Cutting Cataloging Costs: Accepting LC Classification Call Numbers from OCLC Cataloging Copy

Susan A. Massey and S. Michael Malinconico

Cataloging policy at the University of Alabama Libraries allows the acceptance of LC classification call numbers from OCLC cataloging copy into the local database without shelflisting. In this study, we measured error rates for locally unshelflisted samples and a control group of locally assigned and shelflisted call numbers to determine whether this policy produces disarrangement of the local online shelflist. The results show no significant differences between samples, indicating that catalogers' task of local shelflisting is not a cost-effective use of their time. An analysis of the error data suggests that the types of disorder created by shelflisting errors would not impede the retrieval of items while subject browsing, but further study is needed to confirm this.

Local Shelflisting Policy

The University of Alabama (UA) Libraries utilize a national bibliographic utility, the OCLC Online Computer Library Center, Inc. (OCLC), to provide Machine-Readable Cataloging (MARC) records that are downloaded into the local online cataloging system (NOTIS). Since the migration of local online records to a new system in 1990, UA's catalog department has accepted OCLC cataloging copy for monographs without locally shelflisting Library of Congress (LC) classification call numbers already in the records, whether assigned by LC or by a participating library. Regardless of the source of the record, call numbers are not checked against the existing online shelflist or revised to ensure that items are located on the shelf in correct logical order.

Copy catalogers perform a quick visual check of call numbers to make sure there are no obvious problems such as incorrect

Susan A. Massey (smassey@ua1vm.ua.edu) is catalog librarian at the University of Alabama Libraries. S. Michael Malinconico (mmalinco@slism.slis.ua.edu) is EBSCO Professor at the School of Library and Information Studies at the University of Alabama. This research was supported by a grant from the Research Grants Committee of the University of Alabama, and by a Council on Library Resources, "Cooperative Research Grant for Librarians and Faculty Members." The authors wish to thank Usha Midkiff of University of Alabama Systems Development for programming the extraction of the samples from the local NOTIS database. Manuscript received July 25, 1996; accepted for publication November 14, 1996.

or missing subfields or punctuation. They then add an edition date to the call number if it is not already present, along with a lower-case *x* for all call numbers not assigned by LC. Records that lack call numbers or have questionable call numbers (i.e., those that have apparent typographical errors or appear unusual in some way) are routed to original cataloging librarians for review. Call numbers assigned or revised by the catalogers are also shelflisted to fit in the arrangement of the local online database.

Assigning an LC classification system call number consists of both classifying and shelflisting the item. An LC call number is composed of a class number that represents a subject area as designated in the LC classification schedules, a book number that arranges items within a class in a specified order, and any prescribed additional unique identifiers for a particular item. Classification involves choosing the class number. Shelflisting is the process of logically arranging materials in the collection by creating a unique call number for each item through the addition of a cutter number or other identifiers, such as edition dates, to the class number. This activity is achieved in the context of comparing the call numbers to others in the local shelflist, a file of bibliographic records reflecting the order of the materials on the shelves (Library of Congress 1995, G10, 12).

The policies of library cataloging departments may vary widely concerning the extent to which they review call numbers from cataloging copy. While one institution may check whole call numbers from all copy, another may accept class numbers while reviewing cutters only, expecting to find classification errors in the process of shelflisting. Some institutions may accept call numbers only from particular cataloging copy sources without review. These decisions may be based on the library's cataloging philosophy (Taylor 1988, 184), economics, or a combination of factors.

By accepting OCLC cataloging copy containing call numbers assigned by other institutions without shelflisting or checking the class schedules for correct classification, UA's catalog department streamlines workflow and reduces cataloging costs. Class number assignment is not an exact science, however, and shelflist order is shaped by the holdings in the local database, so accepting a call number assigned by another institution may result in local shelflist disarrangement. In the process of shelflisting to assign a new call number, occasionally a section of the UA online database that is cuttered out of order or contains an incorrectly classified item is discovered. The current research was designed to discover whether UA's local shelflisting policy creates extensive disorder in the local database and therefore should be reconsidered.

Serendipity at a Cost

Theoretically, the purpose of local shelflisting is to ensure that an item fits in order by author, title, geographic area, or some other criterion represented by a cutter number within a particular classification or subject area on the library's shelves. The reason for this concern is to enable effective shelf browsing by patrons. If items are out of order to the extent that they are far removed from other similar items, they can only be retrieved if their exact location or call number is known. One of the benefits of the subject arrangement of items on a shelf is depth of access to several full texts when searching for a precise bit of information not reflected in a catalog record (LeBlanc 1995, 296). This serendipitous discovery of information could be lost if call numbers were only locating devices. In the online environment, LC class numbers can be used in searching to increase precision in retrieval and as the basis for broad subject searches (Chan 1989, 531–33). There is less evidence that the correct order of items achieved by cuttering within a classification is a factor in a satisfactory online search. Neither is it clear whether near-perfect shelf arrangement of items within a subject area is necessary for successful stack browsing; moreover, maintaining an exact shelflist order for items may not be essential to information retrieval.

However, the size and nature of a collection may reduce the importance of shelf browsing by classification as a primary subject retrieval strategy. With the variety of information formats increasingly available in libraries, patrons may need to browse several physical locations to obtain a full range of materials in a given subject. This phenomenon increases the patron's dependence on the library catalog as a locating device. In addition, the proliferation of interdisciplinary subjects in recent years and classification of the same topic within a variety of class schedules depending on the discipline emphasized may mean that a single item is classed with only one aspect of its subject matter, and use of the class number as method of subject retrieval becomes impossible to achieve consistently (Taylor 1988, 172).

While not diminishing the importance of shelf browsing, catalog departments must weigh the cost of shelflisting and reviewing classification carefully against its perceived benefits. Assigning call numbers has traditionally been viewed as a duty requiring the expertise of a highly paid professional librarian (Bleil and Renner 1990, 100). Although in some libraries the editing of call numbers may be performed by support staff, it is still a time-intensive and therefore expensive procedure. One recent study indicated that literature items with LC author numbers already established required 3.09 minutes per title to shelflist (LeBlanc 1995, 299). Based on this estimate, if a catalog department the size of UA's were to revise its policy to include shelflisting the 22,000 monographs volumes cataloged annually, its workload would increase by 1,133 hours, requiring an additional 0.58 FTE position. Admittedly, projecting one library's findings for one classification to another library's entire operation may or may not be a reliable method. However, it does provide evidence that notable efficiencies can be achieved by not shelflisting.

As part of the current study, the sample items were manually shelflisted at an average rate of 50 call numbers per hour. This excluded several steps in normal online shelflisting, such as incidental database cleanup. Since the shelflisting was primarily performed from a computer printout, the time per item also did not include online searching and computer response time. This very low time estimate still indicates a required 440 hours to shelflist 22,000 volumes, or close to 0.25 FTE position. It is clear that changing the current policy in order to shelflist all items would be a costly endeavor. The only compelling reason for such a change would be the discovery of a high rate of database disorder resulting from the current policy, in conjunction with the frequent occurrence of types of shelflisting errors that seriously impeded patron browsing by placing titles on the shelf far removed from related volumes.

Tracking the Elusive Error Rate

One reason there has been little definitive research into the benefits of shelflisting may be the difficulty of extracting data to examine. Cataloging policies in the online environment are often fluid, depending on the technology available and the consequent evolution of work flow and procedures. The consistency of shelflisting policy at the UA Libraries and a long-term commitment to the same integrated library system provided a window in time when data were produced that could be sampled with confidence in the validity of the research results.

The main research question of the current study was broken into two parts: How much disorder is created in our local online database by accepting LC call numbers from OCLC cataloging copy without local shelflisting? and Is there a significant difference in the number of shelflisting errors caused by these unrevised call numbers and the error rate of call numbers that have been locally shelflisted? We also looked theoretically at the types of shelflisting errors we discovered in order to determine whether the kind of disorder produced appeared to impede patron browsing seriously. Obvious classification errors that placed an item among others about a different subject were included in the study because it is likely they would be detected in the process of shelflisting online, just as they were indeed found in the course of the study, although correct classification was not checked in the class schedules for every item in the samples.

To answer the main research question, we drew a sample of LC call numbers that had been accepted from OCLC member copy without local shelflisting and compared it to a control group of call numbers that had been locally assigned and shelflisted. Samples represented records added to the catalog between October 1990 and March 1995. The main sample of call numbers from copy cataloging included both MARC 050 (LC-assigned) and 090 (other locally assigned) fields in all classifications, excluding those records containing UA's OCLC symbol as the cataloging or modifying agency. Therefore, this sample included copy cataloging from all participating OCLC institutions except UA. The control group consisted only of our local original cataloging records. The parameter compared was the amount of shelflist disarrangement, measured by the number of shelflisting errors detected. Error for the copy cataloging sample was defined as a call number that placed an item in a different place in the local online shelflist than where a correctly locally assigned call number would normally fall. Error for the local sample was defined as an incorrectly assigned call number. Shelflist disarrangement in all samples included inappropriate classification, incorrectly assigned cutters, and typographical errors that would have been noticed and revised during the shelflisting process.

We also wanted to know whether error rates differed between LC-produced copy and the overall rate of disorder for cataloging copy. This was to provide research data for institutions that accept only LC copy without revision. A sample was drawn of records with call numbers in the MARC 050 field with second indicator 0 (which represents an LC call number assigned only by LC).

We also wondered whether error rates differed between certain classifications, depending on the complexity of the schedule and types of cuttering required. The reason for including this part of the study was that libraries with holdings heavily weighted in a particular subject area might have different shelflist errors than an institution with generalized holdings. Since the subject mix of local libraries can differ considerably, similar error rates across classifications would enhance the possibility of generalizing our results to other institutions with different holdings. For this part of the research, samples were selected from classifications P (literature), Q (science), and T (technology). We expected a wide variation between institutions in local author cutters for literature classes, while the precision and structure of the Q classes would appear to foster greater continuity among shelflists. The technology schedule is typical of classes that include geographic subdivisions, special topic subdivisions, and the same topic addressed in more than one subclass.

Drawing the Data

We developed a sample selection algorithm to choose call numbers from the NOTIS database for each sample. The parameters used a combination of MARC field definitions and NOTIS system fields, taking into consideration UA's local cataloging policies and procedures. Each target call number came from a record that was entered in the database later than the records immediately preceding it and following it in the online shelflist order. This insured that any disorder discovered was due to an error in the sample call number, not another call number entered later out of shelflist order. If the preceding or following call number was found to be in error, the sample call number was discarded. This parameter was dependent on a local system field indicating the date the record was loaded into the system. The study was limited to records entered after October 1, 1990, a date that immediately followed a series of tape loads that occurred when UA changed local systems. All records were given the entry date of the tape loads, making sample selection prior to that date impossible.

The study was limited to monographs records processed in the main library cataloging unit, indicated by a fixed field code (bib lvl "m") and a processing unit code. Nonbook formats at UA are arranged by schemes other than LC classification. Serials are normally locally shelflisted as a matter of policy. Other processing units on campus also have varying shelflisting policies.

Another parameter in the selection algorithm included a check to indicate that the call number had not been added or altered locally after the record was downloaded (the bib record call number must match the local copy holdings call number). This did not control for the possibility that the call number was locally added in OCLC before the record was downloaded, or that it was added to both the local bibliographic record and holdings screen after downloading, which would mean a cataloger had shelflisted the added call number. Therefore, all call numbers in NOTIS bibliographic records from the selected samples (except the control group) were compared to the call numbers in the corresponding OCLC records. Sample call numbers that differed were discarded.

Using the sample selection algorithm, a program was designed to choose the sample call numbers automatically. Since the local NOTIS database is updated in real time, a computer-generated shelflist frozen in time was used as the universe for sample selection. Samples were drawn on March 11, 1995. The programmer modified an inventory program to reproduce an exact replica of the NOTIS online shelflist order. Records were selected with a uniform random number generator, then the sample selection program was applied to each record. To compensate for items that were expected to be discarded for reasons discussed above, larger samples than needed were requested from the system. Then, during the manual processes of checking NOTIS call numbers against OCLC and shelflisting the sample items, call numbers that were discarded were simply replaced with the next available sample item until a total of 200 call numbers per sample was reached.

A report for each sample was printed showing a grouping of the selected call numbers with the call numbers immediately preceding and following them in the NOTIS shelflist order, along with the main entry, title field, and date of entry in the online catalog. These fields were considered basic for a quick visual scan of correct shelflist order. Sample items were printed in the order of random selection and numbered for identification on data analysis coding sheets.

Correct cutter number order and appropriateness of classification were determined by a visual check of the sample printouts. When a possible error was encountered by examining the printout data, or when the basis for cuttering was not immediately apparent from the main entry and title, the target call number was shelflisted in the NOTIS database as it would have been at the time of cataloging. This meant that bibliographic records, holdings screens, and classification schedules were consulted to determine the correct call number. All errors detected in sample call numbers were recorded on a coding sheet.

The types of errors we encountered were: the assignment of an incorrect classification for the subject matter (not when the question was cataloger's discretion, but an obviously wrong class number); the assignment of the wrong cutter number (i.e., it did not match the main entry or follow the cuttering instructions in the class schedule); and call numbers cuttered out of local shelflist order (based on the criterion used for cuttering). No duplicate call numbers were discovered, but these also would have been considered errors, since UA assigns unique call numbers to items. All categories of errors potentially contained typographical errors that could not be identified as a separate category, but in any case, should have been edited during the shelflisting process if the call numbers were checked at the time of cataloging. No obvious typographical errors, such as the transposition of two letters or numbers, were found.

Sample Size Determination and Statistical Analysis

In order to estimate the rate of occurrence of various kinds of shelflisting errors in the database, we drew random samples of entries resulting from different shelflisting policies. The fraction of each sample with incorrectly assigned call numbers was determined. Reduced to simplest terms, we sought to measure a binary variable (correct vs. incorrect call numbers) by examining a random sample of a large population (the entries in the UA catalog). It should be intuitively obvious that there is a correlation between the size of the sample studied and the reliability of the estimate derived from it. However, the amount of effort required to examine a sample increases in direct proportion to the size of that sample. Thus, considerable thought was given to the size of the sample to be drawn.

Optimum Sample Size

Counterintuitively, the size of the population from which a sample is drawn is of no consequence—provided certain reasonable conditions are met. However, the likelihood of occurrence of the condition to be investigated is an important factor. When random samples are drawn from a population and they are examined for a binary variable, it can be shown that the sample means are normally distributed. The sample mean serves as an estimate of the mean of the population and the standard deviation of the population can be estimated from the formula

$$\sigma = \sqrt{\frac{pq}{n}}$$

where p = the observed probability of occurrence of the condition under investigation, $q = 1 - p$, and n = the size of the sample (Hoel 1971, 82–85). The standard deviation has the property that approximately 68% of all measurements will fall within a range that is ± 1 standard deviation from the mean, and 95% of all measurements will fall within ± 1.96 standard deviations from the mean.

We examined preliminary samples in order to get an idea of the probabilities we would be seeking to measure. The preliminary samples were selected in the same manner as the study samples. These samples indicated that the likelihood of shelflisting error in the database was less than 10% but greater than 5%. We consequently developed table 1. Each column and each row corresponds to sample size and probability respectively.

TABLE 1

p/n	50	100	200	400	800	1600
2%	1.94	1.37	0.97	0.69	0.49	0.34
4%	1.36	0.96	0.68	0.48	0.34	0.24
6%	1.10	0.78	0.55	0.39	0.27	0.19
8%	0.94	0.66	0.47	0.33	0.23	0.17
10%	0.83	0.59	0.42	0.29	0.21	0.15
12%	0.75	0.53	0.38	0.27	0.19	0.13

The entry in each cell is 1.96 × the standard deviation that we would obtain with such a sample divided by the assumed probability, i.e., it is a measure of how closely we could expect to estimate the population mean if we were to use a sample of the size represented by that column. As can be seen from table 1, the intersection of a sample of 200 and a probability of 8% yields a ratio of less than 1/2. Doubling the sample size to 400 only reduces the ratio to approximately 1/3 and quadrupling it to 800 only serves to reduce it to 1/4. Thus, we chose 200 as our sample size, as this provides what we judged to be the optimum discrimination relative to the effort required to obtain it, i.e., the measurement uncertainty would be 1/2 or less of the value of the variable we would be seeking to measure. Furthermore, the projected uncertainties are also well within reasonable tolerances for a study such as this. Small differences, e.g., 2%–3%, would not be sufficient by themselves to sway a choice of cataloging policy.

Analysis of the Samples

Table 2 provides a summary of the results of examining call numbers in the samples selected. It is readily apparent that there is considerable overlap among the esti-

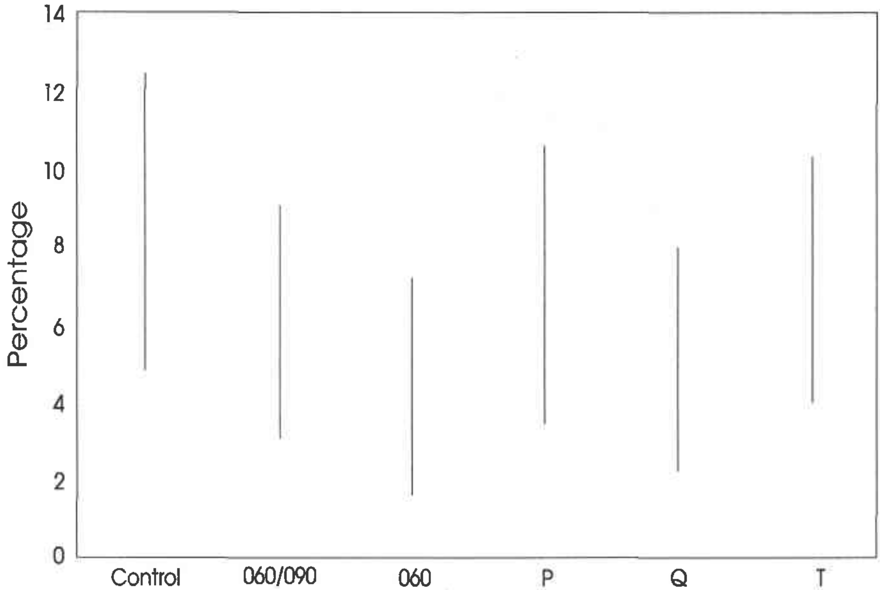

Figure 1. Comparison of Error Samples

mated population means for each of the database subsets studied. This can be clearly illustrated if the data are represented graphically. Each of the lines in figure 1 represents the range within which the population mean falls with 95% probability.

The real question is not so much what the individual error rates are—though they are of professional concern—but rather the *differences between mean error rates under different shelflisting policies*. We can estimate these differences by subtracting the associated sample averages. Clearly, if the differences are large, we can be relatively confident that we have correctly identified a significant difference in the consequences of the shelflisting practices under consideration. If they are very small, our confidence in the significance of that difference is correspondingly weak. Thus, we need a test to assess the significance of differences we measure. Such a test can be established by noting that the difference between the means of two normally distributed variables is likewise normally distributed.

TABLE 2

Sample	Total Errors	p (%)	sd (%)
Original Cataloging (Control Group)	17	8.5	2.0
Copy Cataloging, all sources	12	6.0	1.7
Copy Cataloging, LC	9	4.5	1.5
Copy Cataloging—Class P (all sources)	14	7.0	1.8
Copy Cataloging—Class Q (all sources)	10	5.0	1.5
Copy Cataloging—Class T (all sources)	14	7.0	1.8

Therefore, we can formally state our test by formulating a hypothesis, H_1: *There is a difference between the mean shelflisting error rate, μ_r, for call numbers that have been revised and those that have been accepted from the OCLC shared cataloging database without revision*. It is easier to test the converse of a hypothesis such as the foregoing, the null hypothesis, H_0: *There is no difference in the mean shelflisting error rate for records that have been shelflisted and those that have been accepted from the OCLC shared*

TABLE 3

Sample	dp (%)	p null (%)	z (%)	Reject Null Hypothesis	P(>z) (%)
050/090 - Control	-2.5	7.3	0.96	No	33.5
050 - Control	-4.0	6.5	1.62	No	10.5
P - Control	-1.5	7.8	0.56	No	57.5
Q - Control	-3.5	6.8	1.40	No	16.3
T - Control	-1.5	7.8	0.56	No	57.5

cataloging database without revision. The null hypothesis implies that the average error rates of call numbers in the two samples are consistent with a situation in which all records are selected from the same population, which implies that $\mu_u - \mu_r = 0$.

Attempts to measure this difference will yield sample results that have a normal distribution—we can expect that samples will yield non-zero differences; some differences will be positive and some negative, but their average will be zero. We also expect that large differences will be much less likely to occur than small differences. Thus, if we observe a large difference, we will be inclined to reject the null hypothesis in favor of the experimental hypothesis.

It is known that 95% of all sample means will fall within ± 1.96 standard deviations of the population mean. We do not know a priori the population mean or standard deviation; however, *we can estimate them by assuming, in accordance with the null hypothesis, that both samples were drawn from the same population.* Thus, if the number of shelflisting errors observed in the sample of n_r revised items is e_r, and the number of errors in the sample of n_u unrevised items is e_u, the probability, p, that a call number drawn from the total population will be in error is,

$$p = \frac{e_r + e_u}{n_r + n_u} = \frac{e_r + e_u}{2n}$$

The standard deviation of the difference may also be estimated from the two sample standard deviations,

$$\sigma = \sqrt{\frac{\sigma_r^2}{n_r} + \frac{\sigma_u^2}{n_u}}$$

where σ_r and σ_u are the sample standard deviations of the revised, and unrevised, samples respectively (Hoel 1971, 134–37).

The results of this analysis are presented in table 3, where z is the standard variable computed for each sample.

The standard variable is a construct that simplifies computations and comparisons involving the normal distribution. It is defined as

$$z = \left(\frac{x - \mu}{\sigma}\right)^2$$

As is obvious from the contents of table 3, in none of the cases studied is the standard variable large enough to reject the null hypothesis at a 95% level of confidence. The final column of table 3 gives the probability of observing a standard variable greater than that which was observed in our samples even under the assumption of the null hypothesis.

As we can see from table 3, only in the case of Library of Congress records do we come even close to being able to reject the null hypothesis. Not surprisingly, our sampling data indicates that the incidence of shelflisting error in LC-created records is possibly less than that found in UA-created records. In all other cases we have no evidence that would permit us to reject the null hypothesis, i.e., that there is no difference between the shelflisting error rate observed in unrevised copy cataloging records and original cataloging records.

The Problem in Perspective

Given the small number of errors encountered overall, it would appear that shelflisting every call number is not a cost-effective procedure. Before making that

TABLE 4

Sample	Wrong Class Number		Wrong Cutter Number		Cutter out of Order	
N = 200	No.	%	No.	%	No.	%
Control	2	1.0	7	3.5	8	4.0
050/090	2	1.0	6	3.0	4	2.0
050 (LC)	2	1.0	2	1.0	5	2.5
P class	1	0.5	5	2.5	6	3.0
Q class	1	0.5	2	1.0	6	3.0
T class	1	0.5	6	3.0	7	3.5
Totals	9	4.5	28	14	36	18.0

decision, we wanted to know whether the errors that were found had little meaning for the library consumer browsing the shelves, or if they had major implications for access to the collection. As part of the study, the types of shelflisting errors found were coded and examined. Intuitively, it seems that some kinds of shelflist disorder would have more impact than others on the browsability of a collection. For instance, the assignment of an incorrect classification number could place an item totally out of its subject range on the shelf. The assignment of a wrong cutter number, on the other hand, might result in related works occurring on different shelves within a discipline, but still close enough to be discovered by perusing the spine titles in the general area. Similarly, having cutter numbers out of order that place items on a shelf a few books away from their proper place, or locate an author's works a short distance apart, would have little consequence for retrieval.

Table 4 shows the total numbers of the types of errors that were found in sample call numbers, and the fraction of the sample they represent. The results of this portion of the study show that only a small number of errors involved incorrect classification. Wrong class numbers were recorded when the classification obviously did not reflect the subject matter of the record. The largest number of errors involved cuttering of all types. Wrong cutter numbers included call numbers that did not follow the instructions for cuttering in the classification schedule; call numbers cuttered differently than earlier editions of the same work; or a cutter incorrectly assigned to the main entry. A cutter out of order placed the item out of correct filing sequence on the shelf by main entry, title, etc. This kind of shelflist disorder could be expected to have the least impact on retrieval through browsing. The fact that almost half of the errors involved a cutter out of local shelflist order is not surprising, given the diversity in holdings among institutions and therefore in their shelflists. Further study is necessary to determine the actual effect of item displacement on browsing success, but logically it would appear that local shelflisting has little real value to our patrons.

Another question concerning the interpretation of the data arises from the high error rate in the control group, which consisted of call numbers that had been locally assigned and shelflisted. We would have expected a much lower error rate for these than the rest of the samples, since they had been shelflisted and therefore deliberately placed in order in the database. The individual records for each item were examined to determine possible reasons for the errors and discover whether this phenomenon could be considered a confounding variable in the study. Several of the errors could be traced directly to a project in which a temporary staff member, a cataloging student, was hired to reduce backlog and trained to assign call numbers. Although quality control procedures were implemented, it would not have been cost-effective to check the shelflisting of every item after the student's training phase. During the time period of the study, three new monographs catalogers were also hired in the department. Some errors could be traced to the training periods of these catalogers. A small number of errors could be attributed to individuals by subject expertise, and probably represent random human fallibility. This brought up the question of whether our results can be generalized to other institutions. Although further study would be needed to draw definite conclusions, our circumstances may be typical of other institutions trying to maintain current cataloging workflow while reducing backlogs in cataloging. Different results may have been obtained in a different time period within the same catalog department, as well as from another institution. No cataloger is infallible, and more than likely, no shelflist is perfect. The only way to answer the question of generalizability conclusively would be to replicate this study at other institutions. Given the expense and magnitude of such a project, it is unlikely many similar studies will be undertaken.

In conclusion, the results of this study indicate that local shelflisting is not a cost-effective operation for the University of Alabama libraries, and although it is not certain that the study can be generalized to other institutions, this research should be carefully weighed by other institutions in the process of reviewing local cataloging policy and workflow. The small number of errors detected produced a small amount of shelflist disorder and would therefore be expected to have a low impact on the browsability of the collections. The lack of a difference in disorder created by LC-assigned and member-assigned call numbers argues that differential work flow treatment of call numbers by source of cataloging copy does not significantly improve the quality of the local shelflist.

Works Cited

Bleil, Leslie A., and Charlene E. Renner. 1990. Copy cataloguing and the bibliographic networks. In *Technical services today and tomorrow*, ed. Michael Gorman and others. Englewood, Colo.: Libraries Unlimited.

Chan, Lois Mai. 1989. Library of Congress class numbers in online catalog searching. *RQ* 28: 530–35.

Hoel, Paul G. 1971. *Introduction to mathematical statistics*. 4th ed. New York: Wiley.

LeBlanc, Jim. 1995. Classification and shelflisting as value added: Some remarks on the relative worth and price of predictability, serendipity, and depth of access. *Library resources & technical services* 39: 295–302.

Library of Congress. Cataloging Policy and Support Office. 1995. *Subject cataloging manual: Shelflisting*. Washington, D.C.: Library of Congress.

Taylor, Arlene G. 1988. *Cataloging with copy: A decision-maker's handbook*. Englewood, Colo.: Libraries Unlimited.