 1987–89		
Paper and Journal	Paper's Yearly Average	Paper's Yearly Average	Journal's Yearly Average	Paper's Percent
Cox (1972) *JRSS(B)*	32.0	25.3	545.0	4.6
Kaplan and Meier (1958) *JASA*	21.7	17.3	1120.3	1.5
Duncan (1955) *Biometrics*	0.3	0.3	576.3	0.0
Box and Cox (1964) *JRSS(B)*	16.3	13.7	545.0	2.5
Dempster, Laird, Rubin (1977) *JRSS(B)*	24.0	18.3	545.0	3.4
Efron (1979) *Annals of Statistics*	21.3	18.7	1463.3	1.3

All three analyses of these data were repeated omitting the Cox (1972) citations to gauge the magnitude of its effect. The change was small. Only the export score for *JRSS(B)* was affected appreciably, and the effect there was to lower the score for all citations by 0.06 from −0.06 to −0.12, and to lower the score for older citations by 0.09 from 0.37 to 0.28. Since this is the extreme case, the conclusion is that in general the effect of very small groups of papers upon export scores is not large. This does not alter the fact, however, that the distribution of citations within journals is very far from uniform.

15. CONCLUSION

The analyses presented here have permitted many conclusions of broad generality; that, for example, those aspects of intellectual influence that are represented well by citations tend to flow on balance from more abstract work to more applied work and within specialties in ways represented by the export scores. We find that the size and diversity of journals' exports are considerably more variable than the size and diversity of imports, suggesting great differences among journals in either the efficiency of production or the product itself or both. Most of these conclusions are subject to qualifications, such as that the intellectual influence which is measured by citations, in spite of its multifaceted nature, still does not capture all aspects of that elusive concept. But despite many potential reservations, a sufficiently coherent picture emerges from the analyses to encourage us to believe that something real and fundamental is represented in these numbers.

The recognition of the Cox effect points out the problems of incorporating within one analysis papers playing very different roles in two different constituencies. But even there, the effect is not truly misleading: such papers are actually immensely influential, in Cox's case in both of two major areas, and the citation counts properly interpreted reflect that influence. We can believe that that paper itself was greatly influenced by applied problems *and* at the same time conclude that that influence has been many, many times reciprocated in the ensuing intellectual transactions. There is a message here that should be highly congenial to statisticians: a small investment in statistical theory reaps rewards many times greater, even when examined in terms of journal level aggregates.

APPENDIX—A NOTE ON THE *ANNALS*

The citations to the various *Annals* of the Institute of Mathematical Statistics presented a special problem. After the 1972 volume, *The Annals of Mathematical Statistics* was succeeded by *The Annals of Statistics* and *The Annals of Probability*. To attempt to deal with this and give the new *Annals* appropriate "histories," an elaborate analysis was attempted with the data from 1987–88, where the citations to *The Annals of Mathematical Statistics* were imputed to the successors proportionally to the citations that both received during the years 1973–1978; this proportional allocation was done separately for every source journal where the data permitted. It was found that (as noted above) the propensity for probability and statistics sources to cite each other is so slight that this had no effect except in the analysis of Table 9, the probability journals. Accordingly the subdivision was only retained in that case, and in other analyses the counts of *The Annals of Mathematical Statistics* were aggregated with those of *The Annals of Statistics,* permitting the inclusion of 1989 data in those cases. This is equivalent to assuming in 1973 *The Annals of Mathematical Statistics* simply changed its name to *The Annals of Statistics*, and that is the same year a new journal called *The Annals of Probability* was begun.

ACKNOWLEDGMENTS

This research was supported in part by NSF Grant DMS-92-04504. I am grateful to Claire Friedland, Eugene Garfield, Peter McCullagh, Robert K. Merton, Frederick Mosteller, Richard Posner, Michael Stein, David Wallace and Harriet Zuckerman for comments.

REFERENCES

AGRESTI, A. (1990). *Categorical Data Analysis*. Wiley, New York.

BISHOP, Y. M. M., FIENBERG S. E. and HOLLAND, P. W. (1975). *Discrete Multivariate Analysis*. MIT Press.

BOX, G. E. P. and COX, D. R. (1964). An analysis of transformations. *J. Roy. Statist. Soc. Ser. B* **26** 211–246.

CAUSSINUS, H. (1966). Sur la structure des tableaux de corrélation carrés. *Comptes Rendus de l'Académie des Sciences de Paris* **263** 795–797.

COX, D. R. (1972). Regression models and life tables. *J. Roy. Statist. Soc. Ser. B* **34** 187–220.

DAVID, H. A. (1988). *The Method of Paired Comparisons*, 2nd ed. Griffin, London.

DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B* **39** 1–38.

DUNCAN, D. B. (1955). Multiple range and multiple *F* tests. *Biometrics* **11** 1–42.

EFRON, B. (1979). Bootstrap methods: Another look at the jackknife. *Ann. Statist.* **7** 1–26.

ELKANA, Y., LEDERBERG, J., THACKRAY, A., MERTON, R. K. and ZUCKERMAN, H., eds. (1978). *Toward a Metric of Science: The Advent of Science Indicators*. Wiley, New York.

FIENBERG, S. E. and LARNTZ, K. (1976). Log linear representation for paired and multiple comparisons models. *Biometrika* **63** 245–254.

GARFIELD, E. (1972). Citation analysis as a tool in journal evaluation. *Science* **178** 471–479.

GARFIELD, E. (1979). *Citation Indexing: Its Theory and Application in Science, Technology, and Humanities*. Wiley, New York.

GARFIELD, E., ed. (1987–1989). *SCI Journal Citation Reports: A Bibliometric Analysis of Science Journals in the ISI Data Base*. Institute for Scientific Information, Philadelphia.

GOODMAN, L. A. (1965). On the statistical analysis of mobility tables. *American Journal of Sociology* **70** 564–585.

GOODMAN, L. A. (1979). Simple models for the analysis of association in cross-classifications having ordered categories. *J. Amer. Statist. Assoc.* **74** 537–552.

KAPLAN, E. L. and MEIER, P. (1958). Nonparametric estimation from incomplete observations. *J. Amer. Statist. Assoc.* **53** 457–481.

MCCULLAGH, P. and NELDER, J. A. (1989). *Generalized Linear Models*, 2nd ed. Chapman and Hall, London.

MACROBERTS, M. H., and MACROBERTS, B. R. (1986). Quantitative measures of communication: A study of the formal level. *Social Studies in Science* **16** 151–172.

MERTON, R. K. (1968) *Social Theory and Social Structure*, enlarged ed. Free Press, New York.

MERTON, R. K. (1973). *The Sociology of Science* (N. W. Storer, ed.). Univ. Chicago Press.

PINSKI, G., and NARIN, F. (1976). Citation influence for journal aggregates of scientific publications: Theory, with application to the literature of physics. *Information Processing and Management* **12** 297–312.

STIGLER, G. J. (1968). *The Organization of Industry*. Richard D. Irwin, Homewood, III.

STIGLER, G. J. and STIGLER, S. M. (1992). The journals of economics [paper by G. J. Stigler, with the assistance of C. Friedland, with appendix by S. M. Stigler]. To appear.

STIGLER, S. M. (1978). Laplace's early work: Chronology and citations. *Isis* **69** 234–254.

STIGLER, S. M. (1987). Precise measurement in the face of error: A comment on MacRoberts and MacRoberts. *Social Studies in Science* **17** 332–334.

ZUCKERMAN, H. (1987). Citation analysis and the complex problem of intellectual influence. *Scientometrics* **12** 329–338.