

Better Than Brief Tests: Coverage Power Tests of Collection Strength 1

Howard D. White

Improving on ideas developed in *Brief Tests of Collection Strength*, this paper presents coverage power tests, an empirical method for evaluating collections in all types of libraries by means of ranked holdings counts from OCLC's WorldCat. The new method measures library coverage of subject literatures across levels of the WLN or RLG collection intensity scales that are increasingly difficult to attain. It defines literatures and collections unambiguously, permits objective comparisons of libraries, and is potentially automatable. Results of 38 tests in nine subjects at 30 libraries have high face validity in rating collections. Graphical analysis with the new method also clarifies the bibliometric relation between individual collections and subject literatures.



My 1995 book *Brief Tests of Collection Strength*¹ set forth an empirical method for rating subject collections on the six-level scale devised by the Research Libraries Group (RLG) or the 10-level scale devised by the Western Library Network (WLN). The present article describes an assessment tool that evolved from ideas in that book—a “coverage power” test that once again employs a scale of ranked OCLC holdings counts to rate collections in the RLG or WLN style.² However, the new test leads to much more persuasive ratings. It measures *library coverage of subject literatures*—not as a single percentage but as percentages across levels of coverage that are increasingly difficult to attain.

Items ranked in order of increasing difficulty make this what psychologists call a power test. Although particularly relevant to developers of research collec-

tions, the new method may interest any librarian who uses OCLC's WorldCat. The test defines both collections and literatures in such a way that libraries may be validly compared.

When this article was first submitted, WorldCat users could easily perform coverage power tests as described below. They cannot now: in late 2006, OCLC changed its software, and the key module in WorldCat no longer aggregates holding counts. It remains to be seen whether OCLC will restore this valuable capability, perhaps as part of its Collection Analysis system.³ (That system was not needed to obtain the results reported here.)

The Conspectus Movement

The RLG and WLN scales were central to the Conspectus movement, which in the 1980s and 1990s sought to create a national database of rated collections. Librarians were to define their subject col-

Howard D. White is Professor Emeritus at the College of Information Science and Technology, Drexel University; e-mail: whitehd@drexel.edu.

TABLE 1
Levels of the RLG and WLN Collection Evaluation Scales

	<i>RLG Scale</i>		<i>WLN Scale</i>
0	Out of Scope	0	Out of Scope
1	Minimal Level	1a	Minimal Information Level, Uneven Coverage
2	Basic Information Level	1b	Minimal Information Level, Focused Coverage
3	Instructional Support Level	2a	Basic Information Level, Introductory
4	Research Level	2b	Basic Information Level, Advanced
5	Comprehensive Level	3a	Basic Study or Instructional Support Level
		3b	Intermediate Study or Instructional Support Level
		3c	Advanced Study or Instructional Support Level
		4	Research Level
		5	Comprehensive Level

lections by means of Library of Congress or Dewey class ranges called Conspectus lines and then subjectively assign their holdings in each line a level on the scales seen in table 1. The database in which their ratings were published was to be used by librarians for cooperative collection development and by the public as a reference work on collection strengths.

The project immediately ran into two major difficulties: the Conspectus lines numbered in the thousands, which made rating collections in them very laborious, and, for all the work that went into them, the ratings could not be trusted because of their subjectivity. The scale levels were qualitatively described but never clearly operationalized, which meant that librarians could not know whether they were using the same standards to rate their collections or, indeed, what those standards were. As defined, RLG and WLN levels 1 through 4 are irremediably vague.

RLG sponsored some "verification studies" to test the reliability of the subjective ratings, but these, too, were suspect. They were essentially studies of how well libraries covered various subject literatures, but the literatures were defined by bibliographers' idiosyncratic checklists, which seem to have been drawn up without regard for either the Conspectus lines or the RLG/WLN scale levels. What was

worse, no one ever explained how the "percent covered" figures that emerged from verification studies were to be used to evaluate the subjectively set levels. There was no table that said, in effect, "If you have a percentage score in this range, claim level 4; if in this range, claim level 3" and so on. The verification studies were themselves labor-intensive, both for the bibliographers who compiled them and for the participating librarians who checked their collections against them. However, they did show that librarians were judging their collections inconsistently.

After being offered for years through the interface of RLG's union catalog, the Conspectus database turned out to be not much used, and it was withdrawn from service in 1997. Even so, large fragments of it remain. Academic librarians still use the RLG or WLN scale levels to describe their collections, and the RLG and WLN Conspectus lines have evolved into the "OCLC Conspectus" in the WorldCat Collection Analysis system. In this version, however, it is simply a subject scheme; it does not include members' ratings of their collections on the RLG or WLN scales. "Conspectus" originally meant an overview of collections, not just subjects.

What Are Brief Tests?

The book *Brief Tests*, intended to give li-

brarians a wieldy but data-based method for setting or verifying levels, came along just as the RLG Conspectus project was starting to collapse under its own weight. Brief-test instruments usually consisted of 40 titles chosen by someone with expertise or at least strong interest in a subject. Each selector was given the RLG scale with its qualitative definitions and instructed to choose 10 titles to represent each of the four levels from Minimal through Research. (Level 0, Out of Scope, in which the library collects nothing, and level 5, Comprehensive, in which the library attempts to collect everything, were assumed to be unproblematic.) Selectors were not required to work within Conspectus lines (which would have greatly complicated their task). Rather, they were allowed to pick their own subjects, which ranged from entire disciplines (e.g., Economics, Genetics) to specialties (e.g., Dickens studies, Workplace Democracy, the Holocaust). They picked titles from various LC classes to instantiate these subjects as they saw fit.

In the written instructions, selectors were told: "Your checklist is *not a sample* in the usual sense. That is, it will not be used to estimate, within known margins of error, what percentage of the bibliography of a literature some library holds. Your checklist is a *test*. It is just like a 40-item test for students in which the teacher has tried to include ten easy questions (level 1), ten that are somewhat tougher (level 2), ten that are more demanding still (level 3), and ten that are quite hard (level 4). You might use this notion as a guideline in creating your list—for example, level 4 items are often mathematical or in foreign languages. Level 1 items would be the most likely for a library to hold; level 4 items, the least likely, with the other two levels in between."

Once the 40-item checklists were in hand, OCLC holdings counts for each title were obtained through the old Prism System (WorldCat was still in the future), and the counts were used to rank-order the titles, high to low. This ranking

sometimes moved titles from selectors' levels to levels dictated by the counts, but it made brief tests into power tests, as defined above. The selectors' attention to different levels had the advantage of producing titles with widely varying counts, which heightened their suitability for power tests.

Holdings at various libraries were then checked against the finished lists through lookups in online catalogs (and card catalogs where necessary). If a library had at least 50 percent (5 out of 10) of the items at a given level, it could claim that level for its collection.

The original definitions of the RLG and WLN levels implied (and some definitions now state) that the levels are *cumulative* in nature. Thus, if a library claims an Instructional collection, it should also meet the criteria for Basic and Minimal collections. As power tests whose levels increase in difficulty from Minimal through Research, brief tests embody this notion. They implicitly predict that if a library scores at least 50 percent on a level *more* difficult to attain, it should score at least 50 percent on all levels *less* difficult to attain. In this, they resemble Guttman scales.

The book *Brief Tests* provided massive evidence that the hypothesis of cumulativeness is correct. For example, with few exceptions, libraries that attained Research level on a test also had scores of 50 percent or higher on all lower levels. Table 2 presents corroborating outcomes from a 1990s brief test of Medieval studies collections at Temple University, the University of Pennsylvania, the University of Chicago, and the New York Public Library. The first attains a level 3 (Instructional) collection; the rest have level 4 (Research) collections.

The book was filled with similar tables that showed the ideas behind brief tests to be basically sound and informative. Thomas Twiss⁴ performed a four-part validity check on those ideas, administering his 80-item brief test in Soviet studies to two large academic libraries and two

TABLE 2
Brief Test Results for Medieval Studies

Level	No. of Titles	Percent Covered			
		Temple	Penn	Chicago	NYPL
1	10	80	80	60	100
2	10	90	100	100	90
3	10	80	100	100	100
4	10	10	90	90	70

large public libraries. His results showed the test valid on all four counts—cumulativeness of levels, consistency with expectations for type of library, replicability of previous brief tests conducted by others, and accord with librarians' subjective estimates—a highly intelligible confirmation of theory. Recent data-based analyses by David Lesniaski⁵ and Jay Bernstein⁶ have further substantiated points made in the book.

Problems with Brief Tests

The book discusses problems with brief tests in some detail. Here are three that the new coverage power test solves:

- The convention of having only 40 items, arbitrarily set to make things easy for selectors and catalog-checkers, opens brief tests to a charge of instability. With percentages based on only 10 items per level, each book represents fully 10 percent of that level's score. Thus, reaching or not reaching the 50 percent mark can depend on a single book. Combining tests in a single subject so as to have 80 or 120 items makes for somewhat stabler percentages but also increases the work of test-making and checking.

- Early brief tests were *relative* in matching holdings counts with RLG scale values. That is, the top 10 titles were considered Minimal whatever their counts; the next 10 were Basic whatever their counts, and so forth. Tests made in this way seemed to work well enough, but relative assignments could still lead to undesirable inconsistencies across tests. The book therefore included a newly made

holdings-count scale with *absolute* values, based on a statistical analysis of the counts from 53 existing tests. Future test-makers were advised to work with the values in Table 3 (and many did). For example, a title with a count of 208 would be assigned to Instructional level, whereas a

title with 1,516 would be Minimal level.

- The absolute-value scale did not eliminate a *psychological* problem, which is that people tend to conflate the scale for levels of collections with a conceptually simpler scale for the intellectual levels of works. The latter, existing in everyone's head for textbooks, runs from *elementary* through *basic* and *intermediate* to *advanced*. More broadly, it is a scale of *ease of reading*, with works such as children's books and popularizations at one end and works such as learned treatises and highbrow novels, poems, or plays at the other. The problem arises when the holdings counts of individual titles place them on levels that do not accord with the presumed degree of sophistication needed to read them. Since such discrepancies are discussed at some length in the book,⁷ all that need be said here is that the counts are a radically simple means of defining titles that are appropriate to the levels *on a variety of grounds*. For example, titles held by more than 750 libraries are items of the sort that a library would buy if it collects anything at all in a given subject; they

TABLE 3
Suggested Values for Assigning Titles to Levels

Counts	Levels
751 to highest	1 Minimal
401 to 750	2 Basic
151 to 400	3 Instructional
1 to 150	4 Research

are, in words from the book, “popular, standard, indispensable, classic.” But that does not necessarily mean they are easy to read (e.g., *The Critique of Pure Reason*, *The Sound and the Fury*, *Capital*). At the other extreme, titles held by fewer than 150 libraries have only small, specialized readerships because they may be “advanced, localized, foreign, academic, obscure, old.” But they are not necessarily hard to read (e.g., regional cookbooks, *passé* novels, hobby magazines). The same caveat on “ease of reading” holds for coverage power tests.

Coverage Power: Introduction

The examples of coverage power tests that follow come from an online course in Research Collection Development that I taught for Drexel University’s College of Information Science and Technology during summer 2006. The 24 enrollees, all graduate students, were scattered widely over the United States. Their assignments included making a brief test and administering it at libraries of their choice, and then doing the same thing for a coverage power test.

Working in WorldCat from written instructions, the students first retrieved all titles in Library of Congress classes or class ranges that defined topics of interest to them. LC classes, which are searchable in the Expert mode of WorldCat, are an objective way of defining literatures. (One chooses lc: from a dropdown menu and then enters, for example, HM51, the code for the literature of general sociology.) By default, WorldCat ranks titles in a retrieved literature high to low by their holdings counts. This makes it far easier to create brief tests than in the days of Prism: one simply asks for displays of 100 titles at a time, checks 10 titles with well-spaced holdings counts to represent each of the four levels as in Table 3, and saves the result. (This is also easier than the methods given in Lesniaski, which essentially repeat those of the book.)

After administering their 40-item tests by checking them against the online

catalogs of selected libraries, the students moved on to the coverage power tests with the same library collections and same topical literatures they had used previously. They were told, “Over the next week, you’ll test the *entire collection* of each library against the *entire literature* from which you chose your brief test. No more 40 items; we’re testing coverage in full. That makes the method definitely less frail.” The last remark speaks to the first problem mentioned above: unstable percentages. As to the other problems, coverage power tests are pegged to an absolute scale of holdings counts, and they eliminate perceived mismatches of titles with levels because individual titles are no longer a concern.

Earlier I had suggested that students choose LC classes so as to form literatures in the neighborhood of 1,000 titles for their projects. This advice had to do with my own uncertainties about WorldCat’s technical limitations and my desire not to overburden either the students or the system. It is now apparent that one can work with much larger literatures. The technical problems lie in calling them up. Outside the Collection Analysis (CA) system, WorldCat programming does not readily permit the formation of literatures corresponding to entire Conspectus lines. Within CA, it does not permit coverage power tests of the sort introduced here. Moreover, for serious use beyond exploratory trials, CA is open only to subscribers. The Expert search module, which is open to everyone, makes it easy to form literatures based on one or two LC classes, but those based on broad *ranges* of classes may require elaborate ORing of terms with wild-card truncation. These are by no means permanent obstacles, merely observations on how the state of the art at OCLC affected the coverage power trials in the course.

Whatever the current limitations in coverage-power methodology, it is a boon that literatures can now be defined by explicit, commonly accepted, nonoverlapping LC classes. This fact holds both

for single classes (such as HM51) and for the broad ranges of LC classes that make up Conspectus lines. Moreover, academic library collections can usually be defined as subsets of these same LC classes—for example, “the subset of HM51 titles held by Drexel University”—which gives coverage power tests a conceptual clarity that brief tests lacked. The resulting ratings of collections depend not at all on librarians’ subjective estimates and are strictly comparable across libraries. Tests for any OCLC library can be replicated and checked, which of course was not true of the original RLG Conspectus.

Coverage power tests work equally well with Dewey classes and with literatures formed by LC subject headings, although, in the latter case, there is a complication. Since LC subject headings are not mutually exclusive, titles can be part of multiple literatures. That makes coverage vary with how many subject headings have been assigned to titles. Thus, for no justifiable reason, some

books will count more than others in coverage tests. Nevertheless, the test results will be intelligible.

Methods

The following instructions were given to the students for conducting coverage power tests:

1. Get the OCLC participating library symbol for each of your libraries. These are usually three-letter identifiers, such as DXU for Drexel, YUS for Yale, or CUY for UC Berkeley.
 2. Retrieve the literature defined by your LC class codes (and your other settings, if any).
 3. Get the holdings-count distribution for titles in this literature. You’ll use a holdings-count scale that OCLC itself has devised in WorldCat. Click the “Limit” button at top left on the list of retrieved records and then click “Limit Your Results by Number of Libraries.” Copy and paste the distribution into Excel.
 4. Get the holdings-count distributions for titles in *collections* of this same literature. To do this, add the OCLC symbol for one of your libraries to the search panel, retrieve the literature again, and call up the holdings-count scale again. Repeat this for every collection in which you are interested. Enter the collection distributions into Excel.
 5. Have Excel divide the values of the *collection* distribution by the values of the *literature* distribution and convert them to percentages. This gives each library’s coverage of the literature at each level of the holdings-count scale.
 6. Score the library as to level on my adaptation of the WLN scale. The library can claim the highest level at which it has at least 50 percent coverage.
- Figure 1 shows a literature distributed over WorldCat’s holdings-count scale—the 3,441 titles classed in HM51 (General Sociology) as of July 2006. Two titles are held by

FIGURE 1
The HM51 Literature on the
Holdings-count Scale

Select a Number of Libraries to Limit Your Results	
Search	Clear
<input type="checkbox"/> 1500-1999	2
<input type="checkbox"/> 1000-1499	9
<input type="checkbox"/> 900-999	5
<input type="checkbox"/> 800-899	4
<input type="checkbox"/> 700-799	12
<input type="checkbox"/> 600-699	29
<input type="checkbox"/> 500-599	30
<input type="checkbox"/> 400-499	45
<input type="checkbox"/> 300-399	101
<input type="checkbox"/> 200-299	188
<input type="checkbox"/> 150-199	149
<input type="checkbox"/> 100-149	261
<input type="checkbox"/> 75-99	181
<input type="checkbox"/> 50-74	246
<input type="checkbox"/> 25-49	321
<input type="checkbox"/> 10-24	443
<input type="checkbox"/> 5-9	325
<input type="checkbox"/> 2-4	538
<input type="checkbox"/> 1	506
<input type="checkbox"/> 0	46

more than 1,500 libraries, nine are held by more than 1,000 but fewer than 1,499 libraries, and so on. (The items held by “zero libraries” may be disregarded for present purposes.) The checkbox display allows the searcher to limit a retrieval to items held by some liminal number of libraries. It is designed not for the broad study of coverage, but for title-by-title selection and discovery of unique or near-unique holdings.

The intervals in the WorldCat holdings-count scale are obviously very unequal: they increase by one at the bottom and by 500 at the top, extending over four orders of magnitude (units, tens, hundreds, and thousands). If the holdings count for each title was converted to its logarithm and then these logs were binned along a logarithmic scale, the frequency distribution would be similar. WorldCat’s holdings-count scale in fact approximates a logarithmic scale.⁸ More on this later, because it provides insights into collection development that one need not be a mathematician to grasp.

Figure 2 moves us to a collection analysis for one library. It shows the distribution of the same HM51 literature, this time filtered through the condition that titles must be owned by the University of California at Berkeley.

The final steps in determining a library’s coverage power are to find the percentage of a literature that is covered at each interval on the WorldCat scale and then to read off the highest interval at which a library attains 50 percent. The percentageing is done through a nonstandard operation in Excel. It involves dividing *collection title counts* by *literature title counts* at each interval of the scale and converting the decimal fractions to percentages.⁹ All members of my class did so successfully to produce tables like Table 4. Of course, hand calculations with Excel are needed merely for proof of concept. For serious use by busy librarians, the whole process should be automated.¹⁰

Table 4 displays coverage power tests in the HM51 literature for the Berkeley

FIGURE 2
The HM51 Collection at UC Berkeley

Search	Clear
<input type="checkbox"/> 1500-1999	2
<input type="checkbox"/> 1000-1499	9
<input type="checkbox"/> 900-999	3
<input type="checkbox"/> 800-899	4
<input type="checkbox"/> 700-799	12
<input type="checkbox"/> 600-699	28
<input type="checkbox"/> 500-599	27
<input type="checkbox"/> 400-499	36
<input type="checkbox"/> 300-399	80
<input type="checkbox"/> 200-299	122
<input type="checkbox"/> 150-199	77
<input type="checkbox"/> 100-149	81
<input type="checkbox"/> 75-99	37
<input type="checkbox"/> 50-74	38
<input type="checkbox"/> 25-49	47
<input type="checkbox"/> 10-24	36
<input type="checkbox"/> 5-9	9
<input type="checkbox"/> 2-4	9
<input type="checkbox"/> 1	2

(UCB) and Drexel (DU) collections. It also displays my experimental mapping of WorldCat scale intervals onto those of the WLN scale. (The intervals go higher than 1,500–1,999 for some literatures but not HM51.) Berkeley’s and Drexel’s percentages of coverage appear in the right-most columns. In the top row, Berkeley has 2/2 or 100 percent of the titles in the 1,500–1,999 range, while Drexel has 0/2 or 0 percent; all other calculations follow the same pattern. Berkeley’s percentages exceed 50 percent down through level 3c (“Advanced study or instructional support”), whereas Drexel’s exceed 50 percent only at the 1a level (“Minimal information level, uneven coverage”).

The mapping of WLN and RLG levels onto the intervals of the WorldCat scale is an attempt to devise a useful collection evaluation scheme from existing concepts. Some technical details:

As the WLN scale increases from 1a through 5, librarians have increasingly rejected (or not even considered) the titles of a literature. Its low end thus corresponds to library best sellers, and its high end,

TABLE 4
RLG and WLN Scales Mapped onto WorldCat Holdings-Count Scale, with Results in the HM51 Literature for Collections at Drexel and UC Berkeley

RLG	WLN	WorldCat	Lit	UCB	DU	UCB%	DU%
1 Minimal	1a	1500-1999	2	2	0	100%	0%
1	1a	1000-1499	9	9	9	100%	100%
1	1b	900-999	5	3	1	60%	20%
1	1b	800-899	4	4	1	100%	25%
2 Basic	2a	700-799	12	12	4	100%	33%
2	2a	600-699	29	28	12	97%	41%
2	2b	500-599	30	27	8	90%	27%
2	2b	400-499	45	36	11	80%	24%
3 Instructional	3a	300-399	101	80	7	79%	7%
3	3b	200-299	188	122	16	65%	9%
3	3c	150-199	149	77	5	52%	3%
4 Research	4	100-149	261	81	4	31%	2%
4	4	75-99	181	37	0	20%	0%
4	4	50-74	246	38	1	15%	0%
4	4	25-49	321	47	0	15%	0%
5 Comprehensive	5	10-24	443	36	1	8%	0%
5	5	5-9	325	9	1	3%	0%
5	5	2-4	538	9	1	2%	0%
5	5	1	506	2	0	0%	0%

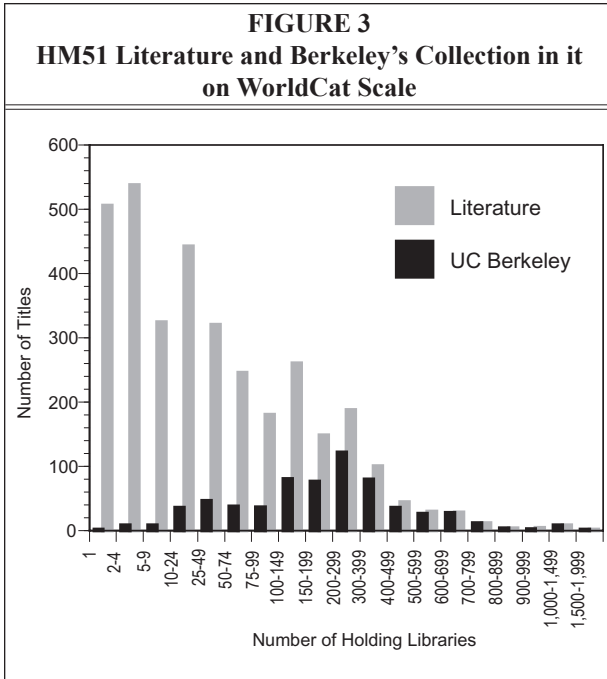
to library worst sellers. In pairing low with high and high with low, the scale is inverse.

Since the WLN scale simply subdivides RLG's into finer levels, the cutpoints of table 3 have been used to make the main divisions from levels 1 through 4. The intervals of the WLN scale, as adapted here, are not equal; in condensed form they group holdings counts in roughly logarithmic fashion as 1: *above 800 holding libraries*, 2: *799 to 400 libraries*, 3: *399 to 150 libraries*, 4: *149 to 25 libraries*, and 5: *below 25 libraries*. Level 1, Minimal, begins at 800 libraries rather than 750 because the WorldCat scale does not have a break at 750. The handling of levels 4 and 5, which WLN does not subdivide, is ad hoc but plausible.

In many cases it makes sense to terminate both RLG and WLN scales at

WorldCat interval 25–49 because collections rarely reach 50 percent at *that* level, let alone any of the Comprehensive ones. Still, the WorldCat intervals at level 5 can register the occasional magnificent collection, and they are necessary when a library buys numerous rarely held foreign-language materials in a subject such as Buddhism or Korean History.

RLG's levels are less elaborate than WLN's and in coverage power tests may be more understandable. As matters stand, using RLG levels involves merging the WorldCat counts at each level and repercentaging. But, again, OCLC should simply automate such operations. On the RLG scale, Berkeley scores Minimal 90%, Basic 89%, Instructional 64%, Research 20%, and Comprehensive 3%. Drexel's much smaller, more fragmented collection scores Minimal 55%, Basic 30%,



Instructional 6%, Research 0.5% and Comprehensive 0.2%. Berkeley's collection is thus Instructional; Drexel's, Minimal.

An Interlude on Distributions

In HM51, Berkeley holds 659 items out of the 3,441-title literature. How do the two distributions compare? How does Berkeley's library, or any library, cover a particular literature? In figure 3, the distribution of Berkeley's collection has been superimposed on that of the HM51 literature as a whole. The literature distribution is non-normal and highly skewed to the right, whereas the collection distribution is roughly normal. (When the latter is plotted separately, it appears more like a bell-shaped curve, as in figure 4b below.) These two outcomes are not quirks. Rather, they are typical of any sizeable literature and any collection of hundreds or thousands of items. (Considerably smaller collections, such as Drexel's, do not look roughly normal if graphed; they look like rubble on a plain.)

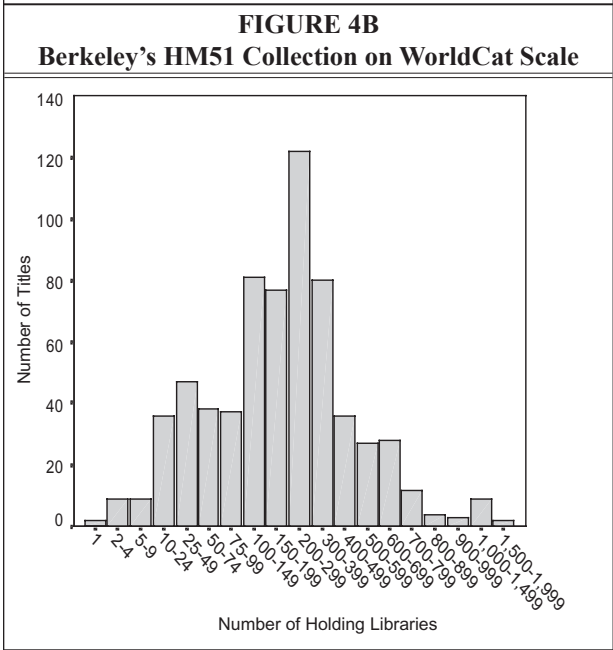
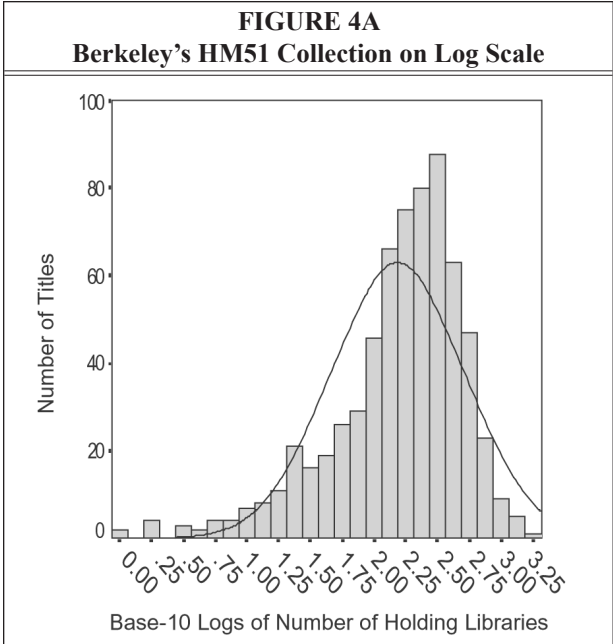
Figure 3 helps to explain the highly skewed *literature* distribution: it comes

from many hundreds of librarians agreeing strongly on the value of titles at the high end of the scale and disagreeing strongly on the value of titles at the low end of the scale. In other words, there are a few handfuls of titles that most librarians want to have (best sellers) and increasing numbers of titles that fewer and fewer librarians want to have (they sell decreasingly well, so to speak).

In the part of the scale that reflects titles held by 400 or more libraries, note the close correspondence between what Berkeley has bought and what there is to buy in

the literature. Of the 10 dozen or so titles available, Berkeley has bought almost all, and they are roughly the same dozens that every other library has bought. At the other end of the scale, Berkeley does not acquire hundreds, let alone thousands, of the low sellers; it again acquires just dozens of them. But these are mostly *different dozens* from what all the other libraries have bought. Everyone getting different dozens creates the part of the literature held by fewer than 150 libraries.

The skewed holdings-count frequencies of the literature distribution in figure 3 form a hyperbolic or "reverse-J" curve that has long been one of bibliometricians' fundamental objects of study. Less abstractly, one may speak of a core-and-scatter distribution reminiscent of the ones S. C. Bradford famously described in the 1930s when he was head of the Science Library in London.¹¹ Here, the *core* of HM51 would comprise titles whose high holdings counts place them at the right of the scale in figure 3. The topmost titles turn out to be classics by sociologists with high name recognition,

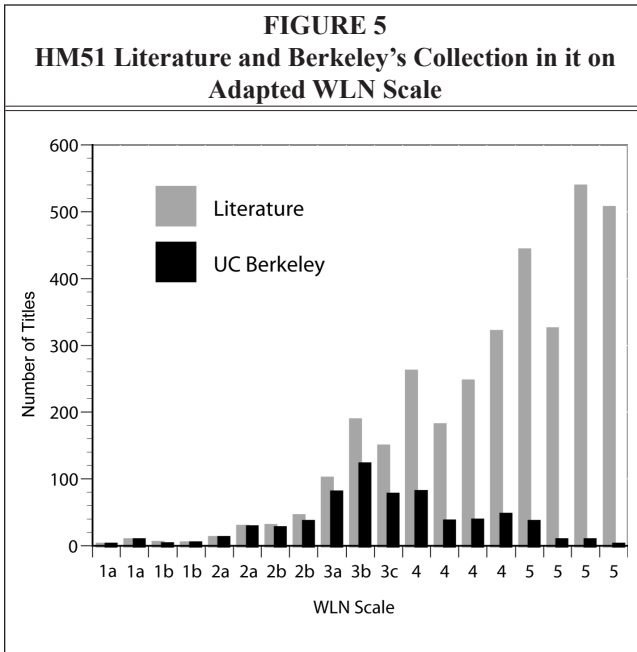


The leftmost point on the scale comprises the long tail of singletons that all Bradford-type data exhibit—in this case they are titles held by only one library each. These unique items are seldom valuable rare books; many appear to be sociological scraps that someone happened to catalog.

To turn to the collection distribution, it was said earlier that WorldCat's holdings-count scale approximates a logarithmic scale. We may now conclude that title counts in any sizeable collection are log-normally distributed. That is, the distribution not of the counts but of the logarithms of the counts is approximately normal. Figure 4a presents the counts for Berkeley's HM51 collection in base-10 logs (0 = 1 holding library, 1 = 10 holding libraries, 2 = 100, 3 = 1,000, etc.). Figure 4b is the same collection along the WorldCat scale. Both figures were made by SPSS. The two distributions are by no means identical (the binning of values is different), but both are roughly normal. (If the title counts of Berkeley's HM51 collection, or any other sizeable collection,

such as Robert Merton, Talcott Parsons, Peter Berger, and Ralf Dahrendorf. The *scatter* begins at some mathematically definable point leftward on the scale as holdings counts diminish and titles and authors decrease in recognizability.

are simply plotted without being binned on the WorldCat scale or logged and binned on a logarithmic scale, they look like the *literature* distribution in figure 3—non-normal and highly skewed to the right.)



Berkeley buys greater absolute numbers of titles, but decreasing proportions of what there is to buy; the black and gray bars become more divergent. Finally, at levels 4 and 5, as the mountain of relatively little-held titles grows, Berkeley's editorial resistance screens out ever more of them, as shown by the black bars shrinking in proportion to the gray bars. So we end up with a roughly symmetrical distribution for this and every other big collection.

The collection distribution comes about, of course, through librarians' behavior. Like editors in publishing houses, librarians accept and reject texts. More precisely, they approve or fail to approve the choices of publishers in disseminating texts. Replacing the WorldCat scale of figure 3 with the WLN scale, its inverse, clarifies this notion.

As figure 5 shows, the replacement simply reverses the layout of figure 3's bars. For collection developers at particular institutions, the WLN scale may now be seen as a scale of *editorial resistance*: it forces titles into groups that reflect librarians' resistance to having bought them, from low to high. The relatively few titles in the literature at 1a and 1b are held by at least 800 libraries, and they are irresistible, so to speak, to Berkeley collection developers as well; the black bars and the gray bars are identical in height. So, too, with the titles in 2a, which are held by 799 to 600 libraries. But in 2b, titles held by 599 to 400 libraries, we begin to see a divergence of black and gray bars; Berkeley buys the great majority of titles but resists some of them. Across 3a, 3b, and 3c,

interesting to learn the characteristics of publications in what we might call "zones of resistibility." The zones would be the RLG levels as seen in table 4, which correspond to changes in the proportions of literatures that librarians buy. *Brief Tests* laid out a number of variables on which the zones might systematically vary: for example, trade vs. academic imprints, print vs. microform (and other nonprint), U.S. vs. foreign publishers, English vs. foreign languages, in-print vs. out-of-print status, many vs. few or no reviews, multiple vs. single editions, nonmathematical vs. mathematical content, reputation of authors (measurable by citation counts), and proportions of various types of libraries among the buyers.¹² As foreshadowed in *Brief Tests*,¹³ this last variable is now being used in OCLC's experimental system for determining Audience Level of titles.¹⁴

The goal of research along these lines would be to establish bundles of characteristics that are associated with the different levels of resistibility of titles. This could lead to analysis of their relative importance in determining coverage of literatures.

TABLE 5
Results from Coverage Power Tests in Nine Subjects

<i>20th Cen French History</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature	1	2	7	8	10	16	28	38	55	73
Drexel U	100%	50%	71%	50%	30%	19%	32%	34%	11%	14%
South Dakota State U	100%	100%	43%	75%	90%	50%	57%	45%	49%	26%
U Montana	100%	50%	86%	75%	70%	69%	61%	50%	36%	19%
U Binghamton	100%	100%	100%	100%	60%	75%	89%	76%	73%	62%
Cornell U	100%	50%	86%	100%	100%	94%	93%	92%	91%	81%
U Michigan	100%	100%	100%	100%	100%	75%	96%	87%	85%	75%
Central Missouri State U	0%	0%	0%	0%	10%	6%	18%	11%	18%	10%
<i>Buddhism</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature	2	1	3	4	9	9	21	28	48	80
Drexel U	100%	0%	0%	0%	11%	33%	5%	7%	15%	6%
Ball State U	50%	100%	33%	75%	78%	67%	48%	36%	33%	28%
Vanderbilt U	50%	100%	33%	75%	78%	56%	67%	54%	42%	59%
U Pittsburgh	50%	0%	100%	50%	78%	67%	67%	57%	46%	53%
Yale U	100%	0%	67%	75%	89%	56%	62%	75%	60%	60%
UCLA	50%	100%	67%	100%	100%	89%	81%	71%	67%	73%
Harvard U	100%	100%	67%	75%	78%	89%	76%	75%	65%	68%
<i>Environmental Policy</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature	1	1	5	7	8	9	15	31	52	127
LA Public Library	100%	100%	100%	100%	75%	78%	60%	61%	46%	24%
UCLA	100%	100%	100%	86%	88%	67%	93%	90%	87%	79%
British Library	0%	100%	100%	86%	75%	44%	80%	71%	77%	66%
<i>Proust Studies</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature		1	5	2	8	11	10	15	30	42
U Delaware		100%	60%	50%	75%	64%	60%	73%	80%	69%
U Washington		100%	80%	50%	75%	91%	70%	93%	87%	62%
Library of Congress		100%	100%	50%	100%	82%	50%	87%	67%	62%
Stanford U		100%	80%	50%	88%	100%	90%	80%	73%	79%
<i>Astronomical Telescopes</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature			1		3	3	3	4	8	15
Brigham Young U			100%		67%	67%	100%	75%	63%	33%
U Arizona			100%		100%	67%	67%	25%	75%	40%
<i>Quantum Theory</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature	1	3	2	7	3	6	19	26	37	98
Drexel U	0%	67%	50%	29%	100%	67%	21%	54%	35%	33%

Coverage Power Tests of Collection Strength 167

TABLE 5
Results from Coverage Power Tests in Nine Subjects

<i>20th Cen French History</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	112	85	221	226	476	1319	2047	1373	1448	934
Drexel U	4%	5%	1%	0%	0%	0%	0%	0%	0%	0%
South Dakota State U	27%	5%	5%	1%	1%	0%	0%	0%	0%	0%
U Montana	20%	9%	5%	2%	2%	1%	0%	0%	0%	0%
U Binghamton	61%	59%	43%	35%	30%	16%	6%	2%	0%	0%
Cornell U	73%	78%	74%	69%	63%	51%	30%	12%	2%	1%
U Michigan	73%	79%	76%	76%	70%	59%	30%	11%	6%	2%
Central Missouri State U	4%	1%	1%	0%	1%	0%	0%	0%	0%	0%
<i>Buddhism</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	193	170	263	215	367	883	4650	4150	6475	7064
Drexel U	2%	2%	2%	0%	0%	0%	0%	0%	0%	0%
Ball State U	20%	12%	10%	7%	4%	1%	0%	0%	0%	0%
Vanderbilt U	54%	46%	35%	25%	16%	7%	0%	0%	0%	0%
U Pittsburgh	52%	46%	30%	22%	17%	18%	8%	5%	2%	0%
Yale U	60%	52%	48%	40%	35%	33%	25%	37%	20%	5%
UCLA	63%	68%	56%	45%	39%	37%	44%	17%	9%	6%
Harvard U	64%	65%	54%	52%	46%	32%	31%	10%	2%	0%
<i>Environmental Policy</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	150	121	176	118	118	211	455	554	1146	1259
LA Public Library	21%	20%	10%	9%	10%					
UCLA	80%	76%	56%	49%	39%					
British Library	57%	54%	61%	41%	31%					
<i>Proust Studies</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	98	96	173	142	182	273	425	231	348	321
U Delaware	60%	52%	27%	20%	13%	4%	3%	1%	0%	1%
U Washington	65%	59%	45%	32%	21%	13%	7%	5%	1%	0%
Library of Congress	61%	49%	51%	44%	37%	25%	15%	11%	6%	2%
Stanford U	70%	77%	68%	62%	50%	29%	11%	4%	1%	0%
<i>Astronomical Telescopes</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	22	13	33	17	48	59	136	146		
Brigham Young U	45%	23%	24%	24%	2%	15%	1%	0%		
U Arizona	50%	62%	42%	59%	46%	25%	14%	9%		
<i>Quantum Theory</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	73	67	42	36	69	133	111			
Drexel U	34%	28%	24%	8%	4%	0%	0%			

TABLE 5
Results from Coverage Power Tests in Nine Subjects

Stanford U	100%	67%	50%	57%	67%	100%	95%	77%	73%	77%
Cal Tech U	0%	100%	50%	71%	67%	67%	74%	85%	84%	78%
U Colorado Boulder	100%	100%	100%	100%	67%	83%	68%	77%	76%	64%
<i>Communicable Diseases</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature		1	1	6	31	15	16	15	31	37
U Pittsburgh		100%	100%	100%	94%	93%	81%	73%	81%	73%
U Washington		100%	100%	100%	90%	80%	75%	93%	68%	73%
Natl Library Medicine		100%	0%	83%	94%	80%	94%	100%	81%	76%
<i>Genealogy</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature		3	4	5	4	11	10	20	27	35
Natl Archives & Records		67%	25%	60%	50%	45%	50%	35%	33%	31%
New York State Library		100%	75%	80%	75%	64%	60%	40%	59%	74%
Rochester Public Library		100%	100%	100%	100%	82%	80%	75%	85%	57%
New Eng Hist Gen Soc		67%	100%	40%	75%	45%	70%	60%	67%	66%
<i>Nazi Germany</i>	1a	1a	1a	1b	1b	2a	2a	2b	2b	3a
Literature		5	15	31	12	20	29	51	64	78
U New Hampshire	60%	80%	61%	50%	45%	62%	45%	49%	33%	32%
U Rhode Island	100%	60%	65%	50%	65%	59%	48%	51%	30%	36%
U Delaware	80%	87%	81%	67%	90%	76%	79%	69%	64%	45%
Harvard U	100%	93%	77%	83%	80%	86%	79%	78%	72%	67%

Test Results

Table 5 gives results from 38 coverage power tests in nine subjects at 30 libraries. Since the Drexel students all had usable results, these nine subjects were chosen simply for diversity. The LC classes that produced them are given in the endnotes, along with credits.¹⁵ Had Conspectus lines been available for testing, they would have produced literatures with larger bases for percentaging, but they would pattern like those in table 5.

Most of the collections tested are at universities, some world-class. The Library of Congress, the National Library of Medicine, and the British Library also appear, along with a sprinkling of public and special libraries.

The latter include the libraries of the U.S. National Archives and Records Administration, New York State, and the New England Historical Genealogical Society.

For compact presentation, the WLN scale has been transposed to run horizontally, as it does in figure 5. Below it, the counts of titles in each literature run horizontally across levels. Below them, again read across, are the percentages of titles held at each level by the libraries—the coverage scores. The rightmost value at which a library attains at least 50 percent coverage is in bold face; the rating may be read off the WLN scale above it. For example, in 20th Century French History at top, the highest level at which Drexel

TABLE 5
Results from Coverage Power Tests in Nine Subjects

Stanford U	70%	45%	48%	25%	10%	7%	2%			
Cal Tech U	71%	58%	48%	22%	17%	6%	1%			
U Colorado Boulder	73%	58%	55%	31%	10%	5%	2%			
<i>Communicable Diseases</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	79	121	217	201	294	212	279			
U Pittsburgh	71%	70%	65%	61%	48%	36%	31%			
U Washington	71%	64%	59%	56%	42%	33%	27%			
Natl Library Medicine	77%	82%	76%	73%	66%	58%	55%			
<i>Genealogy</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	80	86	155	138	296	1230	6100	2764	3926	3356
Natl Archives & Records	21%	10%	3%	6%	3%	0%	0%	0%	0%	0%
New York State Library	43%	36%	24%	32%	18%	12%	29%	4%	0%	0%
Rochester Public Library	49%	40%	24%	14%	12%	5%	1%	0%	0%	0%
New Eng Hist Gen Soc	51%	51%	48%	54%	45%	49%	15%	13%	10%	10%
<i>Nazi Germany</i>	3b	3c	4	4	4	4	5	5	5	5
Literature	125	92	258	250	424	934	1534	1280	1571	1228
U New Hampshire	30%	16%	10%	8%	5%	1%	1%	0%	0%	0%
U Rhode Island	32%	20%	9%	5%	3%	2%	0%	0%	0%	0%
U Delaware	50%	42%	43%	39%	30%	14%	3%	1%	1%	0%
Harvard U	65%	62%	69%	67%	71%	66%	51%	36%	17%	7%

attains 50 percent is the first 1b, Minimal Information, Even Coverage.

Percentages may shift slightly as new holdings reports come into OCLC, and so close misses of 50 percent are not final. Most percentages are based on larger, sometimes much larger, bases than the 10 titles per level of brief tests. The percentages at Minimal level on the WLN scale may have bases as small as a single title, but these bases are not likely to change much, because it is so rare for titles to become library best sellers.

Down each subject in table 5, the libraries appear from lowest-rated to highest. (One exception: in 20th-Century French history, Central Missouri State University is unrated.) Trials in these nine subjects

at smaller academic, public, and school libraries would of course produce many more ratings of Minimal and Basic than appear here.

A library's coverage will sometimes slip below the 50 percent threshold at one level and then meet or exceed it at subsequent levels (for instance, Drexel's in quantum theory). Collections have not been penalized for such dips, but the coverage power test does reveal them.

The raw counts of library holdings are not given but can be computed. For example, in the literature on Nazi Germany at bottom, there are 1,228 titles at the right end of the scale. Harvard, bottom row, holds 7 percent of these. Therefore, it holds $0.07 \times 1,228 = 86$ items.

Literatures vary slightly in how many categories are needed to bracket their most widely held items. The table shows three 1a's, standing for WorldCat categories 2,000–2,499, 1,500–1,999, and 1,000–1,499. In one subject, four 1a's were needed, and the two highest items were combined in one; in other subjects, only one or two are needed. That accounts for the empty columns at left in the table.

Empty columns occur at right because some of the students truncated their tests for titles held by fewer than five libraries. This had no bearing on their results, and so the truncations have been allowed to stand.

Moving to substantive results, the tests have a high face validity. Large research libraries that one would expect to have excellent collections have them (for example, Cornell and the University of Michigan in 20th-century French history). Famous libraries tend to outdo the less famous, although there may be small surprises (at least to outsiders), such as the University of Colorado topping Cal Tech and Stanford in quantum theory.

Cumulativeness of holdings and appropriateness of ratings for type of library are two of the validity checks named in *Brief Tests*. (The other two, replicability of past results and conformity to librarians' subjective ratings, are less salient in this new context.) Confirming *Brief Tests*, the coverage power tests in table 5 leave no doubt that library collection development is cumulative across levels. Research libraries, for instance, do *not* ignore the lower WLN levels to buy almost exclusively in levels 4 and 5. Whatever the type of library, in fact, almost all collections pattern cumulatively. The rare exception stretches over a range of levels without making 50 percent in any—for example, Central Missouri State's collection in 20th-century French history. Oddities like this might be put at 1a, minimal information, uneven coverage.

The ratings also make sense for particular types of libraries. The big national

libraries attain level 4 (in environmental policy and Proust studies) or even 5 (communicable diseases). The two public libraries have good mid-level collections (Los Angeles in environmental policy and Rochester in genealogy), but academic and research libraries are stronger. Drexel, a technological university that does not collect deeply in the humanities, has only Minimal collections in French history and Buddhism but does better in quantum theory.

At least one rating seems low—Yale's 3c in Buddhism. While Yale just misses 50 percent at the first category of level 4, a graver concern is that, in a very large literature with many esoteric items, it has substantial holdings at all stages of levels 4 and 5. Like Harvard and UCLA, its modal acquisitions in absolute numbers are in a category of *level 5* (they include more than a thousand books in Japanese), and its overall coverage of Buddhism (total holdings of 5,366 / total literature of 24,635) is the best of any of the libraries shown. How Yale librarians might handle the 3c is a judgment call; they could simply declare theirs to be a special case of level 4, or take the 3c on the assumption that the collection would evolve to level 4 over time, or buy the relatively few items that would bring them to that level. Alternatively, special coverage-power scoring might be devised to take extraordinary foreign-language holdings such as Yale's into account.

In any case, the coverage power test again reveals facts about an actual collection and its relation to others that were missing from the original Conspectus. Thus, if two collections turn out to have the same rating, their performance across levels can be compared (cf. Lesniaski, 19), probably to one library's advantage. If a library sets a higher-level collection as a goal (aiming to move from 3c to 4, for example), the coverage power test shows the collection developer exactly what needs to be done to reach it. WorldCat can be used to generate the very titles that are needed to reach the goal.

Advocates of the original *Conspectus* hoped that libraries with strong existing collections would assume "primary collecting responsibility" (PCR) for *Conspectus* lines so that other research libraries could benefit from cooperative collection development and adjust their acquisitions accordingly. The 1982 *RLG Conspectus Manual* described PCR thus (1–5): "A library will not be assigned primary collecting responsibility for a subject or geographic area unless it collects at least at level 4, as defined...in this manual. Nor will a library be assigned primary collecting responsibility unless it collects at least 50 percent of what is available."¹⁶ Unfortunately, before WorldCat, librarians had no easy way of knowing whether collections covered at least 50 percent of the available literature, and this may have been one reason many *Conspectus* lines lacked a library with PCR. Table 5 shows that "50 percent of what is available" cannot be taken literally. Every literature defined by LC classes has large residues of titles in the level 5 or Comprehensive part of the distribution (that is, titles held by 24 or fewer libraries), and no library yet seen comes close to holding 50 percent of all of them. (Probably no library should want to; Comprehensive collecting really makes sense only in a sharply limited domain, such as writings by and about a famous author.) But some instances of making the first category of level 5 exist even in this very limited sample: the National Library of Medicine in communicable diseases and Harvard in the history of Nazi Germany. Surely that more than qualifies them for PCRs if this idea is ever revived.

Discussion

While the goals of the *Conspectus* movement remain worthy, the definitions of the RLG and WLN levels currently on the Web (see endnote 2) preserve all the illusions of the project as formerly conceived. For example, the 1996–1997 WLN definitions offered by Columbia University Libraries are full of phrases

that, while not meaningless, can be construed in any number of ways, none authoritative, such as "few selections," "high percentage," "broader and more in-depth array," "extensive collection," "significant numbers of retrospective materials," "basic informational needs," and "core works in the field." Moreover, taken literally, these definitions still imply that librarians will match collections in thousands of subjects against multiple complex criteria and rate them in three separate aspects: current collection level, acquisition commitment, and collection goal. Since nowadays bibliographers will rarely if ever do this, why preserve the myth of careful subjective evaluation? The authors of the definitions would say, "We mean them only as broad guidelines." But guidelines to what? Ratings made by swift guesswork are untrustworthy; ratings made by painstaking criteria-matching are impossible. Unless measurement on the RLG and WLN scales can be systematized and automated—which entails operationalizing the levels—the definitions will remain part of the culture of pseudo-evaluation that has long marked American librarianship.

The advent of WorldCat, which incorporates features of the old OCLC/AMIGOS Collection Analysis CD, makes possible a fresh approach to *Conspectus* ratings. This article has shown that, even now, WorldCat can be "bibliomined"¹⁷ to provide solid ratings of collections in the WLN or RLG style. Coverage power tests may as yet be makeshift, but they put data about collections into a new and highly meaningful order—one that relieves librarians from having to match holdings against complex verbal definitions. The tests automatically partition collections into levels that reflect the RLG or WLN definitions and their conceptual framework. In fact, the information in the holdings-count scale may capture what the definers *meant* better than their words do.

Consider two key sentences from the same paragraph of the Columbia Librar-

ies text (their order has been reversed): "In most instances, each successive level includes the elements, formats and characteristics of the previous levels.*** These are not equal incremental steps, however, since the difference from one level to the next may be measured both in terms of quantity and quality and the amount of materials necessary to move from one level to the next greatly increases as one moves up the scale." The first sentence anticipates cumulateness; the second, a logarithmic scale. Coverage power tests operationalize both notions nicely.

As for qualitative differences of the materials at various levels, only formal surveys of content will tell. But coverage power tests do a good job of modeling these levels substantively. Some indications:

- The tests model level 1, Minimal, as small numbers of selections that are, on librarians' say-so, the ones most worth having for general inquiries. The way they model level 1a ("unsystematic representation of a subject") as opposed to 1b ("systematic representation") may be seen in table 5: compare Drexel's 1a in Buddhism with its 1b in 20th-century French history, and the difference is immediately clear.

- Level 2, Basic, calls for limited collections of monographs and reference works that introduce and define a subject for the general public and college freshmen and sophomores. HM51 is too narrow a slice of sociology to include serials, encyclopedias, periodical indexes, and statistical sources, but Berkeley's collection in it, as seen in table 4, is otherwise apt: it consists of just over 100 items from trade and academic publishers that are introductory textbooks, anthologies of readings, studies of famous theorists, explorations of major sociological ideas, and bibliographic guides. The distinction between 2a and 2b rests on "defined access" to electronic resources, which would be simple to determine outside the test. Berkeley's level 2 collection does have items that OCLC labels "Internet resources."

- At level 3, Instructional, Berkeley's acquisitions in HM51 are highest in absolute numbers (279 items), and they deepen level 2's mix of titles from both trade and academic publishers. Qualitatively, they look much like those in level 2, but one sees the proliferation of both theoretical and applied topics and increased holdings in particular authors' *oeuvres*. Berkeley's holdings at levels 1 through 3 could well support instruction and self-study up through master's-degree work in sociological theory.

- Literatures in level 4, Research, and 5, Comprehensive, have by far the most titles. As anyone can check in WorldCat, these titles have diverse characteristics that no short verbal definitions can pin down. For example, in environmental policy, they are specialized as to countries and regions; in Buddhism, they are Asian in origin and language; in communicable diseases, they include reports of mainly historical interest from international and nongovernmental organizations or foreign governments; in astronomical telescopes, they tend to be very technical in a rarefied field. But in being both numerous and narrow in appeal, they qualify as the kinds of titles only doctoral students, professors, and other serious researchers are likely to want.

All of these considerations make coverage power tests a promising venture. Comparable tests of coverage are something that all librarians grasp but rarely do—certainly not with data of the sort illustrated here. In the past, in-depth evaluations of collections have proved difficult, given all the other demands on librarians' time. Such evaluations must be automated if they are to be a realistic prospect, and that is the ultimate goal of the present research.

When *Brief Tests* appeared, OCLC and RLG ran largely incompatible collection evaluation systems and seemed autonomous giants, barely in communication. About brief tests, both could claim "Not invented here." Now that OCLC, RLG,

and WLN are a single organization, the coverage power test is attractive because it uses OCLC data and software to achieve ratings that are much in the spirit of RLG and WLN evaluation but that were

not technically (and perhaps not politically) feasible before. WorldCat gives us a database in which collections could be automatically assessed—rated RLG- or WLN-style—on a comparative basis.

Notes

1. Howard D. White, *Brief Tests of Collection Strength: A Methodology for All Types of Libraries* (Westport, Conn.: Greenwood Press, 1995).

2. For RLG levels, see the Library of Congress at www.loc.gov/acq/devpol/cpc.html or the University of Texas Libraries at www.lib.utexas.edu/admin/cird/policies/newlevels.html. For WLN levels, see Columbia University Libraries at www.columbia.edu/cu/lweb/services/colldev/collection-depth.html. [Accessed 5 November 2006].

3. For the Collection Analysis system, see www.oclc.org/collectionanalysis/default.htm. [Accessed 5 November 2006].

4. Thomas M. Twiss, "A Validation of Brief Tests of Collection Strength," *Collection Management* 25 (2001): 23–32.

5. David Lesniaski, "Evaluating Collections: A Discussion and Extension of *Brief Tests of Collection Strength*," *College & Undergraduate Libraries* 11 (2004): 11–24.

6. Jay H. Bernstein, "From the Ubiquitous to the Nonexistent: A Demographic Study of OCLC WorldCat," *Library Resources & Technical Services* 50 (2006): 79–90.

7. White, *Brief Tests*, 138–39, 145–47.

8. The upper part of the WorldCat holdings-count scale can be approximated as 25, 50, 100, 200, 400, 800, 1600, and 3200. Converted to base-10 logarithms, these numbers form a scale with equal intervals of 0.3—that is, 1.4, 1.7, 2, 2.3, 2.6, 2.9, 3.2, and 3.5. This scale maps well onto RLG levels 1 through 4. The part of the WorldCat scale below 25 is slightly less regular in base-10 logs, but that does not matter because all values there simply become RLG level 5. The simplifying tactic of using "150 and below" as a cutpoint for level 4 (and now level 5) is explained in *Brief Tests*, 123–37. See also Lesniaski, 16–17.

9. With Excel 2004 on a Macintosh, one simply copies the literature or collection distribution along with the WorldCat scale, pastes it into the spreadsheet, names the literature and collection columns as variables with Insert > Name > Create, uses the Calculator to divide "Collection" by "Literature" (if those are their names), and converts the resulting decimal fractions to percentages with Format > Cells... > Number > Percentage (setting "Decimal places" to zero). With PCs, the process is slightly slower because Excel 2003 will not compute ratios from pasted figures (they are not in numeric format), nor does it have the Calculator. After cutting and pasting the literature or the collection figures from WorldCat into the spreadsheet, the data must be re-entered by hand in columns defined by Format > Cells... > Number > Number. (This should take less than a minute per column.) Using the Formula bar, one then divides the count in the *top cell* of the collection column by the count in the *top cell* of the literature column, selects the result along with the empty cells in the results column that correspond to the filled cells in the two data columns, and chooses Edit > Fill > Down. This fills the cells with decimal fractions that can be converted to percentages as above.

10. White, *Brief Tests*, 147–49.

11. A short introduction to this work is M. Carl Drott, "Bradford's Law: Theory, Empiricism, and the Gaps Between," *Library Trends* 30 (1981): 41–51. A more recent discussion will be found in Stephen J. Bensman and Stanley J. Wilder, "Scientific and Technical Serials Holdings Optimization in an Inefficient Market: An LSU Serials Redesign Project Exercise," *Library Resources and Technical Services* 42 (1998): 147–242. The latter is online at www.lib.lsu.edu/collserv/lrts/. [Accessed 5 November 2006].

12. White, *Brief Tests*, 139–42.

13. *Ibid.*, 142–45.

14. For the Audience Level Prototype, see www.oclc.org/research/researchworks/audience/default.htm. [Accessed 5 November 2006].

15. The Table 5 literatures are defined with LC classes as follows: 20th-century French history: DC361–DC424; Buddhism: BQ7001–BQ9800; environmental policy: GE170–GE190; Proust studies: PQ2631.R63; astronomical telescopes: QB88; quantum theory: QC174.12; communicable diseases: RA643*–RA644*; genealogy: CS42–CS65; Nazi Germany: DD24*. My thanks to the students whose projects these were—respectively, Carol Smith, Regina Hart, Alison Vasquez, Christopher Matson,

Melissa Woolley, Michael Dulock, Bobbi Otte, Nancy McCrave, and Adrienne Johnson.

16. Research Libraries Group, *RLG Collection Development Manual*, 2nd ed. (Stanford, Calif.: RLG, 1982).

17. The coinage is Scott Nicholson's. See, for example, his "Bibliomining for Automated Collection Development in a Digital Library Setting: Using Data Mining to Discover Web-Based Scholarly Research Works," *Journal of the American Society for Information Science and Technology* 54 (2003): 1081-90.

70,000 Publishing Sources. 3.2 Million Book Titles. We're Connected. Are You?

IT'S EASY TO BE CONNECTED WHEN YOU'RE AN EMERY-PRATT CUSTOMER.

At Emery-Pratt, we work extra hard to be closely connected to more than 70,000 publishing sources, allowing you to have access to millions of new book titles. And thanks to our close relationship with these sources, we learn about the hottest upcoming books long before they're released. As a result, we can make sure you and your customer get all the books you want...right when you want them.

Would you like to be connected? It's easy. Just call us at 1.800.248.3887. We'd love to help you anyway we can. Or go online at emery-pratt.com

THE NICEST PEOPLE IN THE BOOK BUSINESS

Emery-Pratt[®]
Book Distributors since 1873

VISIT OUR WEBSITE TO SEARCH
FOR BOOKS, PLACE ORDERS,
AND MUCH MORE.

www.emery-pratt.com

1966 West M-21, Owosso, MI 48867-1397 | Phone (toll-free) 1.800.248.3887
Fax (toll-free) 1.800.523.6379 | E-mail: mail@emery-pratt.com