Incentive and Culture: Shaping Information and Social Dynamics in Online Information Sharing Systems

by

Jiang Yang

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Information)
In The University of Michigan
2011

Doctoral Committee:

Professor Mark S. Ackerman, Co-Chair
Associate Professor Lada A. Adamic, Co-Chair
Professor Yan Chen
Assistant Professor Kristen R. LeFevre

Acknowledgements

Studying the Internet is an exciting, challenging, and rewarding experience. I would first like to gratefully and sincerely thank my great advisor team: Prof. Mark Ackerman, Prof. Lada Adamic, and Prof. George Furnas for bringing me in, supporting me, and guiding the way throughout this adventure. This great advisor team not only offered me the best all-around and insightful training in understanding human-computer interaction in the Internet world, but also were great role models of academic passion, integrity, and expertise. Of course, I have also enjoyed having you as friends for the past 5 years, as well as well on down the road.

Prof. Yan Chen and Prof. Kristen LeFevre deserve special thanks as my thesis committee members. I highly appreciate your constructive and insightful suggestions and discussions, and I appreciate your being flexible and considerate with my dissertation schedule.

I have been lucky to have all passionate, intelligent, and industrious collaborators, including my advisors, Tracy Xiao Liu, Prof. Yan Chen, and all my internship mentors. You made our collaboration very productive and enjoyable. In particular, I want to express my greatest appreciation to my collaborators from industry labs for offering wonderful opportunities and experiences for my research: Zhen Wen and Ching-Yung Lin from IBM T.J. Watson Research Center, Scott Counts, Jaime Teevan, and Meredith Morris from Microsoft Research.

I would thank the University of Michigan School of Information, the place that nurtured me and supported me for five years. I will be proud to think of her as my second home. I thank Suzanne Schuon, Judith Olson, Stephanie Teasley, and Yvonne Perhne for their great help throughout my entire journey. I thank all colleagues (and friends) who have shared with me their expertise, inspiration, and various kinds of help in my studies, research, and life. To name a few: Xiaomu Zhou, Tracy Xiao Liu, Lian Jian, Eytan Bakshy, Liuling Gong, Beth St Jean, Sean Munson, Edwin Teng, Tao Dong, Jun Zhang, Jina Huh, Kevin Nam, Mark Newman, and Eytan Adar.

I am very grateful for the Human-Computer Interaction (HCI) and Social Computing Community, which have been very open, resourceful, and supportive for my academic growth. In particular, I would like to thank ICWSM for providing me invaluable experience as a program committee member and other conferences and journals such as CHI, CSCW, WWW, for letting me be a reviewer and interact with the community. In addition, it was a great honor to have been supported by multiple National Science Foundation grants: NSF 0325347, NSF IIS-0746646, and NSF IIS-0948639.

My deepest love and gratitude goes to my family. I could not make this achievement without the trust, encouragement, and unconditional love of my mother, father, and sister. I also want to thank Watson, my little dog, for being sweet company and bringing endless happiness.

Finally, I want to thank my husband, Xiao Wei. There are simply no words better than to say that I cherish his love and support deep in my heart.

Table of Contents

List of Figures

List of Tables

Abstract

The most revolutionary power of the Internet lies in the way it changes people's collaborative work, aggregating social knowledge at an ever-increasing speed. This trend has been manifested in a variety of social information sharing and augmenting systems, such as Wikipedia, Question-and-Answer (Q&A) forums, and crowdsourcing websites. Understanding the information and social dynamics involved in these systems is crucial to improve their design and truly harness their power. This dissertation is devoted to investigating how two important factors, incentive and culture, significantly and interactively shaped users' information and social behavior in the information-sharing websites I studied.

This dissertation is organized by four interlinked studies, which address incentive design in online information sharing systems from four perspectives: how users learn and adapt their behavior to an incentive design dynamically; how users' adaptation dynamics contribute to a positive feedback mechanism that sustains the community; how culture deeply influences information and social dynamics, even given with very similar virtual point incentive designs and system platforms; and how incentive design can interact with a particular community structure and cultural context in a very comprehensive and complex way, and how the interaction can lead to a co-evolution process between the users and the way users perceive and use the incentive design.

Chapter 1

Overview

One of the most revolutionary potentials of the Internet is how it changes people's collaborative work. And nowhere is this more obvious than in sites that collect intellectual contributions from otherwise disparate and distributed peer users on a massive scale. This trend has manifested itself in various familiar examples such as open source projects, Wikipedia, Question-and-Answer (Q&A) forums, and social tagging sites such as Flickr and Del.icio.us. These collaborative sites are designed to support a variety of knowledge augmentation processes. For example, Wikipedia is building a massive encyclopedia by accumulating knowledge contributions from distributed peer experts. Question-and-Answer (Q&A) forums provide platforms for people to exchange information and knowledge on an ongoing basis. A more recent type of collaborative service is called "crowd-sourcing," in which intellectual tasks are directly outsourced to individual workers through public solicitation (Howe, 2006, 2008; Kleeman et al., 2008). Crowdsourcing sites have been growing fast in number, popularity, and research attention. For example, one of the earliest sites that have been studied, Taskcn.com, is using a competition mechanism to outsource diverse types of tasks, such as designing a company logo or translating a research statement. Amazon's Mechanical Turk (MTurk) is similarly designed to collect human labor to accomplish "human intelligence tasks" (HITs) requested by users, who pay workers a small fee (Mason & Watts, 2009).

More so than regular webpages, collaborative sites can easily flop and fail. Even a site by one of the Internet's most visible companies, Google Answers failed and exists no longer. Butler found that almost half of many social, hobby, and work mailing lists become silent after about 122 days (Butler, 2001). And among those still active, traffic was very low, with a median of one message every 3.6 days (Cummings et al., 2002). To stay alive, collaborative sites need to maintain a reasonable group size and enough participation over time. But that can be more challenging in online communities, where people have few strong ties and less commitment to the group than they do in offline settings. Even active online communities suffer from sparse participation. For example MovieLens (http://www.movielens.org), an online movie recommendation site, more than 22% of the movies on the site obtained fewer than 40 ratings. This defeats the purpose of the site, since it doesn't have enough data to make personalized recommendations and predictions (Chen, Harper, et al., 2010).

A wealth of literature has tried to understand what it takes to motivate people to join online groups or participate more (Joyce & Kraut, 2006). Empirical studies have suggested that there are both extrinsic reasons (such as gaining reputation, education, and money (Hertel et al., 2003; Lakhani & Hippel, 2003)) and intrinsic reasons (like the inherent pleasure of problem solving, altruism, and commitment to a community (Holohan & Garg, 2005; Rossi, 2004)). These people share similar motivations when participating in online Question-and-Answer (Q&A) communities. For example, on Korea's largest Q&A forum Naver, interviewees frequently report altruism, learning, and business motives as reasons for answering others' questions (Nam et al., 2009).

More often, sites build in incentive mechanisms to motivate people's participation. For example, Taskcn and Amazon's Mechanical Turk both allow requesters to pay contributors. Community-based Q&A forums, such as Yahoo!Answers (Adamic et al., 2008), Baidu Knows (Yang et al., 2010), and the now defunct Google Answers (Chen et al., 2010) have used incentive schemes ranging from flat-rate virtual currency in Yahoo! Answers, flexible-rate virtual currency in Baidu Knows and Naver Knowledge-In (Nam et al., 2009), to real-market schemes (in Google Answers). But incentives do not have to be money or points. For example, Chen et al. (2010) did a field experiment with Movielens and found that a social comparison design can also motivate people to contribute.

As one of the primary questions concerning both researchers and designers, whether and how these incentives can motivate more and better contributions has been the major question to assess these systems. For example, field experiments conducted on a series of Q&A sites have indicated that higher awards can induce more answers but yielded mixed results in quality of answers (Chen et al., 2010; Harper et al., 2008). In Amazon's Mechanical Turk, researchers found the consistent result that money increased the *quantity* of contributions, but not the quality (Mason & Watts, 2009). Field studies with Baidu Knows (Yang & Wei, 2009) and Naver Knowledge-In (Nam et al., 2009) also found that incentives increased the amount of answers.

Incentive systems are so prevalent in collaboration sites that it is almost shocking how little research has been done to understand and evaluate these systems. Most studies that have done so are limited because they measure "one-time" transactions: whether an individual award can bring more payoffs and payoffs of higher quality (answers on a Q&A site, for example). But system designers want to know about users' responses to an incentive design in a dynamic perspective, and how to encourage and sustain participation over time. For example, it is important to know whether users learn to adapt to the incentive design, whether the adaptation provides positive or negative feedback to the system, and which design features will encourage continuing participation.

People's learning and adaptive behaviors have often been observed in lab experiments, but very few studies have documented how learning works online. A study on eBay found that users learn over time to snipe, or submit their bids close to the end of the bidding period (Wilcox, 2000). In addition, experienced users are less likely to make multiple bids on the same item (for example increasing their limit once they see that they have been outbid). This is especially true for items that are easy to place an objective dollar value on, as opposed to items that are more based on personal preference.

Joyce and Kraut (2006) first investigated how online newsgroups sustain newcomers, who often face obstacles in their peripheral participation (Wenger, 1998). Across six different online news groups, the results showed that people's initial interaction experience (got-reply) can predict whether newcomers continue participating, while characteristics of initial posts are related to the characteristics of the replies they obtained. A similar study (Arguello et al., 2006) found that the chance of getting a reply varies across different topic groups, and that newcomers and old timers differ in their ability to get replies. This implies that sustainability depends on the particular system and the culture of its users. Lampe and Johnston (2005) explored another kind of online forum, Slashdot, which has a different conversation format and moderation rules. This study also found that newcomers' first interactions with the system (group) were important and that feedback from the system (moderation) is related to the time to post and score of the second post.

These studies all show that initial interactions are crucial to sustain new users, but these initial interaction dynamics vary across different users, domains, and systems, calling for more systematic and controlled investigations. In addition, these studies have all looked at the initial stage of participation—whether someone returns after their first visit. This is not sufficient to describe the variety in individuals' length of stay, level of involvement, and ways of participating. Moreover, systems evolve over time, so how user participation changes depending on whether the system is new or well established.

Arguello et al. (2006) found that interaction dynamics can be very different across online groups. But the interaction between incentive design and national culture has the potential to be much more unpredictable. The 400-million-strong Chinese market is a great example of this. When information systems have simply taken their product more or less "as is" into China, (e.g., eBay, Orkut, and Yahoo!), they have encountered serious challenges. To overcome these challenges and design better systems, we need to understand different cultures.

Cross-cultural psychologists and sociologists have found that Westerners and East Asians have fundamental differences not just in the content of their beliefs and ideas, but in the basic way they view things and process information. Psychologists have called the Western pattern “analytic,” and the East Asian pattern “holistic” (Nisbett et al., 2001). Across dozens of experiments, psychologists have shown that Westerners tend to focus more of their attention on central items rather than the context; they are more narrowly focused; they see items as individuals, rather than as bound up in a situation or in relationships with others; and they tend to use formal logic or rules of non-contradiction. East Asian thought is called “holistic” because they pay more attention to the entire visual field, including the background and the context; they see human behavior as tied up with that context; and they think more intuitively and dialectically, finding meaning and value in contradiction (Nisbett et al., 2001; Varnum et al., 2010). Some psychologists have argued that this difference can be traced back to these cultures’ social orientation: individualism and collectivism (Varnum et al., 2010). Western cultures value independence, individualism, autonomy, and self-achievement (Hofstede, 1980); in contrast, Asian cultures emphasize interdependence, harmony, relatedness, and connection (Hofstede, 1983; Singelis, 1994; Triandis, 1995). These differences interact with and shape deeply cultural things, such as value systems (Aristotelian vs. Confucian intellectual traditions)(Lloyd, 1996; Pye, 1985), languages (Varnum et al., 2010), religions (Dollinger, 1988), economic ideology (Ralston et al., 2007), and industrialization and geographic mobility (Kitayama et al., 2009).

These inherent cultural characteristics can significantly affect how people perceive and use a system/design, and how they interact with others in collaborative sites. In fact, I found that cultural factors predict people’s perception, preferences, and motivations in their social Q&A behavior more than other demographic variables (Yang et al., 2011), which suggests that cultural differences are an important variable to understand to design across groups and contexts.

Fortunately, there is a rapidly growing stock of research on cultural issues in HCI and CSCW (Setlock & Fussell, 2010). These studies have focused mainly on the adoption and usage of collaboration tools. For example, Asian users have been found to prefer multi-party chat, audio-video chat, and emoticons in IM (Kayan et al., 2006), benefit more from rich communication media in negotiation (Veinott et al., 1999), and be less satisfied with asynchronous communication (Massey et al., 2001). Setlock and Fussell (2010) found that Asian participants involve additional considerations when deciding on appropriate communication tools, especially the ability to support social processes.

But because incentive design is so complex—and because studying it across cultures is so difficult—empirical studies are rare. However, understanding this is crucial in designing cross-cultural collaborative systems. Simply put, more in-depth studies on the interaction between incentive and culture are therefore motivated.

Around this general theme, I conducted four interlinked studies of incentive design in online knowledge sharing systems. Each study takes a different perspective: (1) How users learn and adapt their behavior in response to incentives in Witkey websites. Witkey websites are an emerging knowledge market design in which an all-pay auction model is used to crowd-source peer expertise. (2) How adaption dynamics contribute to a positive feedback mechanism that sustains community in a Q&A site, Baidu Knows. (3) How information and social dynamics differ across different cultural contexts, even given with very similar virtual point incentive designs and system platforms. (4) How incentive design can interact with a specific cultural context.

In the first study (Chapter 2, published as Yang et al., 2008 b), I examined the behavior of users on one of the biggest Witkey websites in China, Taskcn. On Taskcn, people post diverse types of tasks (e.g., designing a company logo or translating a research statement) with a monetary reward for the reward to the person who has their solution selected. I found that users were learning and adopt strategies over time. Users tended to select tasks where they were competing against fewer opponents to increase their chances of winning. They also selected tasks with higher expected reward; these tended to be tasks that require a high skill level, but low work-load. There was a small portion of users who had won multiple times and were able to practice the strategies better than others, discovering less competitive tasks and submitting solutions in a later stage. Despite that, overall, users do not increase their chance of winning. But this small core of winners (0.12% of all submitters) managed to increase their win-to-submit ratio over time. Moreover, since most users quit after only a few submissions, this core group proposes nearly 20% of the winning solutions on the site and actually sustains the site behind the high traffic of casual participants. This study revealed patterns of how design causes users to develop strategies over time and how user groups with vastly different patterns of behavior all contribute to a site's dynamic.

The second study (Chapter 3, published as Yang & Wei, 2009) investigates user behavior in a large-scale knowledge sharing community, Baidu Knows. Askers and answerers on Baidu Knows had evolving behavior patterns that sustained the community. In particular, there is a positive feedback cycle: you put in effort, you win, you are rewarded, and you participate more. There is also a core of generalized reciprocity. A large fraction of users are tied through indirect helping relationships, including askers and answerers. In addition, the core group of users who both ask and answer are motivated by the virtual point system and make the majority of contribution to the site. As such, the system has been able to successfully exchange knowledge among distributed experts.

To further assess incentive design in a broader context, I conducted the third study (Chapter 4), a comprehensive analysis of users' activity lifespan and participation pattern across three predominant online knowledge-sharing communities: Yahoo!Answers, Baidu Knows, and Naver Knowledge-IN in English, Chinese, and Korean, respectively (originally published as Yang et al., 2010). These three community-based Question-Answering

(Q&A) sites share very similar virtual point incentive and system design, and are comparable in their history and scale. Extending previous work focusing on initial interactions of new users, I used survival analysis to quantify participation patterns that can be used to predict individual lifespan over the long term. Across all three sites, users who prefer answering tend to stay longer and they are also sensitive to the initial experience on the site. In addition, users' first-month experience can account for a considerable amount of variance of predicted lifespan. In particular, users' self-selection effect (whether a user is active or what type of role he/she likes to play) and performance in the community account for the most variance in the prediction. Despite these similarities, there were very intriguing differences between the sites: answerers tend to be more active in providing answers in Yahoo! Answers than the other two sites, and the question-answering dynamics on Yahoo!Answers tend to be more conversational than Baidu Knows. This might be explained by a complex interaction between the incentive design and cultural contexts. Furthermore, a longitudinal comparison of the communities' evolution between two distinct stages suggests that users' commitment or a site's ability to sustain users can evolve over the different life stages of a system.

Because the interaction between incentive design and culture is so complex, the fourth study (Chapter 5) provides an in-depth picture of the social motivations and dynamics of Mitbbs.com (published as Yang, Ackerman, and Adamic, 2011). Mitbbs is a thriving Chinese online forum for Chinese people located overseas to share information and sustain their virtual bond of common identity. Based on over 4 years of observation and data collection, this study shows how a virtual currency system—not unlike those used by many American websites—has evolved into an essential medium for extremely diverse and culturally specific social exchange activities. The social interactions reflect the traditional Chinese idea of *guanxi*, or interpersonal influence and connectedness, while at the same time incorporating the norms of a new generation of Internet users. This study demonstrates how incentive design can interact with a particular culture and how that pushes along how users perceive and use the incentive design.

These four studies take four different angles to the question of how people's behavior on online information sharing systems is shaped by incentive design and culture. They together provide the first comprehensive and in-depth understanding of the dynamic perspective of these systems: users learn and adapt to incentive designs; this process interacts with their particular cultural characteristics and context; and the way users perceive and use these systems evolve over time. These four studies approach this problem with quantitative measures, comparative study, statistical modeling, and field ethnographic investigation. After explaining the studies, Chapter 5 will address the theoretical and design implications, as well as important future work still being done.

Chapter 2

Learning and Strategic Behavior in a Crowdsourcing Site

Crowdsourcing, or the use of an Internet-scale community to outsource a task, has garnered considerable interest in the popular press. Articles in Wired (Howe, 2006) and Business Week (Hempel, 2006), for example, repeat the same success stories for video, stock photography, and even corporate R&D. However, the media coverage consists primarily of anecdotal evidence in an often relentlessly enthusiastic manner. Empirically-based analytical studies of crowd sourcing sites are, unfortunately, lacking.

This chapter presents one such study. It analyzes use of a Witkey site, Taskcn.com, where users offer monetary awards for solutions to problems. Other users provide solutions in the hopes of winning the awards. Taskcn has 1.7 million registered users. Users have requested solutions for nearly 3,100 tasks, and 543,000 solutions have been proposed all in less than two years.

It might appear that the site should be drowning in newbies and lurkers; yet, the site appears to be quite successful. Askers clearly get solutions. More interestingly, the site appears to be socially stable, there is a core of users who repeatedly propose and win. The large numbers of new users ensure many answers, while also providing new members for the stable core. The data from this study will show that crowdsourcing works, albeit perhaps only as long as it is a popular phenomenon.

In this chapter I focus on three important aspects of this expertise-sharing marketplace. The first is whether tasks are priced according to the expertise and effort level required. The second is the set of factors involved in strategic selection, and whether users learn to better their chances of winning over time. The third is what distinguishes the successful users from the unsuccessful users over time. All of these are important to maintain the site as an ongoing and successful marketplace.

The chapter proceeds as follows: First we introduce the literature background on knowledge market, and Witkeys and Taskcn in particular. We then talk about our data collection. This is followed by a discussion of our findings about pricing, strategic selection, and winners. We conclude with design implications and future work.

1. Literature Background

Knowledge Market

A variety of collaborative sites have been designed to collect intellectual contributions from distributed peer users of large scale. These sites accumulate many types of knowledge, ranging from the pieces of information or knowledge exchanged in Question-and-Answer (Q&A) sites, aggregated meta knowledge about information items such as social tags on Flickr and Del.icio.us and recommendations on Netflix, to structured knowledge repositories like Wikipedia. “Crowd-sourcing” sites collect another type of knowledge in which well-defined tasks are outsourced to individual workers through public solicitation (Howe, 2006, 2008; Kleeman et al., 2008). Crowdsourcing sites have been growing fast in number, popularity, and research attention. For example, one of the earliest sites that have been studied, Taskcn.com, is using a competition mechanism to outsource diverse types of tasks, such as designing a company logo or translating a research statement. Amazon’s Mechanical Turk (MTurk) is similarly designed to collect human labor to accomplish “human intelligence tasks” (HITs) requested by users, who pay workers a small fee (Mason & Watts, 2009).

Unlike Wikipedia or Flickr, crowdsourcing sites are task-driven with arbitrary requirements (or expectations) for completion time, quality, and other features. These more defined "tasks" tend to fall outside of the intrinsic motivation or some form of social reward (Nov et al., 2008), compared to those free-structured and non-defined contribution tasks. Therefore, financial incentives have been increasingly involved in designing crowdsourcing services. For example, Taskcn and Amazon's Mechanical Turk both allow requesters to pay contributors. Q&A forums offer virtual currencies and even real money.

Economists and sociologists are interested in understanding how incentives can be used to motivate contributions to these systems, and several past studies have focused on this question. For example, field experiments conducted on a series of Q&A sites have found that higher awards induce more answers, but yielded mixed results in quality of answers (Chen et al., 2010; Harper et al., 2008). A study on Amazon's Mechanical Turk found that money consistently increased the quantity of contribution, but not the quality (Mason & Watts, 2009). Field studies on Baidu Knows (Yang & Wei, 2009) and Naver Knowledge-In (Nam et al., 2009) found that virtual points can also motivate more of answer contributions.

These studies, however, are limited to evaluating "one-time" transactions. It is also important to know how users dynamically respond to incentive design and how online systems evolve over time.

User Behavior in Knowledge Market

A number of studies have looked into how people participate in online knowledge markets. Hiltz and Turoff (1993) first found that the distribution of contribution is extremely skewed. A few people contribute a lot, while everyone else only contributes a little. This has been found in a large number of computer-mediated communication (CMC) systems and online settings, for example, in Wikipedia (Kittur et al., 2007) and online social media such as Del.icio.us (Golder & Huberman, 2006) and FlickR (Marlow et al., 2006). In Wikipedia, a small subset of authors make a larger fraction of the edits, but their edits have greater longevity (Adler & Alfaro, 2007). In contrast, users who casually contribute content have a higher rate of bad edits that are quickly reversed. Despite these problems, quality remains high. For example, Wikipedia has been shown to be close in accuracy to Encyclopedia Britannica (Giles, 2005). Interestingly, that may change in the future because of a recent trend whereby Wikipedia and Del.icio.us are seeing more and more contributions from non-elite users (Kittur et al., 2007).

Other recent work has examined the dynamics of knowledge-sharing systems, such as technical forums and community-based question-answering sites. For example, a study investigating Java Forum hosted by Sun found that the most active users also tend to be the most expert, and they are likely to answer both newbie questions and technical questions (Zhang et al., 2007). Another study found very diverse knowledge and expertise sharing on Yahoo!Answers, where many questions are prompts for discussion or support, rather than pure information-seeking (Adamic et al., 2008). In addition, in top-level categories, such as science and math, where most questions are of a factual nature, specializing within a subcategory correlates with a higher proportion of "winning" answers.

But real-world studies of crowdsourcing sites are few and far between. Most studies of strategic and learning behaviors have been carried out in the lab. Even among the few real-world studies, most studies have been restricted to online auctions. Researchers have found that eBay users learn over time to snipe, or submit their bids close to the end of the bidding period (Wilcox, 2000). Although the timing of the bid should not matter if all players are rational and submit their true valuation, early bidding can prompt irrational bidders to up their bid. It is therefore advantageous to submit one's bid later, and indeed, 13% of the bids on eBay occur in the last 5 minutes of the auction. Wilcox (2000) also found that more experienced users are less likely to submit multiple bids on the same item, especially for items with a large common value component.

2. Taskcn as A Knowledge Market

A Witkey website is a new type of knowledge market website, in which users offer monetary awards for a question or task, and other users provide solutions to compete for the award. The website plays the role of the trusted third party by collecting the money from the requester and distributing the award to the winner(s), whom the requester chooses. The website takes a small portion of the award as a service fee.

The term "Witkey" was coined by the founder of the website Witkey.com[1] in 2005, and it became the name of a series of similar websites in China. This business pattern has quickly motivated a number of followers: in the last two years, more than 10 Witkey websites have been launched (e.g., Witkey.com, Taskcn.com, zhubajie.com, and k68.cn). Within a relatively short time, the Witkey model has demonstrated its capability to gather people to share knowledge. Taskcn is one of the biggest Witkey websites in China, and we analyze it here. It had 1,691,404 users registered between June 2006 and December 2007. "Witkey" is a very popular phenomenon in China, and there are many more sites in addition to Taskcn and Witkey.com. For example, on k68.com, 936,462 users participated in at least one task from July 2004 to January 2008;[2] and zhubajie.com claims to have added 497,169 users from its launch date of December 2005 to January 2008.[3]

Witkey websites can be seen as harbingers of the freelance markets that were forecast in Malone's "The Future of Work" (Malone, 2004). Witkeys differ from open question answer forums such as Yahoo! Answers, because instead of questions that are answered by other users without payment, the requesters offer awards for completion of tasks they pose. Witkeys also differ from the (now defunct) Google Answers, which while allowing

---

[1] http://www.witkey.com/lfarticle/articledt.asp?aid=20000

[2] http://www.k68.cn/

[3] http://www.zhubajie.com/info/about/

requesters to offer rewards for answered questions, limited the participation of those competing to answer to a few (<500) vetted individuals. In contrast, Witkeys seem to foster a new and completely open way to share more complex knowledge among individuals of distributed expertise (Yang et al., 2008 a). On Witkey websites, tasks usually require particular expertise and include a moderate investment of effort on behalf of the task-solvers. For example, many companies are looking for logo designs, a task that requires solution providers to have particular design expertise. In addition, to complete the tasks the solution providers need to take efforts to learn about the companies. The mechanism provides an incentive: the potential monetary award that encourages people's participation.

Because users submit their work directly, and concurrently with other users competing on the same task, they have little guarantee that their work will receive a monetary reward. This is different from Google Answers, which recruited a small number of expert answerers. Google Answers' answerers would select tasks to complete, and would have an exclusive lock on the task for a period of time. Witkeys also differ from sites such as eLance and TopCoder, where requesters pose a task and anyone can submit their credentials and proposals, but the task is not attempted until after the requester chooses a person or team to complete it.

Witkeys therefore occupy an interesting position in the design space of online knowledge -sharing sites. The tasks must necessarily be of relatively low complexity and effort, since a user has no guarantee of collecting the award before submitting their solution. However, this also poses a distinct advantage to the requester, as they are able to choose among several possible solutions to their task. The financial component also encourages users to contribute expertise beyond simple question answering, as might occur on no-fee sites such as Yahoo! Answers. For instance, some tasks ask for professionals to develop websites: there are relatively fewer individuals who have this particular expertise than there are people who would be able to answer a typical question on Yahoo Answers such as: "what is good facial cleanser for acne?" Therefore,

Figure 1: Home page of Taskcn.com (snapshot on Jan 3, 2008).

the monetary-award competitive mechanism aims to attract people with some expertise, but for relatively non-complex tasks.

Taskcn.com

Taskcn.com, the Witkey website we selected for our study, is one of the biggest Witkey websites in China. Taskcn.com has had slightly over 3,100 posted tasks categorized into 7 different types: design, strategy planning, programming, personal service, website, and “others”. Around half of the all tasks are in the design category: 1412 design tasks out of a total of 3112.

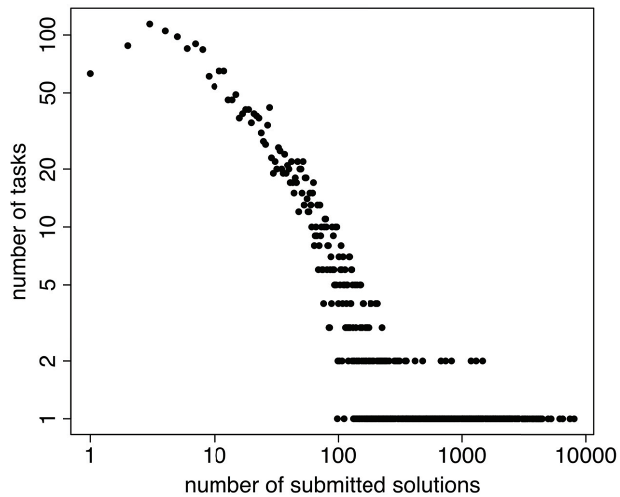

Figure 2: The distribution of the number of solutions posted to the different tasks.

Figure 1 is a snapshot of the front page of the website taken on January 3, 2008. This page contains categorized task lists (by post-date, by number of views, by award etc.), task categories, and help/tutorial information. The tasks vary in topic and amount of award. For example, there are tasks that offer 2 yuan (at 7.5 yuan/dollar) and also at the same time a task that offers 4000 yuan. The task contents range from logo design and website development to business plan writing. On the top of the page, there is an instant status-updater: the number of tasks that have been submitted and announced, the award that has been offered, the number of registered users, and the number of users currently online. These numbers all show that the website is a very active community of knowledge exchange.

Taskcn.com exhibits potentially excessive entry, where decentralized participation creates market inefficiencies since many entrants race to compete for a single prize but ignore the negative externality on other participants of their entry (Dasgupta & Stiglitz, 1980): Out of 185,429 users (11% of all registered users) who have participated at least once, there are only 5953 (3.2%) who have ever won. Yet, for all the completed tasks, each task has an average of 184 competitors, with each user being able to submit only once. Note, however, the distribution of submissions is heavily skewed as shown Figure 2. This variable, namely the number of competitors in a task, is the single strongest factor determining a user's probability of winning, exceeding even the individual's past winning record (Yang et al., 2008 a). This result in particular suggests that intentionally choosing less popular tasks to participate could potentially enhance winning probabilities, even if one's own expertise remains the same. In section 4 of the paper, we will see that, as they gain experience, users in fact do choose less crowded tasks.

3. Data Set and Approaches

The data include all 3112 tasks that were completed (i.e., the task is closed and winner is decided) from June 2006 (Taskcn's launch date) to December 2007. A task has these basic variables:

Table 1: Variables of tasks on Taskcn

Variable Name	Description
Task Start-Time	When a task is posted and competition starts
Task End-Time	The deadline for submission
Task Period	The period between Start-time and End-time, counted by day
Task Type	The type the task has been categorized into
Task Award	The amount of monetary award the task offers
# of Registrations	The number of users who registered to participate in the task
# of Submissions	The number of users who submitted solutions for the task
# of Winners	Some tasks can have more than one winner, this variable is obtained by the actual result of the competition

We also collected activity data on all users who had participated at least once in these tasks and excluded those who had registered on the website but had never contributed.

This yielded a total of 185,429 users with data on the number of submissions made (excluding tasks for which they registered but did not submit a solution), the number of wins, and initial date they registered on the website. Although Taskcn.com requires users to provide their real names, we excluded all identifiable information in this study. We also collected data on the interaction between a user and a task, such as the time of submission and whether or not the user won in the task.

Basic Participation Pattern

To understand the expertise sharing on Taskcn, it is important to understand the strategies users employ in task selection, since users can choose the tasks they want to compete. All ongoing tasks are listed in many ways such that users can browse: e.g., by categories or by award range; and all tasks are listed by recency. In addition, when one is exploring a particular task page, there are several similar tasks listed aside: for example, next to a task requesting a company logo design, there will be several other recent logo design tasks listed.

First, we look at how users behave in participating in a task. Users can view a task, place a task in their profiles, register to participate in the task, and submit a solution to the task.

A higher monetary award will result in significantly more views of the task page. After log transforming the variables (due to the skew in distribution of both award amount and number of views), we find a high positive correlation ($\rho = 0.64$, sig. $<10^{-4}$). Users will also be more likely to register for the task upon viewing it if the award is higher ($\rho = 0.60$, sig. $<10^{-3}$), but there is a lower correlation with the number of solutions submitted ($\rho = 0.43$, sig. $<10^{-4}$). If we consider the difference between the number of registrants and final submissions as an indicator of how often people gave up on their submission, we find that higher money award is correlated with a higher percentage of users who give up ($\rho = 0.37$, sig. $<10^{-4}$). This result suggests that monetary award can draw users' attention and even intent; however, there are other factors that affect the final number of submissions and quality. One possible explanation is in the result from the following section, where we found that higher award tasks tend to require a higher skill level. Users may initially register for the task but give up once they find their skills inadequate.

Tasks, Effort, and Award

In addition to the award offered by each task, it is important to understand the relationship between other task properties and users' participation, particularly with regard to the effects of pricing. To do this, we employed human-coding to rate the implicit properties of tasks. We used two raters (professional designers) to evaluate 157 randomly selected tasks in the design category (10% of the total tasks in the category) in terms of following task dimensions:

Table 2: Human-coded variables for design tasks

Dimension	Definition	Inter-rater reliability
Skill requirement	The lowest professional skill required for completing the task	Spearman's rho = 0.38**
Workload	The time an average person of the required skill level will take to complete the task	Spearman's rho = 0.40**

The raters evaluated the tasks without knowing the amount of reward that had been offered, so did not obtain cues as to the value of the task from the price. For this sample of tasks, the inter-rater reliability is relatively low. However, note that we are then using these scores (the average of the ratings given by the two raters) to correlate e.g. the skill level required to the amount of the reward. The low inter-rater reliability would only introduce noise that would make the correlation with any other variable lower than it potentially is. We therefore report these correlations with the understanding that the effect we are observing is at least of this strength. So for example, we may be underestimating the degree to which skill level correlates with the amount of reward offered, but we are not overestimating it.

There are interesting correlations among the task properties. Task award is a positive indicator of the skill requirement, which means that users tend to offer more money on tasks of higher skill requirement. A combination of interest in both money and design can make those tasks more desirable.

Table 3: Spearman's correlations of task variables

	Award	Skill-Required
Skill-Required	0.493**	
Workload	-0.443**	-0.629**

Interestingly, award and skill-required are negatively correlated with workload. Although at first one may expect that tasks that require more effort in terms of time should be compensated appropriately, note that the raters were instructed to rate the amount of time it would take a person *of an appropriate skill level* to complete the task. Even so, it is interesting that workload should be *negatively* correlated with reward. Anecdotally, users who post a high-quality task and offer more money often just ask for a concise solution. For example, one task offered 2000 yuan for a logo design for a conference organized by a famous magazine. On the other hand, there are also many cases in which tasks requiring a great deal of work come with a tiny money award.

Table 4: Spearman's correlation between task properties and number of submissions to the task

	Award	Skill-Required	Workload
Submission #	0.211	0.253**	-0.242**
** Correlation is significant at the 0.01 level (2-tailed).			

We also find that these task property variables influence participation: Table 4 shows that award and required skill level attract users, while people avoid the tasks that have larger workload or don't offer sufficient award. We will show in Section 4 that these task properties also have different effects on participation for experienced users.

4. Users' Continuing, Learning, and Adaptation

Winning as Incentive to Continue

As we see that the properties of the tasks and the award offered by the tasks can influence people's decision on whether to participate in a competition, we further examine what makes them continue. Merely winning appears to play an important role in contribution. The vast majority of the users on the Witkey websites actually get nothing from their contributions, since the probability of winning is so small. One might therefore expect that a lot of users would leave after a couple of failures. In fact, from 2006 June to May 2007, there were 66,182 users who had one, two, or three submissions during this period and never submitted anything else after May 2007. These users, one third of Taskcn's total participants, disappeared. The high number of registered users who have never attempted a task (89%) suggests that although there are many people interested in participating, they might be hindered by the very likely futility of their efforts.

For those who do elect to participate, the first attempt in the competition can be very important in influencing their subsequent participation in Taskcn. There are 2307 users who won on the first attempt and 169,456 others who failed on the first attempt. Figure 3 shows the portion of users in the winner and loser group who had 2, 3, 4 … j attempts. Both groups have a heavy tailed distribution of attempts: the majority of users have a couple of attempts and a handful of users attempt many tasks. One can observe that, on average, the winners have more attempts than the losers group.

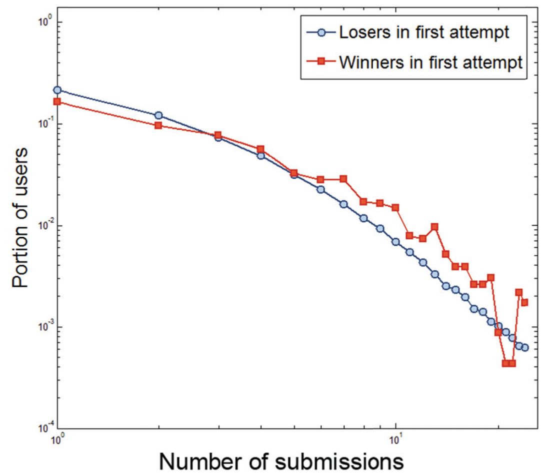

Figure 3: Distribution in the total number of submissions depending on whether the user won or lost on their first attempt.

A Cox proportional hazards analysis shows that users who win on the first attempt have a 19% lower probability (sig.<10^{-4}) of stopping after each subsequent attempt. As Table 5 shows, this translates to approximately 1 additional attempt on average for the winner group. If the first win occurs on the third attempt, there is a smaller difference of 12% in whether the user continues participating. This suggests that the result of a user's first, and subsequent, competitions can be an important factor in later participation behavior: winning encourages users' contribution.

Users Learn to Submit Later

Next we investigate users' participation pattern from a dynamic perspective. That is, how do users adaptively change their behavior over time?

The timing of users' submissions is an important participation dynamic, since users can chose to submit early, or wait to see how many other submissions a task receives. We normalize the time of a user's submission by the task period (the duration from the start time to the end time of a task), so that a user who submits at the very beginning has a submission time of 0, and one who submits at the end has a submission time of 1.

Table 5: Comparison between the number of submissions for first time winners and losers

	Winners in 1st attempt	Losers in 1st attempt
Mean	4.388817	3.20194
Variance	85.02092	25.54748
Observations	2307	169456
P(T<=t) two-tail	8.04E-10	
t Critical two-tail	1.960985	

We find that, for all users, task award correlates with users submitting solutions later. It may be an indication that people are more intent on winning higher awards (ρ = 0.067, sig. $<10^{-4}$) and so either take longer to devise a solution or "sit" on it until they are certain it is their best effort. Interestingly, tasks of longer duration have a slightly later submission time (relative to the overall duration (ρ = .026, sig. $<10^{-4}$). One possible explanation is that users may notice the task after it has started and still have sufficient time to submit a solution.

Furthermore, we find that the number of submissions is negatively correlated with the time when people submit (ρ = -.128, sig. $<10^{-4}$). A simple reason could be that most tasks with many submissions require little effort, and so users can complete and submit solutions sooner. An alternate explanation could be that when people see that many others have participated in the task, they may not want to follow up. This would result in a higher proportion of submissions having an earlier submission time.

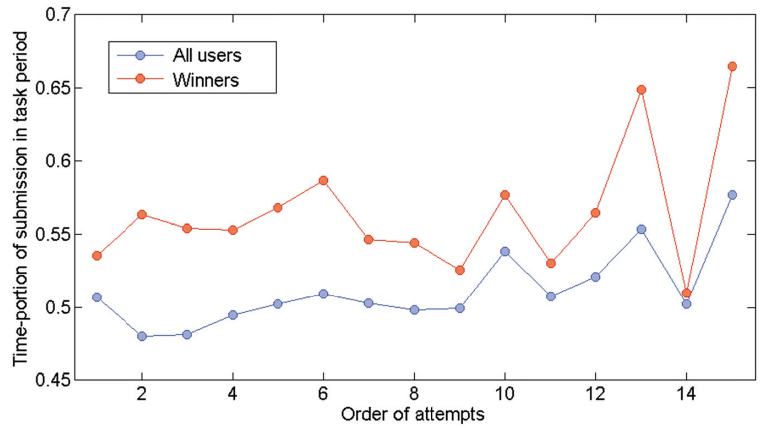

Figure 4: Point of submission in the task period for users who have participated 15 or more times.

Figure 4 shows that users, both those who have won at least once (winners) and non-winners alike, are likely to submit slightly later as they participate over time. We also observe that winners consistently submit later in the time period. Since they cannot see others' submissions, there is little additional information they can gain by waiting. However, the later submissions may be an indication of greater effort expended.

Users Learn to Choose Less Popular Tasks

Since a user's chance of winning largely depends on the number of other users competing in the task (Yang et al., 2008 a), we hypothesized that users would learn to select tasks with fewer competitors, in order to enlarge their winning probability.

In general, there is a learning pattern of users over time: users are more and more likely to choose tasks with fewer competitors (for all users and all participation levels, correlation between the number of competitors and order of attempt: $\rho = -0.23$, sig. $< 10^{-4}$). Note that this is occurring not because there is an overall decline in task popularity over time. In fact, overall task popularity is rising very slightly. Rather, on average, users are choosing the less popular tasks as they gain more experience on the site. Note also that the variance explained by the learning trend is necessarily low, given that much of the variance is due to the difference in popularity of tasks chosen by different users (e.g. on first attempts, mean=2362, standard deviation=1971 submissions).

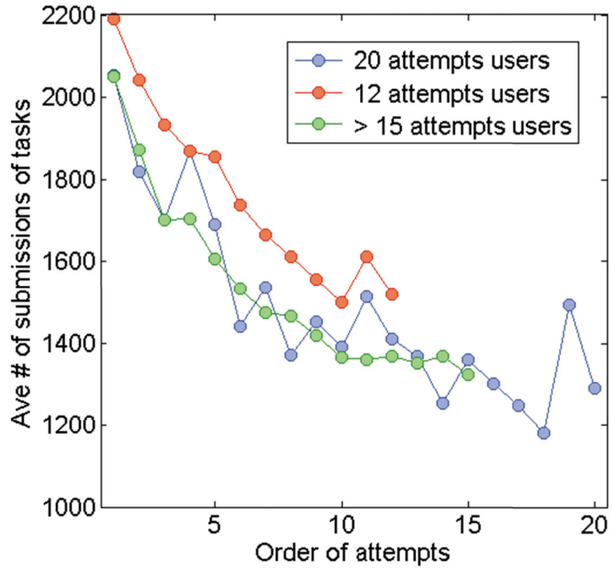

Figure 5: Average size of the competition for users in each task that they attempt.

In order to take a closer look of this trend, we select subsets of users who attempted the same total number of tasks and look at the average characteristics of the task they chose at each attempt. Figure 5 shows the trends of how users selected tasks, considering participants with exactly 20, exactly 12, and more than 15 attempts (the sets have 193, 928, and 3520 users respectively). We plot the average number of the submissions of the tasks in which they chose to compete, for each attempt they make.

Another direct way to see that users adopt a strategy favoring less popular tasks over time is to measure the average experience of the user (given as their total number of submissions before and after the particular task) and compare it with the popularity of the task. Indeed, we observe a negative correlation ($\rho \sim -0.2$, sig.$<10^{-4}$) between both the number of views and number of submissions for a task and the average experience of the user.

We can further run a regression for the submission order of a user, as related to the recency of a task, represented by the order in which it appeared, and the total number of submissions for the task. We find that both variables are significant (sig. $<10^{-4}$), with later submissions by users naturally corresponding to more recently posed tasks, but in addition also corresponding to less popular tasks.

Given that users tend to adopt the same strategy of choosing less popular tasks, it is of little surprise that experienced users find themselves attending to the same tasks. If we select two users at random for each task, we observe a positive correlation (ρ = 0.13, sig. $<10^{-4}$) for the number of submissions by each of the two users. This implies that inexperienced users are more likely to go up against other inexperienced users who are making their first attempts, while the old timers are likely to find themselves in the company of other old-timers.

Beyond simply attempting to increase their odds of winning, we find that the more experienced users have even more interesting selection criteria. Using the human-rated sample of 157 tasks, we find that, on average, experienced users are more likely to participate in tasks with a higher skill requirement (ρ = 0.253, sig. = 0.002). In addition, the higher workload of the task actually hinders experienced users from attempting the task (ρ = -0.242, sig. = 0.003). The result suggests that the serious users of the site have a combination of multiple strategies when choosing the next task to participate in. In addition to selecting tasks of higher winning probability and expected award, they also tend to challenge themselves by participating in tasks requiring greater skill; but they are thrifty with their effort by selecting tasks of lower workload.

Users Learn to Choose Tasks with Higher Winning Odds

Some of the tasks have more than one winner (in most cases, multiple winners will be announced ahead of time) and this can also affect the chance of winning. For example, some tasks need multiple people to complete the task or simply want to attract more people to participate. Thus we define WinChance as the number of winners divided by the number of submissions for the task. Intuitively, this ratio can denote the winning probability of a task in general, without regard to a particular participant. Strategic users might be expected to select the tasks of a higher WinChance. Indeed, we found that the WinChance to be increasing very slightly on average with each subsequent attempt by the user (ρ = 0.19, sig. $<10^{-4}$).

Similarly, we can also compare the task selecting patterns of the three user groups (participants with exactly 20, exactly 12, and more than 15 attempts) separately in terms of WinChance. Figure 6 shows that all three groups present increasing average trends, which means that each group has successfully improved their chances by selecting particular tasks. In addition, the group that tends to stay longer (the blue points representing users having more attempts in total) selects tasks with higher WinChance.

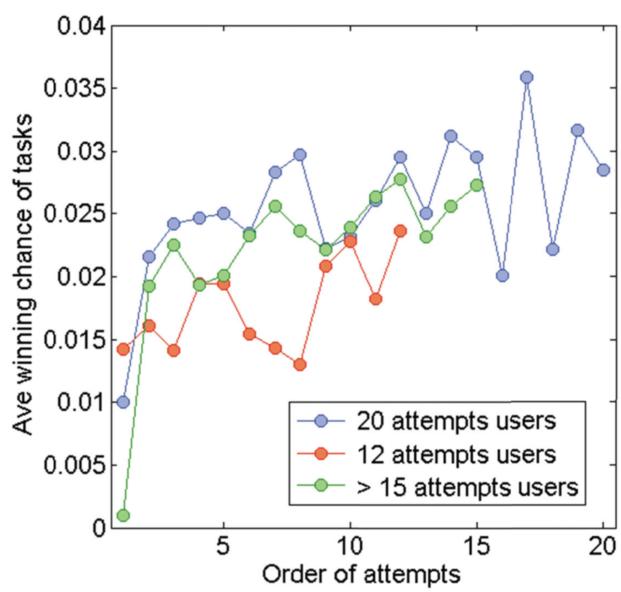

Figure 6: Average chance of winning (# of winners for task)/(# of participants) for each tasks users participate in.

The improved WinChance is due in part to the users selecting less popular tasks, as we described above. The remainder is explained by the actual number of winners of the tasks that users participated in increases over time.

Users Also Raise Their Award Expectation

We just saw that as users gain experience, they tend to enhance their winning probability. Now we ask whether they may also be attempting to combine a higher likelihood of winning with a higher award expectation. To obtain the expected earnings for a winner of a task, we divide the total award by the number of winners. As shown in Figure 7, there is a significant but very weak trend of users increasing their winning-award expectation over time ($\rho = 0.04$, sig. $< 10^{-5}$).

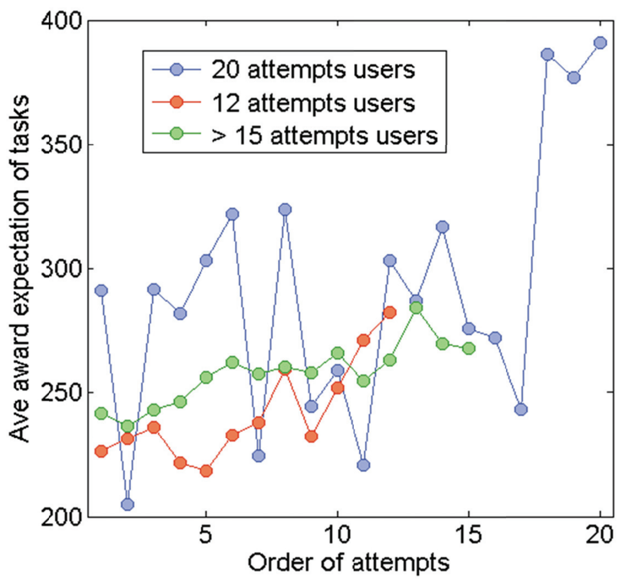

Figure 7: Average expected award (amount of award)/(# of winners for task) for each task users participate in

In summary, these results indicate that those users who remain active on the website appear to be incentivized by the award. They adjust their participation strategy such that they are likely to select less popular tasks, which yield higher odds of winning. Moreover, they improve the winning award expectation by choosing tasks with higher awards and fewer awardees that it is distributed between.

The Paradox: Users Fail to Improve

Unfortunately, for most users this effort is not significantly rewarded: we investigated user groups who had 8, 12, 15, 10, and 25 attempts and there is no emerging trend of improving win rates (defined as the number of winning submissions by those users divided by the total attempted submissions) or increasing the actual money won.

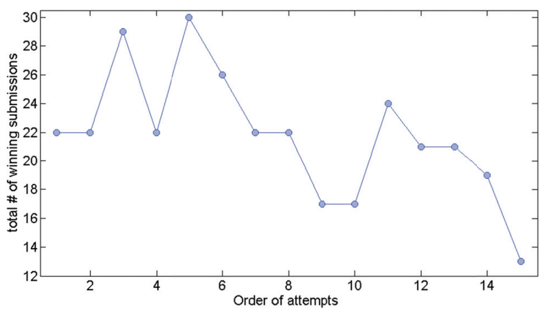

Figure 8: Aggregate number of winning submissions by the order of attempts for the set of users participating at least 15 times in the design category.

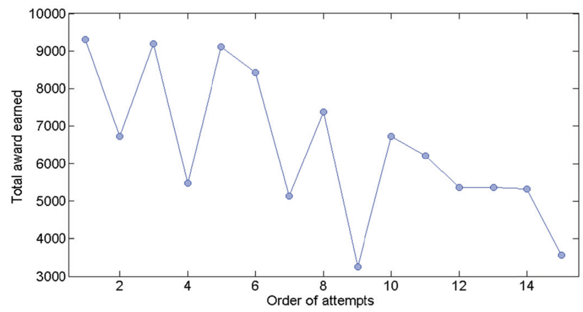

Figure 9: Total award earned by the group who have more than 15 attempts in the design category.

What is worse, although there is no significant trend on all users' performance along time in terms of wining rate or winning award, in the Design and Strategy Planning categories there is actually a very slight downward trend, which means in these two categories, users perform even worse as they spend more time on the site. Figures 8-9 show the performance over time of the users who had more than 15 attempts in the Design category. We can see that both the overall winning rates and earned award declined slightly over time among these users.

Thus, the question arises why users were able to enhance their winning probability by choosing unpopular tasks and even had increasing award expectation, but their performance has grown worse instead. We can answer this by contrasting users who are consistently winning with those who are not. That is the goal of the next section.

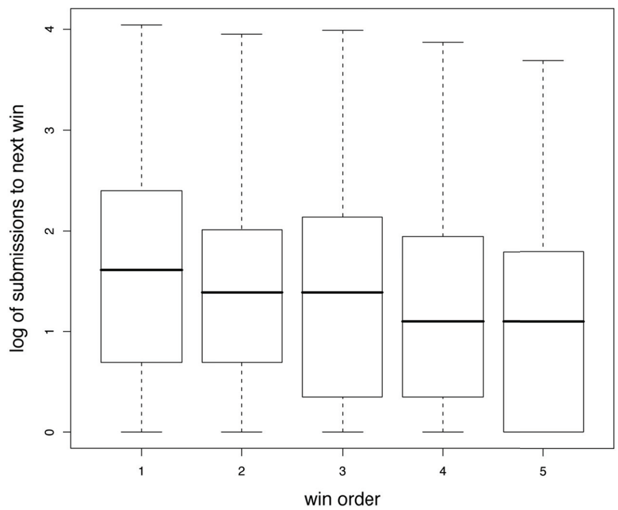

Figure 10: Interval until the next win for users with 5 or more wins.

Winners' Strategy

In our data, we see a significant learning effect for a group of winners, or at least that this group tends to get more efficient at winning over time. We take all users who have won at least 5 times, and observe the interval between wins, i.e. the number of submissions preceding their first win, the interval between their 1st and second win, etc. For their first five tasks, we observe a quickening in the succession of wins (see Figure 10). Each additional win comes 0.68 submissions sooner (ρ = -0.12, sig. $<10^{-4}$) out of a mean of 5.3 submissions between wins (median of 2). So while most users manage to worsen their chances of winning, the winners learn how to improve them.

Even though there were only 231 users winning 5 or more tasks, their wins accounted for a full 19.9% of the total wins on the site. It is therefore especially interesting to observe that this core set of users learned effective strategies for winning.

So what is it that winners do differently? In many ways, they are just like other users. Winners (here defined as users who won at least once) tend to have the same strategy as the rest of the users, in that they participate in less popular tasks over time. However, they tend to take longer to submit the task (note that the content of other submissions is not visible users before the task finishes, so that one does not benefit from seeing others' solutions). All submitters except for winners have a mean of .5039, which indicates these users tend to submit solutions in the middle of the task duration; while the winners, are likely to submit later at 0.6176 (after 61.8% of the task period has elapsed). The difference is statistically significant at p =.001. Similarly, comparing users from the winner group (who won at least once) and the rest of the users, we find that the winners tend to submit later than other users (mean difference = 0.095, sig. = .000). Figure 4 shows that like all users, winners also tend to delay their submission time over the sequence of attempts, but they also have a consistent delay relative to other users: they are always submitting later than others.

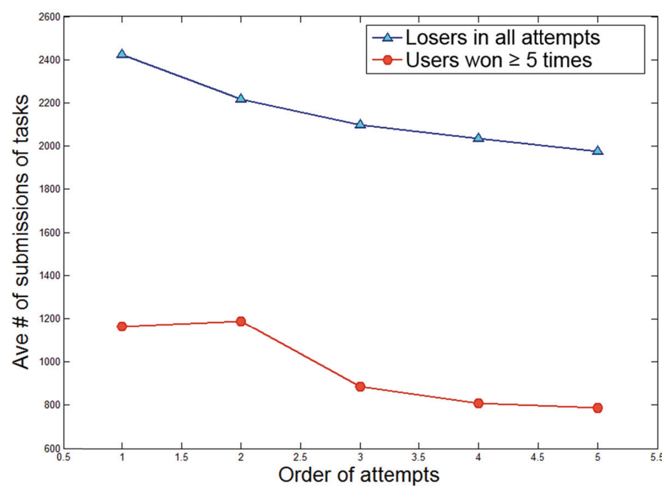

Figure 11: Ave-submissions of the tasks users who had at least 5 attempts participated: losers are those who have never won and winners are those who won at least 5 times.

Similarly, when considering popularity, winners are selecting less and less popular tasks on average. But the winning group has actually been successful in selecting the tasks of even lower popularity, starting with their very first attempt. The difference is significant (sig. < 10^{-4}). As we can see in the Figure 11, the winner group has always a significant lower average number of competitors than the loser group (defined by those users who had never won in all their at least 5 attempts). This result does not directly suggest that these users are more expert, but it does show that they on average more aggressively practice the strategy of choosing unpopular tasks, in order to enhance their winning probability.

In addition, the winner group is often able to find tasks of higher winning chance on average. The difference is also statistically significant (sig. < 10^{-5}). (See Figure 12.)

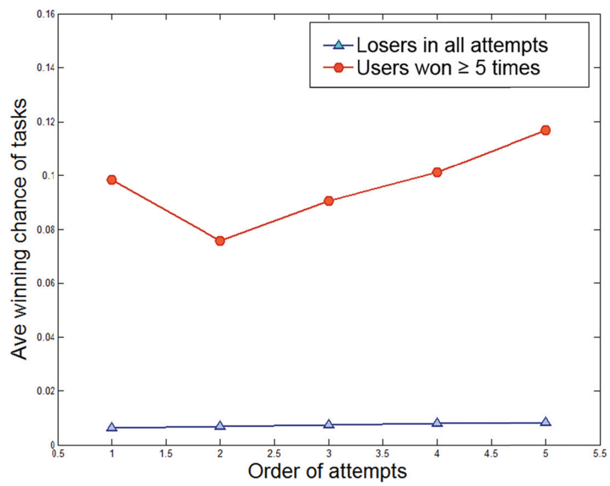

Figure 12: Ave-WinChance of the tasks users who had at least 5 attempts participated: losers are those who have never won and winners are those who won at least 5 times.

However, there is no difference between winners and others in terms of award expectation over time, although averages for the two groups both hint at a slight upward trend.

The comparison between the winner group and others implies winners are better than others at starting and sticking with a strategy that will improve their winning chance. This result is consistent with our previous finding that the best predictor of whether an individual will win is the size of the competition, only then followed by the expertise of the user. Winners are simply better at executing this strategy.

5. Conclusion

We observed several characteristics in users' activity over time on Taskcn that have implications for crowdsourcing and similar phenomena. On Witkey sites, many participants are willing to put their solutions forward in exchange for a chance to win payment. What is more, the requester of the task benefits by being able to choose among different solutions. While some designers perceive such sites as encouraging "spec work", there is little doubt that Witkeys present an open marketplace to match workers with tasks, where it makes sense for the workers to present up-front effort.

The patterns we observe hold clues to both the success of a freelance marketplace and crowdsourcing, and raise interesting design implications for such sites. On Taskcn, most users become inactive after only a few submissions. Others keep attempting tasks. Among those users, we see different behaviors. Over time, users will tend to select tasks where they are competing against fewer opponents, to increase their chances of winning. They will also, perhaps counterproductively, select tasks with higher expected rewards. However, on average, they do not increase their chances of winning, and in some categories of tasks, their chances actually decrease. This does not paint the full picture, however, because there is a very small core of successful users who manage not only to win multiple tasks, but to increase their win-to-submission ratio over time. Whether this is a case of the rich getting richer, since their successful wins give them a reputation that may enhance the chances that their submission is selected, or whether it true evidence of learning, remains unclear.

The design implications of this work are important: it is likely that it will be necessary to incentivize this core group of winners in order to maintain their continued presence on the site. It should be possible to identify and reward these users. For example, one could modify the interface to guide users to less popular tasks, or ones that match their particular expertise based on prior tasks they participated in. It is also likely to be critical to identify promising participants early (perhaps earlier than is currently possible on Taskcn), since many people leave after only a couple of task attempts. Furthermore, given the way that Taskcn works, it is critical to continue to drive large numbers of prospective members towards the site, since those members may over time become a part of small, but highly active core of users that provides 80% of the solutions.

Chapter 3

A Sustainable Mechanism for Baidu Knows

1. Introduction

Across the globe, community-based online Question-and-Answer (Q&A) sites 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 generate both demand and supply for people's knowledge and expertise, accommodating a large number of users of diverse interests and expertise, forming thriving online communities on a tremendous scale. For example, the site we investigate here, Baidu Knows, has answered over 47 million questions since 2005 and receives more than 47,000 questions per day.

Although Baidu Knows and Yahoo! Answers have very similar technical platform, there are several critical differences that are likely to cause different user behavior and site performance. First, Baidu Knows allows askers to award extra points to award the best answer, which gives askers the potential to provide higher and more flexible incentives than flat points. In addition, Baidu Knows intentionally establishes a "sense of community" by enhancing people's social interactions (e.g., providing feedback to answerers and an instant messaging service). It also tries to promote community awareness with its prominent honor title system and by explicitly promoting experts.

Interestingly, Baidu Knows has made a system with mechanism that is sustainable because it addresses people's incentives. On other sites, users can largely be separated into disjoint groups of askers and answerers, but on Baidu Knows a significant portion of users participate as both askers and answerers. These users are most likely motivated by the incentive design. Here, we investigate how users spread over multiple categories: which categories do they ask in and which categories do they answer in? In general, we found that users asked in more categories than they answered in. And users tend to spend their time unevenly, spending a lot of time in a few categories while very little in other categories.

Baidu Knows' incentive system is clearly at work. Askers post higher prices for questions they value more, and these questions bring more answers. We found that askers gradually improved their asking efficiency over time (answers per point given). Answerers, on the other hand, learn over time to discover less competitive questions, give fuller answers, be more focused, which end up winning them more rewards. Finally, we found that the only-answering group—although less active—are seeking more challenging questions and performing better than others. This strongly suggests that they are motivated by the community features, rather than the point system.

This chapter first introduces Baidu Knows and the dataset. Then it examines how the reward mechanism works and how users behave differently according to their activity level. In particular, we look into the participation pattern of the core group in the community that both asks and answers questions. We then discuss our findings and related work, and we conclude with a discussion of design implications and future work.

2. Baidu Knows and Data Set

Founded in 2005, Baidu Knows (BK) is the biggest Chinese Q&A community. About 83 million questions have been submitted and 47 million have been resolved. BK's format is similar to many other Q&A sites: the main page lists recommended topics, current questions, and quick links to meta-categories and common sub-categories; there are also frequently updated knowledge entries. Each sequential page consists of a question and its answers, and the asker can provide further feedback on the answers.

Figure 13. Screenshot of the Baidu Knows Q&A community.
http://zhidao.baidu.com/upf/

The asker can select the best answer or invite other people to vote for the best answer. Each question is closed after 15 days; users can prolong the period for another 3 days by adding award points. If there is no answer, if there is an insufficient number of votes (less than 4 votes), or if the asker is dissatisfied and wants to withdraw the question, 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% of the questions were successfully resolved.

The site has two hierarchical category levels including 24 meta-categories and approximately 300 subcategories. Meta-categories include, for example, Health, Computer/Internet, and Fashion/Life. The Computer/Internet meta-category includes C++, viruses, and downloading sub-categories. Askers assign questions to a category when they posting questions. Of course, some users incorrectly categorize their questions, but the system tries to correct that by suggesting categories after users type keywords. The system also tries to avoid redundant questions by automatically generating relevant solved questions.

Baidu Know's point system works by giving to users when they log onto the site and answer questions. BK also allows askers to offer extra points to the person who provides the best answer. This mechanism encourages more and better answers, and it encourages askers to earn more points in order to be able ask. In addition to gaining points, users can also gain "ranks." BK uses an 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. For example, it selects experts who perform well in particular categories as 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. These promotions give users more incentive to garner points, find acknowledgement, accumulate fame, and contribute on the site over time.

The site consciously builds a sense of community by enhancing the social bonds between participants. First, people often give themselves meaningful user IDs and the titles they earn are attached to those IDs, giving that person a user identity.

For example, CEO "Wind karma wind words" is a user who has answered 3,278 questions and ben chosen for best answer 1,360 times. The ID also links to the profile page with ID picture, personal information, ask/answer statistics, and activity on the Baidu forums site.

CEO "Wind karma wind words" writes that he is a male and has a master's degree, likes sleeping late, and hates smoking. He also lists his favorite books and hobbies. This page is also linked to other social networking services on Baidu.

The site also promotes interaction by letting askers provide feedback to the answerers in the question entry page. While some give terse encouragement such as "thank you!" or "very thoughtful!", others start actual discussions. A chat window is also available for users, and we see a lot of evidence that users actually use it or exchange contact information in the Q&A pages for further interpersonal interactions.

In sum, BK promotes contributions by giving experts recognition and promoting a sense of community. We believe that maintaining a palpable sense of community contributes considerably to the site's success.

Data

The dataset used for the analysis here includes all users' activities over 4.5 months (January to mid-May, 2008). During this period, 9.3 million questions were asked, 5,210,163 were resolved (or otherwise closed), and 2,667,518 unique users participated.

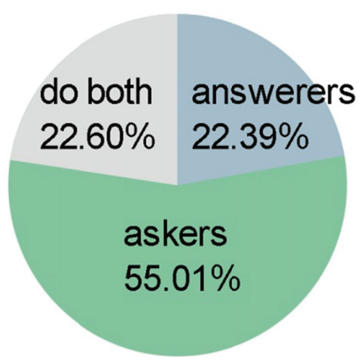

Figure 14: Distribution of three types of participants

In this dataset, only 56% of submitted questions were actually resolved. On average each question received 3.33 replies. This is lower than YA, where the rate is 7.27, but it is better than Naver (1.7). About 55% of users on BK only ask questions, while askers and answerers on YA are fairly balanced. That means that if we calculate the number of questions generated per unit of population, BK users tend to have many more questions than YA users (Adamic et al., 2008). Figure 14 shows the distribution of the three types of users: the percent that both asks and answers is similar to YA, but significantly higher than Naver, where people tend to be either answerers or askers. 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 and offering knowledge and expertise.

3. Incentives

Incentive design is crucial for knowledge-sharing communities to attract contributions. All sites offer some form of explicit incentive. 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 offers real money for the best solutions.

While some field studies have compared different incentive schemes (Chen et al., 2010; Harper et al., 2008), we wanted to know whether virtual points (instead of money) actually encourage contributions, and whether the point incentive would have different efficiencies in different ranges.

We found a correlation between the awarded points and number of answers for all questions (ρ= 0.24, sig.<10^{-4}) and when we limited our analysis to only the questions that offered extra points (ρ = 0.26, sig.< 10^{-4}). Correlations were very consistent among the different meta-categories. Figure 15 shows how many answers questions of different values obtained on average and it is clearly chows a linear trend. In sum, points have a consistent effect on participation.

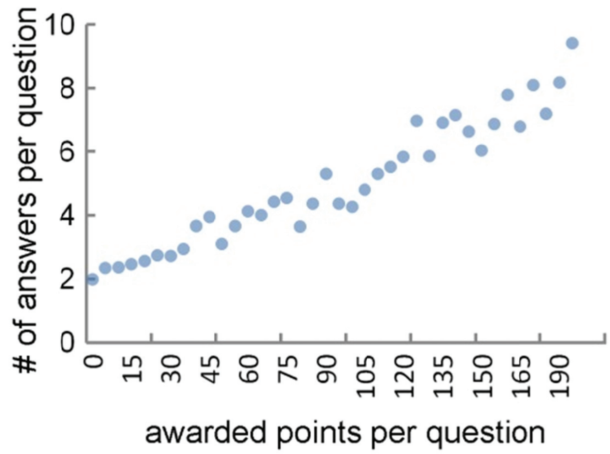

Figure 15: Average number of answers per question by amount of points

In addition, we looked at whether answerers were rewarded more for expending more effort. The results showed that longer answers were rewarded, which is consistent with what we found at YA. A two-sample t-test showed that the best answers were significantly longer than non-best answers (sig.$<10^{-4}$). Best answers were 407 characters on average (sd=1320), while non-best answers were 226 characters on average (sd=986).

Pricing Questions

In order to understand how askers reward answerers, we investigated the distribution of questions' prices as shown in Figure 16. 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 the following section that askers pay more for their first questions, and when people ask fewer questions they also pay higher amounts.

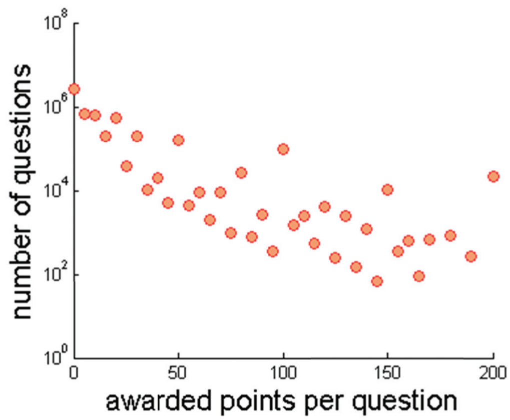

Figure 16: Distribution of awarded points for each question

We found a significant category difference in terms of question pricing. As presented in Figure 17, askers offer more in some categories such as "Music" and "Computer" while they price "Science" and "Brands"[4] the lowest. The price across categories correlates with popularity as measured by the category's total number of questions (ρ =.46, sig. =.025, in the 24 meta-categories). This, however, does not result in more answers per question (sig. = .46). This indicates the complexity in people's pricing behavior.

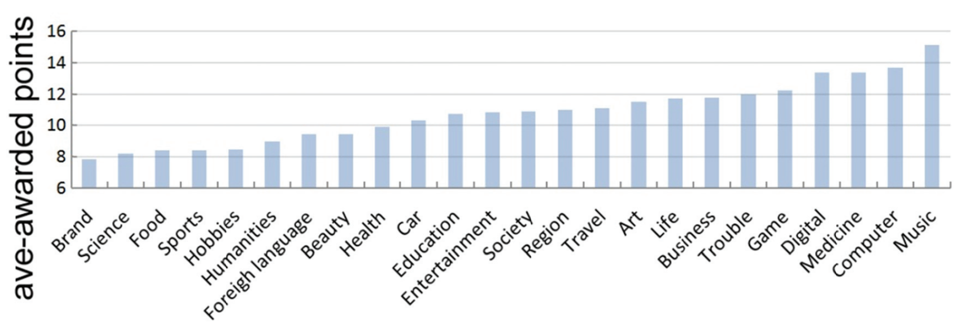

Figure 17: Average awarded points in meta-categories

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

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 (ρ =.47, sig.<0.05). The first question ratio and popularity count for a significant portion of variance of the price (ρ =.63, sig. =.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 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 18 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.

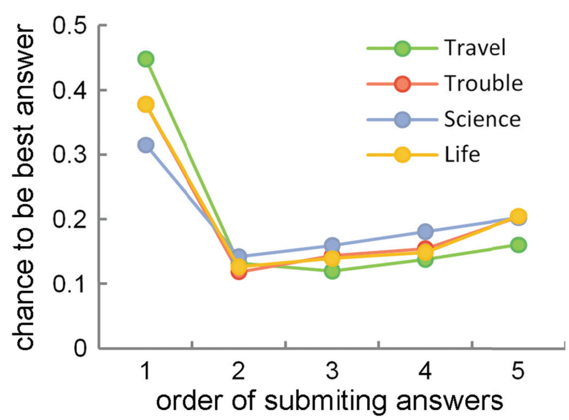

Figure 18: Chance to be selected as the best answer for all question with 5 answers

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 19, 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).

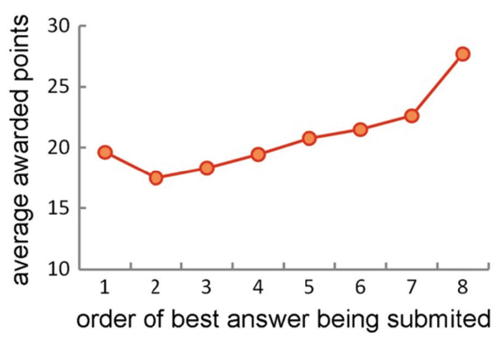

Figure 19: Average awarded points for the questions that chose 1, 2,...8th answer to be the best

4. Users' Behavior Over Time

Users can be differentiated by various dimensions. We previously found on the crowdsourcing site Taskcn that users adapt their behavior over time, and that the behavior of the most successful users is different from the rest. 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 only 35% new active users, therefore we could not capture the initial behavior for the majority of the users. For this reason, we excluded users who participated in the first month of the dataset while being active in later months and count them 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 heaviest answerer 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 questions 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.

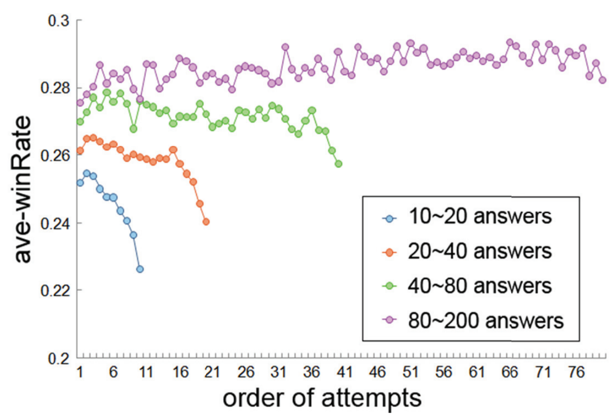

Figure 20: Average winRate at each attempt for the three answerer groups

Wining rate: winRate, a measure of answerer's performance, is defined as the total winning attempts divided by the total number of attempts. Figure 20 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.

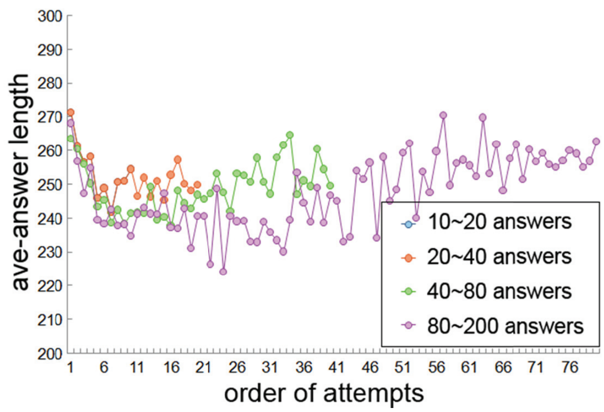

Figure 21: Average answer length at each attempt for the three answerer groups

Answerers' effort: the answer length is a simple metric of answerers' effort. Figure 21 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.

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.

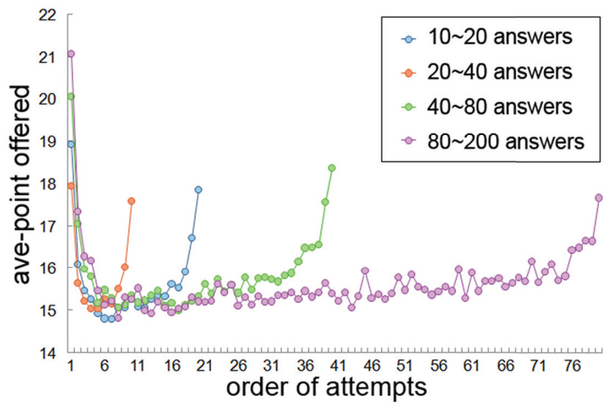

Figure 22: Average point award per question offered at each attempt for the three answerer groups

Experience and performance

As I 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 believe 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 ($\rho = 0.16$, sig.$<10^{-4}$)[5], average award obtained for each question attempted ($\rho = 0.06$, sig.$<10^{-4}$), and Guru Score[6] ($\rho = 0.10$, sig.$<10^{-4}$). Figure 23 and Figure 24 present the collective patterns among answerer groups that answered 40~100, 100~200…questions: more active users consistently perform better than less active ones.

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

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

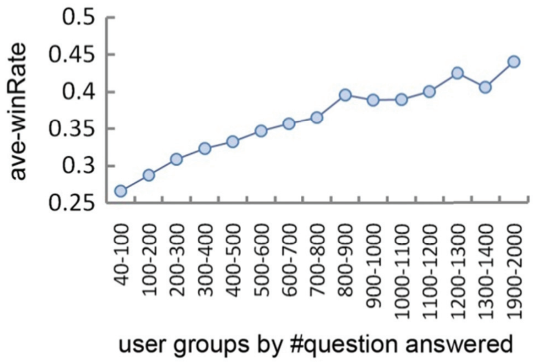

Figure 23: Average winRate by answerer groups

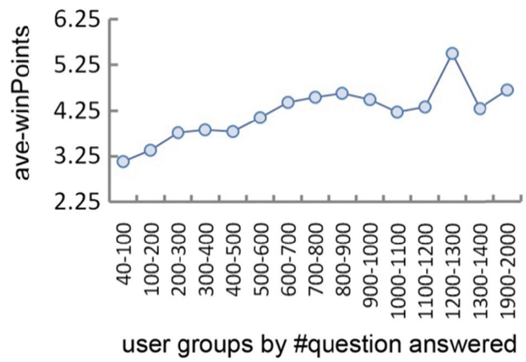

Figure 24: Average points won per attempt by answerer groups

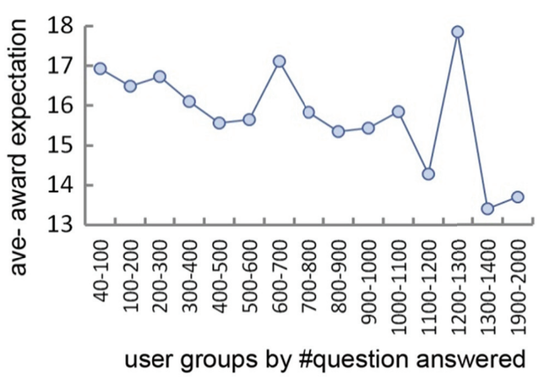

Figure 25: Average award expected per question by answerer groups

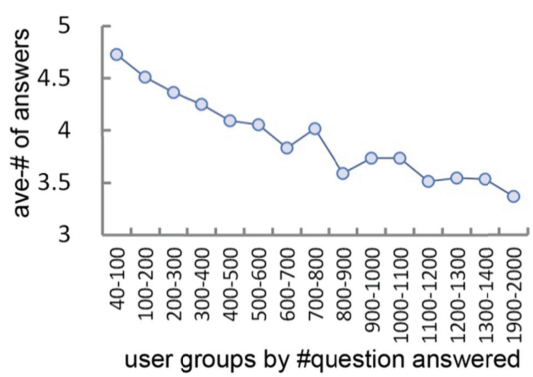

Figure 26: Average number of answers per question by answerer groups

More active answerers put in more effort per answer and are more focused in providing knowledge/expertise. In particular, we use answer length (ρ = .06, sig.$<10^{-4}$, Figure 25) to measure effort and users' entropy (ρ = -.04, sig.$<10^{-4}$, Figure 16) to measure how an answerer is focused on particular domain/s[7].

[7] Entropy: see Adamic et al. (2008)

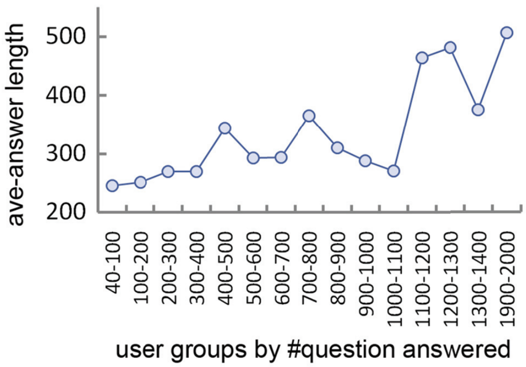

Figure 27: Average answer length per question per user by answerer groups

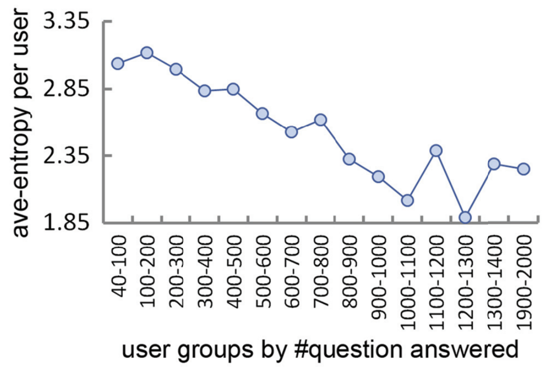

Figure 28: Average entropy per user by answerer groups

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 (ρ = -.73, sig.<10^{-4}) while there is little correlation between the winRate and award per question (sig. =0.104). In addition, answer length and focus also contribute to better performance (ρ = .38, sig.<10^{-4}; ρ = -.18, sig.<10^{-4}).[8]

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

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

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.

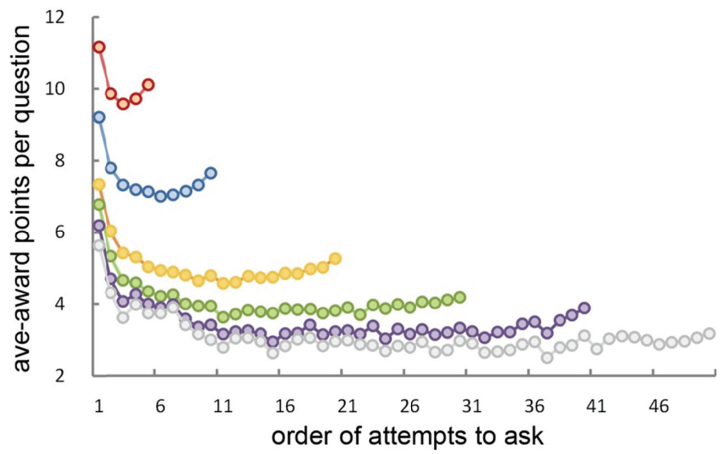

Figure 29: Average point award per question offers at each attempt to ask for different asker groups

Figure 29 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 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 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.

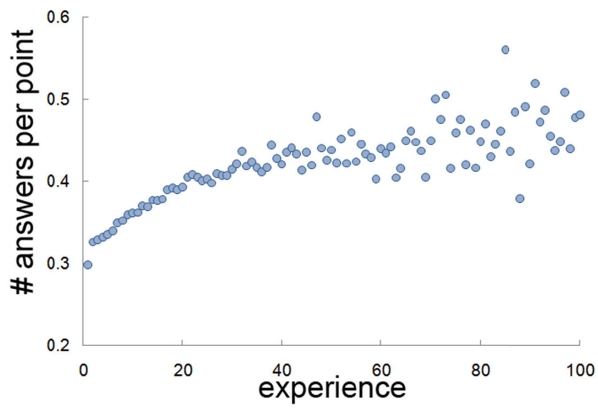

Figure 30: Average number of answers per point offered

In addition, we observe a slight trend that more experienced askers get a higher number of answers per point offered (ρ =0.005, sig. $<10^{-4}$). As shown in Figure 30, although askers offer smaller award, they actually improve the efficiency of each point in terms of buying participation.

5. 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 awards 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, ρ = .26, sig.<10^{-4}) and similarly that those who spend more points also earn more (log#point-earned to log#point-spent, ρ = .19, sig.<10^{-4}); 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 et al. (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.

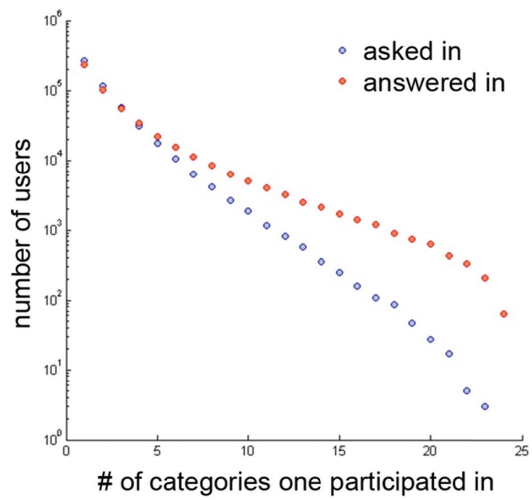

Figure 31: User distribution in terms of the number of categories they asked and answered in

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%.

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 32 and 33), answering patterns present higher concentration in general; and we see more highly focused answerers in each category too.

6. Conclusions and Future work

In this chapter, we studied a large scale Q&A system, Baidu Knows, in order to understand how such a system is sustaining 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 feedback cycle for users to keep participating and improving: you put in more effort, 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 sustainable mechanism.

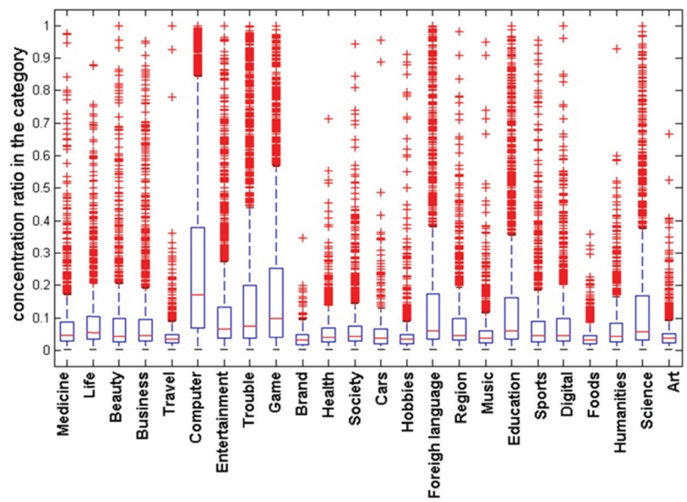

Figure 32: Users ask in categories

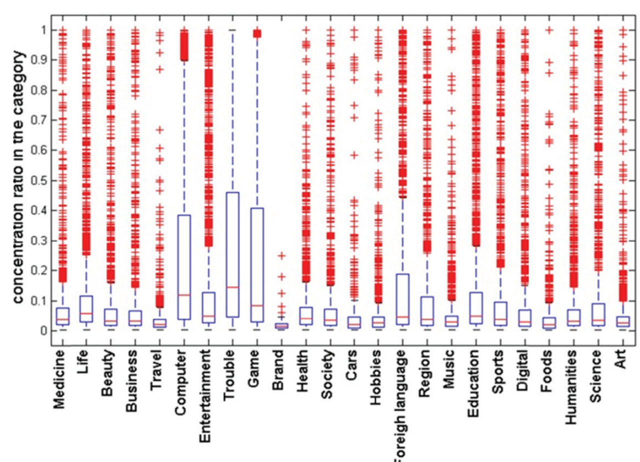

Figure 33: Users answer in categories

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 many online communities, including Yahoo! Answers, Wikipedia (Adler & Alfaro, 2007), Del.icio.us (Golder & Huberman, 2006), and Flickr (Marlow et al., 2006). In addition, most of contributions are made by a small minority of the participants, and this group of users usually have better performance as described by Welser et al. (2007), Adler and Alfaro (2007) and Yang et al. (2008 b). The 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. In addition to monetary incentives such as in previous chapter and in Harper 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 incentivizes contribution would be studied in future work.

In the previous study we investigated how users price the tasks to recruit solutions, and we found that the price correlates with the expertise required for completing the task. In the form of virtual points, askers on BK pay different amounts for different questions: there are category difference and sequence difference in terms of when the question is asked by the asker.

Chapter 4

Survival Patterns in Online Knowledge Sharing Communities

Chapter 3 presents one example of Internet-scale Q&A sites, Baidu Knows, in terms of how the virtual point design encourages users to improve their strategy and performance. We would like to know more about what keeps these kinds of communities going. Due to the sheer size of their populations, Internet-scale Q&A communities might suffer more from sparse social interactions and thus low levels of commitment. However, relatively little is known about motivating people over the long run in online communities. A few studies have focused on what makes help sharing systems sustainable over time (e.g., Ackerman & Palen, 1996), and some work has investigated this problem in Internet-scale communities. Joyce and Kraut (2006) investigated newcomers' retention across six newsgroups and found interrelations among newcomers' initial post properties, reply properties, and the probability of posting again. Arguello et al. (2006) conducted a similar study with eight Usenet newsgroups, comparing the interaction patterns between newcomers and old timers, and found that they differ in their ability to get replies and in the ways they write messages. In an exploration of the online forum Slashdot, Lampe and Johnston (2005) found that how a newcomer's post is rated and moderated affects her probability of returning.

These studies, however, are limited. They largely consider initial participation, and commitment is measured only to the second post. In this chapter, on the other hand, I examine users' participation lifespans to assess how systems might sustain users for the long term. As I will discuss, survival analysis shows that participation patterns and performance factors can account for considerable variance in predicting participation lifespan.

As well, this is a comparison study across three major Q&A sites (the three sites mentioned above). Thus, this is not only the first user retention study on Internet-scaled Q&A sites, but also the first comparison study among these three large communities.

1. Data and Methods

Data Description

As mentioned, I investigated three Internet-scale Q&A sites across languages and cultures: Yahoo! Answers (YA) in English, Baidu Knows (BK) in Chinese, and Naver Knowledge-IN (NK) in Korean. NK, started in 2002, was the earliest, while YA and BK were launched a short time apart in 2005. All three experienced a boom in user population and traffic starting in 2006. The set of sites is well suited to a comparative study, since they are similar in scale, purpose, and basic functionality. They, however, vary in cultural context, incentive structure, and site design, which might potentially influence users' participation patterns and social interactions on each site.

All of these sites presented challenges in collecting data. Below I briefly describe each site and its data collection. We were limited to two years of use at each site for reasons we will explain below.

Yahoo! Answers (YA) is an online knowledge sharing community website launched by Yahoo! in Dec 2005. YA sites exist in various languages, but we limit our analysis to the English site, which is by far the largest.

YA, like NK and BK, allows any user to ask or answer a question and provides a virtual point system to realize its knowledge market. Users pay a flat fee in virtual points to ask a question and can recover some of those points by selecting the best answer among those received in response to the question. If the asker does not select a best answer, it is selected by votes from other users. On the answerer side, a small number of points are awarded for contributing an answer, and a bigger, flat number of points for being selected as best. The total number of points earned, as well as the percentage of a user's answers that were selected as best are displayed in a user's profile.

Using YA APIs, we were able to crawl all questions in each category and the corresponding users within the first year after the launch of the site. However, we could not reach questions posted in the second year using this method. Instead, we used a random sample of 150K users from Adamic et al. (2008) over a period of three months starting a month into YA's second year.

Baidu Knows (BK*)* founded in June 2005, is the largest Chinese Q&A online community. As successful as its co-named search engine in China, BK has garnered a huge user base and traffic. To date, more than 100 million questions have been asked with more than half of them successfully solved (i.e., the best answer was chosen by the asker or voted on by other users).

As YA's Chinese peer, BK shares many features with YA such as using virtual point system as well as some deviations listed in Table 6. Our BK data includes a full history of the first two years of the site, including all undeleted questions and answers (57 million posts in total), with corresponding users.

Table 6: General site comparison

	YA	BK	NK[1]
Founded in	2005/12	2005/6	2002
Incentive for answering	Earn flat point rate	Earn flat point rate+ flexible points	
Incentive for asking	Pay flat rate of points	Earn flat rate of points, but optionally offer extra flexible points	

Table 7: Data description & general characteristics of the sites

Year 1	YA	BK	NK
#sampled users	6841	14,683	6,460
Ave # ques per asker	14.3	4.9	4.64
Ave # ans per answerer	217.1	26.08	7.01
Ave # ans per question	12.1	6	3
%asker	54.1%	54%	21.4%
%answerer	7.1%	10.4%	43.5%
%doBoth	38.8%	35.6%	35.1%
Year 2	**YA**	**BK**	**NK**
#sampled users	72,099	18,871	61,177
Ave # ques per asker	5.32	4.17	1.78
Ave # ans per answerer	51.33	16.61	5.44
Ave # ans per question	12.71	5.1	1.72
%asker	59.3%	58.1%	60.5%
%answerer	6.8%	12.4%	24.7%
%doBoth	33.9%	29.5%	14.8%

[1] To our knowledge, the flexible rate award was not widely used during the observation time for this study.

Naver Knowledge-iN (NK) is the largest online Q&A community in South Korea. The site has over 43 million questions[2]. Since Naver's API access to NK data is restrictive, we manually crawled 2.6 million questions and their answers from 15 categories between 2002 and 2007. The data collection technique we used is described in Nam et al. (2009).

The site differences are summarized in Table 6. Table 7 presents the data description of the sampled users who joined each site during the first three month of each year. We can see that while there are some noticeable discrepancies between the years, the site differences are significant: YA users contributed significantly more in terms of both asking and answering than BK and NK users, which resulted in a higher average number of answers each question obtained. In addition, there are always more doBoth users (who had both asked and answered during the observation period) in YA.

Survival Analysis

Survival analysis (Cox & Oaks, 1984) is the main method in this study to measure the lifespan of users' participation. The technique has been widely used in biological and medical science, engineering, and sociology. It involves modeling of a lifetime against a specific event. In particular, For example, two applications include how many days a cancer patient will survive (against death) and how long a marriage will last (against divorce). Note that survival analysis must deal properly with censored data, or where the event has not occurred before the end of the observation period. In our context, a user "survives" on the site if they keep participating. As demonstrated in the sections below, survival analysis can test the difference in participation lifespan between groups and quantify individual predictors using a Cox proportional-hazards regression model.

Defining Lifespan. In our context, users "survive" on the site if they keep participating. Defining lifespan with respect to participation is tricky because, unlike actual death or

[2] http://kin.naver.com/, retrieved on Dec, 23, 2009

divorce, from which few recover, users can just be inactive for a while and return again. As long as the account is still valid and the site is still running we cannot be sure that a user will not return. Thus we first examine users' intermittent pattern. As Figure 13 shows, the cumulative distribution of the maximum intervals between any two sequential actions for the users is heavy tailed. More than 70% users had no more than 100 days between posts, which implies that most users are unlikely to return if they have left for more than 100 days. Therefore, we calculate the users' lifespan as the duration of active participation from when a user first posts to the forum to her last post, with no gap greater than 100 days. We performed a sensitivity analysis using alternate cutoffs of 50 and 150 days and obtained no major differences in the statistical analyses presented below.

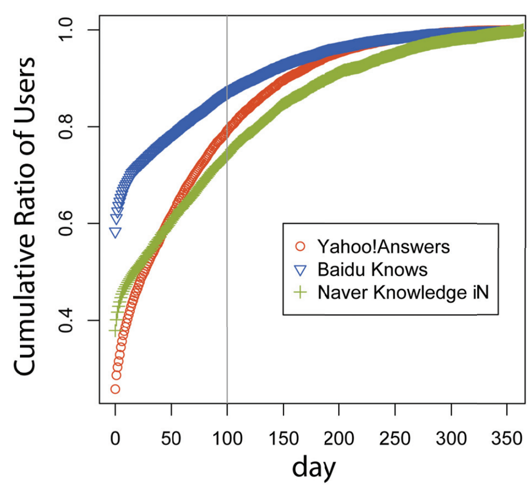

Figure 34: Cumulative distribution of maximum inactivity intervals for users

Characterizing User Lifespan across Sites

To study the sites' ability to retain users, we split the data into two stages: the initial year after the site was launched and a following year, representing a more mature period for each site. In addition, in order to be able to observe the lifespan of a user for a sufficiently long time period following her debut, we further restrict our sample to include only users whose first post occurs in the first three months of either year (as described in Table 7). Therefore, we have a range of 9 to 12 months to observe how long a user continues to participate. At the end of the observation period we then can assign each individual's life status as either dead or censored. A censored user is one who has not exceeded the cutoff interval of inactivity at the conclusion of the observation period and can thus still be considered alive. Survival analysis allows us to properly account for censored data.

Identifying Participation Patterns

We define the following variables which were used in the statistical tests to predict users' lifespans.

Two other variables describe users' asking and answering activity. One is the ***ask/reply ratio*** (***A/R ratio***) representing users' preference between asking and replying to questions (as defined below). The second, ***netPoints***, is defined as the net point balance: points earned minus points expended in asking and answering activity during a period of time.

Table 8: Variables of participation pattern

Asking Variables Description	
# Questions	# questions user asked during a period of time, indicating activity level of asking
# Answers/Q	Ave. # answers obtained per question, indicating ability to get answers
Len_Ans	Ave. length of answers obtained
%Answered (BK)	Whether the question has >= 1 answer
%Solved (BK)	Whether a best answer was selected (either by the asker or by a vote)
# Points	Ave. # points user offered per question
%chosenBest (NK)	Whether best answer is chosen by asker
%userChosenBest(NK)	Whether best answer is chosen by others

Answering Variables Description	
# Questions	# questions answered during a period of time, indicating activity level of answering
# Answers	Ave. # answers for each question, indicating level of competition
Len_Ans	Ave. length of answers
# win	# times that user's answer was selected as best
winRate	# win/#question
Guru	winRate incorporating question competitiveness
Points earned	Ave. # points earned per question answered
Points expected	(Ave. points offered for each question the user answered)/(# questions)
%comment	% of answers that were commented on by the asker

2. Analysis and Results

Comparing User Lifespans across Sites

We first compared the general survival curves across the three sites as shown in Figure 35.

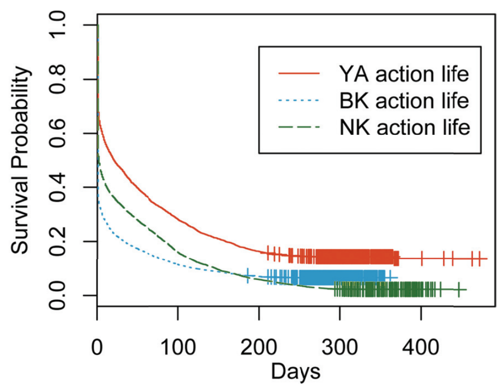

Figure 35: Active user lifespan survival curves of year 1

All sites had a stark initial drop-off, with 30%~70% of users leaving after posting just once. The observation echoes previous studies (e.g., Joyce & Kraut, 2006), which found that the first interaction is critical for sustaining a large number of potential users. Subsequently, the curves for all three sites flatten, suggesting that the longer a user remains active, the more likely they are to remain even longer. Overall, YA users are significantly more likely to remain active than users on the other two sites. NK is better able to retain users during the first 100 days than BK. However, the 10% of BK users who remain active past the first 200 days are likely to stay the full year, while NK users' survival continues to drop to as little as 3%.

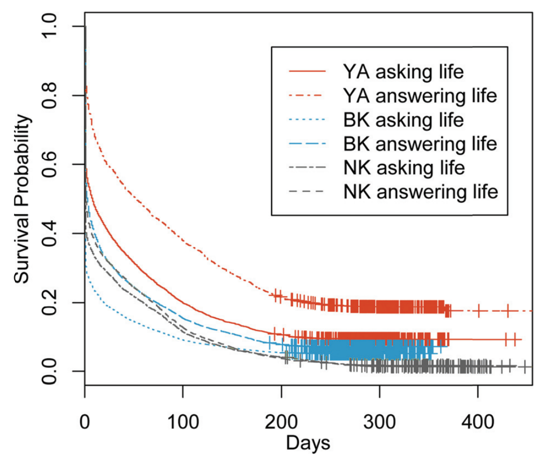

Figure 36: Users' asking lifespan versus answering lifespan (Best seen in color)

So far, our analysis has looked at combined asking and answering activity in aggregate. However, a user need not do both or may participate in one type of activity longer than the other. As shown in Figure 36 answering activity is more likely to persist than asking activity, but this difference is pronounced for just YA and BK. However, even when broken down by the type of activity, the relative difference in survival likelihood remains between sites. Interestingly, answering activity in BK has a similar lifespan to that of NK users, but it is the asking lifespan in BK that is significantly shorter.

The Role You Play: The Life You Have

Activity preference between asking & replying. As asking and replying lifespan patterns of individual users are different, we wanted to measure how users' preference for either of these two roles was related to their continued participation. We define:

$$A/R\_ratio = \frac{\#questionsAsked}{\#questionsAsked + \#questionsAnswered}$$

Activity preference and survival. We then used the A/R ratio to characterize individuals' activity lifespans. Table 9 provides the regression result. Note that the exponentiated coefficient indicates the direction of the effect: when larger than 1, it presents a negative relationship between the variable and estimate of lifespan. It presents a positive relationship when smaller than 1. For example the exp(coef) for YA is 2.497, meaning that when the A/R ratio is higher, then the lifespan is shorter (i.e., users who primarily ask questions tend to leave earlier). Pr(>|z|) indicates the statistical significance; R^2 is the strength of correlation.

Table 9: Correlating survival and A/R ratio

	exp(coef)	Pr(>|z|)	R^2
YA	2.497	<0.001 ***	0.125
BK	3.297	<0.001 ***	0.147
NK	1.591	<0.001 ***	0.040

The results are consistent across all three sites: users who stay longer prefer answering to asking. Also consistent with Figure 36, this difference is more pronounced in YA and BK than NK. This is a corroboration of what one might intuitively assume: Answerers demonstrate much greater commitment to Q&A communities, by contributing more and staying significantly longer.

How First Time Experience Matters

As mentioned, previous work measured the effect of the first interaction experience on the probability of a user's returning to the online forum. Here we extend the analysis past the probability of returning once to quantifying the degree to which the initial interaction correlates with the length of the entire lifespan of a user's participation.

Table 10: First time action preference in year 1 and year 2

1st action	YA1	YA2	BK1	BK2	NK1	NK2
= asking	63.7%	73.8%	73.4%	74%	36.5%	68.4%
= answering	36.3%	26.2%	26.6%	26%	63.5%	31.6%

Since users can take two different initial actions, asking or answering, we examine them separately and predict user' participation lifespan using variables corresponding to asking and answering. Table 10 presents the ratio of users who initially ask to those who initially answer for the three sites. Except for year 1 of NK, all sites present a significantly larger preference for asking as the initial action. Interestingly, BK is rather stable between the two years, and YA and NK gained a larger portion of users who joined by asking first. It may reflect changes in how users initially discover the sites (e.g. by being routed from a search engine while performing a query) or new community-oriented services and designs. It may even reflect initial social instability in new sites.

Predicting activity lifespan by the first question asked. Table 11 presents the results of Cox proportional-hazards regressions on user lifespan using individual variables and the overall multiple regression using all predictor variables relating to the question. The results were statistically significant for both YA and BK, but not for NK.

Table 11: Predicting lifespan by first asking activity

YA				
Individual predictor variables				
	exp(coef)	z	Pr(>|z|)	R^2
#answers	0.9817	-4.979	< .001 ***	0.006
Len_Ans	1	-1.558	0.119	0.001
Len_Ques	0.9994	-6.188	< .001 ***	0.01
Multiple R^2= 0.015, p < .001				
BK				
#answers	0.941	-10.84	< .001 ***	0.013
Len_Ans	0.9999	-5.445	< .001 ***	0.004
Len_Ques	0.9956	-9.817	< .001 ***	0.011
?answered	0.6573	-16.31	< .001 ***	0.026
?solved	0.5648	-20.33	< .001 ***	0.043
# Points	0.9982	-1.546	0.122	0
Multiple R^2= 0.049 , p=0				
NK				
#answers	1.010	0.729	0.466	0
?chosenBest	1.397	0.747	0.455	0
?userChosenBest	1.273	2.266	0.024 *	0.002
Multiple R^2= 0.003, p=0.122				

Table 12: Predicting lifespan by first answering activity

YA				
Individual predictor variables				
	exp(coef)	z	Pr(>|z|)	R2
#Ans	1.01	0.827	0.408	0
?Win	0.8015	-4.07	< .001 ***	0.007
Multiple R^2 = 0.007, p < .001 ***				
BK				
#Ans	1	0.65	0.516	0
?Win	0.9217	-3.097	0.002 **	0.001
earnedPoints	0.9985	-1.824	0.068	0
Len_Ans	1	0.032	0.974	0
?best_Commented	0.9042	-3.377	< .001 ***	0.001
Multiple R^2 = 0.001, p=0.0238				
NK				
#Ans	1.014	1.762	0.078	0.001
%Win	0.98	-0.622	0.534	0
Multiple R^2= 0.001, p=0.214				

In BK, whether the best answer was chosen, either by the asker or by voting, is the most significant factor that results in longer lifespan for the asker. In addition, obtaining more answers (in YA and BK) and longer replies (in BK) for the initial question also encourages askers to stay longer. As indicators of the level of investment on the part of the asker, writing longer questions (on YA and BK) and offering more points (on BK), can not only attract more answers, but are also positively associated with longevity.

Predicting activity lifespan by the first answer. Similarly, we predicted lifespan for those users who started by answering questions (shown in Table 12). The results show very limited prediction power with small R^2 for YA and BK; the variables remain non-significant for NK.

Consistently between YA and BK, having one's answer selected to be the best is a promising sign for a longer lifespan. On BK, earning points also had positive effect, and importantly, getting feedback about the answers from the asker (?best_Commented) also was correlated with users staying longer.

Participation Patterns That Predict Lifespan

Next we look past the first interaction to see how users' continued participation patterns can be used to predict total lifespans on the site. These patterns can be only observed and identified through a period of time. For example, users' performance can be measured as the average points earned per question answered.

Thus, we selected users who had stayed for more than 30 days and used the variables obtained during this period of time to predict how long those users would continue to participate. The results show that users' aggregate participation patterns can yield considerably more predictive power than using just a user's initial experience.

Predicting asking lifespan by the first 30 days. We first predict users' asking lifespans, based on their asking patterns in the first 30 days of participation (shown in Table 13). First, on all three sites, those who asked more questions remained longer after the 30 days, and this accounts for a significant portion of total explained variance. On NK, this difference in activity level accounts for the majority of prediction. This may imply that NK is less capable of sustaining low-use users for the long term, while the other two sites might be able to accommodate users at a variety of activity levels.

Table 13: Predicting asking lifespan by first 30 days

YA				
Individual predictor variables				
	exp(coef)	z	Pr(>|z|)	R^2
# Question	0.972	-10.49	< .001 ***	0.056
Ave# Answer	0.982	-2.857	0.004 **	0.003
Len_Ans	1.000	-0.935	0.35	0
Len_Ques	1.000	-2.455	0.014 *	0.002
A/R ratio	0.598	-10.08	< .001 ***	0.036
Multiple R^2= 0.09, p=0				
BK				
# Question	0.800	-38.12	< .001 ***	0.194
Ave# Answer	0.880	-23.86	< .001 ***	0.058
Len_Ans	1.000	-13.79	< .001 ***	0.022
Len_Ques	0.991	-21.96	< .001 ***	0.052
%answered	0.417	-32.61	< .001 ***	0.089
%solved	0.347	-36.99	< .001 ***	0.116
Ave_offerPoint	1.000	0.202	0.84	0
A/R ratio	3.442	38.03	< .001 ***	0.121
Net_Points	1.080	67.33	< .001 ***	0.292
Multiple R^2= 0.392, p=0				
NK				
# Question	0.7688	-27.34	< .001 ***	0.372
Ave# Answer	0.9681	-2.326	0.0200 *	0.002
%chosenBest	1.277	0.686	0.493	0
%userChosenBest	1.577	3.672	< .001 ***	0.004
A/R ratio	1.432	6.339	< .001 ***	0.012
Net_Points	0.9886	-2.75	< .001 ***	0.002
Multiple R^2= 0.382, p=0				

Table 14: Predicting answering lifespan by first 30 days

YA				
	exp(coef)	z	Pr(>|z|)	R^2
# Question	0.998	-6.452	< .001 ***	0.029
Ave#Ans	0.999	-0.388	0.698	0
#win	0.994	-4.85	< .001 ***	0.016
winRate	0.840	-1.697	0.090	0.002
guru	0.894	-2.472	0.013 *	0.003
A/R ratio	8.181	13.66	< .001 ***	0.075
Multiple R^2= 0.091, p=0				
BK				
# Question	0.977	-19.49	< .001 ***	0.1
Ave# Answer	1.001	5.745	< .001 ***	0.005
#win	0.946	-13.56	< .001 ***	0.053
winRate	0.802	-3.868	< .001 ***	0.003
guru	0.986	-1.684	0.092	0.001
Ave_earned_Point	0.997	-1.549	0.121	0
Ave_expectedPoint	1.002	4.822	< .001 ***	0.004
Len_Ans	1.000	-2.175	0.030 *	0.001
%best_Commented	0.840	-2.351	0.019 *	0.001
A/R ratio	1.519	7.937	< .001 ***	0.01
Net_Points	1.023	47.77	< .001 ***	0.14
Multiple R^2= 0.488, p=0				
NK				
# Question	0.9073	-30.11	< .001 ***	0.335
Ave# Answer	0.9263	-8.385	< .001 ***	0.015
winRate	0.9128	-2.059	0.040 *	0.001
guru	0.965	-2.308	0.021 *	0.001
A/R ratio	0.9384	-0.986	0.324	0
Net_Points	0.9856	-6.737	< .001 ***	0.009

Multiple R^2= 0.371, p=0

Some aspects of the general experience are important. Askers who continuously put in more effort (as measured by the average questions length) will also stay longer. As well, across the three sites, getting more answers each time is correlated with the asker's continued participation. On BK, getting longer answers and the proportion of questions being answered (demonstrating a greater effort on the part of the community) encouraged askers to keep asking. Considering that almost 45% of questions have never been solved on BK, we can imagine many askers being discouraged by obtaining no answer. However on NK, where both the asker and other users can select one of the answers as best, whether the asker made a selection is not statistically significant, while other users selecting the best answer, for reasons unclear to us, actually has a negative effect on lifespan.

Preference for role was mixed. On BK and NK, a user who prefers asking tended to continue asking. However, on YA, she would be less likely to continue asking (Table 13, rows for A/R ratio). This might imply the more marked tendency of YA users to switch roles between asking and answering; there is larger portion of users on YA who have both asked and answered (Table 7).

The incentive structure also had a mixed effect. Net point balance had a different effect size and direction between BK and NK. (We were unable to collect best answer selections, and therefore point balances, for YA.) For BK the factor yields a high predictive power. This might imply that on BK, net point surplus could also yield a negative effect in sustaining users.

Predicting answering lifespan by the first 30 days. Similarly, activity level in answering (#questions in the 30 days) is also an important factor for staying longer (see Table 14). Furthermore, users' performance in answering (as measured by #win, winRate, or guru score) is often positively, but weakly correlated with continued answering. Interestingly for BK, users who self-select for participating in higher reward questions (Ave_expectedPoint) die earlier, which is consistent with the finding in Yang and Wei (2009) and Yang et al. (2008 b) that experienced users tend to adopt a strategy of choosing less well-rewarded and therefore less competitive questions. The willingness to put in more effort in the form of longer answers is also weakly correlated with a higher survival rate. While unpredictable based on their first post, NK users lifespan becomes much more predictable once one accounts for the first 30 days of activity.

Community Evolvement

Next we investigated how users' participation patterns change over time. We conducted a longitudinal comparison between two different periods: the first year after the launch of site and the subsequent year, with the YA sample falling 1.5 month behind the beginning of the second year.

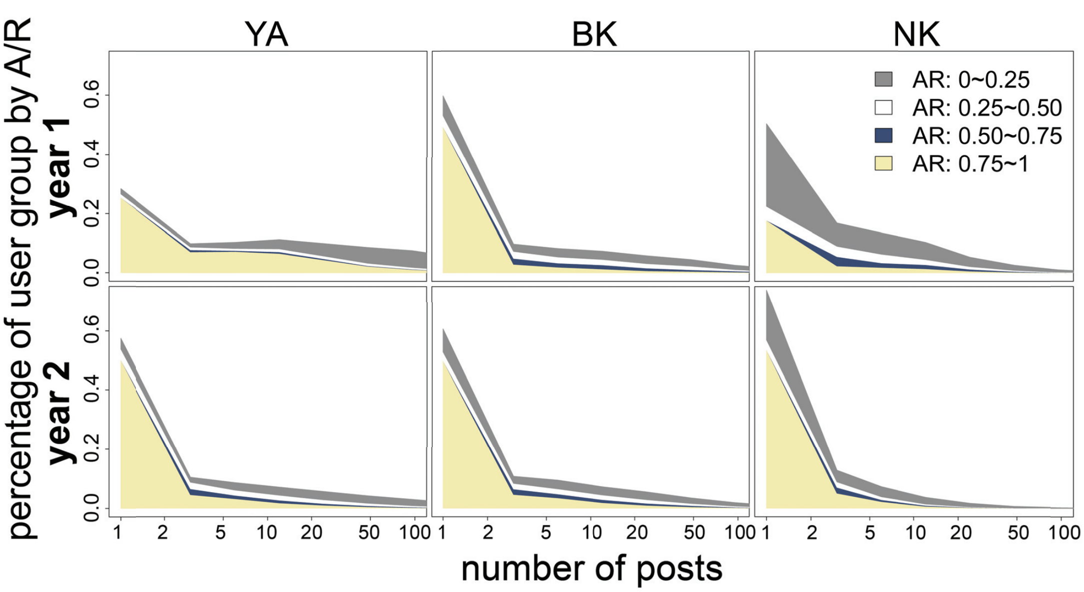

Figure 37: User distribution over #posts and A/R ratio: users are grouped via A/R ratio and presented in different color area; for each level of #post, the length on the Y-axis presents the portion of all users who fall in the combination of #post and A/R ratio

Figure 37 displays the user distributions over two dimensions: #posts and A/R ratio during a year. BK had a rather consistent distribution over A/R ratio and #posts between the years. On the other hand, YA and NK both had a similar shift: users who made a few posts, mainly asking, took over the largest portion of user population in the subsequent year. The shift for NK was more dramatic, consistent with Table 7.

We then compared users' activity lifespans between the two years (Figure 38). Interestingly, all sites presented a decline in survival rate from year 1 to year 2, especially for YA. In the second year, user retention in YA dropped to a similar level as BK, which maintained around 5% users after 250 days. NK suffered more difficulty in sustaining users in the second year as almost no users were left after 250 days. If new users become less committed one year after launch, this might suggest the difficulty of sustaining a quickly expanding population. It may also reflect a difference in enthusiasm between early and later adopters. We do note, however, that YA users were still more active (asking or answering) in both years (Table 7).

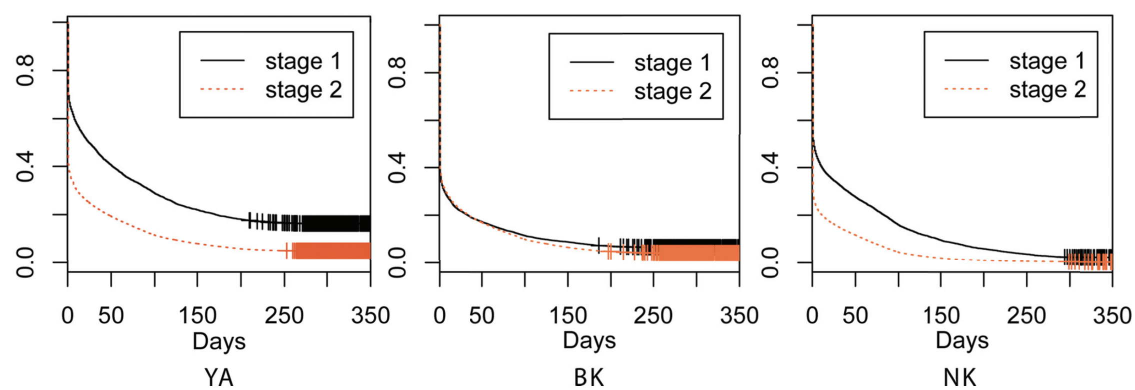

Figure 38: Comparison in user retention between early and established periods for Q&A sites

Lifespan Differences by Category

All three sites have a similar category structure, encompassing topics ranging from science to relationship advice to entertainment. Previous studies found that users have different interaction patterns within categories, according to whether the category is of more informational or conversational nature (Adamic et al., 2008; Harper et al., 2009; Yang & Wei, 2009); While conversational categories usually involve substantial mutual interactions, informational categories consist of more one-directional information seeking and giving (Adamic et al., 2008). Here, we examine whether categories also differ in users' survival patterns, reflecting different commitment to the sub-communities defined by categories.

We found that there is significant and consistent difference in survival patterns among categories on YA and BK, as presented in Figure 39. On both sites, categories like "entertainment" and "trouble/advice," which would be considered as conversational, have higher survival rates over time than informational categories such as "games" on YA where users ask how to access and maneuver different computer games, or "medicine" on BK where users seek medical information.

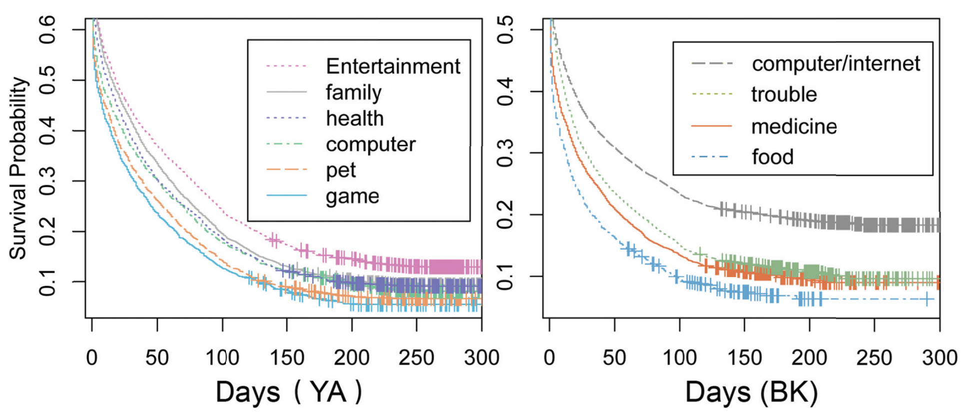

Figure 39: Survival curves in sample categories on YA & BK (Best seen in color)

Note, however, that the category of "computer/internet" on BK is an exception. It is often considered informational but has a significantly higher survival rate than all other categories. We suspect that there is an important factor in this: culturally-based information use.

To examine this further, we conducted human coding to evaluate a random sample of BK and YA questions. To control the variance across categories, we sampled 80 questions from each of two meta-categories of YA and BK: "Entertainment" and "Computer/Internet," which should represent conversational and informational topics respectively. Raters rated each question with all of its answers, in terms of how the question is asked and answered. This integrated question-answering process is evaluated on a 5-point Likert scale from 1 (information seeking and providing objective information) to 5 (social discussion and conversation with subjective opinions and attitude). The Spearman inter-rater correlation was 0.83.

As predicated across sites, "Computer/Internet" categories tend to be more about information seeking and offering, while "entertainment" categories tend to be more social and conversational in terms of both the contents and interaction patterns. In addition, we found that both categories in YA have higher average scores than those of BK, as shown in Figure 40. The matched differences between the sites are statistically significant. This suggests that in general YA's question-answering interactions tend to be more social conversational, consistent with the observation that YA has significantly more answers per question on average, and might also help explain why YA users tend to more frequently switch between asker and answerer roles.

Both the properties of questions being asked and the patterns of answering questions contribute to the Q&A sites' dynamics. BK users ask more questions seeking objective information, prominently questions regarding online resources and computer assistance (e.g., Where can I download XXX?); YA users like to raise discussion topics to garner others' opinions or simply for fun (e.g., What is your favorite website besides this one? Or, have you lived an enchanted life?). On the answering side, we observe that compared to BK users who merely provide answers, YA users tend to add more humor, offer personal opinions, and express sociable statements.

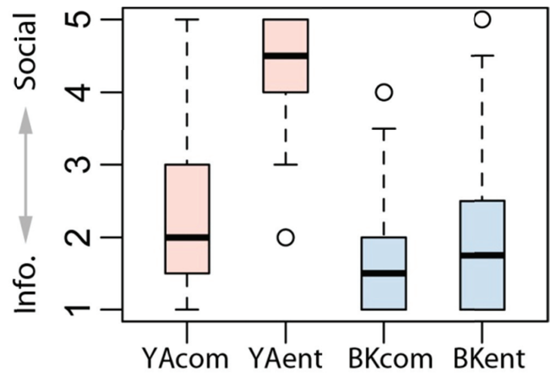

Figure 40: Rating in sample categories on YA & BK

Therefore, we tentatively interpret this significant difference across sites as a consequence of the complicated interactions among: (1) incentive design (2) information needs and (3) cultural difference. First, both BK and NK encourage asking questions by rewarding points while YA deducts points for asking, which may lead to more contribution on the asking side thus leaving many questions without sufficient answers. In addition, the flexible rate of points for answering questions in BK and NK might also create a “transaction” mood rather than a “mutual help” mood, which can jeopardize the intrinsic motivation for social conversation and social bonding. Second, many complicated reasons including Internet censorship have largely limited information availability in China, thus people might have to rely on these human-generated information sources such as BK, especially the information about online resources and computer assistance. The huge demand for such information is also corresponding to the high activity and high commitment level on relevant categories like “Computer/Internet” in BK. Finally, there are significant cultural differences between the East Asians and Westerners: in particular, Western people tend to emphasize individuals’ values, and they are more willing to express their opinions and feelings (Gilbert & Karahalios, 2009; Hofstede, 1980; Yang et al., 2011) Thus the question-answering interactions on YA might be driven by Westerners’ this characteristic to be conversational and interactive. More research, however, will be required to understand this.

3. Discussion

In this chapter we used Survival Analysis to compare users' participation lifespan and social interaction pattern across the three major Q&A sites. First, we examined participation roles, finding consistently across sites that users who preferred answering tend to have a longer, more active life within the site.

While retaining these more committed users helps sustain the Q&A community, generating enough questions for them to answer is also important. As might be expected, askers tend to stay longer if they can get better, more numerous, and longer responses. It is unclear, however, how much of the askers' lifespan is explained by others' responses, and how much of it is explained by the askers' intrinsic motivation. Askers who put in more effort, in terms of the number and average length of questions they write, both get more answers and tend to stay longer.

Answerers who get acknowledgement of their contributions by having answers selected as best or commented on tend to stay longer. This implies a potential need to reinforce the dynamics of information seeking and perhaps offering ways to improve both the askers and answerers' experience (for example, a routing system).

In contrast to earlier studies, which focused on users' initial interactions, we find that such interactions were only very weak predictors of user participation in the long time—and only for two of the sites. Users whose first action is to ask are a bit more predictable than those who first post an answer. This suggests that askers are more sensitive to how their first interaction goes. Thus, in order to help keep new arrivals, it might be useful to help first-time askers by offering help or wizards about how to formulate and post a successful question.

There are also some intriguing differences between the sites. The most noticeable cross-site difference we found was that answerers tend to be more active in providing answers on YA, while askers ask a similar number of questions on all three sites. There is a higher rate of answers per question on YA compared to BK and NK. In addition, YA users tend to stay on the site longer than BK and NK users, showing a stronger commitment to the Q&A community.

We believe that there may be a subtle cause behind this difference: incentive design. Both BK and NK encourage asking by rewarding askers with points, while YA deducts points for asking questions and rewards answers. While incentivizing asking activity may have contributed to the incredible growth in question volume on BK, it might also bring issues such as high drop-off rates, and an insufficient supply of answers or even large numbers of unsolved questions. An additional possible factor for the high volume of questions with few answers, and the low retention of askers on NK and BK is that many casual askers may come to the sites through search portals provided by Baidu and Naver.

By further looking into people's Q&A interaction contents, we found that the interaction patterns on YA tend to be more conversational than BK. This might partially account for the higher answering contribution on YA. But how the sites developed different question-and-answering dynamics can be a complex result of design differences and cultural contexts. First, it might be that the incentive designs and flexible points on BK stamp out people's intrinsic motivation to help and create social bonds. Instead, users may treat the site like a series of transactions. Second, the large amount of informational questions on BK could be because China often lacks comprehensive and up-to-date information sources for information like store hours, bus schedules, and phone numbers. Therefore, people spend more time asking other people these questions. In addition, the difference in the interaction dynamics also hints at potential cultural characteristics, as Westerners tend to be more willing to express their opinions and feelings and involve more interactive discussion in their Q&A dynamics.

Chapter 5

Virtual Currency in Supporting Social Exchange in a Cultural Context

Various information services and systems (e.g. eBay, Orkut, and Yahoo!) have encountered serious challenges when entering China, an emergent and promising market with 400 million Internet users. We argue that in order to successfully localize, such services need not only adequately navigate the current Chinese economic and political landscape, but also need to account for the deeply rooted Chinese culture.

To address the need to understand how Chinese culture interacts with online systems, in this chapter we present a case study of diverse social interactions among Chinese netizens, based on over 4 years of comprehensive data collected from an online bulletin board system (BBS), Mitbbs. Mitbbs is the most frequently used online forum for Chinese nationals who are studying or working abroad, primarily in the United States. Because Mitbbs is hosted in the US, it is less affected by censorship than forums located in China.

Founded in 1998 by volunteers, Mitbbs was later commercialized and is supported through the sale of advertisements. However, in essence, it has been sustained by the hundreds of thousands of Chinese who are seeking both help and a sense of community during their stays abroad. Similar to the experience that most Chinese young people had with their college bulletin board systems (BBSs), i.e. participating in a virtual community and developing social networks, Mitbbs supports a significant part of its users' informational and social life abroad. This could be seen on the *Anniversary* board, created on the 10th anniversary of Mitbbs' launch, where users have posted their experiences and memories of Mitbbs. For example, a user, *anvv*, who had used Mitbbs throughout his ten years abroad, wrote a post titled "an 'unknown'[9] dream for ten years".

A community such as Mitbbs can succeed only if it can motivate users to exchange information and socialize. To this end, Mitbbs introduced "*weibi*," a virtual currency system. For example, posts promoted to the front page earn their authors *weibi*. However, despite the limited and rather unimaginative initial prescribed use of *weibi*, the virtual currency system evolved to be an essential mechanism in Mitbbs because it supports critically important social interactions. It is through the lens of *weibi* that we study social interactions in this community.

As the virtual currency was adopted by Mitbbs' users, its uses evolved in a very culturally specific manner: It quickly began to serve as a mechanism for social exchange activities termed *guanxi* in Chinese. Thus, the online social interactions can reflect real life Chinese social dynamics. *Guanxi* networking (to be explained further below) has been viewed as an "informal aspect of the institutional culture" (Walder, 1986) and a stimulus of social actions (Alston, 1989). Virtual points, through their flexibility and ambiguity, allow users to carry out socially important and culturally nuanced *guanxi* behavior.

This chapter is organized as follows. We first provide the literature background that is crucial to understanding the social interactions of Chinese users in Mitbbs. We then introduce the Mitbbs system and show how its diverse social exchange is supported through the interaction between the virtual point system (*weibi*) and the *guanxi* networking dynamics. Finally, we discuss the implications of our findings and conclude.

[9] "Unknown space" is a nickname for Mitbbs, based on "Unknown BBS," its precursor hosted at Beijing University.

1. Literature Background

We ground our study by describing three general streams of literature: cultural differences between Westerners and East Asians, including studies within CSCW, studies about *guanxi* in China, and virtual points in online communities.

In inter-cultural sociology and cultural psychology, Westerners and East Asians are often categorized as belonging to two differing groups. In terms of this literature, Westerners tend to be labeled as more analytic while East Asians tend to be more holistic; and thus, Westerners are context-independent, more narrowly focused, and use formal logic, while East Asians are field-dependent, broadly focused, situational, and dialectical (Nisbett et al., 2001; Varnum et al., 2010). In terms of social orientation, Western cultures tend to value independence, individualism, autonomy, and self-achievement (Hofstede, 1980); in contrast, Asian cultures emphasize interdependence, harmony, relatedness, and connection (Hofstede, 1983; Singelis, 1994; Triandis, 1995). Cultures of independent-orientation tend to view the self as bounded and separate from others, while interdependent-orientated cultures view the self as interconnected and encompassing important social relationships (Markus & Kitayama, 1991). Thus Asian cultures are "characterized by belonging, mandating the fulfillment of obligations and responsibilities to others" (Heine et al., 1999). In addition, many generic differences associated with each culture may also interact with and shape the culture, such as value systems (Aristotelian vs. Confucian intellectual traditions) (Lloyd, 1996; Pye, 1985), religions (Dollinger, 1988), economic ideology (Ralston et al., 2007), and industrialization and geographic mobility (Kitayama et al., 2009).

CSCW has generally followed this literature, although there is wide concern over its limitations for design. For example, Asian users have been found to prefer multi-party chat, audio-video chat, and emoticons in IM (Kayan et al., 2006), benefit more from rich communication media in negotiation (Veinott et al., 1999), and are less satisfied with asynchronous communication (Massey et al., 2001). Setlock and Fussell (2010) found that

Asians involve additional considerations when deciding on appropriate communication tools, especially on the ability to support social processes. Lindtner et al. (2009), however, points out that these contrast-focused approaches can force a problematic distinction between "here" (the West) and "there" (China). Instead, as Lindtner et al. did in their study of gaming in Chinese Internet cafes, we focus here on specific Chinese practices around social exchange and their re-enactment in digital environments so as to avoid this reification.

Central to our study of social exchange is the notion of *guanxi*. *Guanxi* is a major theme in social interactions in Chinese society. Those familiar with Chinese culture would not doubt its importance, but Westerners find it hard to grasp. Essentially, *guanxi* describes the ties between an individual and others (Jacobs, 1980), fostered through exchanges of favors (Pye, 1982). It is based in a sense of *renqing*, sometimes translated as harmonious relations. As Yang states:

> An important feature of *renqing* principles is the notion of the necessity of reciprocity, obligation, and indebtedness in human relations. What activates reciprocal relations, what imbues relationships with a sense of obligation and indebtedness are the work of relational sentiments and ethics. Concrete expressions of *renqing* are found especially in the gift-giving that goes on at special occasions such as birth, deaths, weddings, and New Year's. (Yang, 1994, p. 122).

Gift exchange plays an important role in establishing and sustaining *guanxi* networks. Yang (1994) and Yang (1996) in their ethnographic studies found two characteristics in *guanxi* to be prominent. First, reciprocal obligations for favors are assumed (Hwang, 1987), and the interactions are designed to cultivate mutual dependence and manufacture obligation and indebtedness (Yang, 1994: p.6)). As Kipnis states (p.307):

> There is ... a congruence between the size of gifts, the burden of obligation, the strength of feeling that either existed or was hoped to develop, the closeness of the *guanxi*, and the dependability of the *guanxi*.

Second, gifts tend to be perceived as equivalent to money, and they can be circulated, calculated, and compared (e.g., cash gifts, gifts recycled to another person).

According to Chiao (1982) and King (1991), *guanxi* is based on and enhanced by shared social experiences among individuals. For people in non-hierarchical or family relationships (i.e., among friends), *guanxi* connections are a primary mechanism of Chinese social life (Farh et al., 1998; Tsui & Farh, 1997). *Guanxi* is essentially "not only instrumentality and rational calculation, but also sociality, morality, intentionality, and personal affection" (Yang, 1994: p.88).

Guanxi is also viewed in the sociology literature as a social mechanism substituting for formal institutions in current Chinese society (Xin & Pearce, 1996). "The structure of social relationships in China rests largely on fluid, person-centered social networks, rather than on fixed social institutions" (Yang, 1994: p.14). As mentioned, *guanxi* is difficult for Westerners to understand. While it can be compared to social capital (Bourdieu, 1986; Portes, 1998) in some ways, *guanxi* is more oriented toward dyadic relationships and is less societally structural, allowing *guanxi* networks to be freely connected and often bridge institutional boundaries (Yang, 1994). Therefore, social capital adheres to and affects (positively or negatively) a social unit, but *guanxi* networks are fluid and autonomous with respect to any institution. It might be noted that despite active research on social capital in Internet contexts, how *guanxi* networks evolve in Chinese online communities is understudied.

Table 15 lists the Chinese terms that will be important to our discussion in this chapter, including a rough translation. The reader is reminded that terms relating to Chinese culture seldom translate precisely to English, and it is important to focus on the Chinese concept rather than its English translation.

Table 15: Important concepts important in Chinese culture

Term	Literal meaning	Free translation
renpin	moral quality	close to the popular use of the term "karma" in the US
RP	same as renpin	abbreviation of renpin (popular online slang)
renqing	human relationships	human feelings, emotions, relationship, favor
guanxi	relationships	ties between individual and others, through the exchanges of favor

The third line of literature concerns virtual points and gifts in online communities. Virtual points are often used to motivated contribution and participation in question-answering (Q&A) forums. Interestingly, while the English site Yahoo! Answers used a fixed-point rate per question and answer, the two major Asian sites (Baidu Knows in China and Naver Knowledge-iN in Korea) allowed users to make variable point offerings to obtain answers. Yang and Wei (2009) found that more points can attract more answers in Baidu, and over time users learned to optimize point expended per answer gained in both Baidu (Yang & Wei, 2009) and Naver (Nam et al., 2009). Yang and Wei (2009) also revealed how Chinese users priced the questions differently based on topics and degrees of importance. Wang and Mainwaring (2008)'s study of virtual currency usage in the largest Chinese social service-Tencent QQ examined the perceived value of the virtual currency and its complex interaction with the currency type and the contexts of obtaining and spending the points. Finally Hjorth (2008) noted the pervasiveness of gift-giving in Korean CyWorld use. These studies all hinted at a very diverse, flexible, and contextualized usage of virtual currency and other gifts to support the complex online social exchange activities in Asian culture, and the current study will address this in depth.

2. Mitbbs: Representative and Unique

Founded in 1998, Mitbbs is the largest and most frequently visited online forum for Chinese people in the US. As of 2004, there were more than 300,000 users visiting the site monthly, and about 20,000 are currently online at any given time. Most users share the common background of being the first generation studying or working in the US and they are typically highly educated. With the arrival of 10,000 new Chinese students to the US each year, Mitbbs continues to enjoy high participation.

Mitbbs shares major design features with other online Chinese and non-Chinese forums, including the format of the homepage, the subdivision of forums by topic, and supplementary blogging and social-networking services. On Mitbbs boards *ChinaNews*, and *Military*, people share and discuss news -- usually Chinese news, which is often hard

Figure 41: Typical post page for Mitbbs: a user is asking "how to review a review-paper" in *Immigration*, where there might be senior people in academia to answer this question.

to access and discuss on other Chinese forums or websites due to censorship. They help one another by sharing information on various boards such as *Job-hunting*, *Immigration*, *Next-generation (child-rearing)*, or *Postdoc*; share money-saving tips or conduct business on *PenneySaver* and *eBiz*; seek a romantic match through *Pie-bridge*[10] and hobby-buddies on forums such as *Movie*, *Tennis*, and *Photography*; meet local people on forums such as *Michigan* or *Seattle*; and find schoolmates on the alumni boards.

When participating in discussions of sensitive topics, users may not want their ID and true identity to be linked. *Majia,* or alternative IDs (see Table 2), are thus frequently used. On the other hand, as in many other webboards, people on Mitbbs post images, including pictures of themselves. Posting a photo is a frequent self-disclosure activity, and the site explicitly encourages this through awarding points and holding various campaigns or contests. For example, people model their dresses on *Fashion*, show off their muscles on *Fitness*, pets in *Pets*, and even body parts on *Sex*. For example on *Pets*, *jackyang* posted some pictures of his son and dog:

> Ally (name of the dog) and his younger brother: taken 5 months ago. [jackyang, 08-04-2010]

Around twenty people replied, including:

> Wow, your son is so tall now, time passed so quickly, in just a second he went from a baby to a handsome boy. [Ted, 08-04-2010]

> I'm admiring… Ally [she] is still too beautiful!!! [magua, 08-04-2010]

> It has been so so so long since I last saw you [xiaoshu, 08-04-2010]

Flame wars are also frequent on Mitbbs. Standard topics for flame wars include "democracy in China," "whether one should return to China," "should one buy a Japanese car or American car," and "should Chinese girls date Americans." These can entice

[10] A legendary bridge where couples meet.

threatening or abusive posts, which might result in users being banned or board masters impeached.

Mitbbs posts are asynchronous, but because of the large user base, interaction can be quick. For example, *little9*'s post to the *Soccer* board obtained 143 replies within 3 minutes when thousands of users gathered on the board during the 2010 World Cup. On the other hand, discussions about controversial topics can potentially last months.

It should be noted that in Mitbbs, users can remove some posts and images, allowing an interaction pattern closer to synchronous systems. For example, as mentioned, a typical post on the Fashion forum is a participant's presenting herself in a favorite dress:

> Just had final(exam)today… I am so bored now I want to show my new dress. Is anyone interested in seeing? [eggpiggie]

After receiving some replies expressing interest, the picture will be posted for minutes to hours and then deleted from the original post. A later reply was regretful:

> Oh, I missed it again! Can you show me one more time? [cocoLily]

While Mitbbs has a number of unique interaction characteristics, those are not the central concern of this chapter. We focus here on the use of weibi for social exchange.

Table 16 lists some of the specific terms that are extensively used in Mitbbs discussions and which will be important in this chapter.

Table 16: Mitbbs Jargon

Term	Literal meaning	Meaning on the forum
weibi	fake currency	virtual money on Mitbbs
baozi	steamed stuffed bun	a pack of 10 weibi[11]
ben	run quickly	1. post pictures of one's self; 2. virtually present in public
majia	vest, shell	alternative ID one might have to post with particular concern

Mitbbs structure and coordination

Similar to other Chinese webboards, Mitbbs is structured with more than 300 sub-forums (called boards). These boards are grouped into 12 large categories, such as "news," "oversea life," "sports," and "alumni." Each board is a space for people to post on a given topic. The centralized homepage provides links to highly ranked threads and boards.

Each board is coordinated by volunteer administrators. There is one "board master" (BM) with several (1 to 5) "vice board masters" who are subordinate. These BMs mark or promote posts; edit (e.g., delete impropriate posts and archive old posts); reward, warn, or ban users; coordinate discussion; manage the board's balance of virtual points; and, organize events (e.g., organize a special event with awards for posts). BMs themselves are organized through a board called the "Family of BMs," where people can propose to initiate new boards, to be a BM, or to complain about a BM. A "station master" sysop can then make decisions based on this information.

[11] Very interestingly, Yang (1994) observed villagers circulating real *baozi* in their gift-exchanging during the Chinese spring festival. In the old time or rural countryside in China, food such as *baozi*, eggs, wine, tea, and cigarettes were popular gifts among people, because they are commonly welcomed and can be recycled as gift for other people.

In addition, new or updated BBS policies are also posted, discussed, and modified through the thread format on boards. In fact, new policies have been invented, discussed, and institutionalized through the history of the BBS. For example, some users do not like their posts getting promoted onto web site's front page (which may bring too much attention to the concerned parties in the post), thus there is a new policy that an author or BM can prevent this promotion from happening by adding a tag in the post's title. Sometimes the policies are local to a board. For example, on the *Military* forum, a frequent topic for a flame war is "Chinese girls dating Americans." In July 2010, the *Military* BM posted the policy "whoever raises this topic again will be banned for 3 days."

Mitbbs' coordination system reflects a mixture of hierarchy and autonomy. In many ways, the administration is similar to any webboard: Users are moderated by BMs, there is a subordinate relationship between chief and vice-BMs, and stationmaster's authority overrides any decisions. In addition, users can voice their opinions on any issue, which can affect the administrators in that they need to satisfy users to the extent that they will return. On the other hand, there is a deference to central authority, or at least an acknowledgement of it, that is unusual on Western sites.

Finally, a point system or virtual currency was introduced to incentivize participation. As we will demonstrate, points have been freely used in a variety of social interactions for diverse purposes and across contexts beyond their original intended use. We will describe the use of these virtual points following a brief description of the study.

3. About the Study

The author was a casual user of the Mitbbs site for over 4 years. Mitbbs forums have been a significant part of her life: providing political or entertainment news, offering experience and advice about problems with living in the US, as well as random surfing.

To further analyze Mitbbs behavior, we first read more than 2,400 threads from a period of 10 months. Some of these threads were translated by the author for subsequent analysis and to explain how the board was used to other co-authors. From the same time range, we selected more than 600 exemplars that were representative instances of Mitbbs social interactions around *weibi*. These threads were then translated, analyzed, and categorized in discussion with the other authors. We further scanned each board by querying for the keywords *baozi* and *weibi* in the thread titles, to obtain a basic sense of how participants use virtual points on each board. We identified a set of popular boards that extensively involved point exchanges, such as *Ebiz* (e-business), *NextGeneration* (child-rearing), and *Fashion*. These were compared to boards with few point interactions, such as *Military*, *Returnee*, and *Family* (which is mainly about controversial family issues, for example, "I fought with my father-in-law because he smoked indoor").

In addition, to gain the participants' perspective about their interactions, we conducted 1.5~2 hour interviews with 13 Mitbbs users. Nine interviewees were from the author's personal social network, including friends of friends. Four additional volunteers were recruited via messages sent to forum participants. We also interviewed a board-master, who provided insider information on how the point mechanisms work. Most interviews were conducted through Google Talk, since using IM rather than other mechanisms allowed the interview subjects to release only a handle which could not be connected with their Mitbbs ID or real name.

In this chapter, we have pseudo-anonymized all names and removed identifying detail.

4. Points on Mitbbs

Mitbbs launched its virtual point (*weibi*) system in 2006. The system was intended to reward contribution and administrative jobs. To mimic a real-world currency system, the site allows users to deposit (to get interest), transfer, and exchange these virtual points. *Weibi* are often exchanged as *baozi*, or units of 10 *weibi*.

Users can earn points through writing a post on one of the boards (0.1 *weibi*), having their post selected or promoted (10~100), posting pictures (10~20), gambling profits, and receiving them from other users.

Points can be used in a number of ways. The primary uses of points are modifying avatars, gambling, giving to a board account, and social exchange. Each is covered in turn.

Avatar fashion. About 10% of Mitbbs users who post display avatars next to their posts. These users must "clothe" their avatars on a regular basis, and this requires points. There are hundreds of items (e.g., jackets, purses, pets, facial expressions, hairstyles) available, priced from zero to 50 *weibi*. New items expire in 45 days after purchase, while second-hand items will continue only for the remainder of the 45 days since the first purchase. If the user does not clothe his avatar, the avatar will be shown in underwear. One of our interviewees tells us that she needs *weibi* to buy some clothing, "when I want to post something, ... I don't want to *ben* in underwear.... I need a fig leaf. ... [The] cheapest or second-hand works" [I2].

The expiration was intended to encourage people to continue to earn and spend money for their avatar. The BBS also operates regular contests for best dressed avatars based on the votes from other users. However, as mentioned, only 10% posts are by users who have avatars in their profiles[12].

Gambling. Gambling can be a relatively fast way to earn or spend *weibi*. Gambling can be run by an individual, but mostly BMs run gambling as the banker in the name of a board. Gambling themes are diverse: people bet on stock values, soccer matches, birth dates of children, exchange rates for the RMB (Chinese currency), and even when the

[12] People can change their display image from a picture to an avatar, or the reverse. For example, during the World Cup 2010, when users could buy their favorite team's uniforms, we observed relatively more users using avatars.

Amazon.com website will crash. Gambling on stock values is often an on-going activity, while others like soccer scores are seasonal or event-driven.

Board accounts. Each board maintains its own *weibi* balance. BMs reward users through board accounts for high quality posts or for participating in posting or presenting campaigns. Board accounts can also fund gambling. This will be discussed further below.

Finally, users can give points away to other users. This is discussed at length next.

5. Virtual Points in Use

Weibi was intended originally to motivate contribution. However transferring points among individuals soon became the medium for a diverse set of social interactions, prevalent in many of the boards.

Value of *Weibi*

It is important to note that *weibi* is not officially convertible to real currency, and its real value is ambiguous at best. For example, one interviewee assigned a small monetary value to *weibi*. He recalled how he had used 100 *weibi* to buy a "15-off-75" Staples coupon (i.e., a coupon to save 15 $US on a 75 $US purchase), which he thought was a good deal: "You need real money to buy coupon on eBay.... Once I saw some people use US dollars to buy *weibi* at 150:1, which means I paid less than 1 dollar for that coupon" [I10] (Occasionally, users in need of *weibi* will post to a board asking to buy them from other users). Others do not perceive any monetary value for *weibi*. An interviewee who gave 1000 *weibi* to his friend for gambling, said, "*Weibi* is worth nothing [in real life]" [I12].

Instrumental Uses of Weibi

Rewarding and incentivizing others' contributions is frequently observed on Mitbbs. Often, users post questions seeking serious and professional answers with a promise of *weibi*. For example, users have needed to know about house closing costs, formatting green card application letters, medical symptoms, lowest possible prices for computer equipment, or even how to find a Dell customer support "phone number for a living person." Offering points in hopes of obtaining better answers is similar to the use of points in question-answer forums (Yang & Wei, 2009; Yang et al., 2010).

People also use *weibi* to gather others' attention. *Baozi* may be offered to people to promote a post onto more prominent positions (especially to the Mitbbs homepage). For example, one user offered *baozi* to those who would reply to her post about Jian-lian Yi, a Chinese star player in the NBA, in the hope of having it promoted to the front page. Another user showed his loyalty to a national soccer team by offering a *baozi* "reimbursement" to those users who would purchase Argentine uniforms for their avatars to support the team during the World Cup. *Baozi* has been used quite frequently in donation campaigns, e.g., two users offered 1 *baozi* to each of up to 50 users who would support (by replying) a post calling for donations for Qinghai (China) earthquake victims. Another user promised to give away all his *baozi* for votes in the "Chase Community Giving" campaign contest on Facebook.

Less instrumentally, people also award others for a good post they encounter. For example, when one user enjoyed reading a post, she sent a *baozi* to the author and also posted the reply, "Hey, I really like your post, I will give you a *baozi*." These "afterward" *baozi* gifts act to further social interaction and one's *guanxi* network. One of our interviewees has posted several times on "how to apply makeup" with her photos demonstrating different techniques. She got a lot of compliments in the thread, but also, she received several messages with *baozi* attached. She said, "*baozi* is useless for me. I share for fun… [but] I am happy to get these [messages]" [I13]. Another interviewee also received *baozi* with questions regarding his post about job seeking, "I replied in detail… and I think he added me as friend [on the site]" [I4].

Purchasing Favors

Exchanging favors is often done between pairs of people who have good *guanxi*. *Weibi* can facilitate these kinds of exchanges among strangers on the BBS. Since users on the BBS share the common identity of Chinese students and very similar life experience (e.g., studied in the same schools or worked in the same places before and after coming to the US), the community formed a "small world network" in which people are closely connected with one another. Thus these strangers, although outside of one's preexisting social (and *guanxi*) network, may be only a couple of degrees away and are very likely to develop new *guanxi* in the future in such a small-world community. These indirect *guanxi* links are very important in Chinese culture as indicating potential of new *guanxi* development (Hwang, 1987; Yang, 1994). Most of these exchanges will not be done without this kind of connection, and *weibi* serves as "indirect payment" when one seeks favors from outside of one's preexisting networks (Yang, 1996).

Illegal copies: One interviewee [I5] told us she spent hundreds of *weibi* to buy books in PDF format from other users on the site. For example, she bought an unlicensed copy of "Career Cup," a book on answering technical interview questions, from another user for 100 *weibi* (sold as a legal electronic copy for 29 $US online). This was a deal off the thread and board, but we could infer many such transactions from people's relevant conversations.

Review referral: A user posted a request on the *Faculty* board, asking for an opportunity to review journal or conference papers. One's review record is crucial when applying for an EB-1 green card (for US permanent residency), and there is considerable discussion about this in relevant boards such as *Immigration* or *Faculty* (who are very likely to apply for this type of green card):

```
Journal: 20 baozi

Conference PC(Program Committee): 5 baozi
```

```
My research direction is data mining in computational
biology. And later I switched to information retrieval on
mobile devices. Send me BBS mail, and I will send you my
vita.[happyLife]
```

Similarly, *horseJean* asked for code that could be used to compute a "tight-binding model in NanoST". Another user wanted a sample reference letter for faculty job applications, and she offered 5 *baozi*.

On the *Automobile* board, *baozi* is frequently used for checking a car record with a given VIN number. Some people purchase the Autocheck or Carfax service when checking the record of a used car, and they usually can look up the records for more than one car. This favor is done for free as the person has already spent the money, and people started to use *baozi* in exchange for this help. This has recently started to change, as people have begun to ask for money to do the checking.

Collecting Renpin

Collecting *renpin* (often "RP") is a very common use of *weibi* on Mitbbs. *Renpin* in Chinese was originally used to describe one's character or moral quality. Online, its meaning has shifted to more commonly connote something akin to "karma in present life". Positive actions or deeds can accumulate *renpin* and result in later good luck.

Renpin reflects a mixture of karma from Buddhism and the norm of reciprocity in people's *guanxi* networks. However, *renpin* is a kind of karma that will pay out within one's present life. In addition, Chinese people practice *guanxi* networking by cultivating mutual dependence and exchange of favors (Hwang, 1987; Yang, 1994), and people who do not follow the rule of equity will lose *renpin* and be considered untrustworthy (Alston, 1989).

Netizens have even derived a "Law of Conservation of Character[13]": one needs to spend certain amount of *renpin* in order to get good luck in a particular situation, and if one gets bad luck that is because he has used up his *renpin*. It has been observed that *renpin* often drives people to "do good" in social interactions on the Internet in the absence of other social norms and religion. While some interviewees were skeptical of *renpin*, many more expressed sentiments like "I don't exclude the possibility to collect RP by distributing *baozi*."

On Mitbbs, *weibi* plays an important role as the medium of collecting *renpin*. People distribute *baozi* to accumulate "blessings" from other users when they are hoping for a good outcome, for example, for a pregnancy, a parent's visa interview, a spouse's upcoming job interview, or a new romance. As another example, a user, *oke*, said in her post that:

> Baozi on the Baby board (NextGeneration) was really effective. I had a very smooth delivery of my baby, after I sent 66 baozi when it was 4 days overdue.
>
> Now I want to distribute 66 baozi again, asking for blessings on our 2-month overdue greencard approval. [oke,03-24-2010]

Another user gave away "double-filling" (20 point) *baozi* in order to get rid of his "bad luck":

> I bought baozi specifically for this. I was really unlucky in May, Lost a package around 3000 bucks Got a 800 bucks ticket and got my car back with another 200 bucks and got a very bad negative feedback14 and threw away my contact lens as trash. Everyone please give me some luck… [IamLegend,06-22-2010]

This post obtained more than 150 replies although it only offered *baozi* to the first 20 people who sent blessings. Many people replied with posts such as "really bad fortune-

[13] http://baike.baidu.com/view/1586.htm

[14] As a seller in eBay

loss," "serious blessing," and "endless good luck." The posting also received numerous suggestions such as to wear "something fortunate" such as crystals.

One of our interviewees described his experience of distributing *baozi:* "I requested blessings for my doctoral dissertation …It was a complicated situation, and I got [my thesis] signed by my committee on the last day …The process, anyway, was very tricky." He sent *baozi* to each of the 100 users who replied, because "it would show my sincerity." He believed he had collected RP for his thesis process: "Eventually I was surprisingly lucky… It passed and I believe the *baozi* worked" [I3].

Banquets of Baozi

Banquets are one of the most prominent social instruments to sustain *guanxi* networks (Yang, 1994). On Mitbbs, people frequently hold "banquets of *baozi*" to celebrate various events. According to the "theory of *renpin*," one needs to re-accumulate *renpin* as one "redeems" a portion as good luck. Thus people need to "do something good" by gifting back to the community to keep the "*renpin* balance." Akin to food and drink given by hosts in real-life banquets or *hongbao* (a red small envelope containing cash) (Yang, 1994), the *baozi* that is given on Mitbbs is considered a carrier of luck, thus rewarding the community.

A few examples of banquets of *baozi* include a successful delivery of one's baby, a successfully obtained visa, an approved green card application, an accepted offer on a house purchase, and getting job offers. The success can be smaller as well: the successful sale of one's used computer, finding a good deal on a purchase, or even celebrating Spain's 2010 World Cup championship.

Baozi are not the only kind of gift item used for this. As we observed on the *Job-hunting* board, many users provide *mianjing*[15] to reward the community. For example, one interviewee asked for blessings before her husband's interview, and she urged him to post his *mianjing* onto the board as she promised.

> I did not have many points so I did not give. I felt an obligation to share mianjing …I think there are some people who blessed us because they want [mianjing]… it can somehow help people...It is returning the favor.[I2]

Often, people give both *baozi* and useful information together as a reward to the community. For example, one user gave away *baozi* for receiving her Eb1a card (the highest priority greencard) and shared comprehensive information about her background.

Gifting W*eibi* to Friends

Weibi exchange is not limited to simply enhancing one's *guanxi* online, but can be exchanged as a gift among friends to enhance both online and offline relationships. However, unlike other *weibi* uses on Mitbbs, *weibi* transfers among offline friends is often invisible. Despite this, we saw many cases where people indicate their transfers of *weibi* in discussion threads. For example, a user transferred 500 *weibi* to the *eBiz* board "sponsoring *Xiongxiong* to distribute [his/her] *baozi*."

Our interviewee who distributed 100 *baozi* for his dissertation also got many *weibi* from his offline friends: "I had accumulated a few by myself…and I know a couple of rich guys. I asked many from them, hah hah!" [I3]. Another interviewee said he had given 1000 *weibi* to a friend from college, "*baozi* is worthwhile… 1k *weibi* can make [my friend] very, very happy, why not?" [I12].

[15] Literally translated as "scripture of interview," this is where people write about the experience they had with the interview, especially "what kind of interview questions one has been asked and how he answered."

Individual-group interactions

Guanxi exists not only between two individuals, but also between an individual and the group. Correspondingly, *weibi* can also be utilized to reinforce relationships between an individual and a community (often within a board). For example, a user, *yue*, felt sad to see some users leave the *Connecticut* board, and offered them *baozi* to stay in touch:

> Several key people of our board – kekeLee, catFish, and yunQi are leaving! Sigh, I just got to know people here. I am moving too, but it is good that I will stay in Connecticut…
>
> I will give baozi to all tongxu[16] who are moving, welcome whoever is coming and [say] farewell to whoever is leaving. …I wish all you happiness anywhere. …Those who are leaving please come back to chat when you get time, friendship is forever…
>
> There is a condition for eating [my] baozi: people moving out need to tell where they are going then I can find you later. People moving in also please tell me where you live, we can take care of each other. [yue, 07-21-2010]

People often donate to a board they liked as the reward to its community. For example, *appleSky* donated 500 *weibi* to the *eBiz* board, and he/she said in the post: "I just donated 500 to *eBiz*, come on, let's donate *baozi*, accumulate RP, and build our board together." It is also quite common that people donate part of their *weibi* to the board account and ask the BM to help distribute the rest. A mutually beneficial interaction can be realized through this process: the donor can show kindness to the BM by offering points to both the board and people on the board, while the BM can help distribute *weibi*, bridging both the donor and other community members in *guanxi*.

These community-rewarding activities often take place on boards where mutual help is appreciated and community is cherished. For example, on *NextGeneration*, people share knowledge and experience, and support one another going through the process of

[16] Classmates or schoolmates, general names for young people who are likely in school.

becoming new parents, while on *Pets*, people not only share experiences of their loved pets, but often defend against outside pet-haters. By donating *weibi*, people show their appreciation and good wishes to a board and its community. One user wrote "thanks to the *Baby* board" when donating *weibi* to the board. Another user, *maya*, wrote on his return to China:

> I have not figured out how to 'cross the border[17]' onto Mitbbs, so I give away all my baozi I have…and, I want to give my best wish to the board, being still thriving and pure, in this superficial society. Farewell… [maya, 04-02-2010]

Another kind of individual-group interaction through *weibi* involves sponsorships for on-board activities. For example, *Soccer* has an approximately 2 year old tradition of bidding for a board-logo sponsorship with a rate of 100 *weibi* per day. This allows fans to promote their beloved players or teams, as the winner will get his team logo on the right-hand side of the thread list. As well, one interviewee stated that people donate *weibi* as a form of registration fee or title sponsorship in food contests or online game competitions.

6. Discussion

Weibi is of token value, and the system has support for them – this has fostered their use as an important resource for the Mitbbs users. Indeed, *weibi* serves as a critical lubricant for a wide range of Chinese customs and norms. We have shown above how people on Mitbbs use *weibi* to serve a number of purposes they find valuable: The Mitbbs users, Chinese students and workers in the US, value *guanxi*, or their networks of reciprocal obligations, as would any person in China. *Weibi* can also serve to foster *renpin*, or karma conservation. It is hard to overstate the fit to Chinese culture that the uses of weibi have on Mitbbs.

[17] Mitbbs.com is forbidden in mainland China, although people can gain access, for example, by using a VPN service from abroad.

Weibi works for these purposes on Mitbbs because the virtual points are ambiguously valued, can be given away in flexible ways, fit socially valued goals, and are visible. Its use, however, is more explicit than perhaps would be desired.

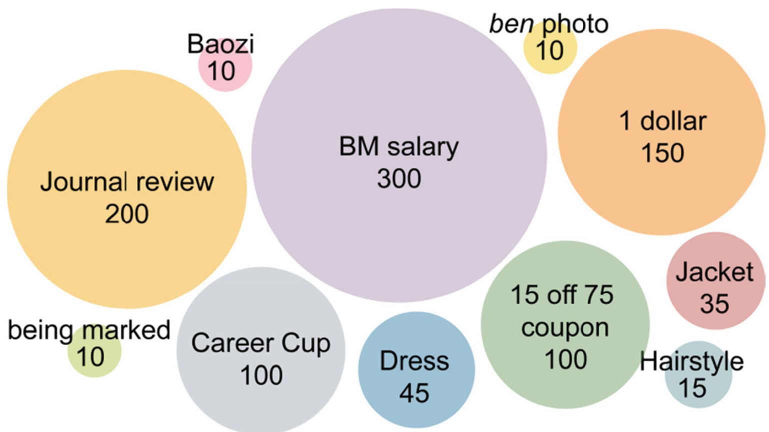

Figure 42: Examples of *weibi* exchange rates

Sweet Spot between Nothing and Value

Because of the vague value of *weibi*, it serves well as the sweet spot between gifting nothing and providing a small gift (such as a token amount of real money). A classic economics experiment serves as a counter-claim: a small amount of monetary incentive can be more detrimental than paying nothing (Gneezy & Rustichini, 2000). *Weibi* is **not** real money, especially when disguised by the name of *baozi*. Instead it is the carrier for kindness, blessings, good luck, and social obligations. In a Chinese setting, however, *weibi* appears to work well for a large range of social interactions, even though its value is ambiguous, if not nil. In fact, *weibi* has somehow blurred the boundary between extrinsic motivation and intrinsic motivation, and works in this capacity because of its fit to the fluid and flexible *guanxi* and *guanxi* relationships.

Boosting Social Interactions

Since *weibi* can circulate, social interactions are thus boosted through it (but probably resulting in less tension than with real money). Contributions on Mitbbs appear to be encouraged through the use of *weibi*, and positive social interactions also appear to be encouraged. For example, users are willing to gift *weibi* to show their appreciation for a posting, and they post their pictures and valuable information. As well, users broadcast and celebrate their personal news through distributing *weibi* to other users, while obtaining plenty of admiration, praise, and blessings.

Since the value of *weibi* is contextual and perceived differently by different people, it can thus be flexible and substitutable as an instrument in these reciprocal social exchanges. For example, one user had posted pictures of real *baozi* she made, to thank the board for helping her settle in Seattle and to introduce herself to the community.

Karma Conservation

Weibi in Mitbbs facilitates both the processes of collecting and returning *renpin* between individuals and the community, in order to maintain a "balanced karma." During the interactions about *renpin*, *weibi* can act as the token of the social debt (Parry & Bloch, 1989). In addition, due to its casual nature, it is easy to collect from and give to the crowd, thus it can enact the idea of karma circulating through the social system. *Weibi*, *by fitting in as a resource in socially valued ways,* can thereby add to a sense of generalized social obligation in a very Chinese culturally-specific manner.

Visibility of Use

The uses of *weibi* are visible, reflexively reinforcing users' desires to create and maintain their *guanxi* networks through *weibi*. When seeing many people giving away *weibi* to celebrate, users come to understand that they should also do so. As well, they may be told do so, if they have something that could be celebrated or when they need blessings.

For example on *Soccer* during the World Cup, a user (*Tevez*) was asked to give away *baozi*, when the same-named soccer player scored a goal in an important match:

    tevez(should) give us baozi lah! [shanren2]

In another case, *mirror* was asked to give away *baozi* for his good luck in shopping. One replier posted mirror's account info showing he is "rich" with 3800 *weibi*:

    Attention everybody, stand in line, mirror has 380 baozi [avatar10]

Another reply stated simply:

    [one] has to give baozi for this [kind of good luck]. [Inception]

Note that this reflects a social norm in Chinese culture that higher status people should contribute more to the community or to society, and in a *guanxi* network, weaker parties should be often favored in the relationship (Alston, 1989).

Explicitness

Bluntness in social interaction is often required online. Similarly, growing and utilizing one's *guanxi* networks online by requesting favors and calling for action results in less subtlety than one might use offline. On the other hand, *weibi* is not money, and people are thus free to be less explicit in their requests.

The degree to which people explicitly use *weibi* to exchange favor varies across people and boards. For example, one interviewee felt uncomfortable when seeing people ask for *baozi* when offering information or answering others' questions in *PennySaver*, "I never saw this on other boards. It is a very happy thing to give away *baozi* in many other boards, to get blessing or celebrate,… I don't go to *PennySaver* often so that might be why I am not used to it… I come from other boards where people just answer questions to help others. I help people too, without asking for *baozi*." [I10]

7. Conclusion

In this chapter, we discussed a thriving and devoted online community, Mitbbs. It is an online community where the users display behaviors typical of Chinese – the social uses of personal networks of reciprocal obligation called *guanxi*.

Although Mitbbs is a single site, and any generalizations must necessarily be limited, social interactions on Mitbbs appear to be fostered by a small design feature. This design feature, however, is one that turned out to be critical for Mitbbs' Chinese users -- a virtual currency. This virtual currency, or *weibi*, has little real value. Because its use is not structured, but is flexible, visible, ambiguously valued, and fits Chinese social practices, virtual currency is a valuable resource for Mitbbs users for a wide variety of their own purposes -- all in a very Chinese manner.

As well, the virtual points are intensively used to practice and enhance a new social norm for the netizen generation: "karma conservation," which evolves from mixing Buddhism and the *guanxi* networking philosophy. This norm acts not only as an additional basis for social reciprocity, but it is also a significant motivation for contribution in this online community. *Weibi* allows *renpin* to serve as an important mechanism for peer contribution, again showing *weibi*'s Chinese culturally-specific design value.

Chapter 6

Conclusion and Future Work

The previous four chapters used different methods to provide the first comprehensive and in-depth understanding of the dynamics of user behavior and system evolution, and how incentive and culture shape their complex dynamics. The four studies examined three types of information-sharing systems that vary in purpose, scale, and mechanism. But they present a variety of common characteristics that are important for understanding and designing other systems.

The studies revealed the multifaceted characteristics of users' behaviors in these information-sharing systems. Users' behaviors are adaptive, diverse, and complex. First, incentive design—in addition to affecting individual transactions—can change users' long-term behavior as users adapt to the system. Overall, users learned to get the most out of the system. For example, Taskcn participants discovered less-competitive and more-expensive tasks, and answerers in Baidu Knows learned how to improve answer quality and answering performance over time. On Mitbbs.com, users learned to use points by imitating others and inventing new ways to use points according to their specific need and context. These behavior patterns suggest that it is important to understand how users respond to incentive designs, what motivates them, and what the potential effects of users' adaptive behaviors.

Individuals have adaptive trajectories over time, but so do the sites. Sites' information and social dynamics evolve over time as well. For example, I observed that Q&A sites evolved over time in terms of the distribution of contributions and their ability to sustain users. In general, earlier adopters tend to be more committed to the site, which might indicate that they are more intrinsically motivated to use the site and that they have an advantage in accumulating reputation and experience. Often, the change in the composition of contributors responds to the changes in the overall survival rate of the users on a site. In Mitbbs.com, the variation and scope of the virtual currency usage expanded over time, and how the users perceive the value and meaning of the virtual currency has been evolving and diversifying too. All these findings suggest there are different stages of a site or a design (e.g., virtual currency in Mitbbs.com). Thus I might need to evaluate the status of a site or design by measuring multiple dimensions—participation, composition of participation, and survival rate, for example. In future work, I would like to investigate the evolution of a site's dynamics for a longer time, and develop comprehensive metrics to identify the different life stages of a site.

Users are diverse, and they vary on all kinds of dimensions: what motivates them, what strategies they use, and their expertise. On Taskcn, in particular, a large number of casual users contributed to its high traffic, while a small core of users provided the winning solutions, continually improving their performance. Casual users were less strategic, and thus tended to lose more, which made them less motivated than the core winning users. This suggests that in order to sustain the website, it is important to incentivize a core of users, as well as attract potential users to this core group from the high volume of attempters. In addition, it is also crucial to continue to drive large number of prospective members towards a site, to enhance a site's publicity, sustain challenging competitions (for the task requesters), and bring new blood to the small, but highly active core of contributors.

The second study demonstrated how Baidu Knows is composed of different user groups who differ in their motivation, contribution, participation, and strategies. For example, the users can be grouped by their levels of contribution and expertise, but they also switched roles between asker and answerer. Between different roles they play, users were differently motivated and had different levels of commitment to the service. In Mitbbs, the usage and perception of the virtual currency *weibi* varied dramatically across topic-boards too. The analytical method developed in these studies will be applied to examining more complex and diverse user behaviors in future work. In addition, I will develop systematic metrics to differentiate and evaluate each type of users and recommend schemes to motivate them respectively.

Users' behaviors can be very complex. The third study used Survival Analysis to quantify users' participation lifespan, which reflects their commitment to a site and the site's ability to sustain users over time. Consistently across sites, users who prefer answering tended to stay longer and were less sensitive to their initial experience. In addition, users' first-month experience accounted for a considerable amount of variance in predicting lifespan. In particular, users' self-selection effect (e.g., whether a user is active or what type of role one likes to play) and performance in the community accounted for most of the variance in their behavior. This suggests that intrinsic motivation (e.g., whether a user enjoys answering questions) is the key factor behind bringing in and sustaining users; however, sites can still motivate users by improving their asking and answering experiences. For example, sites can direct users to the questions that might be interesting and proper for them to answer. Similarly in Mitbbs.com, when the virtual currency *weibi* evolved to be a blended medium of social support and exchange, as well as good luck and karma, it blurred the boundary between intrinsic and extrinsic motivation, encouraging more intensive social interactions.

Culture comes into the picture and makes users' behavior even more complex. For instance, Taskcn's success at getting so many might partially result from China's large surplus of human labor. The third study provided the first large-scale comparison study among the three countries' popular Q&A. Although the three sites have similar incentive schemes, system designs, and scales, they present significant differences: users of

Yahoo!Answers tended to stay longer on the site, and the answerers tended to be more active in providing answers than the users of the two other sites. By analyzing the contents of sample questions and answers, we found that the question-answering dynamics on Yahoo!Answers tended to be more conversational than Baidu Knows. This difference might be explained by a complex interaction between the small discrepancy of the incentive designs and culture of the participants. This interaction changes both what people ask and how people answer. Therefore, in order to evaluate how efficiently the sites exchange information and social supports, it will be necessary to conduct field experiments across sites in future work.

The fourth study is a deeper investigation into how incentive design and culture can interact and co-evolve in a very complex way in a Chinese information-sharing system. *Weibi*, the virtual currency used on Mitbbs, was designed to motivate contribution not unlike ones employed by many US-based websites. However, the ambiguity of its value ended up perfectly supporting the crucial social dynamics of Chinese communities — *guanxi*, which requires fluid networking, calculated reciprocity, and contextual *renqing* interactions. In addition, users helped *weibi* evolve to incorporate the norms of a new generation of Internet users: *renpin*, or karma. This created many new purposes and uses of *weibi*, boosting the social interactions and contributions in the community. Furthermore, people attached new meanings to *weibi*, such as kindness, blessing, and good luck. Thus *weibi* can motivate users both extrinsically and intrinsically.

Mitbbs.com provided a specific instance of how incentive design can interact with a particular community structure and culture in complex ways, and how the interaction can lead to a co-evolutionary process between the design and the way users perceive and use the incentive design. This suggests that sites will face stiff challenges when they cross cultures. There are two areas I would like to pursue in future work: 1) I would like to explore what cultural considerations have been taken when designing information-sharing systems in different cultures. For example, how and why did Yahoo!Answers and Baidu Knows design different incentive schemes? How and why did Twitter and Sina

Weibo (a Chinese micro-blogging service) develop different designs? (2) I would like to identify which dimensions are particularly influenced by culture and how to design information-sharing systems to best fit different cultures.

References

Ackerman, S. M., & Palen, L. (1996). *The Zephyr Help Instance: promoting ongoing activity in a CSCW system*. Proceedings of the SIGCHI conference on Human factors in computing systems: common ground, Vancouver, Canada: 268-275.

Adamic, L. A., Zhang, J., Bakshy, E., & Ackerman, M. S. (2008). *Everyone knows something: Examining knowledge sharing on Yahoo Answers*. Proceeding of the 17th international conference on World Wide Web, Beijing: 665-674.

Adler, B. T., & Alfaro, L. d. (2007). *A Content-Driven Reputation System for the Wikipedia*. Proceeding of the 16th international conference on World Wide Web Beijing, China: 261-270.

Alston, J. P. (1989). Wa, Guanxi, and inhwa: managerial principles in Japan, China, and Korea. *Business Horizons, 32*(2), 26-31.

Arguello, J., Butler, B., Elisabeth, J., Kraut, R., Ling, K. S., Rose, C., et al. (2006). *Talk to Me: Foundations for Successful Individual-Group Interactions in Online Communities*. Proceeding of Human-Computer Interactions, Montréal, Québec, Canada.: 959-968.

Bourdieu, P. (1986). The Forms of Capital. In J. Richardson (Ed.), *Handbook of Theory and Research for the Sociology of Education*. (pp. 241-258). New York: Greenwood.

Butler, B. (2001). Membership Size, Communication Activity, and Sustainability: A Resource-Based Model of Online Social Structures. *Information Systems Research, 12*(4), 346-362.

Chen, Y., Harper, M., Konstan, J., & Li, S. X. (2010). Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens. *American Economic Review, 100*(4), 1358–1398.

Chen, Y., Ho, T.-H., & Kim, Y.-M. (2010). Knowledge Market Design: A Field Experiment on Google Answers. *Journal of Public Economics Theory, 12*(4), 641-664.

Chiao, C. (1982). Guanxi: A Preliminary Conceptualization *The Sinicization of Social and Behavioral Science Research in China*. (pp. 345-360). Taipei: Academia Sinica.

Cummings, J., Butler, B., & Kraut, R. (2002). The quality of online social relationships. *Communications of the ACM, 45*(7), 103-108.

Dasgupta, P., & Stiglitz, J. (1980). Industrial structure and the nature of innovative activity. *The Economic Journal 90*(358), 266-293.

Dollinger, M. J. (1988). Confucian ethics and Japanese management practices. *Journal of Business Ethics, 7*, 575–584.

Farh, J.-L., Tsui, A. S., Xin, K., & Cheng, B.-S. (1998). The Influence of Relational Demography and Guanxi: The Chinese Case. *Organization Science, 9*(4), 471-488

Gilbert, E., & Karahalios, K. (2009). *Predicting tie strength with social media.* Proceeding of the 27th international conference on Human factors in computing Boston, MA, USA: 211-220.

Giles, J. (2005). Special Report Internet encyclopaedias go head to head. *Nature, 438*, 900-901.

Gneezy, U., & Rustichini, A. (2000). Pay Enough Or Don'T Pay At All. *The Quarterly Journal of Economics, 115*(3), 791-810.

Golder, S. A., & Huberman, B. A. (2006). Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science, 32*(2), 198-208.

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, Florence, Italy: 865-874.

Harper, F. M., Moy, D., & Konstan, J. A. (2009). *Facts or Friends? Distinguishing Informational and Conversational Questions in Social Q&A Sites.* Proceeding of the 27th international conference on Human factors in computing systems, Boston, MA, USA. : 759-768.

Heine, S. J., Lehman, D. R., Markus, H. R., & Kitayama, S. (1999). Is There a Universal Need for Positive Self- Regard? *Psychological Review, 106*(October), 766-794.

Hempel, J. (2006). Crowdsourcing: Milk the masses for inspiration. *Business Week*.

Hertel, G., Niedner, S., & Herrmann, S. (2003). Motivation of software developers in open source projects: An Internet-based survey of contributors to the Linux kernel. *Research Policy, 32*(7), 1159-1177.

Hiltz, S. R., & Turoff, M. (1993). *The Network Nation*: The MIT Press.

Hjorth, L. (2008). Home and Away: A Case Study of the Use of Cyworld Mini-hompy by Korean Students Studying in Australia. *Asian Studies Review, 31*(4), 397-407.

Hofstede, G. (1983). Dimensions of national cultures in fifty countries and three regions. In J. Deregowski, S. Dzuirawiec & R. Annis (Eds.), *Expiscations in Cross-Cultural Psychology*. (pp. 335-355). Lisse NL: Swets and Zeitlinger.

Hofstede, G. H. (1980). *Culture's consequences: international differences in work-related values*: SAGE.

Holohan, A., & Garg, A. (2005). Collaboration online: The example of distributed computing. *Journal of Computer-Mediated Communication, 10*(4).

Howe, J. (2006). The Rise of Crowdsourcing. *Wired, 14*.

Howe, J. (2008). *Crowdsourcing: Why the Power of the Crowd Is Driving the Future of Business*. New York: Crown Business.

Hwang, K.-k. (1987). Face and Favor: The Chinese Power Game. *American Journal of Sociology, 92*, 944-974.

Jacobs, J. B. (1980). The Concept of Guanxi and Local Politics in a Rural Chinese Cultural Setting. In S. L. Greenblatt, R. W. Wilson & A. A. Wilson (Eds.), *Social Interaction in Chinese Society*. (pp. 209-236). New York: Praeger.

Joyce, E., & Kraut, R. E. (2006). Loyalty: Predicting continued participation in newsgroups. *Journal of Computer Mediated Communication, 11*(3), 723-747.

Kayan, S., Fussell, S. R., & Setlock, L. D. (2006). *Cultural differences in the use of instant messaging in Asia and North America*. Proceeding of the 20th conference on computer supported cooperative work, Banff, Alberta, Canada: 525-528.

King, A. Y.-c. (1991). Kuan-hsi and Network Building: A Sociological Interpretation. *Daedalus, 120*(2), 63-84.

Kitayama, S., Park, H., Sevincer, A. T., Karasawa, M., & Uskul, A. K. (2009). A cultural task analysis of implicit independence: Comparing North America, Western Europe, and East Asia. *Journal of Personality and Social Psychology, 97*(2), 236-255.

Kittur, A., Chi, E. H., Pendleton, B. A., Suh, B., & Mytkowicz, T. (2007). *Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie*. Proceeding of Alt.CHI at CHI San Jose, CA.

Kleeman, F., Voss, G. G., & Rieder, K. (2008). Un(der)paid Innovators: The Commercial Utilization of Consumer Work through Crowdsourcing. *Science, Technology & Innovation Studies, 4*(1), 5-26.

Lakhani, K. R., & Hippel, E. V. (2003). How open source software works: "Free" user to user assistance. *Research Policy, 32*(6), 923-943.

Lampe, C., & Johnston, E. (2005). *Follow the (slash) dot: Effects of feedback on new members in an online community*. Proceeding of International ACM SIGGroup conference on supporting group work, New York, NY: 11-20.

Lindtner, S., Mainwaring, S., Dourish, P., & Wang, Y. (2009). *Situating Productive Play: Online Gaming Practices and Guanxi in China*. Proceeding of the 12th IFIP TC 13 International Conference on Human-Computer Interaction, Uppsala, Sweden.: 328-341.

Lloyd, G. E. R. (1996). Science in antiquity: The Greek and Chinese cases and their relevance to the problems of culture and cognition. In D. R. Olson & N. Torrance (Eds.), *Modes of thought: Explorations in culture and cognition*. Cambridge, UK: Cambridge University Press.

Malone, T. W. (2004). *The Future of Work: How the New Order of Business Will Shape Your Organization, Your Management Style, and Your Life*. Boston, MA: Harvard Business School Press.

Markus, H. R., & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review, 98*, 224–253.

Marlow, C., Naaman, M., boyd, d., & Davis, M. (2006). *HT06, Tagging Paper, Taxonomy, Flickr, Academic Article, ToRead*. Proceeding of the Seventeenth Conference on Hypertext and Hypermedia: 31-40.

Mason, W., & Watts, D. J. (2009). *Financial incentives and the "performance of crowds"*. Proceeding of the ACM SIGKDD Workshop on Human Computation, Paris, France: 77-85.

Massey, A. P., Hung, Y.-T. C., Montoya-Weiss, M., & Ramesh, V. (2001). *When culture and style aren't about clothes: perceptions of task-technology "fit" in global virtual teams*. Proceeding of the International ACM Conference on Supporting Group Work, Boulder, Colorado, USA: 207-213.

Nam, K., Adamic, L. A., & Ackerman, M. S. (2009). *Questions in, Knowledge iN? A Study of Naver's Question Answering Community*. Proceeding of the twenty-seventh annual SIGCHI conference on Human factors in computing systems, Boston, MA: 779-788.

Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic vs. analytic cognition. *Psychological Review, 108*, 291–310.

Nov, O., Naaman, M., & Ye, C. (2008). *What drives content tagging: the case of photos on Flickr*. Proceeding of the twenty-sixth SIGCHI conference on Human factors in computing systems, Italy: 1097-1100.

Parry, J. P., & Bloch, M. (1989). *Money and the morality of exchange*: Cambridge University Press.

Portes, A. (1998). Social Capital: Its Origins and Applications in Modern Sociology. *Annual Review of Sociology, 24*, 1-24.

Pye, L. W. (1982). *Chinese Commercial Negotiating Style*: Oelgeschlager, Gunn & Hain.

Pye, L. W. (1985). *Asian power and politics: The cultural dimensions of authority*. Cambridge, MA: The Belknap Press of Harvard University.

Ralston, D. A., Holt, D. H., Terpstra, R. H., & Kai-Cheng, Y. (2007). The impact of national culture and economic ideology on managerial work values: A study of the United States, Russia, Japan, and China. *Journal of International Business Studies*, 1-19.

Rossi, M. A. (2004). Decoding the "Free/Open Source(F/OSS) Software Puzzle" a survey of theoretical and empirical contributions. *Department of Economics University of Siena 424, Department of Economics, University of Siena.*

Setlock, L. D., & Fussell, S. R. (2010). *What's it worth to you?: the costs and affordances of CMC tools to asian and american users*. Proceeding of the conference on Computer supported cooperative work, Savannah, Georgia, USA: 341-350.

Singelis, T. M. (1994). The Measurement of Independent and Interdependent Self-Construals. *Personality and Social Psychology Bulletin, 20*(October), 580-591.

Triandis, H. C. (1995). *Individualism and Collectivism*. Boulder, CO: Westview.

Tsui, A. S., & Farh, J. L. (1997). Where Guanxi Matters: Relational Demography and Guanxi in the Chinese Context. *Work and Occupations, 24*, 56-79.

Varnum, M. E. W., Grossmann, I., Kitayama, S., & Nisbett, R. E. (2010). The Origin of Cultural Differences in Cognition: The Social Orientation Hypothesis. *Current Directions in Psychological Science, 10*(1), 9-13.

Veinott, E. S., Olson, J., Olson, G. M., & Fu, X. (1999). *Video helps remote work: speakers who need to negotiate common ground benefit from seeing each other*. Proceeding of the SIGCHI conference on Human factors in computing systems, Pittsburgh, Pennsylvania, United States: 302-309.

Walder, A. G. (1986). *Communist neo-traditionalism: work and authority in Chinese industry*. Berkeley: University of California Press.

Wang, Y., & Mainwaring, S. D. (2008). *Human-Currency Interaction: learning from virtual currency use in China*. Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems, Florence, Italy: 25-28.

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

Wilcox, R. T. (2000). Experts and Amateurs: The Role of Experience in Internet Auctions. *Marketing Letters 11*(4), 363-374.

Xin, K. R., & Pearce, J. L. (1996). Guanxi: Connections as Substitutes for Formal Institutional Support. *The Academy of Management Journal, 39*(6), 1641-1658

Yang, J., Ackerman, M. S., & Adamic, L. A. (2011). *Virtual gifts and guanxi: supporting social exchange in a chinese online community*. Proceeding of the ACM conference on Computer supported cooperative work, Hangzhou, China: 45-54.

Yang, J., Adamic, L. A., & Ackerman, M. S. (2008 a). *Competing to Share Expertise: the Taskcn Knowledge Sharing Community*. Proceeding of the International AAAI Conference on Weblogs and Social Media, Seattle.

Yang, J., Adamic, L. A., & Ackerman, M. S. (2008 b). *Crowdsourcing and Knowledge Sharing: Strategic User Behavior on Taskcn*. Proceeding of the ACM Conference on Electronic Commerce, Chicago.

Yang, J., Morris, M., Teevan, J., Adami, L. A., & Ackerman, M. S. (2011). *Culture matters: A survey study of social Q&A behavior*. Proceeding of the International AAAI Conference on Weblogs and Social Media, Spain.

Yang, J., & Wei, X. (2009). *Seeking and Offering Expertise across Categories: A Sustainable Mechanism Works for Baidu Knows*. Proceeding of the International AAAI Conference on Weblogs and Social Media, San Jose.

Yang, J., Wei, X., Adamic, L., & Ackerman, M. (2010). *Activity Lifespan: An Analysis of User Survival Patterns in Online Knowledge Sharing Communities*. Proceeding of the International AAAI Conference on Weblogs and Social Media, Washington DC.

Yang, M. M.-H. (1994). *Gifts, favors, and banquets: the art of social relationships in China*: Cornell University Press.

Yang, Y. (1996). *The Flow of Gifts: Reciprocity and Social Networks in a Chinese Village*: Stanford University Press.

Zhang, J., Ackerman, M. S., & Adamic, L. A. (2007). *Expertise networks in online communities: structure and algorithms*. Proceeding of the 16th international conference on World Wide Web, Banff, Alberta, Canada: 221-230.