

How I Stopped Worrying and Learned to Love the Usage Data

Michael P. Hughes

The digital environment has transformed how data on library material use is collected and reported, providing librarians with more information about usage but less clarity about how to interpret it. This article discusses current approaches to reporting and assessing library book use, addresses the question of what qualifies as a worthwhile use of library materials, and presents an analysis of four years of COUNTER 4 BR2 ebook reports at a single research institution to explore the reliability of page view-level usage data for collection assessment. It reveals ways assessment theory and practice fail to capture the value of library materials throughout the research lifecycle, and argues for an inclusive view of collections use.

Introduction¹

With the proliferation of electronic resources, and ebooks in particular, there have arisen a multitude of ways to record how library patrons use library materials. Early on, collections librarians relied on idiosyncratic reports that varied from vendor to vendor, which made comparisons between different packages from different vendors impossible. The work of Project COUNTER to develop a standard for reporting electronic resource usage, and corresponding reports reflecting that standard has improved things greatly by encouraging consistency across platforms.² Despite some variation from vendor to vendor in how to interpret the standards, the collection manager is currently in a much better position to understand the use of electronic resources. However, there remains uncertainty about how we are to understand some of the data provided, including which reports are more likely to provide reliably actionable insights into the use of our collections. Perhaps the one most likely to be questioned is the “successful section request” reported in COUNTER 4 BR2, which reports the total number of successful user requests for the smallest section of a book provided by the vendor. Some vendors provide page-level reporting, while others serve books by chapter or section. This article takes the uncertainty engendered by the BR2, and particularly by reports that record use at the page level, as the starting point for an investigation of collection use theory both as practiced and as implied by common practices. At issue are 1) what qualifies as a worthwhile use of our materials, which is to say what makes a use significant enough to record; 2) the ways our theory and practice fail to capture the value of library materials throughout the research lifecycle; and 3) how we might move forward using the data we already have. After a discussion of the history

*Michael P. Hughes is Collection Management Librarian at Brooklyn College, City University of New York; email: michael.hughes@brooklyn.cuny.edu. ©2020 Michael P. Hughes, Attribution-NonCommercial (<https://creativecommons.org/licenses/by-nc/4.0/>) CC BY-NC.

and theory of monograph use analysis, the article presents an analysis of four years of Ebook Central³ COUNTER Book Report 2 (BR2) reports at a single research institution to support an inclusive view of collections use.

What Is a Use?

Prior to the tracking of networked electronic resource use, data available in studies of print use were most often limited to circulations (that is, cases where a patron made the commitment to check out a book from the library). In such studies, we could have no record of any use prior to the initial circulation by the patron, nor any insight into how the patron used the book during the loan period. And although librarians have long been aware of the limitations of circulation data in representing collection use,⁴ and despite some attempts to measure in-house use,⁵ the idea that real use was measured by such a commitment on the part of the patron has taken hold and has influenced the view of use data to this day. This influence is seen most clearly in the ways assessment confers more value on use types that seem to reflect a greater level of engagement by the user, such as circulation or full-text access or download of electronic books and articles. Discussing librarians' attitudes toward ebook use, Robert Slater cites "a widely held but empirically unsupported belief that an access of an e-book represents a less-thorough use than if a book is checked out."⁶ One of the indisputable advantages of the circulation approach is that the data for such uses are easily quantifiable and allow us to make at least some sense of an obscure phenomenon: what are our users actually doing with our materials? However, it is not self-evident that such extended or committed usage reliably reflects more important use, nor that materials used in this way are of greater import than others.

George S. Bonn was the first to provide a tool to analyze raw circulation numbers into something actionable through his introduction of "use factor" in his still-relevant 1974 article "Evaluation of the Collection."⁷ As the name suggests, use factor is a self-referential measure, where the number generated, in this case through a ratio of circulations or other use to holdings, is a factor of the holdings number. For example, a collection with 400 volumes and 600 circulations in a year has a use factor of 600:400, which is to say, 1.5. This metric appears under different names in the literature, such as Percentage of Expected Use (PEU)⁸ and Performance.⁹ Circulation studies lend themselves to such quantitative metrics as rate of use to holdings or the rate of use to acquisition, which, while illuminating, have allowed the false impression to take root that library books are commodities akin to items on a supermarket shelf and that the goal of the library is to circulate those items (in a manner of speaking, to "sell" them). While it is true that books on the market are commodities, they are no longer so once they are *library books*, and to think of them in those terms obscures the values of these materials. Lacking data on more ephemeral uses, which would give a richer view of collections use (such as in-house use or brief consultation in the stacks), collection managers have found their views shaped by the limitations of the data presented to them. However, what librarians have been able to identify as "use" does not capture the wide range of productive and essential uses of books and other resources in the research process, and we are therefore at a disadvantage in assessing the value of our collections to our current users.

Circulation data became the standard because circulation is measurable and easy to understand; however, it has also given rise to a limited conception regarding what constitutes a use of a book. Relying on print circulation as the standard for use not only misses important

cases of use of library materials but has also given rise to some confusion about what we mean when we talk about patrons using library materials. The difficulty arises, at least in part, because when we talk about “use data” we are prone to focus on the “data” and can easily forget that “use” refers to what our users actually do with our books, not simply to what we can record, and further that use of library materials is not essentially transactional but is involved in knowledge production. The use of library materials is crucial at each point of the research lifecycle, and the nature and extent of that use varies across that lifecycle. Recorded circulation reflects only a portion of use at a particular point in that lifecycle. In the examples that follow, “book” is simply a rhetorical stand-in for any library materials.

Library Materials and the Research Lifecycle

Imagine a researcher setting out on a new project. In a much-simplified form, this is what the process might look like in terms of library books. With the initial research topic in mind, say “action theory in moral philosophy of the eighteenth century,” the researcher may check the library catalog (or discovery system) for the topic “action theory” and retrieve a list of records of books held in the stacks. Taking down one or more of the call numbers, the researcher heads to the stacks, finds the appropriate range, and walks toward the books with the target call numbers. It is not controversial to say that most researchers of this type will look at the books in the same subject area in the stacks, without limiting the time in the stacks to retrieving the initially identified items. Suppose, then, that our researcher sees a book on the stacks with a title like “Action Theory and Group Action.” Our researcher is not an expert in action theory, and is curious about this subtopic, so takes the book from the shelf, peruses it briefly, contemplates whether the book is relevant to the topic at hand, decides it is not and places it back on the shelf. Likewise, a researcher who knows more about action theory may simply look at the title on the spine, consider whether the book is relevant, and decide against it without opening the book at all. What is to be made of this? Has the researcher used the book? Is this an example of research? What was the goal of opening the book? What was the effect of perusing it? Not only has the researcher used the book, but through this use the topic has become clearer and the research more focused. Regardless of how short or cursory the interaction was between the researcher and the book, this is a case of productive research. In the early stages of research, this is a crucial use of library materials. As with the term “book,” the terms “research” and “researcher” are used here for the sake of convenience. The example above applies just as aptly to students gathering material for a term paper at whatever level, and equally to a public library patron who, interested in learning to knit, realizes that the perused book teaches a previously unknown style of knitting in which our patron, upon learning of it, has no interest.

Truncating the process considerably, imagine that, by the time the researcher locates the originally sought books, the parameters of the topic have been set and a certain subset of books has been selected as appropriate for deeper use. In some cases, there may only be a single chapter of interest, and our researcher will read the chapter in the library or make a scan without checking out the book. In other cases, the researcher will check out the books and the collection management librarian finally has evidence of use.

While the researcher has the books checked out, of course, the collection management librarian has no insight into how the books are being used, but since the researcher has committed to carrying the books home, this is our classic use case. Common sense and personal

experience tell us that some of the books will be used extensively, the researcher poring over hundreds of pages, reading and rereading sections, perhaps scanning chapters at home to read on a portable device when away from home, whereas other borrowed books may never be opened or might be opened once and then rejected as being irrelevant to the project. The latter is a use akin to the perusal in the stacks, but the former is not a use at all (that is to say, the use really only took place in the stacks, prior to checkout, and the recorded circulation does not correspond to any greater or additional use).

Over the course of the project, the researcher will likely repeat these steps as the project takes shape, but at some point the books will be returned. From the circulation perspective, this is the end of the use. As every reference librarian knows, however, even this late in the research process, researchers are not done with the books. Often, a researcher will need to check a book for a citation, or a quote, or simply reread a passage quickly to clarify matters. And, just as often, the researcher will choose not to check the books out for this purpose but will go to the stacks or rely on the reference librarians to check a citation without the researcher engaging with the physical book again. Clearly, these are also cases of uses necessary to the successful research process. Clearly again, these sorts of uses are not recorded. What's more, these cursory uses are not limited to the same project (for which, perhaps, the researcher did check out the book) but can carry over from project to project as researchers build their own sets of notes and citations. A book checked out for one paper may be used only cursorily for another, leaving no record of use for that project.

Since circulation data alone do not reflect all of the important uses of a given title, suppose instead that we were able to look at a different metric to capture use.

- What if, in addition to the initial checkout, we counted all the times the borrower used the book during the loan period?
- What if we counted the times someone sat with a book and read from it in the library, but did not check it out and take it home?
- What if we counted the times someone browsed the contents of the book while standing in the stacks, but took it no further?
- What if we counted the times someone paused in the stacks and considered the information printed on the spine before moving on to other books?
- What if we were able to identify all the checked-out books that were never used by the borrower?

If we were able to collect and analyze the data types presented above, we would have a much richer understanding of library collection use. We might even think that this would provide us with clarity around library use. Much of what is described in the list above is simply the tangible equivalent of data we can extract from many ebook usage reports, and yet the meaning of our data seems to become no clearer. In some reports, we can look at the number of times a user looked at a record, but not the book. In some cases, we can see when an ebook has been downloaded; other times when there has been a "successful section request" indicating that the user consulted a portion of the book. Many of us have been involved in conversations about what constitutes the use of an ebook, with the result often being a lot of head-shaking and hand-wringing. Anecdotally, in many cases a single download is, explicitly or not, seen as a "real use," while the value of 12 successful section requests remains murky. This is not to say that librarians have not tried to clarify matters. In her method developed to compare page-level ebook usage data with print circulations, Cathy Goodwin identifies a threshold of 11 for

the minimum number of page views in a single session to be considered “substantive use” to capture more in-depth use of ebooks reported this way.¹⁰ While Michael Levine-Clark, Kari Paulson, and Paul Moeller note the variety of uses of ebooks and the danger of favoring one sort over another, they, too, prioritize some uses over others in terms of the value provided by the resource: “If a library were to only consider page views or session information, it would be missing the fact that when a history book gets used the user interacts deeply with it, and thus it provides good value.”¹¹ Implicitly, at least, we treat some uses as better than others.

Are Some Uses Better Than Others?

Although Bonn’s introduction of use factor was instrumental in allowing collection management librarians to measure and assess the use of their collections, it is fundamentally limited as a metric for collection use. Use factor compares the rate of holding of books in a certain area (often LC subclass or other subject area) to rates of circulation, with the arbitrary ideal being a 1:1 ratio, so that a library with 400 books on knitting should hope to see 400 circulations in a given period. If the number of circulations is below 400, then the collection management librarian is alerted to the possibility that the intensity of collection development in this area has been excessive (although the disparity could just as easily signal that greater outreach to local knitting circles is appropriate); likewise, if the number of circulations exceeds 400, it may be the case that additional resources should be devoted to collecting books on knitting. As we have seen, circulation numbers do not reflect all uses, so the data have only limited power to illuminate (and certainly not enough power to justify a data-driven approach to collection management). Use factor, however limited it may be, is not itself fundamentally flawed; however, it has provided the basis for the importation and development of faulty models of collection assessment based on a commodity model of library materials. For example, modifying the use factor model to compare rate of use with rate of acquisition over a certain period is supposed to provide the theoretical basis for collection development decisions to reduce acquisitions in areas with lower circulation than acquisition. Although this appears to be a helpful and illuminating metric, it is instead a good example of the misuse of commodity thinking in collection management. In fact, a major stumbling block to interpreting the use of our collections is the importation of business models, which can obscure matters and confuse budgetary with service standards.

To understand the fundamental mismatch between the commodity model and the library, it is fruitful to examine where the commodity model is appropriate. Take, for example, a supermarket. The supermarket acts as a go-between for producers and consumers. The mission of the producers is to sell food and other products to supermarkets; the mission of the supermarket is to sell those products to consumers. That is why a supermarket purchases the products. This is all very clear, and in such a case it makes perfect sense for the supermarket manager to compare the number of cans of peas ordered with the number of cans sold, and to reduce the number of orders in response to poor sales. To do otherwise would be foolish on the part of the manager. Things are different, however, from the customer’s perspective. Imagine a shopper in the canned vegetable aisle of a supermarket. This shopper is in the market for peas. While walking down the aisle to the desired brand of peas, our shopper stops to look at another brand. Perhaps the shopper takes the can of peas off the shelf, reads the label, and considers purchasing this new brand. Our shopper, however, puts the can back on the shelf, proceeds as originally intended to the desired brand, takes a can off the shelf, purchases it

and later consumes its contents. It is obvious that the purchase of the peas constitutes a use of the peas that fulfills the missions of the producer, the supermarket, and the shopper. But what of the interaction with the rejected can of peas? As with the book on action theory and group action in the earlier example, the consultation of the label of the new brand of peas seems to have clarified matters for our shopper. Maybe there was something in the ingredients list that was undesirable; maybe the shopper read a recipe on the rejected can, committed it to memory, and prepared it later that day using the competing brand of peas. Whatever the case, we can reasonably say that this interaction with the rejected peas helped to fulfill the shopper's mission, being part of the information-gathering involved with pea purchasing and ultimate meal preparation. It does not, however, fulfill the mission of the producer of those rejected peas, nor the mission of the supermarket, as the shopper did not purchase more than initially planned due to this interaction, and the supermarket would have benefited equally had the customer never consulted the rejected can of peas. No matter how much more certain the shopper is of the decision to purchase the original brand, no matter how much more expertly the shopper prepares the peas, nothing in the interaction with the rejected peas fulfills the missions of the producer or supermarket. So we can say that the missions of the producer and supermarket, taken separately or together, do not correspond exactly to the mission of the shopper, and that there are ways that the shopper can use the products of the producer and supermarket to fulfill the shopping mission while being contrary to the missions of the producer and supermarket (their missions are never fulfilled when something is not sold).

Returning to the library, we find that, unlike in a supermarket, where a customer's passing interest in a product ultimately rejected does not fulfill the mission either of the supermarket or the producer, when a patron examines a library book, or any other library resource in the course of research, even the recognition that the book in question is not relevant to the project, fulfills the purpose of the library. **The fact is, there is no category of use that runs contrary to the library's mission.** This is not to say that collection managers should devote limited funds willy-nilly based on cursory use; nor does this mean that we should treat the numerous and undifferentiated uses reported in a BR2 the same way we do with title-level BR1 or circulation reports. If we recognize, however, that all uses are in keeping with the mission of the library, then distinguishing among different types of use takes on a different, less decisive character. We make better decisions when we understand the limitations and biases of our data and models.

To this point, the focus has been on what the data do not show us, specifically in the case of circulation, where the checking out of a book gives us no insight into how the book is used or what other books not selected for borrowing were useful to the researcher. At the other extreme are reports like COUNTER BR2, which gives counts for "successful section requests" and appears at first blush to provide an inflated and, perhaps, unreliable accounting of use.

What Do We Want the Data to Show Us?

When determining whether a certain dataset is relevant or illuminating to a given purpose, we are looking at information quality (IQ). Information quality is context-dependent, meaning that what determines the quality of the information is the project itself, not an abstract ideal of IQ. Nor is it the case that the more granular the data, the higher the IQ. Rather, it depends on what we want to know. In the case of circulation data in a project evaluating patron use of the collection, IQ is quite low due to the limited nature of the data, which show us only those

cases where a user checks out a book and misses all of the other uses. However, if what we want to know is what percentage of our collection circulates at any given time to make stacks management decisions, then the same dataset has high IQ. The IQ does not increase for a project if the detail or completeness of the information exceeds the requirements of the project. In many cases, details like the number of unique sessions, total number of full-text downloads, or even the total number of successful section requests for individual items do not increase IQ over aggregate data. When looking at a collection, or a subset of a collection (grouped, say, by subject area), there is a point at which too much detail brings no greater insight. From a collection development perspective, what we really want to know is not primarily which individual books are being used, or which books are subjected to extremely high successful section requests, but more what categories of titles are being used. Analogously, if we want to know which serials subscriptions to renew, we don't look at lists of all of the articles used but instead at the aggregate use across the title. In cases where the object really is to identify the individual titles with the highest and lowest use, it is easy enough to find that information without doing any sort of analysis or significant manipulation of the data.

In this paper, it is argued that a use of library materials, however measured, fulfills the mission of the library and is thus reaffirming of the choice of the title for inclusion in the collection. This appears to be undermined somewhat by inflation when something as granular as a single page view counts as a use of the item, as in the case of some COUNTER BR2 reports. The COUNTER BR2 report gathers the number of "successful section requests" for any ebook, the definition of which varies by vendor and platform. If a platform offers ebooks at the chapter level, then requesting and receiving a chapter counts as one successful section request, regardless of how many pages are read. If a platform offers ebooks at the page level, as is the case with Ebook Central, then each page viewed counts as a successful section request. It is immediately clear how this disparity can lead to confusion among collection managers and others. This issue of the varying definitions of a section in BR2 reporting has been discussed by Karin Byström¹² and Jonathan H. Harwell and Erin Gallagher.¹³ According to the COUNTER 5 documentation, the current TR_B1 does not place more prescriptive requirements on reporting of section requests, but provides "comparable statistics" to the COUNTER 4 BR1 and BR2.¹⁴ Since the purpose of the current study is not to assess or criticize the COUNTER approach to reporting data, the supersession of the BR2 report is not relevant to the discussion at hand. Taking COUNTER 4 BR2 reports as our example, when a page is viewed in Ebook Central, it counts as a use; when the same page is printed, it counts as a use: one page, two recorded uses. This sort of thing appears to lead to wildly inflated numbers that can't be easily reconciled with our traditional view of use as circulation, and, as a result, when talking among themselves, many librarians suspect the BR2 of somehow gaming the system to make it appear that use is higher than it is. If this is the case, then the IQ of the BR2 would, like standard circulation data in a print collection usage study, be low. In the case of the BR2, rather than excluding real uses, the abundance of data points would potentially be masking some real use with duplicative and irrelevant data.

Methodology

The approach to use data analysis presented below arises from work in collection development, where the concerns are not primarily budgetary but are directed toward creating a collection to support subject-specific research. This subject-level approach is of primary importance

to collection assessment. Further, when we look at our collections from a purely budgetary perspective, we are no longer assessing the collection as an intellectual resource, but as a financial liability. The argument here is not that we should refrain from assessing our collections for financial viability in the context of our materials budgets, but that assessing them primarily financially can cause us to lose sight of the objects we are assessing, which have (or lack) value that cannot be reduced to economics. From this perspective, it makes sense to look at aggregate use by field or subject (such as by LC subclass or LC topic). This allows the collection manager to see which areas are well-used, perhaps requiring additional resources, and which areas are less well-used. With print circulation, we know that our circulation data underreports use for many items, as discussed above; for ebooks, we seem to run into the opposite problem of inflation of use reporting. For example, in the collection analyzed below, there were 981 Ebook Central books classified in LC basic class N (the arts in general); for the month of December 2015 alone, there were 3,300 successful section requests from those books (a use factor of 3.4). By contrast, the print holdings in the main library at the same institution in LC basic class N numbered 40,200 at that time with a circulation of 3,380 for the entire fall 2015 semester (a use factor of 0.08 over a longer period). These two sets appear to be irreconcilable, and the ebook data appears to be unreliable. While it is not the purpose of this paper to compare print and ebook usage, interested readers can find studies of such comparisons by Karin Kohn,¹⁵ Justin Littman and Lynn Silipigni Connaway,¹⁶ Steven Knowlton,¹⁷ and others. The question here is simply: how can we make sense of the ebook data?

The study here begins with the following hypothesis: if we value all use equally, then we can gain an illuminating view of the data by transforming the title-level data in a binary fashion (used/not-used) and then running the same analysis as performed on the raw data. This transformation to used/not-used is employed in the three studies comparing print and ebook usage cited above because, in addition to facilitating comparison between ebook and print circulation data, it “alleviates the problem of inconsistency in COUNTER reports.”¹⁸ In statistics, this sort of data transformation is called censoring—and, more precisely, right-censoring. Right-censoring sets a maximum observable value, C . According to Joseph M. Hilbe, “this value in the data actually means ‘greater than or equal to C .’ If $C = 15$, then all response values greater than 15 are revalued to 15.”¹⁹ Because in our case $C = 1$, this transformation can also be called binary-censoring. In many cases, censoring takes place because the data are incomplete and not all values are known; in other cases, as is the case here, censoring is done deliberately to limit the values for one reason or another. In their study of methods for analyzing different levels of time-series data, James E. Alt, Gary King, and Curtis S. Signorino make the clearest argument for transformations of this kind: “We do not want the form in which the data happen to be collected to determine the substantive ideas which we can explore.”²⁰ In this study, we are transforming the original event count dependent variable into a binary, or dichotomous, dependent variable for the same times series.

That is, if we have use data for each title by month, instead of treating the 213 successful section requests of *Etymologies of Isidore of Seville* in July of 2012 as 213 uses, we give *Etymologies of Isidore of Seville* a use number of 1 for that period. Likewise, for a title that had only one successful section request during that period, we give that title the same value of 1. That is, we give a title with 213 successful section requests during our sampling period the same weight as a title with only a single successful section request, and we transform the raw number of successful section requests per title into a use-month with a value of either 0 or 1. *Prima facie*

this seems to be absurd, as it appears to be obvious that reducing the high numbers of uses will skew the aggregate data into meaninglessness. In an attempt to counteract the inflation of our raw data, are we not throwing out 212 babies with the bathwater (even if we suspect that some of those babies are just dolls and not real babies at all)? Are we not stripping our use data of all of its richness and usefulness? The argument presented below is that not only is this transformation legitimate even in the absence of print comparisons, but also that it tells us something interesting about how our materials are being used, and that it is illuminating to the study of use analysis in general.

A Study of Ebook Central Use

The data analyzed here, derived from monthly COUNTER BR2 reports with added Library of Congress call numbers for each title, were collected each month during 48 months from January 2012 through December 2015 at a R1 institution. Holdings data are taken from January 2016. Library of Congress call numbers are included in the holdings report, but not in the BR2 report. To associate the call numbers with the correct titles in the BR2 report, it was necessary to match the titles in each report by the vendor DocID, which is included in both reports, and then essentially to create a custom BR2 report including the call number field. Once the call numbers were associated with the titles in the BR2, the data were analyzed by LC basic class and subclass using custom scripts written by the author for collections holdings and use analysis.²¹ Then the dataset was copied and the use numbers were reduced in the manner described above: any use total greater than 0 for a title in a given month was transformed to 1; months without use of that title remained at 0. The transformed dataset was then processed in the same way as the raw data for side-by-side analysis. The initial hypothesis of the study was that analysis of the transformed dataset would present a significantly different view of the use of these materials by correcting what appears to be inflation. By removing the excessive, presumably duplicative, *intensive* use reflected in page views, prints, and downloads, it was expected that falsely identified high-use areas would be exposed as misleading and that other areas would rise to the top based on *extensive* use. As it turns out, however, the areas with intensively used titles (that is to say, those titles with many successful section requests each month) are also the most extensively used (that is, the areas with the most use-months over the life of the study). Furthermore, the relative use by LC subclass reflected in the binary-reduced data (in other words, extensive use) is strikingly similar to the relative use reflected in the original data (that is, intensive use).

In the original dataset, a total of 12,314,975 successful section requests across the entire collection in Ebook Central were reported. While nearly all of the titles in the analyzed collection could be matched to call numbers, a small number of titles did not have associated call numbers and therefore are not included in this study. Reducing the data from the number of successful section requests overall to the number of use-months for each title in the study produced a total of 311,557 use-months across the entire collection in Ebook Central during this period. This is a reduction of 97.47 percent from the original data. As discussed in more detail below when looking at particularly well-used subclasses, this reduction is typical across subclasses. To arrive at a meaningful number, any LC subclass (such as HQ or KFA) with fewer than 100 successful section requests over the course of the 48 months was excluded. The purpose was to remove subclasses that have insufficient use to warrant much attention to avoid skewing the data one way or another. Many of the low-use subclasses saw no re-

duction whatsoever. The mean reduction percentage at the LC subclass level over 48 months in this study is 96.43 percent; the median is 97 percent and the coefficient of variation for all subclasses with more than 99 successful section requests is 2.18 percent, indicating very little dispersion across subclasses. The consistency across the collection of the ratio of successful section requests to use-months is striking (see figure 1).

FIGURE 1
Distribution Curves for 48 Months of Successful Section Requests Compared to the Same Data Measured in Use-Months

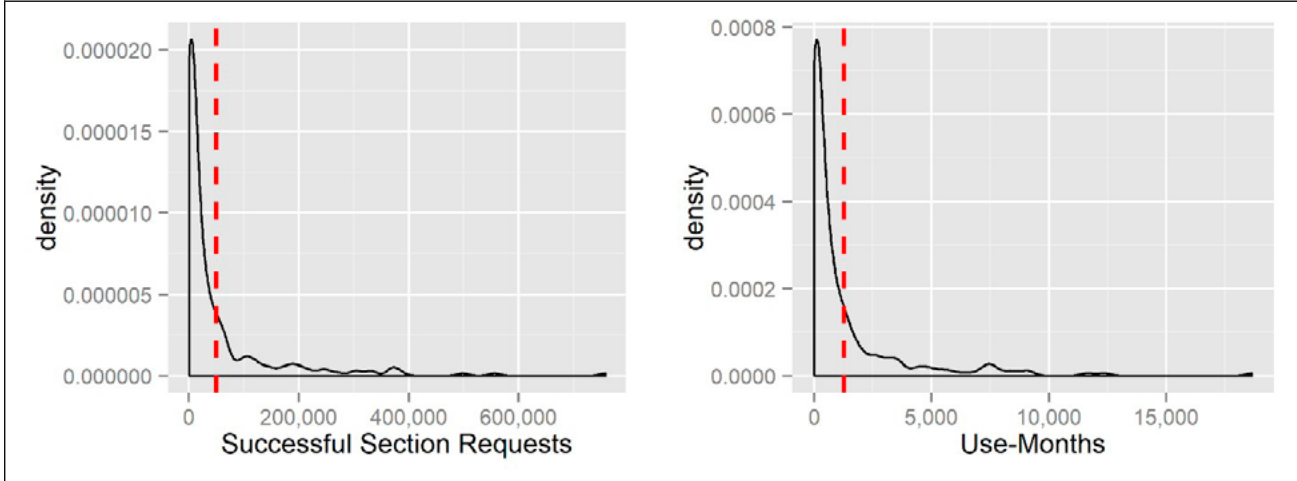
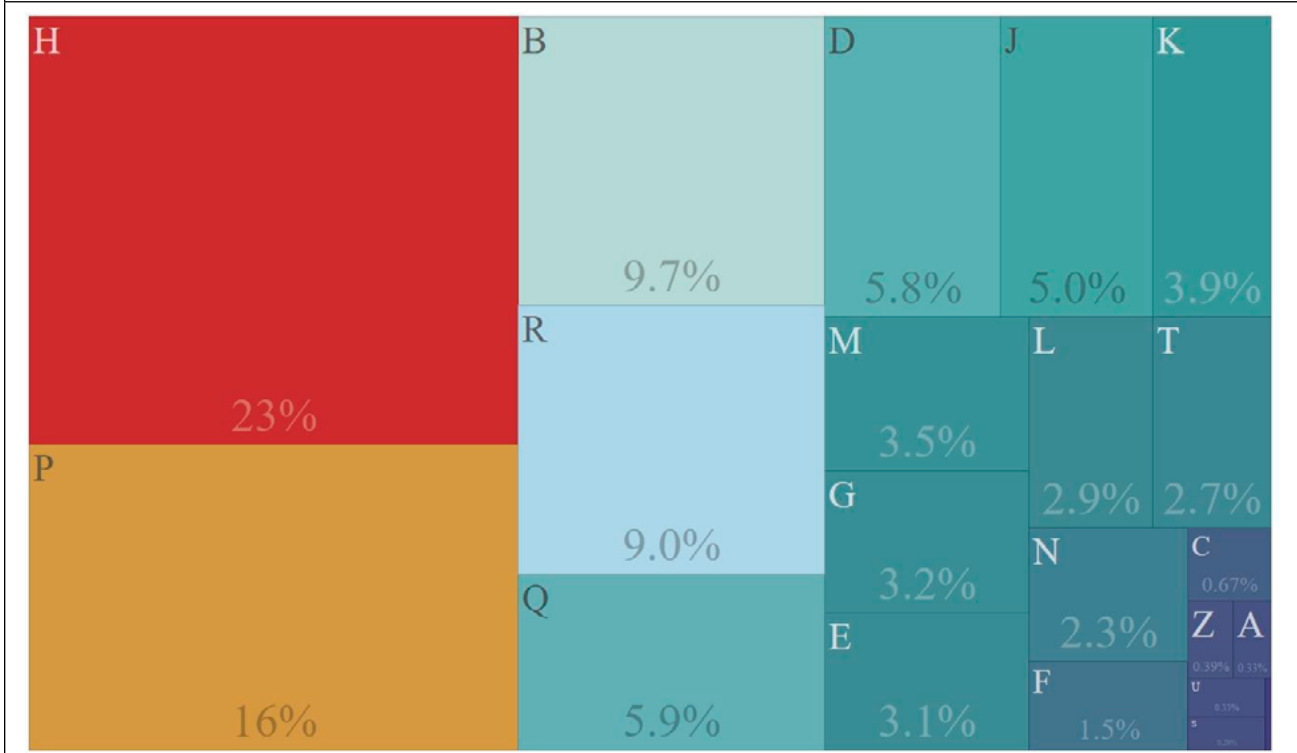


FIGURE 2
Ebook Central Successful Section Requests per Month January 2012 through December 2015



Looking at monthly use data at the basic class level, LC class H consistently makes up the largest percentage of successful section requests (see figure 2). Across the entire four years of this study, we see that there were 2.79 million successful section requests for ebooks classed in H, making up 23 percent of the total number of successful section requests during this period. Class P makes up the second largest group by use with a 16 percent share due to the 1.99 million successful section requests, compared to 35,400 titles held, accounting for 13 percent of the total.

If we subject our monthly sampled data to binary reduction to derive the use-month value, where each positive value for a title counts only as one, and then those monthly totals are added up across the 48 months of the study, LC class H use counts for only 63,100 uses. Binary reduction of successful section requests has resulted in a reduction of almost 98 percent. In the reduced data, class H continues to be the most-used class in the transformed dataset, making up 21 percent of uses (compared to 23 percent in the original data) to 17 percent for P (compared to 16 percent in the original data) (see figures 2 and 3). Looking more closely at the subclasses, we see a similarly slight change. In the raw data, subclass HD accounts for 555,000 successful section requests during these 48 months, making up 20 percent of total uses in the H-subclasses (see figure 4). Binary reduction of the HD data results in 11,600 uses, which accounts for 18 percent of the total (see figure 5). As with the higher-level view, the relative values of the subclasses remain largely the same.

FIGURE 3
Ebook Central Successful Section Requests after Binary Transformation to Use-Months
from January 2012 through December 2015

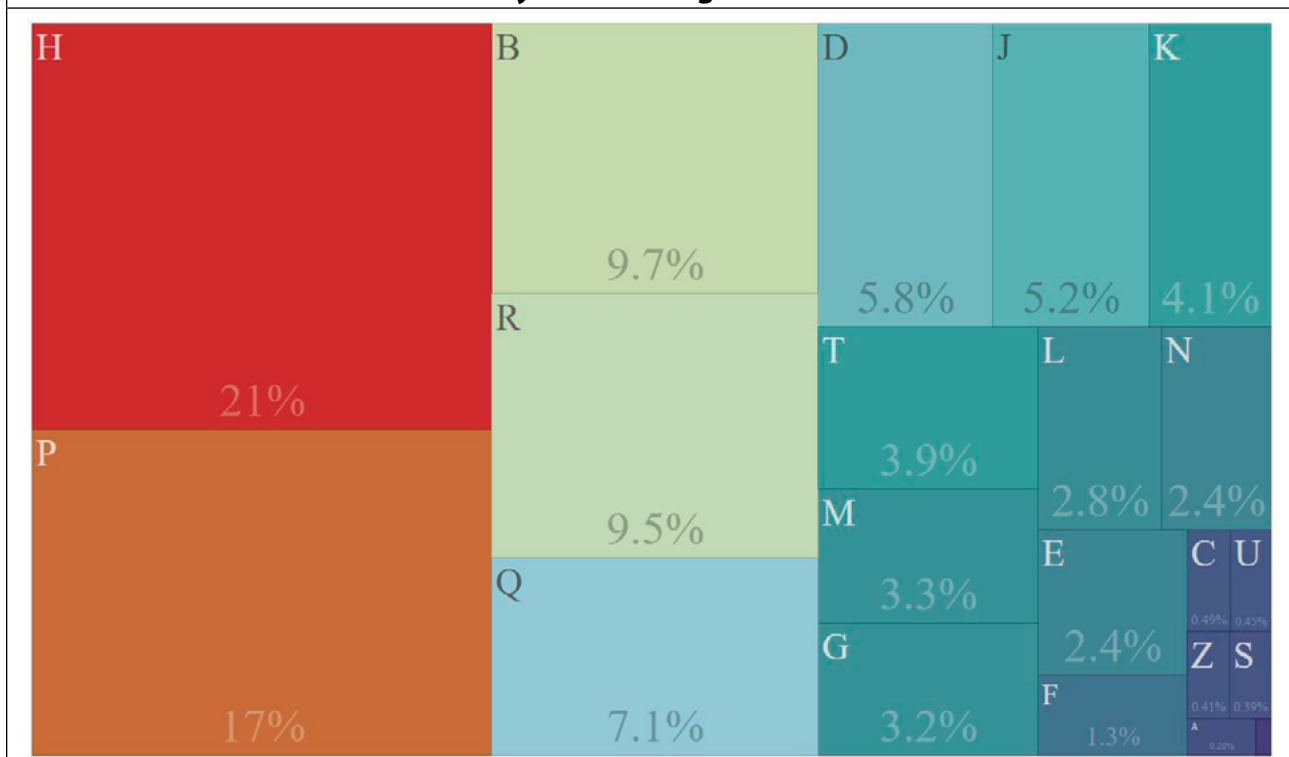


FIGURE 4
H-Subclasses in Ebook Central, Successful Section Requests Sampled Monthly from January 2012 through December 2015

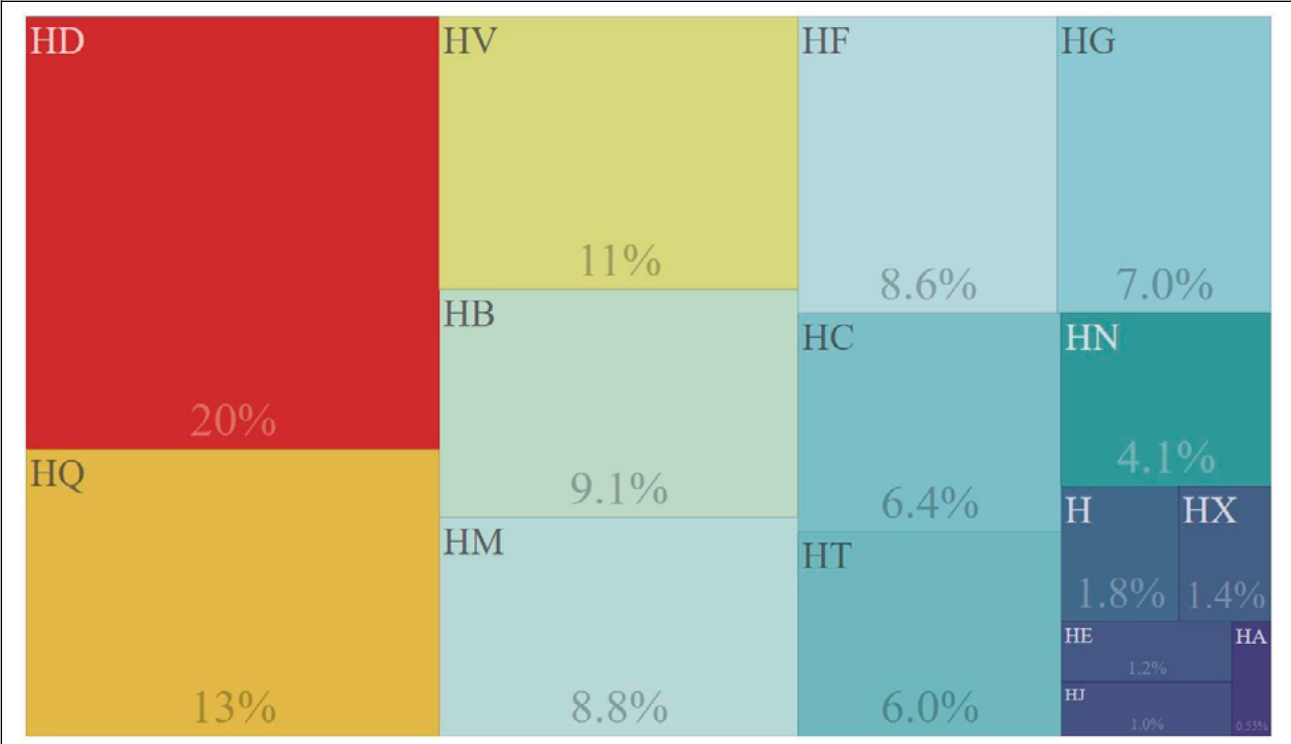
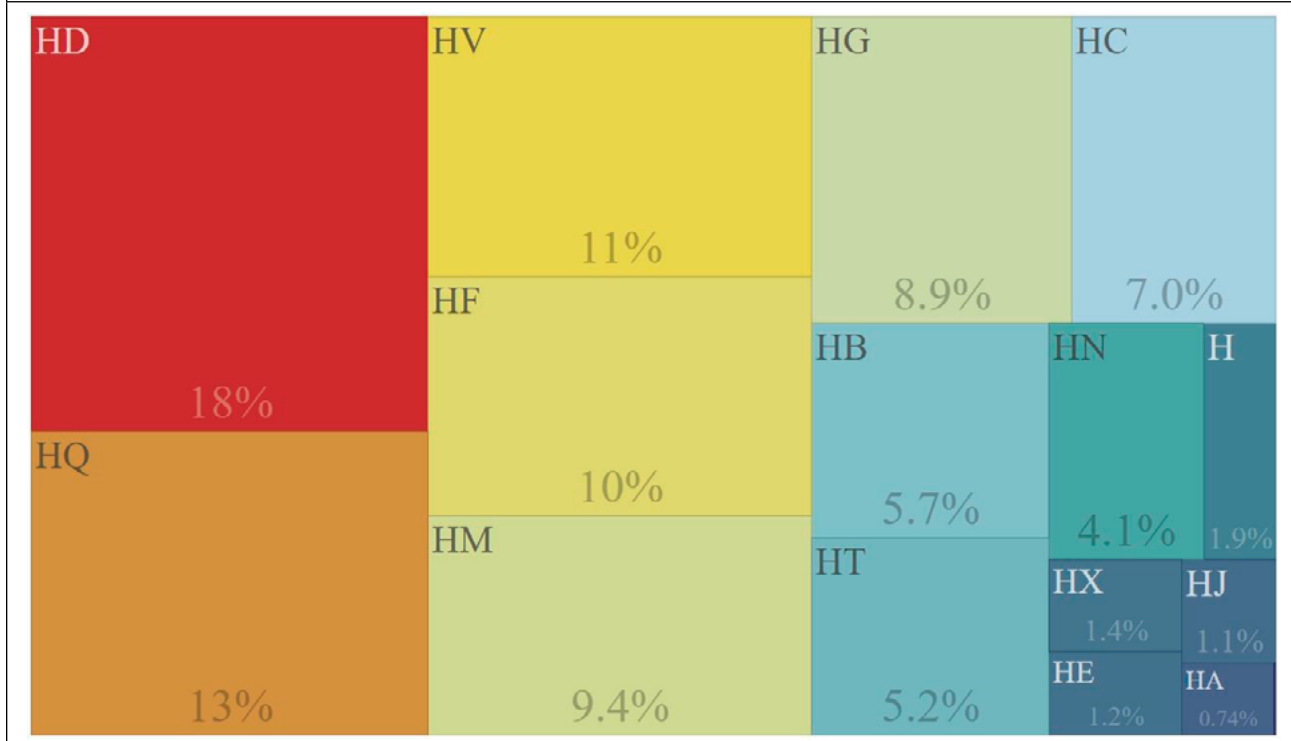


FIGURE 5
H-Subclasses in Ebook Central, Use-Months from January 2012 through December 2015

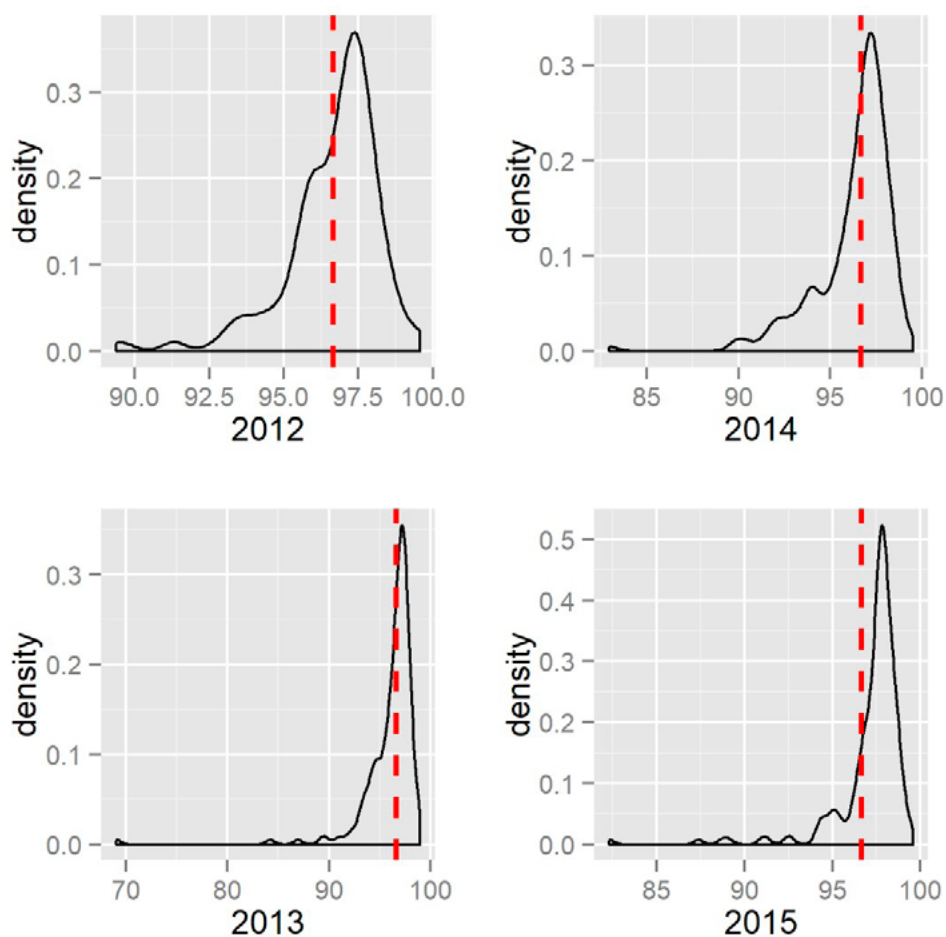


Of course, we don't usually look at use data during a four-year period; instead, we generally collect our statistics annually. By subjecting the annual data from the four years in this study, from 2012 through 2015, to the same transformation as above, and excluding subclasses with fewer than 100 successful section requests per year, we see a similarly consistent relationship between the intensive and extensive uses. During the full course of the study, the mean reduction from successful section requests to number of use months per title is 96.44 percent, the median 97 percent. The distribution is very tight, with a coefficient of variation of only 2.18 percent. For 2015, the mean is 97.19 percent, the median 97.7 percent, and the coefficient of variation 2.06 percent. Other years in the study are remarkably similar (see table 1 and figure 6).

TABLE 1
The Percentage Reduction from Successful Section Request Data to Use-Month, by Year and during the Life of the Study

Percentage Reduction by Year and Overall															
	2012			2013			2014			2015			2012–2015		
	Mean	Median	CV	Mean	Median	CV	Mean	Median	CV	Mean	Median	CV	Mean	Median	CV
Percentage Reduction	96.66%	97.03%	1.64%	96.08%	96.85%	2.92%	96.29%	96.90%	2.16%	97.19%	97.70%	2.06%	96.44%	97.00%	2.18%

FIGURE 6
Distribution Curves by Year of the Percentage Reduction from Successful Section Requests to Use-Months, with the Mean Indicated by the Red Dotted Line (The Mean Value is Given in Table 1)



Discussion

As noted above, it was expected that subjecting the raw COUNTER BR2 data to binary reduction would offer a much different view of the data, correcting for inflation due to single users consulting the same book repeatedly, double-dipping due to users printing or downloading the same pages they have already read, or a book assigned for a class, all of which can be expected to drive up the use numbers each month. LC class H includes titles such as *Multiple Regression with Discrete Dependent Variables*, with 322 successful section requests in the raw data, of which 296 were recorded in a single month. If we count those 296 as one for that month and do the same for books across the LC class, it can reasonably be expected that the transformed data will display a different pattern of use. That was the initial expectation in this study and, in fact, what led to this approach: the expectation of some more enlightening view of the data, stripped of the inflationary uses. However, contrary to the initial hypothesis, the data presented here suggest that even the most granular reporting of successful section request data does not introduce worrisome inflation of use and that the heavily intensive use of certain titles for classes (which, presumably, drives up the number of successful section requests in a given month) is not extreme enough to skew the relative extensive use when viewed in aggregate (by, for instance, LC subclass).

The results of this method on this dataset suggest that there is a stable relationship between intensive and extensive use when looking at periods of a year or more with monthly data. It can be expected that the variance will increase with smaller datasets and shorter time periods. While most single months in this study show similar patterns between raw data and the transformed data, with percentage share at the major class level varying by no more than a handful of percentage points, there are also months like October 2015. Looking at the raw data from that month, we see the typical pattern of LC class H accounting for the largest number of successful section requests (95,800) and the highest percentage share overall at 24 percent. The transformed data, which at this level show us how many unique titles were subject to successful section requests during October 2015, tell a different story. In the raw data for October 2015, there are 32,100 successful section requests for items classed in R, equivalent to only 34 percent of the number of successful section requests in class H. In the transformed data, however, class R accounted for 2,120 use-months across the collection, a 67 percent increase over H during that time. That is, in October 2015, subclass R saw more extensive use than did subclass H, while undergoing significantly less intensive use.

At the title level, the stable ratio of successful section requests to use-months seen elsewhere no longer holds, even during the course of 48 months. The maximum number of uses in a reduced dataset at the title level is the number of sampling periods, in this case 48 months. In the raw data, *New Jim Crow: Mass Incarceration in the Age of Colorblindness* is the title with the tenth highest number of successful section requests with 23,860. At least one successful section request was reported for 27 of those 48 months, which results in a reduced rank of 225th. Rather than casting doubt on this approach, such a disparity at the item level only reflects the different demands of such granular assessment. Even at the title level, there is a certain threshold after which it makes no difference how many times an item was used. If we could say with certainty that a print book was used in the library, with or without being checked out and without knowing how many pages were read or what sorts of roles the book played in research, once a month, we would be hard-pressed to argue for deaccessioning that title. Put plainly, knowing that a title was used, say, at least 27 times in 48 months tells us just as

much from a collection assessment perspective as does knowing that the same title was subject to 23,860 successful section requests.

Of course, we do not suddenly see use clearly by reducing the data as done here. Rather, by weighing all uses equally and establishing a rough but stable ratio of successful section requests to use-months during a sufficiently long period of time (minimum one year), and thereby being free to reduce them to a binary used/not-used scheme, we can be confident that the original data are not skewed as feared. What is interesting about the transformed data is that, by and large, it gives us the same view of relative use that we see when looking at the raw data aggregated in the same ways (such as by month and LC subclass). This is particularly true when we sample monthly, as is the standard practice, and view over a longer scale, such as annually or longer. It is to be expected that the exact ratio of successful section requests to use months will vary by collection, institution, and the definition of a section in BR2.

Conclusion

This study began with a misconception, namely that the page-level data reporting of the BR2 for Ebook Central provided a misleading view of collections use. That misconception was related to the common belief that certain uses of library collections are better than others and particularly that a certain level of user commitment by borrowing or downloading an item was necessary before the use became worthy of recording. It is argued that this belief is bolstered by the historical standard of circulation as the exemplary use of library materials, which itself is tied to another misconception that users typically read all or most of the books borrowed and, further, that ebooks are used more for quick consultation, while users read print books more thoroughly. As Slater puts it, "The belief that checking out or purchasing a print book indicates a user will read all (or even most) of it is not supported by the research, since academic users engaged in a well-established pattern of reading only a small percentage of print books and journals they consulted."²² As Levine-Clark et al. show, some users spend extended time on a few pages, while others view many pages quickly, indicating again the variety of ways our users use our collections.²³

The binary reduction method presented here is not meant to supplant the existing COUNTER approach. Rather, by uncovering a stable ratio of successful section requests (intensive use) to use-months (extensive use) over four years in a large research library (expressed here as a percentage reduction), this study suggests that the intensive and extensive uses are simply different views of ebook use, telling the collections librarian the same thing. There is no wild inflation of use, even when the section requests are reported at the page level, once the data are aggregated rather than approached on a title-by-title basis. Further study is warranted, not only of the relationship between successful section requests and use-months, but of the different metrics used to record collections use to discover if there are other similarly stable relationships. Whether there are or are not, a deeper and more nuanced understanding of the data and, more important, how our users use our collections, will help all librarians involved in collections provide the materials our users need in the way they need them.

Notes

1. The author would like to thank Cecilia Feilla, Laura McCann, and Jennifer Meyer-Stearns for their feedback at various points during the writing of this article.

2. According to the Project COUNTER website, "COUNTER is a non-profit organization supported by a global community of library, publisher and vendor members, who contribute to the development of the Code

of Practice through working groups and outreach." "About COUNTER," Project Counter, <https://www.project-counter.org/about/> [accessed 21 April 2019]. The Code of Practice for COUNTER version 5 outlines a number of different types of reports, including platform report (PR), database reports (DR), title reports (TR), and item reports (IR). Some vendors have yet to issue their reports in COUNTER 5 format, so any longitudinal study like the one presented here will likely rely on COUNTER 4 reports for some time. Relevant for this article is the Book Report 2 (BR2) from the previous standard, COUNTER 4.

3. At the time the data were collected, Ebook Central was branded "Ebrary." The name Ebook Central is used throughout to avoid confusion.

4. See, for example, Karen Kohn, "Using Logistic Regression to Examine Multiple Factors Related to E-Book Use," *Library Resources & Technical Services* 62, no. 2 (April 2018): 56, <https://doi.org/10.5860/lrts.62n2.54>, and Steven A. Knowlton, "A Two-Step Model for Assessing Relative Interest in E-Books Compared to Print," *College & Research Libraries* 77, no. 1 (January 2016): 21, <https://doi.org/10.5860/crl.77.1.20> for illuminating discussions of the shortcomings of print usage data.

5. There have been some attempts to measure in-house use, such as Lorraine M. Lane, "The Relationship between Loans and In-House Use of Books in Determining a Use-Factor for Budget Allocation," *Library Acquisitions: Practice & Theory* 11, no. 2 (1987): 95–102, [https://doi.org/10.1016/0364-6408\(87\)90046-9](https://doi.org/10.1016/0364-6408(87)90046-9); C. Puvogel, "Stack Attack! In-House Book Usage in a Small College Environment," *College and Undergraduate Libraries* 8, no. 2 (1998): 11–22; Victoria H. Wagner, "Phantom Use: Quantifying in-Library Browsing of Circulating Materials," *Journal of Access Services* 5, no. 1 (2007): 173–79; Michael Hughes, "A Long-Term Study of Collection Use Based on Detailed Library of Congress Classification, a Statistical Tool for Collection Management Decisions," *Collection Management* 41, no. 3 (July 2, 2016): 152–67, <https://doi.org/10.1080/01462679.2016.1169964>.

6. Robert Slater, "Why Aren't E-Books Gaining More Ground in Academic Libraries? E-Book Use and Perceptions: A Review of Published Literature and Research," *Journal of Web Librarianship* 4, no. 4 (2010): 307, <https://doi.org/10.1080/19322909.2010.525419>.

7. George S. Bonn, "Evaluation of the Collection," *Library Trends* 22 (January 1974): 273.

8. Terry R. Mills, "The University of Illinois Film Center Collection Use Study," June 1982, www.eric.ed.gov/ERICWebPortal/contentdelivery/servlet/ERICServlet?accno=ED227821.

9. Michael Levine-Clark, Kari Paulson, and Paul Moeller, "10,000 Libraries, 4 Years: A Large-Scale Study of E-Book Usage and How You Can Use the Data to Move Forward," *Serials Librarian* 68, no. 1/4 (May 19, 2015): 262–68, <https://doi.org/10.1080/0361526X.2015.1017709>.

10. Cathy Goodwin, "The E-Duke Scholarly Collection: E-Book v. Print Use," *Collection Building* 33, no. 4 (November 2014): 103.

11. Levine-Clark, Paulson, and Moeller, "10,000 Libraries, 4 Years," 266.

12. Karin Bystrom, "Everything That's Wrong with E-Book Statistics: A Comparison of E-Book Packages," in *Accentuate the Positive: Charleston Conference 2012* (Charleston Conference, Purdue University Press, 2013), <https://doi.org/10.5703/1288284315105>.

13. Jonathan H. Harwell and Erin Gallagher, "The Secret Lives of Ebooks: A Paratextual Analysis Illuminates a Veil of Usage Statistics," *Convergence* (January 9, 2018), <https://doi.org/10.1177/1354856517751379>.

14. Tasha Mellins-Cohen, "The Friendly Guide to Release 5 for Librarians" (COUNTER, 2018), 20, https://www.projectcounter.org/wp-content/uploads/2018/03/Release5_Librarians_PDFX_20180307.pdf.

15. Kohn, "Using Logistic Regression to Examine Multiple Factors Related to E-Book Use."

16. Justin Littman and Lynn Silipigni Connaway, "A Circulation Analysis of Print Books and E-Books in an Academic Research Library," *Library Resources & Technical Services* 48, no. 4 (October 2004): 256–62.

17. Knowlton, "A Two-Step Model for Assessing Relative Interest in E-Books Compared to Print."

18. Kohn, "Using Logistic Regression to Examine Multiple Factors Related to E-Book Use," 56.

19. Joseph M. Hilbe, *Negative Binomial Regression* (Cambridge, UK: Cambridge University Press, 2011), 395, <http://ebookcentral.proquest.com/lib/brooklyn-ebooks/detail.action?docID=667619>.

20. James E. Alt, Gary King, and Curtis S. Signorino, "Aggregation among Binary, Count, and Duration Models: Estimating the Same Quantities from Different Levels of Data," *Political Analysis* 9, no. 1 (2001): 22.

21. For an extensive discussion of this method and the scripts used, see Michael Hughes, "Assessing the Collection through Use Data: An Automated Collection Assessment Tool," *Collection Management* 37, no. 2 (2012): 110–26, <https://doi.org/10.1080/01462679.2012.653777>.

22. Slater, "Why Aren't E-Books Gaining More Ground in Academic Libraries?" 308–09.

23. Levine-Clark, Paulson, and Moeller, "10,000 Libraries, 4 Years," 266.