Seeking and Offering Expertise across Categories: A Sustainable Mechanism Works for Baidu Knows

Jiang Yang, Xiao Wei

School of Information
University of Michigan,1085 South University Ave. 304 West Hall, Ann Arbor, MI 48109-1107, USA
(yangjian, xiaowei)@umich.edu

Abstract

This paper presents the first comprehensive exploration of the largest Chinese online knowledge sharing community-Baidu Knows. With analyzing 5.2 millions questions and 2.7 million users participated during 4.5 months on the site in 2008, we investigate how users adjust initial attempts and behave differently according to the level of participation; in particular, there is a positive dynamic for answerers to input more, be more focused, win more, and thus be rewarded more. As the result, a core user group forms to actively participate in both asking and answering across categories, thus maintaining a self-sufficient community. In addition, a prominent "sense of community" would enhance the social bonds within the community, especially for the contributors who can offer expertise but can rarely learn from others. The study suggests Baibu Knows as a successful design instance for further studies.

Introduction

Across the globe, peer-based online Question-Answer (Q&A) communities have been rapidly accumulating knowledge and expertise to serve as vast knowledge repositories. Examples include Yahoo! Answers in English, Naver in Korean, and Baidu Knows in Chinese. All these sites share similar point-based systems thus demonstrating the capacity of garnering tremendous popularity and knowledge growth through non-monetary incentives. For example, Baidu Knows, the site we investigate here, has answered over 47 million questions since 2005 and receives more than 47,000 questions per day.

Although sharing a very similar technical platform of maintaining the QA communities with its cousins, Baidu Knows is featured with several critical albeit small deviations in terms of system design that would probably cause different user behavior and site performance. First, it allows askers to provide extra points to award the best answer which potentially gives higher and flexible incentive than flat points. In addition, the site intentionally establishes the "sense of community" by enhancing people's social interactions (providing feedback to answerers and instant messaging) and community awareness (employing a prominent honor title system and explicitly promoting experts). Furthermore, paralleling to other ongoing studies on many English QA sites like Yahoo! Answers or Google Answers, we also expect to learn cultural reflections in terms of system design through comparable large scale analysis on Baidu Knows.

Very interestingly, this site presents a sustainable mechanism where people's incentives can be successfully addressed within the system: unlike some other sites where users can be largely identified between askers and answerers, on Baidu Knows, a significant portion of users play as both asker and answerer; they answer questions and ask somewhere else using the points they earned by answering. We explore how users spread over multiple categories: users asked in more categories than answered; and in some categories, we observe considerably more concentrating users while in others, users tend to participate for short visit.

To the asker end, points are allocated among questions and those valued more important were awarded higher; and consequently obtaining more answers. In addition, askers can gradually improve the efficiency of per point in terms of buying participation. To the answerer end, there is a positive dynamic for answerers to input more, be more focused, win more, and thus be rewarded more. Finally, we observe the only-answering group although less active, seeking more challenging questions and performing better thus suggesting other non-point community features could be complementary to incentivize those users.

This paper first introduces Baidu Knows and the dataset; then examines how the reward mechanism works and users behave differently according to their activity level; in particular, we look into the participation pattern of the user group who both asked and answered which makes the main core of the community; we discuss our findings with related work and conclude the design implications and future work in the end.

Baidu Knows and Data Set

Baidu Knows (BK), founded in 2005, is the biggest Chinese Q&A community; where approximately 83 million questions have been submitted and 47 millions been resolved. BK's format is similar to many other Q&A sites: the main page lists recommended topics and newly asked

Figure 1. Screenshot of the Baidu Knows Q&A community. http://zhidao.baidu.com/upf/

or being voted questions, and quick links to meta-categories and their frequent sub categories; there are also frequently updated knowledge entries. The sequential pages consist of a question and its answers, and the asker can provide further feedback on the answers.

The asker can select a best answer or invite other people to vote for the best answer. Each question is closed after 15 days; one can prolong the period for another 3 days by adding award points. If there is no answer, an insufficient number of votes (less than 4 votes), or the asker wants to withdraw the question because she is not satisfied, the question will be closed as unsolved. In addition, we believe the site may also delete politically sensitive questions and answers. Thus, not all questions are answered; in our dataset, 56% questions were successfully resolved.

The site has two hierarchical category levels including 24 meta-categories and approximately 300 subcategories. Meta-categories include, i.e., Health, Computer/Internet, Fashion/Life. As a further example, the Computer/Internet meta-category includes C++, virus, and downloading subcategories. The category for a question is assigned by its asker when posting the question. This might be problematic in terms of correctly categorizing a question. However, the system offers technical assistance. When one inputs any keywords, the system will attempt to generate relevant solved questions to avoid potentially redundant questions; and no matter if the question has been answered or not, the system provides suggestions for the category.

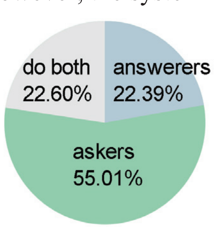

Figure 2: Distribution of three types of participants

BK has a point system, and users actively earn points by logging on the site and answering questions. BK also allows askers to offer extra points to the person who provides the best answer. This mechanism encourages more and better answers, but also, it encourages askers to earn more points in order to ask. In addition to using points when asking questions, users can gain "levels." BK employs a very tempting honor-title system that includes five different themes: business titles (e.g., from trainee to CEO), traditional Chinese imperial examination titles, magical titles, knight-errant titles, and traditional Chinese military titles. The site also explicitly promotes outstanding contributors, encouraging participation. For example, it selects experts who perform well in particular categories to be the "knowledge master" and "star of Knows" each week, and provides links to their profile pages from the portal or category index pages. BK also publicizes users who have been newly promoted. All these means provide additional incentives to garner points find acknowledgement, accumulate fame, and contribute on the site over time.

The site consciously builds up a sense of community using multiple strategies to establish and enhance the participants' social bonds. First, the combination of one's honor-title and the usually meaningful user ID would present a particular user identity. For example, CEO "Wind karma wind words" is a user who answered 3,278 questions and was chosen for best answer 1,360 times. The ID also links to the profile page with ID picture, personal information, ask/answer statistics, and her activities on Baidu forums site. For the same example, CEO "Wind karma wind words" claims he is a male and has a master degree, he likes sleeping late while hates smoking and he has also provided his favorite books and hobbies. This page is also linked to other social networking services on Baidu. In addition, an asker can provide feedback to the answerers in the question entry page, establishing post-question social interactions between the asker and answerers. While some give terse encouragement such as "thank you!" or "very thoughtful!", others initiate further discussion. Instance online chatting client is also available for users' inter-person communication, and we see many evidences that users use it or exchange contact in QA pages for further interpersonal interactions. As mentioned above, BK promotes contributors regularly thereby enhancing awareness of experts and the community overall. We believe successfully maintaining such a more perceptible sense of community considerably contributes to the prosperous dynamic on the site.

Data

The dataset used for the analysis here includes all users' activities over 4.5 months (January to mid-May, 2008). During this period, 9,300,000 questions were asked and 5,210,163 were resolved (otherwise been closed). 2,667,518 unique users participated.

In this dataset, as mentioned, only 56% submitted questions were finally resolved. For each question, there are only 3.33 replies on average. This is lower than YA, where the rate is 7.27, but it is better than Naver which has a mean of 1.7 answers. Since 55% of users on BK ask questions while askers and answerers on YA are fairly balanced, if computing the number of questions generated by a unit of population, BK users tend to have many more questions than YA users (Adamic et al. 2008). Figure 2 shows the distribution of the three types of users: the user group who both ask and answer presents a similar ratio to YA, while significantly higher than Naver, where people tend to play only answerer or asker. As we will discuss below, the group of users who both ask and answer forms the core of the practicing community, and they actively participate across categories, seeking as well as offering knowledge and expertise.

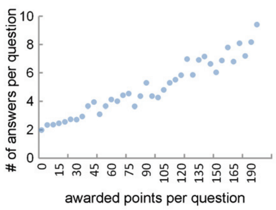

Figure 3: Average number of answers per question by amount of points

Incentives

Incentive design is crucial to knowledge sharing communities for them to be sustainable. Various sites we have observed all provide different explicit incentives. YA gives users points for answering and more points for being chosen as the best answer, BK and Naver allow askers to award extra points from their own account, and taskcn.com offers real money for the best solutions.

While some field studies compared different incentive schemes (Chen et al. 2007; Harper et al. 2008), we wanted to know whether virtual points instead of money can be a positive incentive for contribution, and whether the point incentive would have different effects in different ranges.

We found a correlation between the number of answers and awarded points for all questions ($R = 0.24$, $p<0.0001$) and for only the questions that offered extra points ($R = 0.26$, $p<0.0001$). This rate is very consistent among the different meta-categories, too. Figure 3 shows by average how many answers a question obtained and we can see that the trend is very linear, which suggests that points have a consistent effect to incentivize participation.

In addition, we considered whether answerers would be rewarded for more effort. Simply, longer answers are encouraged, which is consistent with our findings for YA. A two-sample t-test finds that best answers differ from non-best answers ($p<0.0001$). Best answers have a mean of 407 (sd=1320) while non-best answers have a mean of 226 characters (sd=986).

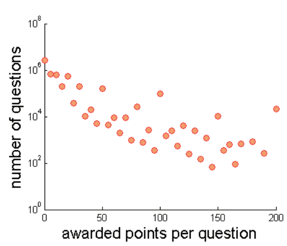

Figure 4: distribution of awarded points for each question

Pricing Questions

In order to understand how askers reward answerers, we investigated the distribution of questions' prices as shown in Figure 4. Although most questions did not offer any extra points for the best answers, on average, each question paid 11.6 extra points to the best answer. The system offers 2 points for submitting an answer and another 20 points for being selected as the best answer, and a user needs to obtain 100 points to be promoted for the first time (i.e., to get a title). Thus, compared to this scale, an incentive of 11.6 extra points seems rather considerable.

In addition, we hypothesize that users value questions differently, which can be partially represented by the award they are willing to offer. In fact, we will show in following section that askers pay more for their first questions and when people ask fewer questions they also pay higher amounts.

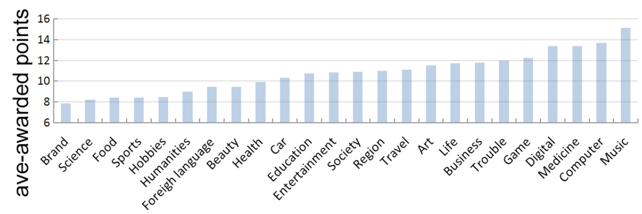

Figure 5: average awarded points in meta-categories

We found significant category difference in terms of question pricing. As presented in Figure 5, askers offer more in some categories such as "Music" and "Computer" while they price "Science" and "Brands"[1] the lowest. The price across categories correlates with popularity as measured by the category's total number of questions (R=.46, p =.025, in the 24 meta-categories). This, however, does not result in more answers per question (p=.46). This indicates the complexity in people's pricing behavior.

As we will show below, people place more value on their earlier questions. We calculated the ratio of users' first questions in each category and found this ratio is positively correlated with the average price of questions (R=.47, $p<0.05$.) The first question ratio and popularity count for a significant portion of variance of the price (R=.63, p=.005; and there was no correlation between them). This would suggest that some categories like Travel, although not necessarily popular, contain questions that trigger people to use the site and are valued higher.

Best Answer Selection

Interestingly, we also found consistent patterns in best-answer selection in terms of answering order (i.e., chronological sequence of answers). People mostly tend to choose the first posted

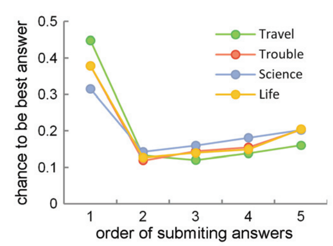

Figure 6: chance to be selected as the best answer for all question with 5 answers

[1] The meta-category "Brands" is about particular product brands; i.e., Adidas, KFC, and Philips.

answer as the best and secondly like to choose the last answer. From the second answer, the chance of being selected as best increases gradually (Figure 6 presents all questions which got 5 answers in 4 example meta-categories). We might expect by intuition, that answers would improve sequentially or at least the answer of the best quality would be random in order, since otherwise people would have less incentive to continue solving the question. In fact, according to our sample set, no answer of any order is necessarily better than others.

We believe some askers reward very prompt answerers. Although overall, the first answer is the most probable to be selected as the best answer, the actual selection of best answer is related to the amount of awarded points. As shown in Figure 7, the questions which selected best answer from different sequence order actually have awarded different amounts of points too. Questions that selected the last answer as the best offered the highest average award, and this value decreases backwards. This indicates that when askers offer fewer points, they tend to reward prompt answerers; otherwise, they may consider answer quality more. Higher awards should attract more participation, and this pattern suggests that askers want to compensate prompt responses when they offer smaller award. This behavior by users, if the case, would encourage contributions, as it provides a buffer between the highly popular questions (with high awards) and unpopular questions (with low awards).

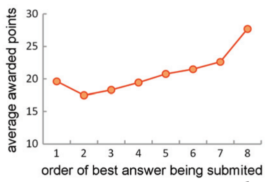

Figure 7: average awarded points for the questions that chose 1, 2,…8th answer to be the best

Users' behavior over time

Users can be differentiated by various dimensions. We previously found on the crowdsourcing site Taskcn.com that users adapt their behavior over time, and that the behavior of the most successful users is different from the rest (Yang, Adamic, and Ackerman 2008). Looking for a similar effect, we examine users' adaptive behaviors upon two dimensions: by role (answerer or asker) and by activity level and we find significant variance among subgroups of users.

The data set includes 35% new active users such that we could not capture the initial behavior for the majority of the users; thus alternatively, we excluded users who participated in the first month of the dataset while being active in later months and count them mostly as new users. However as we will see below, users make greater adjustments in the first several attempts during the period and reach rather stable status; suggesting that the new users would significantly count for these initial adjustments.

Answerers' Activity Level

For all users who have ever answered questions, each has answered 12 questions on average; however like many other online communities, the distribution of contribution is highly skewed and the highest answerer even has answered 18,301 questions during the 4.5-month period. In order to distinguish users of different activity levels, we group them by the number of answers they have answered: groups of answerers who answered 10~20, 20~40, 40~80 and 80~200 times with 132,670, 77,811, 40,010, and 21,305 users respectively.

Wining rate: winRate, a measure of answerer's performance, is defined as the total winning attempts divided by the total number of attempts. Figure 8 shows the average winRate by each group in order of attempts. First, all groups increase their winRate in their first 3 answers, after which their winRate stabilizes or drops. Second, the more active groups tend to have a higher winRate from the start and present smaller declining trend, pointing to a successful self-selection of good answerers.

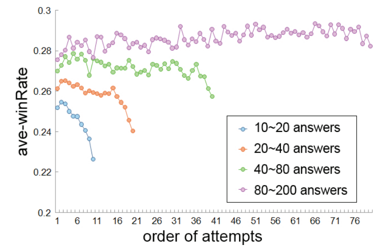

Figure 8: average winRate at each attempt for the three answerer groups

Answerers' effort: the answer length is a simple metric of answerers' effort. Figure 9 shows that answerers provide the longest answers initially (270 characters), but each subsequent answer is shorter, with the sixth answer being 240 characters long on average. From the sixth answer onward the answers gradually lengthen once more. Groups at all levels of activity present a similar pattern– suggesting that users learn to be more efficient in their answers.

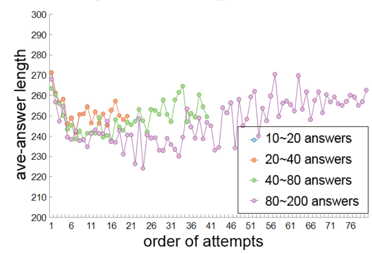

Figure 9: average answer length at each attempt for the three answerer groups

Award expectation: answerers may weigh the points offered for a best answer to a question against their probability of providing the best answer. We observe across activity levels a quick dive in the points a user attempts to gain from the average of 20 points on the first attempt to a lower but stable 15 points by the 5-6th attempt.

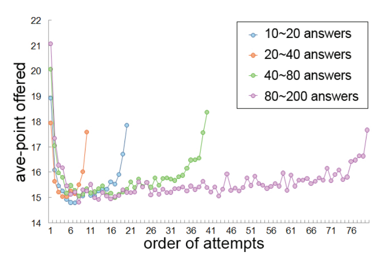

Figure 10: average point award per question offered at each attempt for the three answerer groups

Experience and performance

As we discussed above, there may be a positive reinforcement between experience (the number of answers provided by a user) and their performance (winRate). If users perform well, this might encourage them to participate more, which results in gaining more experience. We consider that a successful system should be able to sustain this kind of positive reinforcement processes for contributors to gain reward, experience, and expertise over time. Indeed, we find that more active answerers perform better: in terms of winRate (R=0.16, p<0.0001)[2], average award obtained for each question attempted (R=0.06, p<0.0001), and Guru Score[3] (R=0.10, p<0.0001). Figure 11 and Figure 12 present the collective patterns among answerer groups that answered 40~100, 100~200…questions: more active users consistently perform better than less active ones.

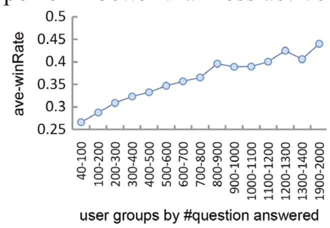

Figure 11: average winRate by answerer groups

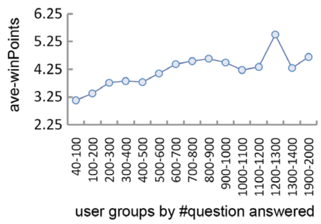

Figure 12: average points won per attempt by answerer groups

Part of the reason experienced users achieve a higher winRate is by selecting questions that offering fewer points (R=0.02, p<0.0001, Figure 13) and they thus face less competition (R=0.09, p<0.0001, Figure 14).

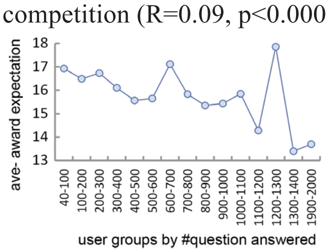

Figure 13: average award expected per question by answerer groups

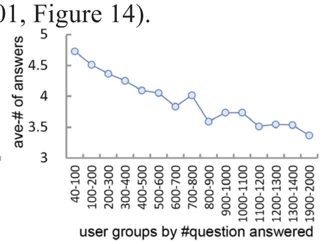

Figure 14: average number of answers per question by answerer groups

More active answerers put in more effort per answer and are more focused in providing knowledge/expertise. In particular, we use answer length (R=.06, p<0.0001, Figure

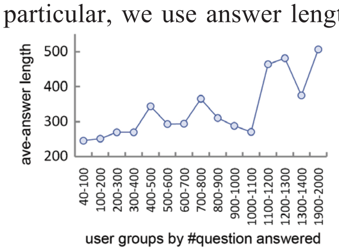

Figure 15: average answer length per question per user by answerer groups

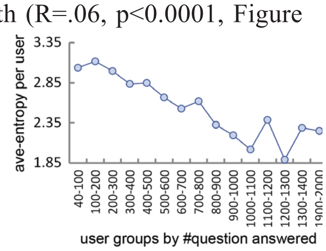

Figure 16: average entropy per user by answerer groups

[2] Correlations are calculated on answerers who have answered at least 40 questions during the period of time.

15) to measure effort and users' entropy (R=-.04, p<0.0001, Figure 16) to measure how an answerer is focused on particular domain/s[4].

Predicting Answerers' Performance

Based on above knowledge, we anticipate predicting answerers' performance by combing all the aforementioned behavior metrics. We found that the price of the questions that an answerer chose (+), competitiveness of the questions (-), answer length (+), and the focus across categories (+) can account for around half of the variance of one's performance. The prediction power slightly increases for more frequent answerers (e.g., for answerers of 100~200 questions, R^2=.54; and for answerers of 500~1000 questions, R^2=.60). In particular, the ability to choose less competitive questions is directly related to the performance (R=-.73, p<0.0001) while there is little correlation between the winRate and award per question (p=.104). In addition, answer length and focus also contribute to better performance (R=.38, p<0.0001; R=-.18, p<0.0001).[5]

Diversity of Askers

It is also important to know how askers ask questions as we hope they can continually contribute questions of good quality. Unlike answerers who answer 12 questions per user, asking activity is more spread out over a larger asker population: on average, each asker has only asked 2.4 questions and the most frequent asker has asked 1033 questions during the period. Similarly, we group askers into different groups according to the number of questions they have asked.

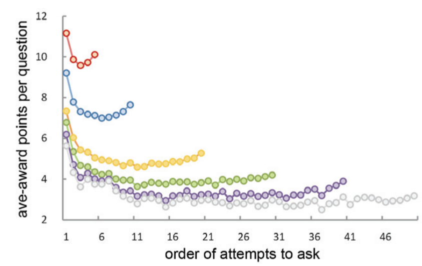

Figure 17: average point award per question offers at each attempt to ask for different asker groups

Figure 17 shows different asker groups (asked more than 5, 10, ..., 50 questions) change the average amount of points awarded for each question by asking order. Askers pay high for the first question and the price drops quickly within the first three questions. There could be two implications here: first, this is an adjustment process where askers learn about a proper price for asking a question; secondly, the first questions may be the trigger for people

[3] See Nam et al. (2009). The Guru score takes into account the odds of winning the best answer

[4] Entropy: see Adamic et al. 2008

[5] Note: correlations are based on answerers of 500-1000; other groups show similar pattern.

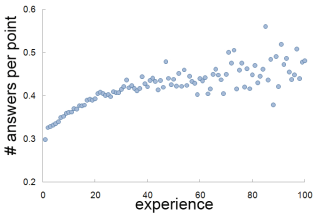

Figure 18: average number of answers per point offered

to start using the site, when people are urgently looking for answer for a particular question.

In addition, frequent askers pay less per question than less-frequent askers. This pattern is consistent; people who only ask a couple of questions they pay on average 14 points per question and for those who ask more than 50 questions, the average price becomes less than 4 points per question. Since this result might be confounded because newcomers join and old-timers leave during the time, we elicit a small portion of users and try to exclude new joiners as much as possible: we looked at the first 50,000 and 100,000 users who asked in the dataset and compare between the subgroups: the one only asked once and others during the whole period. Similarly, the one-time askers pay significantly higher than other users. However, this group of askers has stopped asking not because they did not get enough answers (actually they obtained more answers); thereby implying that askers all have various expectations and incentives of using the site: some only come to ask important questions and are willing to pay higher while some like to hang around more and more actively participate in the community.

In addition, we observe a slight trend that more experienced askers get a higher number of answers per point offered (R=0.005, p<0.0001). As shown in Figure 18, although askers offer smaller award, they actually improve the efficiency of each point in terms of buying participation.

Core Users, Who both Asked and Answered

The Most Active Group

Now we turn to the most active user population on the site: users who both ask and answer questions. This group of 597,297 users comprises 22.6% of the total users who participated on the site during the period of the dataset. And we call them DoBoth users.

- DoBoth users are more active than users who only ask or only answer: they asked almost half of the total questions with an average of 4.1 questions per user; which is significantly more than the group who only asked. In addition, they answered more than the group who only answered (the mean of DoBoth is 15.3 while purely answering users have a mean of 9.2).
- DoBoth users offer higher award when asking (the mean of the award points is 12.3, which is significantly higher than average). They share the same trend in terms of paying points for each question with the general askers; however, they pay higher each time.
- DoBoth users' winRate falls below that of users who only answer; their answers are shorter (mean=258; compared to 296) and they choose less challenging questions (award and number of competing answers for the question) (mean=3.8; compared to 3.9). This suggests that users who only answer may on average be selective in the questions they choose to answer.

From the observation that those who ask more tend to answer more (log#ask to log#answer, R=.26, p<0.0001) and similarly that those who spend more points also earn more (log#point-earned to log#point-spent, R=.19, p<0.0001); we may surmise that DoBoth users are incentivized to answer questions by the fact that they also need points to ask them. This group of users participates intensively and forms a sustainable core dynamic of traders in expertise.

Community across Categories

Consequently, it is important to examine how this dynamic takes place. We construct a users' social network by the help links from asker to answerer and we employ Bowtie analysis (Broder et al. 2000) to learn how users are connected through asking and answering interactions. The large strongly connected component (LSCC) presents the biggest subgroup of users who can reach one another through directed help links. For all pairs (A, B) of users in the LSCC, even if A did not directly help B, A helped someone, who helped someone, … who helped B. For all users on the site the LSCC is 16%, which is similar to the online Java forum community as observed in Zhang, Ackerman, and Adamic, (2007). This suggests that even without an explicit platform for threaded community interactions (e.g., in online forums, users can discuss and reply to one another back and forth), BK presents a connected community where people interact socially through asking and answering questions. In particular, the DoBoth user group contributes the most to maintaining the core of the community.

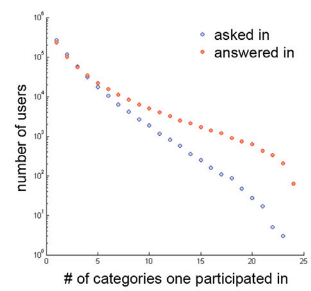

Figure 19: user distribution in terms of the number of categories they asked and answered in

However, Bowtie analysis on individual categories presents much smaller LSCCs ranging from 0.05% to 7.7%. This suggests that rather than only asking and answering in the same category, users participate across categories. They may answer in categories where they have expertise and ask in those where they don't. In general, DoBoth users answered more than asked, and so covered a greater number of categories by answering (mean=2.9) than by asking (mean=2.1). However, if we normalize the number of categories by the number of questions they have asked or answered, this relationship reverses: users cover a mean of 0.81 categories per question, and 0.56 categories per answer given. Finally, for the subset of 24,094 users who asked exactly as often as they answered; the averages are (0.76 versus 0.64). This all points to there being more subjects that individuals need help on, than subjects where they are expert. The power of the Q&A forums is that *collectively*, the users have expertise in all areas.

Category Concentration

Given that many users participate in multiple categories, we were interested in whether some categories more focused users than others. We use "concentration ratio" which is defined as the number of questions in one category divided by all questions one has asked/answered. For example, if a user asked 10 questions in "food" and she has asked 100 questions in total, then her ratio for asking in this category is 10%.

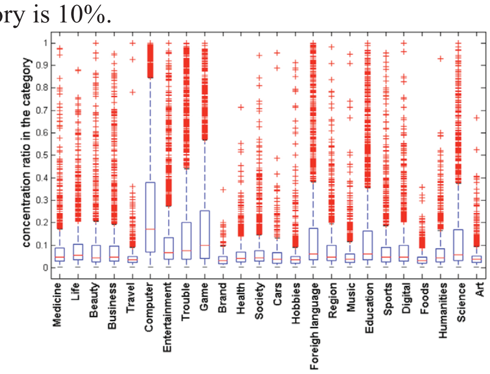

Figure 20: users ask in categories

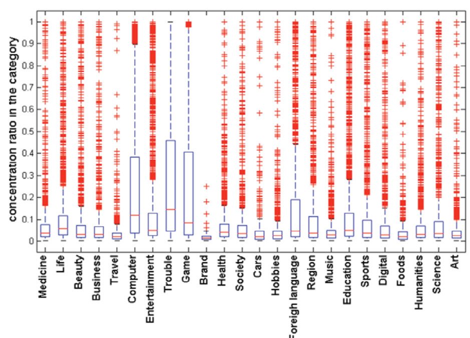

Figure 21: users answer in categories

Overall, users have highly skewed distribution in each category as many other sites. We can also see a difference among categories: for example, the "computer" and "game" categories gather the highest concentration and a few users only ask/answer within these categories; while in "travel" and "food" users tend to just visit shortly. This implies people's various information needs and where user would largely interact with similar people and where they would potentially meet more diverse others.

Comparing concentration distributions for asking and answering (Figures 20 and 21), answering patterns present higher concentration in general; and we see more highly focused answerers in each category too.

Conclusions and future work

In this paper, we studied a large scale Q&A system, Baidu Knows, in order to determine how such a system is maintaining and thriving. We find that the system has successfully accommodated people's various information needs and multiple levels of participation. In particular, there is a positive incentivizing cycle for users to keep participating and improving: you put more efforts, you win, you are rewarded, and you learn. There is also a core of generalized reciprocity-a large fraction of users are tied through indirect helping relationships, and these ties cross categories. As such, the system has been able to successfully exploit the idea of "exchanging" knowledge among distributed experts and the "sense of community" reinforces people's social bonds on the site; thus demonstrating a sustainably working mechanism.

The growing popularity of peer-based knowledge sites has attracted considerable research interests in recent years. Studies have found that users' participation and contribution is highly skewed on various instants of online communities such like Yahoo! Answers, Wikipedia (Adler and Alfaro 2007), Del.icio.us (Golder and Huberman 2006), and Flickr (Marlow et al. 2006). In addition, a large portion of contribution is made by a small minority of the participants and this group of users usually have better performance as in Welser et al. (2007), Adler and Alfaro (2007) and Yang et al. (2008). Participation structure on BK also shares this pattern in terms of skewness. However, there is a core user group who is not extreme on either asking or answering, nor do they necessarily perform better, contributes the most to the site.

This group of users is essentially motivated by the need of points to ask questions. Consistent with previous studies on monetary incentives, e.g., in Harper et al. (2008), Yang et al. (2008), we found that the virtual points can significantly incentivize answerers too. We also attribute this in part to the importance of having a high titled identity in the community, which can be achieved through accumulating points; especially as we see the only-answerers seek to answer high-awarded questions. How this title system plays the role of incentivizing contribution would be further studied in the future work.

In our previous study we investigated how users price the tasks to recruit solutions and we found that the price correlates expertise required for completing the task (Yang et al. 2008). In the form of virtual points, askers on BK pay differently on different questions: there are category difference and sequence difference in terms of when the question is asked by the asker. We would further explore the properties that affect pricing in the future.

Another crucial mission for QA studies is finding experts and understanding users' behavior patterns. This has been a long line of literatures from discovering experts in organizational knowledge systems: e.g., in Streeter & Lochbaum (1988), Krulwich and Burkey (1996), McDonald and Ackerman (2000), to various semantic or graphic-based expertise inference algorithms on the Internet (Kautz et al. 1997; Campbell, et al. 2003; Zhang et al. 2007). However, contributors' behavior pattern has been less explored in online QA communities. Wenger (1998) discussed different roles in community and Welser et al. (2007) used "structural signature" to distinguish "answerers" in online discussion forums; Holloway et al. (2007) and Viégas, et al. (2004) examined various contribution and collaboration patterns in Wikipedia. More recently on QA communities, Nam et al.'s (2009) study on Naver explored top answerers' motivations and intermittent participation patterns, and suggested that higher levels of participation correlate with better performance; our previous study on Taskcn.com presents users' interesting learning patterns and difference among groups of users (Yang et al. 2008). In current paper, we also observe users' initial adaptive behaviors and we find that users present different behavior patterns according to their activity level. Here we should note that we only captured "approximate" initial behaviors, as the dataset is not from the beginning of the site. Furthermore, we examine the special core user group-DoBoth group, which has been little investigated in literature to our knowledge.

In our future work, we also hope to infer cultural difference concerning QA system designs. Baidu Knows presents another successful QA instance with its own interesting characteristics. For example, we would suspect how such a title system would work on an English site; and whether askers would compensate prompt answerers when offering smaller award: all these indicate interesting directions of future studies.

Acknowledgements

We would like to thank Lada Adamic and Mark Ackerman for insightful discussion and review. This work was funded in part by Intel Research, the Army Research Institute, and the National Science Foundation (grant NSF 0325347, NSF IIS-0746646).

References

Adamic, L. A., Zhang, J., Bakshy, E., & Ackerman, M. S. 2008. Everyone knows something: Examining knowledge sharing on Yahoo Answers, Proceedings of the 17th international conference on World Wide Web. Beijing.

Adler, B. T., & Alfaro, L. d. 2007. A Content-Driven Reputation System for the Wikipedia, International World Wide Web Conference (pp. 8-12).

Broder, A., Kumar, R., Maghoul, F., Raghavan, P., Rajagopalan, S., Stata, R., et al. 2000. Graph structure in the Web. Computer Networks, 33, 309-320.

Campbell, C. S., Maglio, P. P., Cozzi, A., & Dom, B. 2003. Expertise identification using email communications, Proceedings of the twelfth international conference on Information and knowledge management (pp. 528-531). New Orleans, LA, USA: ACM.

Chen, Y., Ho, T., & Kim, Y. 2007. Knowledge Market Design: A Field Experiment at Google Answers.

Golder, S. A., & Huberman, B. A. 2006. Usage patterns of collaborative tagging systems. Journal of Information Science, 32, 198.

Harper, F. M., Daphne, R., Sheizaf, R., & Joseph, A. K. 2008. Predictors of answer quality in online Q&A sites. Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems.

Holloway, T., Bozicevic, M., & Börner, K. 2007. Analyzing and visualizing the semantic coverage of Wikipedia and its authors: Research Articles. Complex., 12, 30-40.

Kautz, H., Selman, B., & Shah, M. 1997. Referral Web: combining social networks and collaborative filtering. Commun. ACM, 40, 63-65.

Krulwich, B., & Burkey, C. 1996. ContactFinder agent: answering bulletin board questions with referrals, The 13 th National Conference on Artificial Intelligence (Vol. Part 1, pp. 10-15): AAAI 96.

Marlow, C., Naaman, M., Boyd, D., & Davis, M. 2006. HT06, Tagging Paper, Taxonomy, Flickr, Academic Article, ToRead, the seventeenth Conference on Hypertext and Hypermedia (pp. 31-40).

McDonald, D. W., & Ackerman, M. S. 2000. Expertise recommender: a flexible recommendation system and architecture, ACM conference on Computer supported cooperative work (pp. 231-140).

Nam, K., Adamic, L. A., & Ackerman, M. S. 2009. Questions in, Knowledge iN? A Study of Naver's Question Answering Community, the twenty-seventh annual SIGCHI conference on Human factors in computing systems. Boston: ACM.

Streeter, L. A., & Lochbaum, K. E. 1988. An expert/expert-locating system based on automatic representation of semantic structure. Proceedings of the Fourth Conference on Artificial Intelligence Applications.

Viégas, F. B., Wattenberg, M., & Dave, K. 2004. Studying cooperation and conflict between authors with history flow visualizations, SIGCHI conference on Human factors in computing systems (pp. 575-582). Vienna, Austria.

Welser, H. T., Gleave, E., Fisher, D., & Smith, M. 2007. Visualizing the Signatures of Social Roles in Online Discussion Groups. Journal of Social Structure.

Wenger, E. 1998. Communities of Practice: Learning, Meaning, and Identity. Cambridge: Cambridge University Press.

Yang, J., Adamic, L. A., & Ackerman, M. S. 2008. Crowdsourcing and Knowledge Sharing: Strategic User Behavior on Taskcn. Proceedings of the 8th ACM conference on Electronic commerce.

Zhang, J., Ackerman, M. S., & Adamic, L. A. 2007. Expertise networks in online communities: structure and algorithms. Proceedings of the 16th international conference on World Wide Web.

Zlatic, V., Bozicevic, M., Stefancic, H., & Domazet, M. 2006. Wikipedias: Collaborative web-based encyclopedias as complex networks. Physical Review E, 74.