Taking a Fresh Look

Reviewing and Classifying Reference Statistics for Data-Driven Decision Making

Sarah LeMire, Lorelei Rutledge, and Amy Brunvand

Sarah LeMire (slemire@library
.tamu.edu) is First Year Experience
and Outreach Librarian, Texas A&M
University Libraries. Lorelei Rutledge
(Lorelei.Rutledge@utah.edu)
is Research and User Services
Librarian, University of Utah.
Amy Brunvand (Amy.Brunvand@
utah.edu) is Digital Scholarship
Lab Librarian, University of Utah.

This article describes the results of an extensive review of reference transactions from multiple service points at the University of Utah's J. Willard Marriott Library. The review enabled us to better understand the types of questions asked at our service points and resulted in a new set of codes for categorizing reference transactions that focus on recording the kinds of expertise needed to answer each question. We describe the differences between our model and other scales for collecting reference questions. Our method for reviewing reference transactions and developing new codes may be useful to other libraries interested in updating how they collect reference statistics.

In this paper we develop a strategy to evaluate public service points in an academic library based on the expertise sought by library patrons.[1] Although there is already an enormous body of literature about reference desk statistics, *SPEC Kit 268: Reference Service Statistics and Assessment* identified deep dissatisfaction with the current practice of reference statistics as a tool for evaluation and assessment.[2] Part of this dissatisfaction derives from the failure of conventional statistics to assess changing service models. Recent trends have resulted in major changes: Many libraries have eliminated subject-specific reference desks and adopted "one-stop-shopping" service desks in spaces rebranded as "Information Commons" or "Knowledge Commons." Some libraries have taken librarians off point-of-need service points altogether in favor of offering office hours and research consultations. At the same time, library public services have expanded into virtual space with synchronous "chat" and asynchronous email service points. These physical and virtual public service desks are staffed by some combination of professional librarians, IT personnel, paraprofessionals and part-time staff available to respond to ever-changing patron needs. Given these changes in reference desks, the growth of virtual reference services, and changes in patron needs, reference work is changing as well. Patrons now require reference support in a variety of formats such as face to face, chat, and SMS, in multiple places both on and off campus, and on topics ranging from research to technology troubleshooting. Collecting and employing useful reference statistics for data-driven decision making is more important than ever.

In fall 2013, a team of librarians at the University of Utah's J. Willard Marriott Library embarked upon a project to gather, code, and analyze statistics

from the library's in-person, chat, and email reference services. Although it is common for libraries to evaluate their various reference services, we wanted to try something a little different—to evaluate reference desk, email, and chat reference transactions as three cohesive elements of a single library reference service. To evaluate the library's reference service holistically, we decided to collect the questions we received at all reference service points, code them using grounded theory, and then compare them. We based our initial idea on several studies in which researchers performed grounded theory coding of reference transactions.[3] This project was intended to give us a better understanding of why patrons accessed the library's reference service, and how reference desk, chat, and email reference service points worked in concert, in order to enable us to make data-driven decisions when allocating staff and resources to meet patron needs.

LITERATURE REVIEW

Why Not Use Existing Scales?

There is already an enormous body of literature about reference statistics, including a number of efforts to develop objective scales for reference service assessment. So why develop a new scale? First, many existing scales fall into a trap of pre-assigning value to different question types. These scales privilege a certain type of in-depth, subject-oriented reference question as being the most valuable since they require the expertise of highly trained professionals. However, even "simple" question types can give patrons valuable help and can turn into complex information searches. Secondly, as libraries diversify to offer services such as open-access publishing, maker spaces, technology support, digital scholarship, and other innovative services, service desks may be asked to provide support in ways that are not easily represented in traditional reference desk statistics scales. We considered three commonly-used scales discussed below.

READ Scale

The Reference Effort Assessment Data (READ) scale classifies questions in terms of difficulty. It focuses on "recording vital supplemental qualitative statistics gathered when reference librarians assist users with their inquiries or research-related activities by placing an emphasis on recording the skills, knowledge, techniques, and tools utilized by the librarian during a reference transaction."[4] The READ scale assigns each question a number between one and six based on difficulty as defined by the library staff person. Questions assigned a rating of one "require the least amount of effort and no specialized knowledge, skills or expertise" and generally no consultation of resources.[5] Questions assigned a score of six require in-depth consultation of resources and a great deal of time. The difficulty with this scale is that, although it does take into account the tools and skills necessary to answer questions, it is not necessarily clear what the difference is, for example, between a three and a four. For our purposes, we wanted to simplify our system to reduce the number of decisions our staff needed to make about how to categorize a question. In addition, assigning questions by difficulty is complicated at desks where staff with multiple kinds of expertise answer the same questions, since a question that requires no effort for one might be difficult for another.

Warner Scale

Debra Warner also created a reference collection system. Warner examined how the Eastern Carolina University Health Sciences Library combined its circulation and reference desks and then updated its reference transaction tracking system in order to better identify which questions could be answered by a library technician and which needed to be passed on to a librarian.[6] The Warner scale codes questions into four levels. At the first level are the questions typically referred to as directional, or questions that do not require resources to answer. The second level requires demonstration of a task or skill, while the third level encompasses questions that require a specific use of resources and search strategy. The fourth level is reserved for questions where "the librarian will often have to research recommendations or prepare reports for consultation work."[7] However, this scale has clear problems for our library's reference service model. For instance, by presupposing that all directional questions are easy, this scale obscures times when such questions might become complex. Also, a complex question may be "easy" because the librarian has the knowledge or skills to answer the question, not because the question is inherently easy. Thus, using this system might present falsely most questions at the desk as easy, even when they are not. Second, rating questions by level of difficulty obscures the type of expertise needed for each kind of question.

Katz Scale

Katz offers yet another scale for analyzing types of questions, dividing them into Direction, Ready Reference, Specific-Search, and Research questions.[8] Katz notes that "most [research questions] involve trial-and-error searching or browsing, primarily because (a) the average researcher may have a vague notion of the question but usually cannot be specific; (b) the answer to the yet-to-be-completely formulated question depends on what the researcher is able to find."[9] In contrast, specific-search questions involve locating existing resources. In practice, however, this model did not fit our service point model because we are answering many other kinds of questions, such as technology questions, that are not addressed in this scheme.

Each of these scales attempts to describe query types in terms of difficulty. However, the information they record about the perceived difficulty of each question, whether

based on resource use or the other factors, fails to account for different levels of expertise that would lead to different ratings of the same question by different staff. We also felt that relying on such a difficulty scale would lead us to dismiss the importance of "easy" query types because they do not necessarily require the expertise of a professional librarian. For example, Ryan argued that because most questions are easy, it would be more cost effective to staff the desk with generalists rather than those with a high level of expertise.[10] However, a follow-up study found that reference transactions had significantly declined at a desk with no librarian.[11] Bishop and Bartlett conducted a reference transaction analysis designed to inform staffing decisions at multiple points in the UK Libraries and found that 83.7 percent of the questions were location-based and could be answered by staff rather than librarians.[12] They also explained, however, that these simple questions have a tendency to become more complex and that "training helps staff clarify a user's question and reduces the likelihood of providing inappropriate information in response to the user's original, often ambiguous query."[13] These two case studies demonstrated the benefits of having highly trained staff available, and also encouraged our decision to focus our statistics collection on category types and time spent rather than difficulty ratings.

Marriott Library Service Points

The Marriott Library offers three general-purpose reference service points: a physical location, a chat system, and an email system. The physical Knowledge Commons is a shared service point comprised of the Knowledge Commons Desk as well as the adjacent Student Computing Services (SCS) desk. This service point is staffed up to 111 hours per week by more than seventy staff members, including librarians, staff, and student workers (table 1). Patrons are invited to ask at the Knowledge Commons for everything from releasing print jobs and circulating cables or headphones to in-depth research and technology questions. All in-person transactions, which include telephone transactions, are recorded in the commercial reference statistics system DeskStats. Because the Knowledge Commons was originally conceived as a single service point and many staff members work at both the Knowledge Commons desk and the SCS desk, our DeskStats configuration doesn't distinguish between the two service desks located in the Knowledge Commons space. Many transactions can be completed at either desk, with the one notable exception being the circulation of materials, which is only available at the SCS Desk. Staff members record each statistic in a category based upon question type and duration of the transaction.

Online information service is provided via the commercial system Kayako. Librarians manage and respond to most email reference questions, which are automatically recorded in Kayako. Online reference statistics are not separately entered into DeskStats to avoid having staff enter additional statistics, especially since machine-generated statistics are

Table 1. Fall 2014 Knowledge Commons Staffing

Staffing Type	Weekly Hours Available
Student Workers	111
Staff	13
Librarians	51

more reliable than self-reported statistics.[14] Chat reference is provided by librarians during normal business hours and is supplemented by SCS employees in early morning and evening hours, with a combined total of more than thirty library employees providing chat reference support on a weekly basis. Chat reference transactions are also automatically recorded in Kayako and are not separately entered into DeskStats.

METHOD

To evaluate how the Marriott Library's reference service was being used by patrons, we collected and analyzed data from each service point. We collected self-reported, in-person reference statistics from the Knowledge Commons service point as well as automatically-generated statistics from Kayako, the software used for chat and email reference.

The Data Sets

To create profiles to compare the function of the three service points, we required a sufficiently large sample of questions to ensure that even fairly rare question types were well-represented. Because each of the three service points has very different levels of activity, we were not able to use data sets with exactly the same time parameters and instead selected samples for online services that approximate the volume of one week at the Knowledge Commons Desk. The Knowledge Commons data set contained 1,766 reference queries recorded at the Knowledge Commons one week of the Fall 2013 semester and one week of the Spring 2014 semester. During the two mid-semester sample weeks, November 19–25, 2013, and March 2–8, 2014, all service desk staff were asked to record descriptive comments along with their statistics. Each statistic was entered into the DeskStats system according to both query type and time spent. Staff also had the option to select the category "Other" if they were not sure how to categorize a query. We used this nonstatistical sample to gather a convenience sample, relying on the assumption that the items collected during those weeks were similar in type and quantity to questions received throughout the rest of the year.

Knowledge Commons reference statistics are self-reported, with the accepted limitation that self-report statistics can be inaccurate and are generally undercounted.[15] Still, the Knowledge Commons data set offers broad-based evaluation of question types and difficulty by staff with many

Table 2. Code Book Categories

Library Information and Policy

Circulation/Borrowing/Reserves

Research and Reference

Locate Materials

SFX/EZProxy/Off-Campus Access

Technology

Print/Scan/Copy/Duplication

Feedback

Other

levels of expertise and training regarding the types of queries they handled and how much time they thought they spent answering these queries. The data set also offers a useful estimate of the proportion of each query type. Thus, even recognizing that the data are not a strictly accurate count of all queries received, using the Knowledge Commons data set still enables us to make predictions about what expertise patrons seek at the Knowledge Commons.

The chat and email reference data sets represent questions received via an "Ask the Library" link on the library website. The chat reference data set contains 673 questions received between October 1, 2013 and March 31, 2014. Although Kayako records entire chat transactions, we opted to code the questions based upon the initial query entered by the patron into the chat reference system. Similarly, the email data set contains 1,187 questions received between January 11, 2013, and May 12, 2014, and we coded email reference data based on the initial email question rather than the complete transaction correspondence.

Coding Process

Our first step was to look closely at what kinds of questions we received. We extracted a sample from the data set and, using grounded theory, worked individually to assign codes to describe what type of help patrons were seeking. Then, we compared the themes we found and drafted an initial code book. We then divided our data sets so that two of us looked at and coded each reference transaction and we used those findings to refine the code book. To get a stronger reliability, we then identified the items where our codes disagreed and discussed the codes until we came to a consensus, using this to refine our code definitions.

The Final Code Book

Based on this process, we developed a code book (table 2) to analyze service desk transactions. The code book consisted of nine broad categories that reflected the most common types of questions answered across our reference service. We used this final code book to completely code our three data sets and to generate a table that illustrates the number of questions per category received at each of the three service points (table 3). This data was then formatted into a pie-chart that illustrates the proportion of question types asked via our three reference service points (figure 1).

FINDINGS

Knowledge Commons Pattern Obscured by Low-Complexity Transactions

We anticipated that our three methods of providing reference service, in-person, chat, and email, would have different transaction patterns. Because the Knowledge Commons is a shared service point designed to serve as a sort of one-stop-shopping experience for many patrons, Knowledge Commons staff must answer a wide variety of questions, including many questions that are not reference questions. The Knowledge Commons staff answers high volumes of Print/Scan/Copy/Duplication, Circulation/Borrowing/Reserves, and Library Information and Policy questions. Some of these question types are high-volume because they require intermediation—for example, patrons cannot check out a cable or

Table 3. Questions per Category Received at Service Points

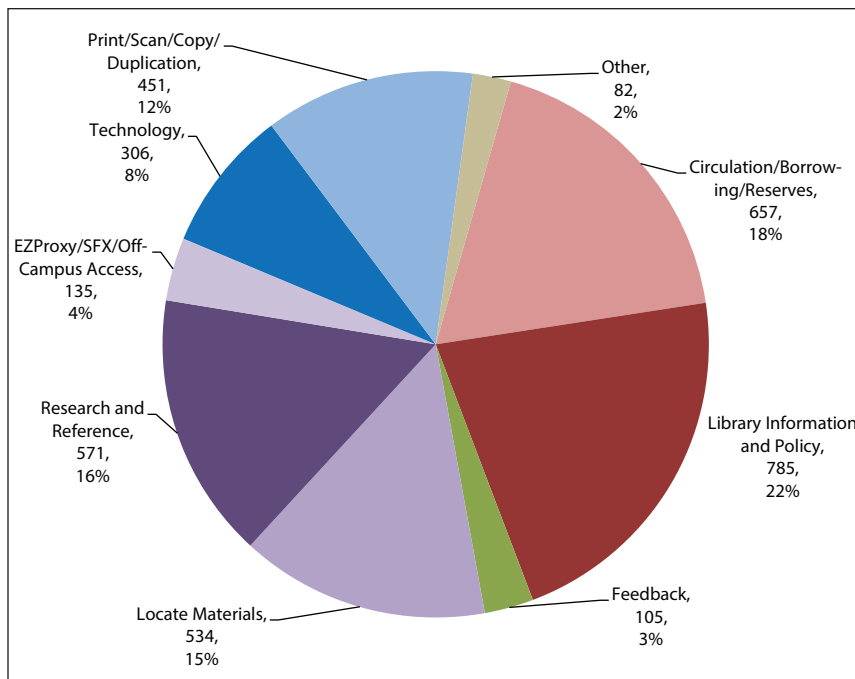| Category	Knowledge Commons	Email Reference	Chat Reference
Circulation/Borrowing/Reserves	479	112	66
Library Information and Policy	318	291	176
Print/Scan/Copy/Duplication	432	7	12
Feedback	8	86	11
Locate Materials	132	250	152
Technology	213	53	40
Other	29	15	38
EZProxy/SFX/Off-Campus Access	6	75	54
Research and Reference	149	298	124
Totals:	1,766	1,187	673

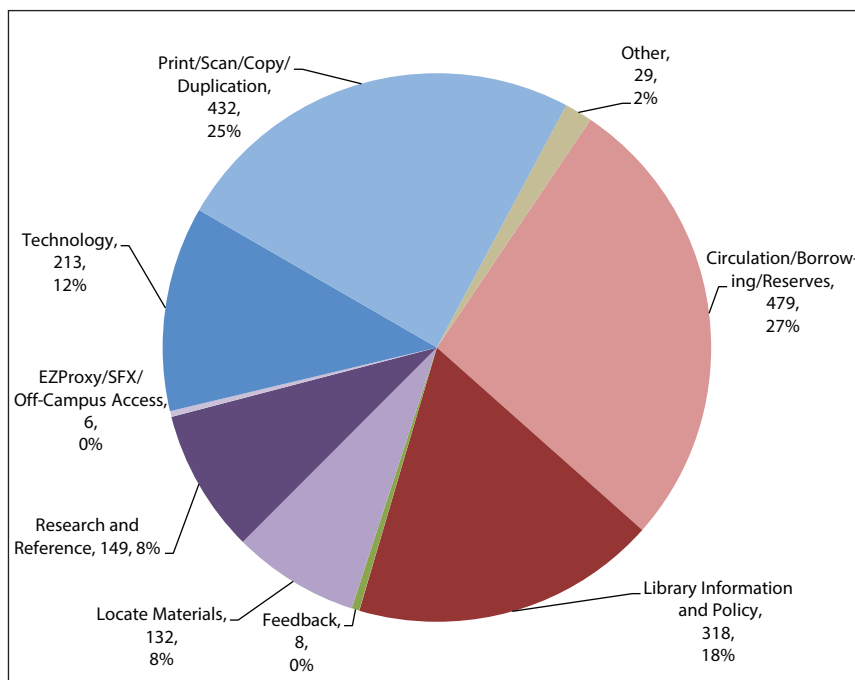**Figure 1.** Combined Chat, Email, and Knowledge Commons Reference Statistics by Code

Figure 2. Knowledge Commons Reference Statistics

Print/Scan/Copy/Duplication questions (which contain many print release requests) and Circulation/Borrowing/Reserves questions (which contain many check-in/check-out requests) were removed from the picture, the breakdown of questions at the Knowledge Commons desk closely resembled the breakdown of queries that we received via email and chat reference (figure 3).

Knowledge Commons Queries Based On Time Spent

While all questions are important to the person who asks, the Knowledge Commons data, which include both question categories and approximate duration, remind us that not all questions are equal. Figure 4 represents the amount of time typically required to answer each category of question. In the Knowledge Commons, Circulation/Borrowing/Reserves questions take less than one minute 86 percent of the time. Assuming, as is established in the literature,[16] that question duration serves as a reasonable proxy for question complexity, our data indicates that circulation-related questions are the least complex type answered in the Knowledge Commons, followed by Library Information and Policy questions and Print/Scan/Copy/Duplication questions. These three question groups make up 71 percent of the questions answered at the Knowledge Commons, which indicates that a considerable volume of question types fielded by Knowledge Commons staff are usually not complex and could be reasonably answered by student workers and staff rather than librarians.

However, some groups of questions tended to be more complex than we anticipated. We expected Locate Materials questions, which are known-item searches, to be a relatively low-skill question, but we discovered that Locate Materials questions take longer than one minute to answer more than 68 percent of the time, suggesting that known-item

a set of headphones without assistance from a staff member. When examining our data set, we discovered that the high volume of low-complexity questions in the Knowledge Commons obscured the fact that a significant number of complex Research and Reference questions were still being asked at the Knowledge Commons desk (figure 2). Indeed, when

searches are frequently more complex than we anticipated. Our data also showed us that 26 percent of Research and Reference questions require more than 15 minutes to answer and that Research and Reference questions are answered in less than one minute only 8 percent of the time. This data confirms our understanding that Research and Reference
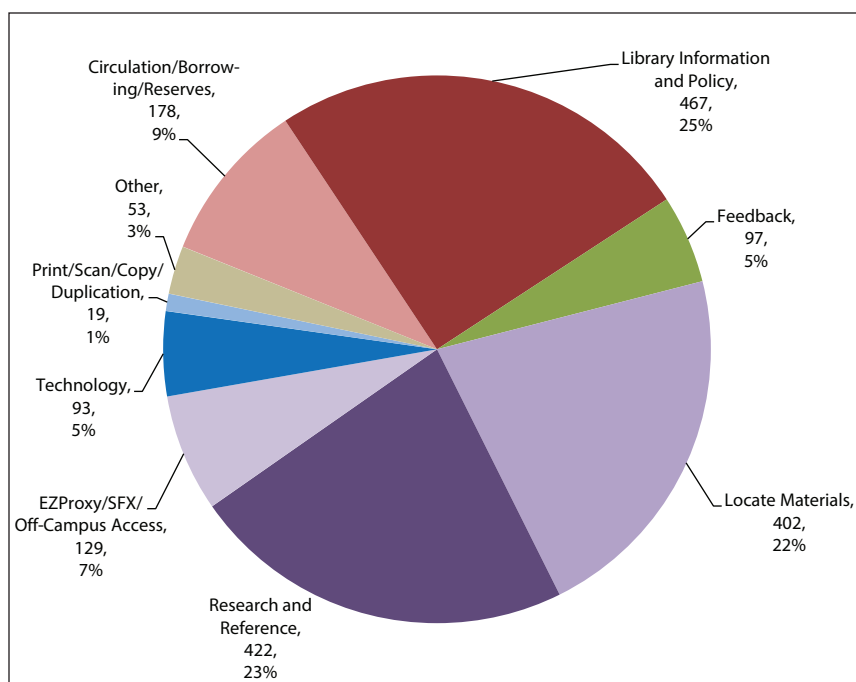
Figure 3. Chat and Email Reference Statistics

tracking how much intervention our EZ-Proxy system requires.

Role of Online Reference in Technology Troubleshooting

While the distinction between Remote Access and EZProxy/SFX/Off-Campus Access is an important one for Knowledge Commons staff to comprehend, we also learned that the majority of the EZProxy/SFX/Off-Campus Access questions are received via chat or email rather than at the in-person service desk. Off-Campus Access questions compose 6 percent of the library's email reference questions and 8 percent of the chat reference questions, suggesting that patrons are running into difficulty with our databases while off-campus and are reaching out for help at the point of need. These results inform us that online reference has a valuable role to play in troubleshooting problems with off-campus access to the library's digital materials. Our in-person reference questions indicated that Off-Campus Access questions were possibly one of the most complex types of questions, requiring more than fifteen minutes 33 percent of the time. Assuming that Off-Campus Access questions retain the same level of complexity when answered via chat or email, this would suggest that there is a real need for higher levels of expertise via chat and email in order to help patrons successfully navigate our EZProxy and SFX systems while off-campus.

questions, which make up 9 percent of the questions answered in the Knowledge Commons, are likely to be complex and require a higher level of skill and training to answer. This suggests that, while the Knowledge Commons answers many more low-complexity questions than Research and Reference ones, patrons still approach this location with Research and Reference questions that are much more likely to be complex and to require a higher level of skill and expertise. This finding demonstrates the need to have highly skilled staff readily available in the Knowledge Commons to answer these types of questions.

Off-Campus Access versus Remote Access

Our current reference statistics system, DeskStats, includes a category entitled "Remote Access/Database/eJournal help," which contains any question relating to patrons accessing services from outside the library. During the grounded theory process, we recognized that our reference statistics conflated two separate types of patron inquiries—those requesting help using software remotely (termed Remote Access) and patrons requesting help accessing articles, journals, and databases through our Off-Campus Access system. While the distinction between these questions is opaque to patrons, library staff needs to recognize the difference because Remote Access queries are essentially a Technology question that is typically answered by an IT staff member, while Off-Campus Access-related questions are EZProxy and SFX-related questions which are routed through Collection Development and our Electronic Resources Manager. Because these two categories are conflated, we have not been

Questions Don't Always Fit Neatly in One Category

This process also taught us much about how we should collect and analyze data. As we refined our codes, we realized that, as with all qualitative data analysis, the ways in which we coded the items were subjective. Although we found, for the most part, that we assigned our codes consistently, there were multiple interactions that we coded differently. We opted to sit down as a group and attempt to reconcile these disparities. While we discovered that most of the differences in coding were simple user errors (e.g., accidentally marking Printing for a Technology question), we also found that there was a small number of records that we could not assign to a single coded category (table 4).

We evaluated these records against our coding scheme in an attempt to determine whether or not there were significant gaps, but upon closer examination we determined that these items were evidence of the multifaceted, multi-step questions that are common at our shared service point. For example, a patron might ask for help finding research on
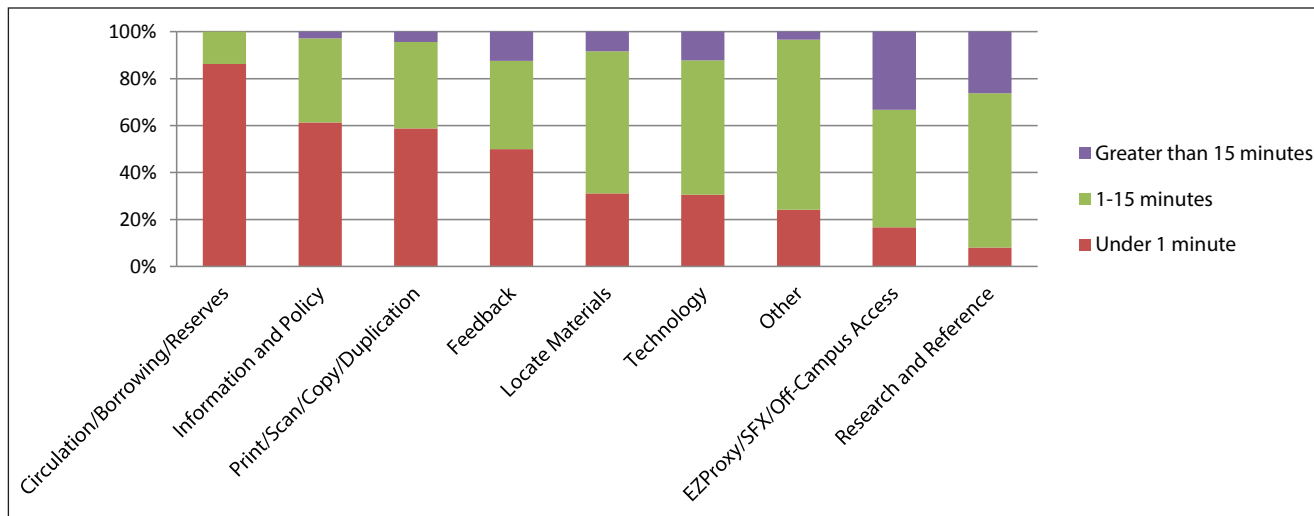
Figure 4. Knowledge Commons Desk Question Complexity by Category

a topic, finding the appropriate place to pick up requested materials, and printing online materials in a single reference desk transaction. In DeskStats, the service desk staff are asked to select the category that best fits that patron's question, leaving secondary elements of the patron's question uncaptured. As we identified records that reflected the complexity of patron questions at a shared service point, we began to question whether a system that requires staff to select a single category was obscuring some of our picture of what is happening at the service point.

THE VALUE OF MIXED-METHOD ANALYSIS

The mixed-method approach we took to this study, which combined qualitative and quantitative analysis, proved to be a very effective way of evaluating our service points. For example, using a quantitative approach to examine the number of questions we received yielded specific kinds of insights, such as the realization that a large segment of our questions were about printing. This suggested to us that we needed to examine our printing procedures for friction points and usage barriers. On the other hand, this quantitative approach treated all questions as equal, which was not always helpful since some question types are more complicated and require more expertise to answer. In this case, coding our data thematically taught us more about how our service points actually functioned. For instance, analyzing the comments we received helped us understand that we were answering multiple queries per transaction and that the kinds of questions we answered were not adequately reflected in our own coding scheme. In addition, considering themes enabled us to see some types of questions that have a tendency to become complex, a trend we would have missed had we looked only at quantitative data. By incorporating both qualitative and quantitative analyses into our research, we were able to

Table 4. Number of Questions with Multiple Codes

Service Point	Number of Questions	Questions with Multiple Codes
Knowledge Commons	1,766	11
Email Reference	1,187	14
Chat Reference	674	8

gain a holistic view of the way our service points are being accessed and a better understanding of how we can improve both service and efficiency.

RECOMMENDATIONS AND FUTURE RESEARCH

This study served to provide us with a number of important insights into our current library reference service, however, it also unveiled additional questions that could be explored in future research. We would like to look more closely at questions to which we assigned multiple codes as these questions demonstrate areas that require multiple types of expertise. At a combined service desk, this may suggest a need for increased cross-training for staff to develop expertise in multiple domains or for the colocation of staff from multiple service areas with expertise in each domain. By examining these kinds of questions in more detail, we would hope to learn which areas of expertise overlap most frequently.

We also hope to take a closer look at chat and email reference questions. Since our staff do not record time spent for these interactions, it is difficult to know whether we spend significantly longer on questions asked via chat and email than we do on questions asked in person. We assumed for this paper that the time spent was comparable,

but examining the data could support or disprove that assumption. Likewise, the number of questions we received via online services asking about locating online materials or troubleshooting access problems suggests the need for online reference providers to offer specific kinds of expertise. We also hope to explore our online reference questions in more detail in order to ensure that we identify the types of training our online staff require in order to develop the expertise to effectively answer these questions.

We would also like to test our codebook against data from other library service points not included in this study such as our Special Collections desk, the data from which was not collected in DeskStats. We look forward to having a wide variety of service points in our library experiment with our coding scheme and help to refine it. In this process, we hope to develop detailed profiles about the types of questions that each service area receives in order to learn more about the types of expertise needed to staff that area effectively.

Finally, our findings show a continued need for research into how to best categorize, record, and use reference statistics. Although many information professionals have evaluated and created methods for recording statistics, it is clear to us that we need to continue this evaluation as traditional reference desks continue to combine with other services and as we continue to see patrons asking complex technology-oriented questions at these service points. Libraries need to take a fresh look at how their service points are being used to find the best methods for capturing the different categories of questions they receive and to use their findings to make data driven decisions about important issues such as staff allocation. We hope that others will use and improve upon the coding scheme we developed.

CONCLUSIONS

We learned a great deal about the kinds of reference service we provide, the types of questions we answer, and the way we collect reference statistics. Both our virtual and physical reference points function as a combined service that requires mixed expertise; however, we answer the bulk of our printing questions in person, while our virtual reference plays a major role in supporting our patrons' access to online materials. Many of the inquiries we receive at the Knowledge Commons desk are requests for help with printing, which suggests that, although we think of our printing, scanning, and copying services as unmediated, they in fact require significant mediation. Because of the high volume of printing questions we receive at the Knowledge Commons desk, which require little time and expertise to answer, we have recommended that the Knowledge Commons leadership explore alternate printing solutions. One such solution could be designating the SCS desk, which is primarily staffed by student workers, as a print release station.

We also found that our virtual reference plays an important role in answering questions about how to access online materials. While only 0.3 percent of the questions answered at the Knowledge Commons desk are related to EZProxy/SFX/Off-Campus Access, such questions occur far more frequently at online service points. Based upon this information, our best practices for online services recommend that staff follow up after patron interactions by reporting broken links, problems authenticating into secured materials, and difficulty viewing online resources.

We also determined that we wanted to modify our criteria for collecting data to reflect our new coding scheme. We found that employees using our configuration of DeskStats had varying ideas of what kinds of questions belonged in each category, and many felt the number of categories in the system made it difficult to use. Some of the categories combined reference transactions that are reportable to the Association of Research Libraries (ARL) with transactions that did not fit the ARL definition of reference transactions, making the process of reporting more difficult. We also wondered if our current coding scheme obscured some types of reference transactions. Based upon these findings, we have been working with the library's Application Development department to build an in-house statistics gathering system that will allow us to record reference statistics according to the categories in our codebook. This system would incorporate several new features based upon our findings. For example, in the new system, Remote Access questions are folded into the Technology category, while Off-Campus Access/EZProxy/SFX questions are designated as a separate category. This will enable us to track questions based upon the training type required to answer each type of question. We also requested that the system allow us to track both individual transactions (i.e., number of patrons helped) and the actual number of questions asked. Should a patron ask questions that fall into multiple categories, we could record the different elements of a single question and ensure that we were capturing how complex some patrons' questions truly are. We plan to continue evaluating the development of the new system as well as the new coding scheme to ensure that they capture the varied and complex nature of questions received at the Knowledge Commons desk.

Our experience suggests that regularly and systematically reviewing reference statistics and how they are collected can be a valuable assessment strategy for any library. It would be particularly useful to do a reference question review if service desk staff report needing more training, which may indicate a disconnect between library assumptions and the types of questions that are actually being asked. Another good time to review reference questions would be prior to any major service reorganization. As our patrons' questions change, the kind of reference support we provide them will likewise need to change. By periodically reviewing reference statistics and coding those statistics using grounded theory, we can keep abreast of the types of questions we are answering in our libraries and ensure that we are collecting accurate and informative data. By ensuring that the data we collect reflects the ever-changing types of questions patrons

ask, we can use that information to make data-driven decisions about important issues that will enable us to better meet patron needs.

References

1. Michelene T. H. Chi, Robert Glaser, and Marshall J. Farr, *The Nature of Expertise* (Hillsdale, NJ: Lawrence Erlbaum Associates, 1988).
2. Eric Novotny, *SPEC Kit 268: Reference Service Statistics and Assessment* (Washington, DC: Association of Research Libraries, 2002).
3. Ellie Buckley, Kornelia Tancheva, and Xin Li, "Systematic Quantitative and Qualitative Reference Transaction Assessment: An Approach for Service Improvements," in *Proceedings of the 2008 Library Assessment Conference: Building Effective, Sustainable, Practical Assessment*, edited Steve Hiller et al. (Seattle, WA: Association of Research Libraries, 2008), 375–85; Amy Brunvand, "Ask the Expert: Using Expertise Domains for Library Service Assessment.," in *Proceedings of the 2010 Library Assessment Conference: Building Effective, Sustainable, Practical Assessment*, edited by Steve Hiller et al. (Baltimore, MD: Association of Research Libraries, 2010), 571–79; Gabrielle K. W. Wong, "Information Commons Help Desk Transactions Study," *Journal of Academic Librarianship* 36, no.3 (2010): 235–41.
4. Bella Karr Gerlich and Edward Whatley, "Using the READ Scale for Staffing Strategies: The Georgia College and State University Experience," *Library Leadership & Management* 23, no. 1 (2009): 26.
5. Ibid., 27.
6. Debra G. Warner, "A New Classification for Reference Statistics," *Reference & User Services Quarterly* 41 no. 1 (Fall 2001): 51–55.
7. Ibid., 53.
8. William A Katz, *Introduction to Reference Work. Vol. 1, Basic Information Services*, 8th ed. (Boston: McGraw Hill, 2002).
9. Ibid., 18.
10. Susan M. Ryan, "Reference Transactions Analysis: The Cost-Effectiveness of Staffing a Traditional Academic Reference Desk," *Journal of Academic Librarianship* 34, no. 5 (2008): 389–99.
11. Debbi Dinkins and Susan M. Ryan, "Measuring Referrals: The Use of Paraprofessionals at the Reference Desk," *Journal of Academic Librarianship* 36 no. 4 (2010): 279–86.
12. Bradley Wade Bishop and Jennifer A. Bartlett, "Where Do We Go from Here? Informing Academic Library Staffing through Reference Transaction Analysis," *College & Research Libraries* 74 no. 5 (2013): 494.
13. Ibid., 498.
14. Alison Graber et al., "Evaluating Reference Data Accuracy: A Mixed Methods Study," *Reference Services Review* 41, no. 2 (2013): 298–312; Sarah M. Philips, "The Search for Accuracy in Reference Desk Statistics," *Community & Junior College Libraries* 12, no. 3 (2004): 49–60.
15. Ibid.
16. Matthew L. Saxton and John V. Richardson Jr., *Understanding Reference Transactions: Transforming an Art into a Science* (San Diego, CA: Academic Press, 2002); Thomas Childers, Cynthia Lopata, and Brian Stafford, "Measuring the Difficulty of Reference Questions," *RQ* 31, no. 2 (Winter 1991): 237–43.