

# Learning to be a Better Q'er in Social Q&A Sites: Social Norms and Information Artifacts

June Ahn<sup>1</sup>, Brian S. Butler<sup>1</sup>, Cindy Weng<sup>2</sup>, Sarah Webster<sup>1</sup>

<sup>1</sup>University of Maryland, College Park  
4110 Hornbake Building South Wing  
College Park, MD 20742  
{juneahn, bsbutler, websters}@umd.edu

<sup>2</sup>Northwestern University  
2240 Campus Drive  
Evanston, IL 60208  
weng.cindy@u.northwestern.edu

## ABSTRACT

Social question and answer sites (SQAs) are increasingly popular knowledge sharing platforms. In this paper, we outline how an SQA site functions as a *social learning community*. The success of an SQA site depends not only on effectively organizing and delivering information, but also on whether it can provide the cues needed by community members to successfully learn to be productive contributors. We explore this learning process in four different SQAs that utilize the Stack Exchange platform: Science Fiction & Fantasy, Seasoned Advice (cooking), Database Administrators, and Android Enthusiasts. Using longitudinal fixed effects models, we examine whether users learn to be better question askers over time, and how user interface features and community norms affect the cultivation of this critical skill. The study offers design implications by highlighting factors that help users develop into literate and productive community contributors in social platforms.

## Keywords

Q&A; online communities; human information interaction.

## INTRODUCTION

Social question and answer sites (SQAs) are an increasingly important source of knowledge and information. These sites allow individuals to pose questions and receive answers from peers and other members. Numerous SQAs exist, differing with respect to the topics they cover, the types of dialogue they support, and the populations that they engage. Some platforms focus on opinions and discussion of social issues, while others focus on factual exchanges (Adamic, Zhang, Bakshy, & Ackerman, 2008). The popularity of SQAs has led to a growing body of research concerned with understanding Q&A behavior and improving the functioning of these systems (Gazan, 2011). Much of the prior work has focused on issues such as motivating

individuals to contribute, classifying questions and answers to facilitate better information transfer, and identifying experts and novices (Hanrahan, Convertino, & Nelson, 2012; Harper, Moy, & Konstan, 2009; Logie, Weinberg, Harper, & Konstan, 2011; Oktay, Taylor, & Jensen, 2010). An underlying thrust of this literature is to understand how to design and implement SQAs as vibrant information sharing communities.

In this paper, we seek to make two contributions. First, we advance an alternative analytical framework for understanding how SQAs function. Prior empirical studies of SQAs have implicitly worked from an information transfer perspective, focusing on how to efficiently classify and deliver content (questions and answers). Instead we draw from online communities research and sociocultural learning theory to argue for a social learning perspective on SQAs. Rather than focusing on the quality or management of information within a site, a social learning approach focuses on questions about how the affordances of successful SQA platforms help individuals transfer community norms (social information) and learn to be skilled community members.

Second, we present a longitudinal study of social learning in SQAs analyzing data from four communities which use the Stack Exchange platform: Science Fiction & Fantasy, Seasoned Advice (cooking), Android Enthusiasts, and Database Administrators. Using this data we examine the following research question:

- How do the various feedback affordances in these Stack Exchange sites correlate with members' learning to be better question askers over time?

The empirical findings show that, depending on a respective online community, feedback features such as past answers, comments, votes, and favorites influence the quality of future questions that members post. The results suggest that the features of SQA platforms provide implicit social feedback that help members learn and develop into more effective users of SQA sites, but that more design work is needed to better leverage social computing artifacts to cultivate members' question asking skills and abilities.

## RELATED WORK

There are several SQA sites that focus on knowledge creation to socialize members. One of the most prominent is Stack Exchange, which is a platform for creating and operating SQA sites. Originally developed to support Stack Overflow, a successful and influential SQA sites for programmers, Stack Exchange now supports nearly 100 SQA sites (Anderson, Huttenlocher, Kleinberg, & Leskovec, 2012; Hanrahan et al., 2012; Mamykina, Manoim, Mittal, Hripscak, & Hartmann, 2011). Each of the Stack Exchange sites utilizes the same underlying platform, but serves a unique population of experts and enthusiasts interested in different content areas, ranging from cooking to database administration.

Members of these communities engage in a group knowledge-creation process by asking and answering questions, editing the questions and answers posted by others, and voting on the quality of the resulting content. The end result is a collectively curated collection of information that is useful not only to the original questioner but as a reference for all site users, present and future, who are interested in the topic.

Responding to the growing popularity of SQA platforms such as Yahoo Answers and Stack Overflow, researchers have begun to examine different aspects of these online communities. Some researchers have focused on understanding the community-wide dynamics of SQAs. For example, prior research suggests that diverse topic areas in Yahoo Answers are related to user behavior (Adamic et al., 2008). Technical forums tend to have fewer replies but longer posts, while discussion based topics typically lead to more replies. Logie and colleagues found that Answerbag's members asked more social, subjective questions, Metafilter's community focused on more objective questions, and Yahoo Answers exhibited a wide variation in question types (Logie et al., 2011).

Mamykina and colleagues argue that the overwhelming success of Stack Overflow is related to features of the platform and how the community developed over time (Mamykina et al., 2011). Affordances such as productive competition through points and badges, and tight interaction with a core expert and developer group, drove productive Q&A norms on the site. Overall, prior work has found that there is significant variation in how well SQAs function, variation that is at least in part related to the social and technical features of the SQA platform.

Arising from roots in knowledge management and information retrieval, SQAs are naturally conceptualized as information sources. Many studies have focused on understanding how to improve the quality of information available and the process of information transfer in SQAs (Gazan, 2011; Rosenbaum & Shachaf, 2010).

Focusing on question processing, some prior work has explored the use of various techniques including human coding, statistical models, and machine learning to classify

the characteristics of questions and answers. Researchers have been able to classify questions in SQAs as conversational versus informational, or along rhetorical characteristics (Harper et al., 2009; Harper, Weinberg, Logie, & Konstan, 2010). Other studies have explored how to characterize difficult questions in Stack Exchange, in order to better direct those queries to expert users (Hanrahan et al., 2012).

Research into other aspects of SQA sites has also focused on direct facilitation of information provision. Work that examines user motivations, such as whether a person is participating for social conversation or for factual information seeking (Gardelli & Weber, 2012), helps better direct individuals to relevant content. Other studies are concerned with assessing SQA answer quality (Shah & Pomerantz, 2010) and accuracy (Fichman, 2011).

While there are different empirical emphases, prior work has largely focused on improving an SQA community's ability to provide high quality and accurate answers (Fichman, 2011; Shah & Pomerantz, 2010). An inherent assumption in the literature is a focus on information provision. If the goal is better information transfer, it is useful to characterize question types or answer quality to enhance the community's ability to process and respond to questions.

## A SOCIAL LEARNING PERSPECTIVE ON SQA DESIGN

Within the substantial literature on information behavior and information provision, there are two general approaches to facilitating high-quality information transfer. The first approach, most commonly seen in the information retrieval literature, focuses on facilitating information transfer by changing the system which processes queries and provides information. This approach is reflected in the information processing perspective taken by prior SQA research, with its emphasis on improving SQA sites ability to respond with relevant, accurate answers in a timely fashion.

The second approach, most commonly seen in discussions of information literacy and education, focuses on facilitating information transfer by enhancing the users' ability to construct queries and use the system effectively. Applying this perspective to SQAs, suggests that improving SQA effectiveness is likely to also be dependent on the ways that SQA participants learn *how to be* a member of the community. This focus on literacy implies that SQA platforms may be successful because they provide affordances, which support the transfer of social norms that enable members to become better question askers over time. In the remainder of this section, we draw from online communities research and sociocultural learning theory to elaborate the nature and implications of a social learning perspective on SQA platform design and evaluation.

SQAs are sociotechnical systems. In a sociotechnical system, the design of an environment constrains and shapes what kinds of behaviors people can enact. However, as

individuals use a system, their social behaviors and routines create new structures that shape community behavior (Rosenbaum & Shachaf, 2010). Communities of practice develop as particular social routines and norms are accepted and used to acculturate new members (Lave & Wenger, 1991; Wenger, 1999). Within SQAs, platform capabilities, routines and norms affect both how individuals respond to questions and how people come to understand how to ask a question that is likely to receive a desirable response.

To illustrate, researchers attribute the success of Stack Overflow at least partially to the fact that the community has successfully implemented norms of only considering objective, fact-based questions (Mamykina et al., 2011). Members are given explicit instructions on the FAQ page that only practical, answerable questions are allowed on the site (“Frequently Asked Questions - Stack Overflow,” n.d.). Similarly, the FAQ page instructs users to learn how to create better questions if they experience the problem of no answers to their queries. Norm-setting using FAQ and other pages is a distinct feature of online communities (Maloney-Krichmar & Preece, 2005).

Stack Overflow provides structures for members to reinforce these norms. For example, members can vote questions up or down based on their assessment of its quality. Users can designate questions as favorites, or leave comments on others’ posts. These features structure how individuals interact on the site, but it is only through the uptake of these features that a social norm (e.g. only post practical, answerable questions) becomes stable (Butler & Wang, 2012). For example, if members vote up questions that are more discussion-based, they can set an alternative norm regardless of the official statements on the FAQ page. As a result, learning how to ask a question that is likely to receive useful responses from an SQA community is not simply a matter of reading the formal policy documents. It necessarily involves a complex process of learning both in and about the demands of the sociotechnical system that is the SQA site.

SQAs then are a distinct form of online learning community where social interaction, mediated by the features and design of the platform, creates both the need and opportunities for learning (Bruckman, 2006). Sociocultural learning theory offers a useful lens for understanding social learning behavior in SQA sites because it highlights three aspects of learning that are particularly relevant to SQAs: that knowledge is necessarily situated in a particular community, learning is mediated by symbols, and social interaction is an important source for individual development (John-Steiner & Mahn, 1996).

From a sociocultural perspective, knowledge is not a discrete object or piece of information that one obtains independent of context. Instead, knowledge is situated in a particular community, in which community members negotiate (and renegotiate) social norms, meaning, and definitions of quality (John-Steiner & Mahn, 1996; Lave &

Wenger, 1991). This conceptualization of learning is reflected in the varying definitions of question and answer quality present in different SQA sites, and even between sub-communities within a single SQA site (Kim & Oh, 2009; Raban, 2009).

Different SQA communities, sub-communities, and individuals have different criteria for defining good questions and answers (Logie et al., 2011). Features and incentives built into SQA platforms influence how members define quality and what types of information they value (Nam, Ackerman, & Adamic, 2009). For example, in one design experiment, introducing market pressures (e.g. money) resulted in more efficient answering behavior, but also less focus on developing social bonds between users (Hsieh & Counts, 2009). A respective SQA community defines its own definition of quality. The social, cultural, and technological constraints of each community influence these definitions. Whatever knowledge an individual learns about how to ask questions is inextricably embedded in the social norms, practices, and meanings of the particular SQA community.

Sociocultural learning theory also highlights how learning is mediated by language and symbols (John-Steiner & Mahn, 1996). The prior example of FAQ pages in Stack Exchange sites nicely illustrates how language plays a role in defining community. Two of the communities we explore in this study – Database Administrators and Science Fiction & Fantasy – use very different language to situate their members. In Database Administrators, members are guided to ask technical questions and avoid topics such as career advice. Conversely, in Science Fiction & Fantasy, members are encouraged to ask questions about plot and historical context of sci-fi works, but avoid factual questions that can easily be answered by reference sites.

Other symbols play major roles in Stack Exchange sites. For example, badges signify expertise and motivate members towards particular activities (Mamykina et al., 2011). Up and down votes on questions determine their quality score and influence where on the interface it appears for others. Features such as rewards (called bounties in Stack Exchange) and comments that express thanks or seek clarification, are related to the social rating of answers in SQA sites (Raban, 2009). In SQAs, the learning that individuals do about how to ask effective questions is necessarily mediated by a complex system of language and symbolic artifacts (features of the site).

Finally, social behaviors are a critical source of learning for the individual, providing feedback mechanisms that enable members to better learn the community’s norms and definitions of quality. Learning how to work with a sociotechnical system, such as an SQA, requires that individuals get feedback about their attempted actions. In an SQA this feedback arises in the context of social interactions that are structured by the design features of the SQA platform. Mechanisms, such as votes, comments, and

favorite indicators, allow individuals to interact and easily provide particular types of feedback. A social learning perspective on SQA platform design suggests that inclusion (or exclusion) of these capabilities are likely to be critical in the success of an SQA because of the role that they play in supporting an individual's efforts to learn how to ask effective, acceptable questions.

**HYPOTHESES**

While sociocultural learning theory provides a general approach to understanding how SQA platforms might affect the way an individual learns how to ask good questions, it doesn't provide specific indications of what types of feedback and social experience are most likely to affect and enhance this learning process. In this section we outline several ways that design features of Stack Exchange sites can act as a source of social learning, and these examples motivate the specific hypotheses we examine in the study.

In asking how members might learn to be better question askers over time, we considered several features of Stack Exchange sites that provide social feedback about this task. How would a user know that their question is acceptable to their respective SQA community? One feature of Stack Exchange sites is the ability to vote up or vote down questions and answers. As noted above, from a sociocultural lens the definition of question quality is negotiated by the community through interactions such as up and down vote capabilities. No matter what the direction of the vote, all votes provide information about how the SQA community has responded to a prior question, thereby providing social feedback to a user about what types and characteristics of questions are seen as valuable by the community. Thus we expect that:

H1: The number of votes (up and down) a user has received on a prior question will be positively correlated with the quality of their subsequent question.

Effective questions receive answers, with some receiving much more community response than others. Thus, the volume of answer behavior can also act as a social signal to a user about their question quality and help them learn how to construct better questions of value in the future.

H2: The number of answers a user has received on a prior question will be positively correlated with the quality of their subsequent question.

Members in Stack Exchange sites can leave comments on users' questions. These comments may help the user refine or edit their question, or explicitly provide feedback about whether the question is suitable for the community.

H3: The number of comments a user has received on a prior question will be positively correlated with the quality of their subsequent question.

Users can also choose an answer that represented the best quality answer to a given question (an *accepted answer* in Stack Exchange). This affordance is an example of sociocultural learning. The process of choosing a best answer requires a member to reflect on the question, its intent, and assess the range of answers to ascertain which answer best addressed those goals. This suggests that question askers who go through the process of choosing an accepted answer may create better mental models about what kinds of questions will elicit quality answers.

H4: If a user was able to choose an accepted answer on a prior question, this will be positively correlated with the quality of their subsequent question.

In addition to choosing an accepted answer, Stack Exchange members can designate questions as a personal favorite. Knowing how many other people designated one's question as a favorite may also serve as a strong social signal that can help a member learn what makes for a quality question.

Record Type	Description
Comments	Logs each comment's text, score, post id of the post the comment belongs to, and the user id of the commenter
Post History	Logs all changes users make to posts, title or content. In this study, we use posts that represent questions asked
Posts	A log of all questions and answers on the site, their creation date, accepted answer id (for questions), score, view count (questions), title and text, user id of contributor, last edits, last activity date, tags, answer count, comment count, favorite count, parent id (for answers) and close date if question is closed
Users	Snapshot of user profiles at the time the data dump was taken; includes reputation, creation date, display name, last access date, about me, profile views, upvotes, downvotes, and demographic information
Votes	Log of all votes and actions made on a post, to which post the votes belong, and the timestamp. Vote types include: accepted answers, upvotes, downvotes, favorites, offensive content flag, close, reopen, bounty information, post deletion, undeletion, spam flag.

**Table 1: Description of Data Types in Stack Exchange Datasets**

H5: The number of favorites a user has received on a prior question will be positively correlated with the quality of their subsequent question.

Learning in a social and cultural community is not passive. A Stack Exchange user might not only learn by receiving social feedback on their posts. They may also learn by actively participating in the community. The act of posting answers, comments, and votes also requires that an individual to reflect on their own developing definitions of quality. Thus, we expect that active participation in the SQA community will influence a person’s learning and ability to ask high quality questions:

H6: The number of answers, comments, and favorites a user has previously posted will be positively correlated with the quality of subsequent questions.

Finally, since sociocultural learning theory posits that knowledge and learning are embedded in particular contexts and communities, it is necessary to consider the possibility that the role and impact of feedback mechanisms might vary from context to context. To examine this idea, we chose four Stack Exchange communities for this study. The communities were similar in their membership and activity levels, but differed along a spectrum of technically focused to social-discussion focused SQAs. Based on sociocultural learning theory and prior work in SQA sites and online communities, we expect that cultural norms, and hence the impact of feedback and experience, would play out differently in these respective settings (Maloney-Krichmar & Preece, 2005). However, since prior work and theory provides little basis for hypothesizing the specific nature of these differences we examine an additional exploratory question:

R1: How does the impact of feedback and experience on individuals’ question quality vary among SQA sites that have different social norms and community aims?

Taken together, we are positing and examining a specific model of technology mediated social participation. Preece & Shneiderman (Preece & Shneiderman, 2009) suggest that users of social computing platforms can be characterized along a spectrum of readers to leaders. The social learning framework and hypotheses we pose here seeks to explain how the affordances of a particular social platform, Stack Exchange, are related to the development and cultivation of SQA users over time.

## METHODS

### Study Sample and Initial Dataset

This study uses data from the April 2012 Stack Exchange data release. In this data release, there were a total of 64 sites from the Stack Exchange network. Each site’s data included XML data files with information about comments, post history, posts, users, and votes. The structure of the Stack Exchange data is outlined in Table 1.

From the 64 Stack Exchange sites for which data was available, we chose four that were similar in membership and activity levels: Android Enthusiasts, Database Administrators, Seasoned Advice (which we will refer to as Cooking), and Science Fiction & Fantasy. Aggregate statistics for the four communities are shown in Table 2. Because each Stack Exchange site is treated as a separate entity, user IDs in the raw data are not consistent across sites. User 1 in Stack Overflow is not the same as user 1 in Android Enthusiasts. To avoid problems arising from non-independence of observations, the dataset for the four communities were created and analyzed separately.

### Measures

From the raw Stack Exchange data, we developed a longitudinal, panel dataset for each Stack Exchange community. The panel data is organized by user ID and the order of questions posted by each user (from earliest to latest). Within each SQA site, each question was assigned a relative question number ( $QuestionNumber_{iq}$ ), such that the first question asked by user  $i$  in that SQA site is question 1, the second is 2, and so on. Each record in the dataset represents a user ( $i$ )/question ( $q$ ) pair and includes values for each measure for that pair.

### Dependent Variable

As a measure for question quality, we use Stack Exchange Question Score ( $QuestionScore_{iq}$ ), which indicates the number of up votes minus the number of down votes question  $q$  posted by user  $i$  received. This number is visible to users of Stack Exchange, and represents the community’s collaborative assessment of quality.

### Independent Variables

Hypothesis 1 (H1) examines the influence of prior up votes and down votes. To test this hypothesis we used the count of up votes ( $UpVotes_{i,q-1}$ ) and down votes ( $DownVotes_{i,q-1}$ ) received in response to user  $i$ ’s previous question.

Community	Users	Questions	Answers	Comments	Votes	Badges Earned
Android Enthusiasts	11,685	7,020	10,755	20,155	66,180	17,968
Database Administrators	7,295	5,134	9,466	21,682	53,756	13,521
Cooking	7,771	5,402	15,495	26,402	93,091	18,138
Science Fiction	5,195	3,698	7,613	21,597	103,137	15,337

**Table 2: Aggregate Activity in the Stack Exchange Communities**

For H2, we use the lagged value of answers received by a user's previous question.  $AnswersReceived_{iq-1}$  is the total number of answers contributed in response to user i's previous question.

For H3, we use the lagged value of comments received by a user's previous question.  $CommentsReceived_{iq-1}$  is the total number of comments contributed in response to user i's previous question.

For H4, we use the lagged value of Accepted Answer on a user's previous question.  $AcceptedAnswer_{iq-1}$  is a binary variable indicating whether user i selected an accepted answer for their previous question.

For H5, we use the lagged value of Favorites received by a user's previous question.  $FavoritesReceived_{iq-1}$  is the total number of answers contributed in response to user i's previous question.

In addition, we calculated the actions a user took themselves in between posting questions (H6).  $AnswersPosted_{iq-1}$  is the total number of answers contributed by a user prior to posting their next question.  $CommentsPosted_{iq-1}$  is the total number of comments a user contributed to other members' posts, prior to posting their next question.  $FavoritesPosted_{iq-1}$  are the total number of favorites a user designated on others' posts prior to posting their next question.

Users' previous voting activity was not included because the available raw data does not record the user ID of

individuals who contribute up and down votes.

The statistical models also included several control variables.  $Views_{iq}$  and  $Favorites_{iq}$  are the total number of views and favorites for question q posted by user i. Question number ( $QuestionNumber_{iq}$ ), which also indicates the total number of questions user i has asked, was included as a measure of the user's experience with the particular SQA site. The descriptive statistics for all the described measures are outlined in Table 3.

Construction of the measures described below requires matching user IDs, questions, and responses (i.e. votes, answers, comments, and favorites). In approximately 4% of the cases these relationships were invalid because of missing user and question data. These cases were dropped from the raw dataset and were not included in the construction of the measures.

In this study we are using longitudinal, individual fixed effects regressions to examine how users learn to be better question askers over time. The measures and models describe individual change over time and thus members who only contributed one question to an SQA site are excluded (since there is no change over time to model). As has been found in prior studies of online communities and SQA sites, a small minority of members account for the majority of activity. In our data, only 6.5%-9.5% of users posted multiple times and so all interpretations are limited to this subset of members. However, these users account for approximately 46%-75% of all question contributions to

	Android Enthusiasts	Database Administrators	Cooking	Science Fiction
Question Score	3.45 (4.92)	3.12 (3.59)	5.50 (5.66)	9.44 (7.84)
Up votes	3.56 (4.90)	3.17 (3.58)	5.59 (5.67)	9.67 (7.85)
Down votes	0.11 (0.39)	0.05 (0.32)	0.09 (0.39)	0.23 (0.72)
Answers received	1.65 (1.57)	1.87 (1.34)	2.94 (2.63)	2.11 (1.65)
Comments received	1.28 (1.82)	1.46 (2.24)	1.62 (2.31)	2.23 (2.83)
Favorites received	0.58 (2.16)	0.47 (1.32)	0.53 (1.61)	0.65 (1.18)
Accepted Answer	0.58 (0.49)	0.67 (0.47)	0.75 (0.43)	0.79 (0.41)
Answers posted	1.30 (8.31)	0.93 (9.59)	1.99 (8.69)	1.06 (8.76)
Comments posted	3.36 (16.40)	2.82 (25.05)	4.56 (22.73)	3.87 (27.45)
Favorites posted	0.41 (2.50)	0.26 (1.25)	0.26 (0.96)	0.53 (2.45)
Question Number (number of questions posted)	6.16 (9.19)	6.26 (9.79)	9.62 (15.08)	35.06 (53.78)
Views	951.61 (3292.03)	321.21 (584.82)	659.31 (1246.54)	396.75 (749.99)
Percentage of Users who posted more than one question	8.2%	9.5%	9.5%	6.5%
Percentage of Site's Question Activity	48.0%	46.3%	59.7%	74.9%

**Table 3: Descriptive Statistics (for users included in the analysis). Mean (Standard Deviation)**

these sites. So while they may be a minority of the user population they account for a significant proportion of activity within the studied SQAs.

### ANALYSIS

We use longitudinal fixed effects regression models in the analysis. This quasi-experimental strategy offers several advantages (Oktay et al., 2010). Each individual is treated as a fixed effect ( $\alpha_i$ ) allowing us to model changes over time within each person. In addition to modeling time-order, the analysis accounts for unobserved variable bias. Any factors related to individuals that are stable over time (e.g. gender or stable measures of dispositions, motivations etc.) are accounted for in the error structure, effectively controlling for unobserved individual differences. To test the proposed hypotheses we estimate the following model with the data from each of the four SQA sites:

$$\begin{aligned} \text{Question Score}_{iq} = & B_1 \text{UpVotes}_{iq-1} + B_2 \text{DownVotes}_{iq-1} + \\ & B_3 \text{AnswersReceived}_{iq-1} + B_4 \text{CommentsReceived}_{iq-1} + \\ & B_5 \text{FavoritesReceived}_{iq-1} + B_6 \text{AcceptedAnswer}_{iq-1} + \\ & B_7 \text{AnswersPosted}_{iq-1} + B_8 \text{CommentsPosted}_{iq-1} + \\ & B_9 \text{FavoritesPosted}_{iq-1} + B_{10} \text{QuestionNumber}_{iq} + B_{11} \text{Views}_{iq} \\ & + B_{12} \text{Favorites}_{iq} + \alpha_i + u_{iq} \end{aligned}$$

This model allows us to examine whether different types of social feedback on prior questions from the community (votes, comments, favorites etc) and past personal experience (answers, comments, favorites posted) influences the score a user receives on their subsequent question (indicating the user is better at posting quality

questions).

### FINDINGS

*(H1) Are prior up votes and down votes correlated with future quality score? (Yes)* The findings (Table 4) show a consistent relationship between prior up votes and down votes received and the quality score of a user’s current question. Positive social feedback is particularly helpful for users to create future quality questions. Receiving more up votes on a previous question is positively correlated with the quality score of a user’s current question. Interestingly, negative social feedback on a prior question in the form of down votes, relates to a lower quality score on a future question. The results suggest that positive feedback mechanisms are more beneficial to users than negative feedback as they learn about the norms of an online community.

*(H2) Do answers received on a prior question relate to future quality score? (Little Support)* We found less support that prior answer feedback acts as a social learning mechanism in these Stack Exchange sites. In three of the communities (Android, Cooking, and Sci-Fi), the number of answers received on a prior question did not correlate with future question quality. In the Database Administrators site, receiving more answers on a question was negatively related to future question quality. This finding may be explained by the intended norms of the Database Administrators community, which is focused on objective, answerable questions. A quality question in such a context would ideally elicit fewer, clearer answers. It seems that in

	Android Enthusiasts	Database Administrators	Cooking	Sci Fi & Fantasy
Up Votes (lagged)	<b>0.06**</b>	<b>0.10**</b>	<b>0.09**</b>	<b>0.11**</b>
Down Votes (lagged)	<b>-0.06**</b>	-0.03	<b>-0.03*</b>	<b>-0.07**</b>
Answers Received (lagged)	-0.01	<b>-0.05**</b>	-0.03	-0.01
Comments Received (lagged)	-0.02	-0.03	-0.01	-0.01
Favorites Received (lagged)	<b>0.03*</b>	-0.01	<b>0.05**</b>	<b>0.03*</b>
Accepted Answer (lagged)	-0.001	-0.001	-0.02	<b>0.03*</b>
Answers Posted (lagged)	-0.03	-0.03	<b>0.07*</b>	-0.01
Comments Posted (lagged)	-0.01	-0.04	-0.06	-0.01
Favorites Posted (lagged)	-0.01	-0.02	-0.01	<b>-0.04*</b>
Question Number	<b>-0.05**</b>	<b>-0.17**</b>	<b>-0.07**</b>	<b>-0.15**</b>
Control Variables				
Current Views	<b>0.08**</b>	<b>0.13**</b>	<b>0.22**</b>	<b>0.39**</b>
Current Favorites	<b>0.74**</b>	<b>0.49**</b>	<b>0.56**</b>	<b>0.39**</b>
N	3,373	2,376	3,227	2,768
Number of Users	956	696	740	339
R <sup>2</sup>	0.70	0.39	0.47	0.51

Normalized beta coefficients. \* p < 0.05, \*\* p < 0.01

**Table 4: Longitudinal, Individual Fixed Effects, Regression Results Predicting Question Score**

this particular site, eliciting more answers in a prior question is a negative influence on a user, providing little social feedback in constructing future questions that will receive higher quality scores from the community.

*(H3) Do comments received on previous questions correlate with future quality score? (No)* The findings show that receiving more comments on a prior question had no relationship to future question quality. One potential interpretation of this result is that the volume of comments does not serve as a strong learning mechanism, but rather the content and nature of comments is more important. Unfortunately, we could not ascertain the nature and content of comments in the dataset. Future work might consider content analysis and machine learning strategies to code comments to better explore whether types of comments relate to better community building in these SQA sites.

*(H4) Does identifying an accepted answer on a previous question relate to future quality score? (Little Support)* We find that the Accepted Answer feature does not serve as a social learning mechanism for three of the Stack Exchange communities (Android, Database, and Cooking). In the Sci-Fi & Fantasy community, identifying an accepted answer on a previous question was positively correlated with future quality score. The results suggest that finding an Accepted Answer serves as at least one indicator for social learning in this one community, where the open-ended, social discussion is the focus. However, the relationship was fairly small (an increase of 0.03 standard deviations in quality score for an increase in Accepted Answer).

*(H5) Are prior favorites received correlated with future quality score? (Yes)* The results show that receiving favorites on a prior question is positively correlated with future quality score in three of the Stack Exchange communities (Android, Cooking, and Sci-Fi). This suggests that the design mechanism of designating favorites serves as a positive feedback mechanism that helps users craft higher quality questions in the future. Interestingly, Favorites had a non-significant relationship to question quality in the most technically focused SQA site (Database Administrators).

*(H6) Does a user's personal activity relate to their future question quality score? (Little Support)* We found little overall support that a user's own past actions help them learn to craft higher quality questions in the future. However, in the Cooking Community, past answers posted was positively correlated with the user's future question quality. Posting answers to cooking-focused questions helps users in this community craft future questions that receive higher community scores. In the Sci-Fi & Fantasy community, designating other posts as favorites was negatively associated with future quality score. One potential interpretation is that in more open-ended SQA communities, the act of favorite-ing posts may be more idiosyncratic to each user, and thus there is less consistent

social feedback that cultivates members towards the community's question and answer norms.

*(R1) Do the same Stack Exchange features have different influences on social learning in each respective SQA site? (Yes)* In this study we analyzed four different online communities that use the same underlying platform, Stack Exchange. Thus, we are able to show that the same feature set exerts different social learning feedback based on the online community. An interesting pattern emerges if one thinks of the communities from most to least technically oriented. Database Administrators represents the most technical, objective-information focused community. Android Enthusiasts and Cooking range in the middle of this spectrum. Sci-Fi & Fantasy represents the most open-ended, social community.

In Database Administrators, almost none of the social features were correlated with users' ability to craft higher quality questions over time. The only feature that appears to provide strong social feedback is prior up votes. In the most open-ended site, Sci-Fi & Fantasy, more of the community features exert social feedback to help users craft higher quality questions. In Sci-Fi, up votes received, down votes received, favorites received, favorites posts, and accepted answer were significantly correlated with future question score. Finally, the Android and Cooking communities fell in the middle of this range. These patterns offer exploratory support to show that factors such as community norms influence user behavior. In this study, we show how social features might exert differential feedback mechanisms to help individuals learn, based on the norms of a particular SQA community.

## **DISCUSSION AND DESIGN CONSIDERATIONS**

The findings in this study offer several contributions and design considerations. First, we demonstrate how analytic frameworks from research in online learning communities and sociocultural theories of learning can inform our understanding of how users develop in SQA platforms. Sustaining vibrant knowledge communities in SQA sites might not only depend on effectively organizing information systems, but also in developing individuals into more skilled members over time. Past research in SQAs implicitly focus on developing more effective systems for knowledge sharing. In this study we contribute a complementary and equally challenging focus on helping members learn how to be better knowledge sharers.

Our analyses show some support that the design features of the Stack Exchange platform – votes, favorites, answers etc. – provide implicit social feedback that is related to members' ability to craft higher quality questions in the future. These social features simultaneously act as markers for question quality and influence members' future actions. However, we note that despite finding positive correlations between design affordances and future question-asking ability, many of these relationships were small (Table 4). Furthermore, it is interesting to note that the Question



Number variable had a small, but negative relationship to question quality. The results hint that, after controlling for other factors, users create slightly lower quality questions over time.

These results suggest that future work should consider design experiments that better enhance the social learning of members in SQA communities. For example, the social features of Stack Exchange sites (votes, favorites etc.) serve as implicit markers for question quality, but small changes to the user interface could make these markers more explicit to the user. A compelling question is whether helping members be more cognizant of the community's feedback on their question, might help them craft better contributions in the future. For example, users could be made aware that one prior question received unanimous positive support, while another received mixed votes, and then prompted to consider the reasons behind these differences in community response. Prior work that has uses various modeling techniques to characterize question types could be used to provide dynamic feedback to users about the qualities of their questions, with prompts to consider the potential success in garnering responses. The use of community feedback mechanisms in SQAs creates an ideal environment for social learning, but making these features explicit and providing scaffolded learning opportunities for members to reflect on these mechanisms may improve the cultivation of members over time.

This study also shows how differences in norms and culture of online communities, interact with the same feature set of a platform (Stack Exchange), to structure varying learning mechanisms. For example, the impact of designating posts as favorites had a differential effect in Cooking versus Database Administrators. The results show that factors identified in previous research such as topic area or use of policy pages to set community norms (Adamic et al., 2008; Maloney-Krichmar & Preece, 2005) complement or interact with the social features available in a platform. Perhaps in open-ended communities such as Sci Fi & Fantasy, the presence of more social feedback mechanisms is needed to better structure and acculturate members. In more defined communities such as Database Administrators, lesser social feedback features are needed. Such questions are intriguing for future work.

#### **LIMITATIONS**

Several limitations of our analysis should be noted to aid in interpreting the empirical findings and highlighting opportunities for future work. First, this analysis only examined four SQAs that use the Stack Exchange platform. Our findings highlight how different communities may enact very different norms and feedback mechanisms, and thus the results of this study are necessarily bounded and not generalizable to all SQA sites. Future work is needed that can tease apart how community norms, structure, and the feature set of social computing platforms influence the development of members over time. There is a ripe

opportunity to identify new online community designs that can maximize the social learning of members and efficiently help new members become skilled members.

Our analysis examined particular features in Stack Exchange (e.g. answers, comments, favorites etc.). We did not consider other important elements of the Stack Exchange platform, such as the role of reputation points and badges. An intriguing future question is whether social feedback mechanisms have differential effects when they interact with badges or reputation points. For example, in Stack Exchange, members earn reputation points via increased participation in the community. The more reputation points a user earns, the more functions of the site become available to that user. The opportunity to advance within the reputation system motivates some users to contribute more to the community, and perhaps social feedback mechanisms are enhanced for those members who actively seek out reputation points or badges.

Finally, our analysis considered the volume of feedback, but could not examine the nature of feedback. One example is the non-significant relationship between volume of comments received and future question quality. It may be likely that volume is not important for certain feedback mechanisms, but rather the content of the feedback is important. Stack Exchange comments that recommend that a question be removed or a comment that offers constructive recommendations for editing, may likely have a differential influence on the social learning of members. Future work that can combine classification or content analysis techniques of posts with longitudinal models will make a needed contribution in this area.

#### **CONCLUSION**

SQA sites represent an increasingly popular method for knowledge sharing. The social features of platforms such as Stack Exchange, and the traces of user behavior they allow us to collect, provide researchers an unprecedented opportunity to empirically observe how social learning occurs in online communities. In this study, we demonstrate one way to quantitatively model social learning and examine how the design affordances of SQA communities influence members' literacy skills over time.

The findings highlight how different design features of Stack Exchange help individuals learn vital literacy skills – e.g. how to ask good questions that provide value to the broader community. This study also articulates a learning-based framework that reorients one to consider whether SQAs and similar online communities are successful in cultivating members over time. A social learning perspective of SQAs complements existing information processing frameworks and contributes to other perspectives in HCI research, such as the reader-to-leader framework (Preece & Shneiderman, 2009), that are concerned with understanding how individuals develop over time into productive contributors to social platforms.

## REFERENCES

- Adamic, L. A., Zhang, J., Bakshy, E., & Ackerman, M. S. (2008). Knowledge sharing and yahoo answers: Everyone knows something. In *Proceedings of the 17th international conference on World Wide Web* (pp. 665–674). New York, NY, USA: ACM. doi:10.1145/1367497.1367587
- Anderson, A., Huttenlocher, D., Kleinberg, J., & Leskovec, J. (2012). Discovering value from community activity on focused question answering sites: A case study of stack overflow. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 850–858). New York, NY, USA: ACM. doi:10.1145/2339530.2339665
- Bruckman, A. (2006). Learning in online communities. *The Cambridge handbook of the learning sciences*, 461–472.
- Butler, B. S., & Wang, X. (2012). The cross-purposes of cross-posting: Boundary reshaping behavior in online discussion communities. *Information Systems Research*, 23(3-Part-2), 993–1010.
- Fichman, P. (2011). A comparative assessment of answer quality on four question answering sites. *Journal of Information Science*, 37(5), 476–486.
- Frequently Asked Questions - Stack Overflow. (n.d.). Retrieved April 4, 2013, from <http://stackoverflow.com/faq>
- Gardelli, G., & Weber, I. (2012). Why do you ask this? In *Proceedings of the 21st international conference companion on World Wide Web* (pp. 815–822). New York, NY, USA: ACM. doi:10.1145/2187980.2188205
- Gazan, R. (2011). Social Q&A. *Journal of the American Society for Information Science and Technology*, 62(12), 2301–2312.
- Hanrahan, B. V., Convertino, G., & Nelson, L. (2012). Modeling problem difficulty and expertise in StackOverflow. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work Companion* (pp. 91–94).
- Harper, F. M., Moy, D., & Konstan, J. A. (2009). Facts or friends?: Distinguishing informational and conversational questions in social Q&A sites. In *Proceedings of the 27th international conference on Human factors in computing systems* (pp. 759–768).
- Harper, F. M., Weinberg, J., Logie, J., & Konstan, J. A. (2010). Question types in social Q&A sites. *First Monday*, 15(7). Retrieved from <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/viewArticle/2913>
- Hsieh, G., & Counts, S. (2009). mimir: A market-based real-time question and answer service. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 769–778).
- John-Steiner, V., & Mahn, H. (1996). Sociocultural approaches to learning and development: A Vygotskian framework. *Educational psychologist*, 31(3-4), 191–206.
- Kim, S., & Oh, S. (2009). Users' relevance criteria for evaluating answers in a social Q&A site. *Journal of the American Society for Information Science and Technology*, 60(4), 716–727.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, UK: Cambridge University Press.
- Logie, J., Weinberg, J., Harper, F. M., & Konstan, J. A. (2011). Asked and Answered: On Qualities and Quantities of Answers in Online Q&A Sites. In *Fifth International AAAI Conference on Weblogs and Social Media*.
- Maloney-Krichmar, D., & Preece, J. (2005). A multilevel analysis of sociability, usability, and community dynamics in an online health community. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 12(2), 201–232.
- Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G., & Hartmann, B. (2011). Design lessons from the fastest q&a site in the west. In *Proceedings of the 2011 annual conference on Human factors in computing systems* (pp. 2857–2866).
- Nam, K. K., Ackerman, M. S., & Adamic, L. A. (2009). Questions in, knowledge in?: A study of naver's question answering community. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 779–788).
- Oktay, H., Taylor, B. J., & Jensen, D. D. (2010). Causal discovery in social media using quasi-experimental designs. In *Proceedings of the First Workshop on Social Media Analytics* (pp. 1–9).
- Preece, J., & Shneiderman, B. (2009). The reader-to-leader framework: Motivating technology-mediated social participation. *AIS Transactions on Human-Computer Interaction*, 1(1), 13–32.
- Raban, D. R. (2009). Self-presentation and the value of information in Q&A websites. *Journal of the American society for information science and technology*, 60(12), 2465–2473.
- Rosenbaum, H., & Shachaf, P. (2010). A structuration approach to online communities of practice: The case of Q&A communities. *Journal of the American Society for Information Science and Technology*, 61(9), 1933–1944.
- Shah, C., & Pomerantz, J. (2010). Evaluating and predicting answer quality in community QA. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval* (pp. 411–418).
- Wenger, E. (1999). *Communities of practice: Learning, meaning, and identity*. Cambridge, UK: Cambridge University Press.