

Competing to Share Expertise: the Taskcn Knowledge Sharing Community

Jiang Yang¹, Lada A. Adamic¹, Mark S. Ackerman^{1,2}

¹School of Information, ²Department of EECS
University of Michigan, 1085 South University Ave. 304 West Hall, Ann Arbor, MI 48109-1107, USA
(yangjian, ladamic, ackerm)@umich.edu

Abstract

"Witkeys" are websites in China that form a rapidly growing web-based knowledge market. A user who posts a task also offers a small fee, and many other users submit their answers to compete. The Witkey sites fall in-between aspects of the now-defunct Google Answers (vetted experts answer questions for a fee) and Yahoo Answers (anyone can answer or ask a question). As such, these sites promise new possibilities for knowledge-sharing online communities, perhaps fostering the freelance marketplace of the future. In this paper, we investigate one of the biggest Witkey websites in China, Taskcn.com. In particular, we apply social network prestige measures to a novel construction of user and task networks based on competitive outcomes to discover the underlying properties of both users and tasks. Our results demonstrate the power of this approach: Our analysis allows us to infer relative expertise of the users and provides an understanding of the participation structure in Taskcn. The results suggest challenges and opportunities for this kind of knowledge sharing medium.

Introduction

The online community knowledge market has been rapidly gathering popularity in recent years. People can easily access and share information and expertise through the Web without time and geographic constraints. For example, Yahoo! Answers, launched in December 2005, has obtained 80 million unique users worldwide and approximately 23 million resolved questions.¹ Similarly, in China, the biggest Internet portal, Baidu.com (the most used site in China according to Alexa.com), also provides an online community based question-answer (QA) platform. Since it began in June 2006, over 25 million questions have been answered on Baidu.² There are many points in the design space of these knowledge sharing communities (Wenger 1998; Ackerman et al. 2003). Yahoo! Answers allows users to ask questions about a variety of topics, and anyone can answer (Adamic et al. 2008). Google Answers, on the other hand, had vetted experts who would answer a user's question for a fee

set by the user. (Although Google Answers is now defunct, similar sites still exist, for example, justanswer.com.)

Witkey websites are a new type of knowledge sharing community in China occupying a new point in the knowledge sharing design space. The Witkey websites have gathered thousands of tasks, hundreds of thousands of participants, and millions of users. Unlike Yahoo! Answers, but similar to Google Answers, users offer a monetary award in return for help and expertise. However, unlike Google Answers, other users, not vetted experts, answer the question and provide potential solutions to a task.

Another feature differentiating Witkey sites from standard question-answer forums is that the tasks (the problems or requests presented by users) appear to be often more complex than those posted to other popular knowledge sharing communities. For example, one might ask for the mockup of a website front page or a graphic design. Because of the complexity of the questions/tasks, contributors may need a non-trivial incentive (e.g., in the form of reputation that may be transferred elsewhere or simple monetary reward). Furthermore, since the tasks can be rather specific, and the answer potentially only of interest to the asker, a monetary reward would seem fair in exchange for effort consumed. For example, while answering a question about a software library or a particular car problem can benefit many individuals, in addition to the asker, designing a logo is normally only interesting to the requester himself. The Witkey websites, therefore, may be harbingers of the freelance markets that have been forecast (Malone 2004).

Crucial to the success of this new medium is whether users have enough incentive to participate and consequently whether tasks will attract solutions from users of a sufficient level of expertise. In order to answer these questions, we studied one of the biggest Witkey websites in China -- Taskcn.com. It has garnered more than 4,000 posted tasks and 1.2 million registered users since its launch date in June 2006.

To our knowledge, these sites have not been studied previously, and given that Witkey sites are a new kind of knowledge sharing community, this important design point should be explored. We would like to understand how these sites work. The first, and most obvious question is whether the monetary award can attract good answerers and

¹ <http://answers.yahoo.com/>

² <http://zhidao.baidu.com/>

buy good solutions, and whether the contributors can obtain appropriate rewards by participating. We find that monetary reward is *not* a significant incentive for user participation in a task, as many users compete in tasks with low monetary reward. This may indicate a healthy, participatory dynamic in the site. Regardless, we also need to understand the distribution of expertise that allows problems to be solved on Taskcn. Therefore, we would like to know whether users can be differentiated according to expertise level.

In order to do so, we apply social network analysis to network representations of the participation structures, identifying both the prestige of the users and of the tasks.

We find that a task's prestige on Taskcn can slightly hinder people's participation, hinting that tasks where even expert users lose are attempted by fewer users on average. On the other hand, tasks where many users participate (making the task central in the network) tend to have lower average participant expertise levels. These two observations imply that "peripheral" tasks attract more prestigious participants, while popular tasks attract on average less discriminating, and less expert, participants. However this does not mean that popular tasks are doomed to attract only mediocre solutions. While the average participant may have a lower expertise level, having more solutions submitted and in particular, having solutions submitted by winners in other tasks, significantly improves the chance that the *winner* will be more expert in Taskcn. Finally, we find a user's chance of winning depends on the number of submitters for the task, her track record and her prestige level.

Taskcn.com

Witkey Websites

As mentioned, a Witkey website is a new type of knowledge market website, in which a user offers a monetary award for a question or task and other users compete for the award. When an asker posts her task requirement and offers money, this money has to be deposited with the website. Upon the deadline the requester can choose the winner(s); the majority of the money will be sent by the website to the winner(s), and the website takes a service fee.

The term Witkey was coined by the founder of the first website Witkey.com³, and became the name of a series of similar websites in China. In the last two years, more than 10 Witkey websites have been launched (e.g., Witkey.com, Taskcn.com, and K68.com), and they have gathered millions of users. On Taskcn.com, one of the biggest Witkey websites in China, and the one analyzed here, 1,258,109 users have registered since June 2006.⁴

Because of the monetary incentive and the relatively small number of task posts, each task can get a considerable

³ <http://www.witkey.com/lfarticle/articledt.asp?aid=20000>

⁴ As claimed by Taskcn.com in November 2007.

number of contributors. For example, in the completed 2,437 tasks in Taskcn.com, more than 150,000 users competed. Since usually one task has only one winner, most contributors do not get any financial reward even though they must invest effort. Therefore, the mechanism potentially benefits requesters more than contributors.

Taskcn.com

Taskcn.com, the Witkey community we selected for our study, has more than 4,000 posted tasks categorized into 7 different types. Our data include the 2,437 tasks that were completed before November 3, 2007. These 2,437 tasks had 158,290 participating users who submitted at least one submission to a task. Although there is an order of magnitude more registered users, most of the registered users did not contribute to any task and were omitted from our study.

Basic Participation Characteristics

Users' Participation. As in the majority of online communities, users' participation is unevenly distributed.




Figure 1: Users' submissions and won tasks.

Figure 1 shows the frequency distribution of the number of users' task submissions and wins. We can see that the majority of users submitted to or won a very small number of tasks, while a handful of users submitted to or won multiple times. So although one may have expected monetary incentives to lead to broader participation, the distributions match the standard skewed distributions of other activity in online spaces such as news groups (Fisher et al. 2006), wikis (Holloway, Bozicevic, and Börner 2007), online dating communities (Holme, Edling, and Liljeros 2004) collaborative tagging systems (Golder and Huberman 2006) and question answer forums (Zhang, Ackerman, and Adamic 2007, Adamic et al. 2008). In the end, people participate in few tasks, and even fewer individuals actually receive monetary rewards for participating.

Table 1. Categories of tasks.

Category	Number of Tasks	Number of Participants
Design	1130	41838
Strategic planning	297	123166
Programming	75	554
Others	185	36804
Personal service	222	9028
Website	436	9049
Writing	92	14480

Category Difference. All tasks are categorized on the site into 7 groups as shown in Table 1. The tasks are unevenly distributed in the categories, with the Design and Website categories being by far the most popular. In addition, the average in the amount of award, number of submitters, view times, and number of votes of the tasks differs among categories. For example:

- The Design category has half of the total number of tasks, and it also has a higher mean award ($m = 408.13$ yuan) relative to the other 6 categories.
- The Strategic Planning and Writing categories have the highest numbers of views (i.e., people more often view the tasks).
- The Strategic Planning category received many more submissions per task (mean=979.51), followed by the "others" category (mean = 239.61).




Figure 2. Task numbers in the 7 categories.

Figure 2 presents the growth of tasks in the categories from May 2006 through September 2007. The Design category had a surge in September, 2006 and stayed as the largest category even while the other categories have also grown in size. This pattern implies that task requesters have found design problems to be a better match with the Witkey system mechanisms than tasks in other categories. We suspect that there are some determinant properties that resulted in this match. Design tasks can often be completely handed off to independent workers, whereas planning a marketing strategy requires much more contextual information and feedback. In addition, the expertise on Design (or Programming) is distributed unevenly (only a few people have considerable expertise) and most requesters don't need a close relationship with the designer (one may need a designer only once for the company logo); thus, an online platform provides a perfect place to establish these rarely occurring connections. This can also explain why other categories like Strategic Planning often have more people submitting and viewing: They do not necessarily have easy tasks, but they do have a lower threshold of participation. For example, most Taskcn users would have the ability to offer suggestions for naming a newborn baby or giving advice for planning or running a new online business.

User Network and Task Network

Users' Structural Prestige

Researchers have widely employed social network analysis to detect interaction patterns. The best known is PageRank for ranking web pages (Page et al. 1998), where an inbound link can be regarded as a positive reference from the source page and the importance of a page recursively depends on that of its referrers. Social network analysis can also provide an informative visualization facility to help people to understand the complex dynamics in online communities. For example, Turner et al. (2005) have mapped people's interactions in Usenet Newsgroups and Fisher, Smith, and Welser (2006) were able to categorize users into different roles through examining the interaction structure in the website. In addition, Kou and Zhang (2003)'s analysis of the replying network of a bulletin board system discovered that the distribution of people's interest spaces is embedded in their replying interactions. Graph-based ranking algorithms can be used to quantitatively discover people's expertise distribution (or the distribution of other properties), for examples, see Kautz, Selman, and Shah (1997), Campbell, et al. (2003), and Dom et al. (2003). Most recently, Zhang, Ackerman and Adamic (2007) explored various graph-based algorithms to find users' expertise in QA online communities.

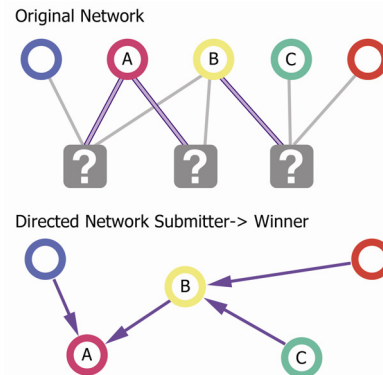


Figure 3: Users are denoted by colored circles and tasks are denoted as question marks. Purple double lines are "wins"; grey lines are "non-winning" submissions. The user prestige network is derived as follows: if users A and B participate in the same task and A wins, then we add an edge from B to A.

Figure 3 shows how we construct the users' prestige network. The directed edges not only imply common interest and expertise level (by participating in the same tasks), but also indicate a competitive relationship. Participating in a task is also participating in a competition, in which the winner beats the other participants. We find that one user winning over another is consistent across tasks. For example, of the approximately 2,000 pairs of users with two wins between them on two tasks in which they both participated, 77% of the time it was the same user who won. Since winning is not random (by chance, the same user would win twice only 50% of the time), this

gives justification to treating the edges as directed. Similarly, of the 384 instances where two users competed 3 times, one would expect the same user to win all three times in only $\frac{1}{4}$ of the cases. Instead we observe the same user winning all 3 tasks 56% of the time. This method of constructing a network where a link from A to B implies "B is more expert than A," is essentially a variation of the *community expertise network* (CEN) (Zhang, Ackerman, and Adamic 2007) in which people's expertise can be measured by structural prestige (Wasserman and Faust 1994).

We use several prestige measures; they have slightly different implications and limitations:

WinRate is a statistic which is computed by the number of tasks a user won over the number of that user's submissions. In general, a user should have higher WinRate when she has higher expertise; however, it could also be largely influenced by the strategy the user takes in choosing tasks.

Indegree of a node in the network represents how many references the node gets from others. It basically shows how many other users one has beaten in tasks. This measure can be used for evaluating a user's expertise level; however, it potentially suffers from the sensitivity to the number of times one has submitted work.

Closeness is a measure of average proximity of a node to all others, treating edges as undirected. A node with high closeness will be located in the "middle" of the network, while one with low closeness will be situated on the periphery.

Betweenness reflects how many shortest paths a node lies on, that is, between how many pairs of other nodes the node is situated. In addition to being correlated with degree, a high betweenness user would participate in a diverse set of tasks, and a high betweenness task would be attempted by a diverse set of users who compete in other tasks.

PageRank is a widely used ranking measure for network structure, which takes into account the recommender's prestige. PageRank corresponds to the principal eigenvector of the adjacency matrix of the directed graph (Berkhin 2005)⁵, which essentially indicates how much a node is recommended by all others in the network, directly or indirectly. By this measure, one should get a higher PageRank by winning over a higher-PageRank user than by winning over a lower-PageRank user.

Tasks' Structural Prestige

We are also interested in finding the underlying properties of tasks. There are many possible properties that a task can have, such as the level of difficulty, which will turn people away; or level of interestingness, which would invite people to join, even if not a lot of money is at stake. All

these will influence how many people finally participate in the task and who actually wins the task.

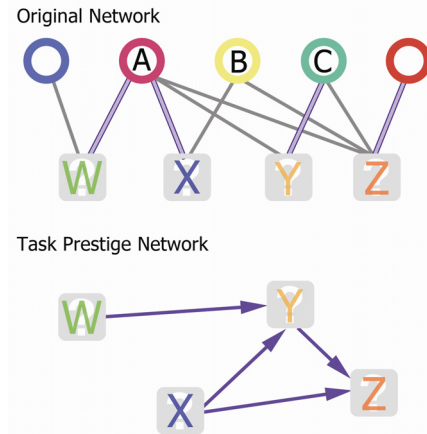


Figure 4: Task prestige network. Users participate in tasks W, X, Y, and Z. If user A wins in task X but fails in task Y, then a directed link is built from X to Y, meaning task Y is more prestigious than task X.

We construct the *task prestige network* according to the method detailed in Figure 4. The idea is similar to the tennis open series: some championships are harder to win due to the participation of many highly ranked players and winning them imparts more credit to the winners. We can describe this kind of property of the championships as their *prestige*.

We also calculate prestige rank (indegree, PageRank) and closeness with regard to the structural interactions for the task network, as we did for the user expertise network.

Characteristics of the Networks

The Bow Tie Structure of the User Network. The bow tie structure analysis was first developed by Broder et al. (2000) to capture the macroscopic picture of the Web. Table 2 shows just how different the bowtie structure of the Taskcn community is from hyperlinks on the Web. The different parts of the bow tie indicate whether nodes can be reached from one another by following directed paths (e.g. A won over B who won over C).

In our community expertise network, there is a tiny core containing users where any user can be reached from any other by following such directed paths. The large IN set, pointing to the core, is mostly comprised of users who submit work but never win. The very small OUT set has users who have won over the users in the Core set. Tendrils contain users who have lost to the OUT set. This means that there are very few reciprocal loops among users and the majority of the users participate without winning.

Table 2 Comparison of bow tie structures.

	LSCC	In	Out	Tendrils	Others
Web	27.7%	21.2%	21.2%	21.5%	8.4%
Design	1.57%	76.17%	0.56%	20.48%	1.22%
Program	0.18%	2.89%	0	73.10%	23.83%
Website	3.27%	28.49%	9.08%	55.35%	3.80%
Service	1.20%	40.05%	3.25%	50.34%	5.13%
Writing	0.41%	59.11%	0.74%	39.55%	0.18%

⁵ For computing PageRank, we used PageRank bias $\alpha = 0.15$ by convention in this paper, which means the teleportation probability of a node is set to be 0.15.

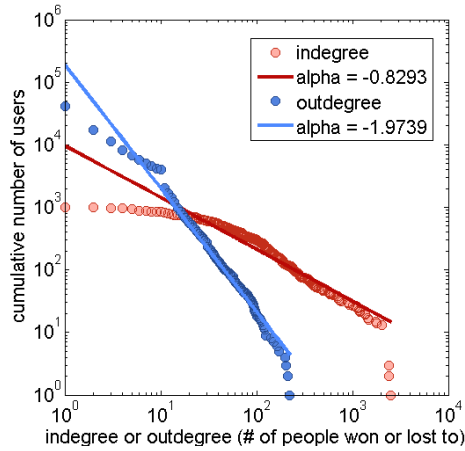


Figure 5: *Indegree and outdegree distribution of user nodes in the prestige network in design category.*

Degree Distribution. The degree distribution further describes the uneven participation and outcomes for the users. Figure 5 displays the distribution of users' indegree (number of people one has won over) and outdegree (number of people one has lost to) in the network. Similar to many other networks, the users' prestige network also demonstrates an evident scale-free nature (Barabasi and Albert 1999). This fact indicates the uneven interactions among the users: The majority of users participates in a few tasks and wins fewer, while there is an extremely small group of users who have had many submissions or have successfully won money.

Network Visualization. Figure 6 shows part of the users' prestige network for the Design category. This subnetwork contains the most recently active 800 users and their interactions through jointly attempted tasks.




Figure 6: *Part of the user prestige network in the Design category. The users who have won at least once are denoted by the blue nodes and the size of a node is proportional to its PageRank. Other users are denoted by smaller green nodes.*

In this subnetwork, a winner is usually surrounded by many submitters, and there are only 28 winners out of the total 800 users. In addition, the numbers of submitters are different among tasks. In the middle of the graph, the high interaction indicates the users actively participate in the same tasks, but far from the center, the users have participated only in a couple of tasks.

Motif Analysis. Motif analysis supplies a finer grained, local view into the networks of users and tasks. Table 3 presents all the frequencies of diadic and triadic motifs in the two networks, compared to expected frequencies in randomized versions of the networks (Milo et al. 2002; Milo et al. 2004; Wernicke and Rasche 2006). In the table, frequencies that are significantly different from a random network are shown in bolded numbers.

These frequencies can inform us about the social interactions and orderings, by reflecting reciprocity or hierarchical structures. Looking at pairwise relationships first, we find only a tiny portion (0.6%) of edges to be reciprocal. This is due to the large number of users who never win a task, and if they do, it is unlikely to be a win against someone who has bested them before. In the task network, there is a slightly higher proportion of bidirectional edges (4.1%) between two tasks, meaning that the winners in both tasks participated in the other task. Since winners are a smaller subset of the users, with at least a minimum level of achievement, it is more likely, although still fairly unusual, that their relationship would be reciprocal, which is reflected in the reciprocal edges among tasks.

The triad motifs represent interactions among sets of 3 nodes. Motif 36⁶ (in which two nodes refer to a third) accounts for 93.86% of the triad motifs, significantly more frequently than in the randomized networks. However, this is due to the way the user network is constructed; usually, a task has many submitters but only one winner. In the task network, the motif is not statistically significant.




Figure 7: *Motif profiles of users and tasks in the Design category. The Y-axis is the normalized z-score of the frequency deviation from the random network.*

Both networks show a statistically significant frequency for several other kinds of loops (e.g., triad 102, 140, 174, 238). The frequencies are quite tiny but they all significantly differ from randomized networks. This implies that although two nodes rarely have direct reciprocal references, they are linked through a layered structure, one in which nodes (users or tasks) of possibly similar prestige level

⁶ The Motif ID, shown in Figure 7, is assigned for each possible structure in Milo et al. (2004).

Table 2 Distribution of motif structures.

Motif Structure														
Motif ID			164	166	102	36	140	14	78	174	238	6	38	46
User network	99.4%	0.6%	1.23%	0.04%	0.01%	93.86%	0.004%	0.100%	0.02%	0.01%	0.002%	1.01%	0.15%	0.01%
Task network	95.9%	4.1%	4.68%	0.79%	0.43%	36.14%	0.326%	3.183%	0.299%	0.17%	0.010%	23.17%	7.27%	0.48%

indirectly refer to each other. Interestingly, both networks have significant symmetrical structures (e.g., triad 166, 238, 46). This also implies that some users or tasks may be on the same prestige level.

Understanding Participation Structure

Armed with the network representations of user participation in tasks, we are prepared to answer several questions concerning the dynamics of Taskcn. First, we correlate the network prestige and monetary reward of a task to the number and expertise level of the participants. We then characterize the expertise level of the winner, and predict the likelihood that any given user will win a task based on both user and task attributes. To answer these questions, we chose the Design category since its network is larger and denser than those of other categories.

Incentives to Participate

The first goal for a person posing a task is to attract high quality solutions from as many participants as possible. From the submitter's perspective, the decision to participate in a task may depend on both potential monetary reward and the task difficulty. Intuitively, tasks offering higher rewards should provide higher incentive for users to participate.

We find that higher rewards do attract more views ($R=0.383, p<0.001$). However, counter to intuition, but in agreement with a study of Google Answers (Chen, Ho, and Kim 2007), we find that task reward is uncorrelated with the number of submissions. This implies that users are attracted by a high-paying task, but perhaps seeing the greater amount of effort or skill required, are not more likely to attempt it. It also implies that many users are willing to give some time and expertise for small financial rewards. We do find a very small but significant negative correlation with task PageRank ($R=-0.080, \text{Sig.}=0.007$), indicating that more difficult tasks (those where users who have won previous tasks end up losing) attract slightly fewer attempts. There might be many other factors, such as enjoyment, wording politeness, potential achievability of the task, that determine the task's popularity.

Average Expertise of All Users of a Task

It is not just the number of submitters that counts, but also the expertise level of those users. A person posing a task would want people with sufficient expertise submitting potential solutions.

Table 3 Correlations between task prestige and users' average expertise level *

		Task Total Degree	Task Outdegree	Task Closeness	Task Betweenness
AvePagerank	Pearson Correlation	-.293(**)	-.491(**)	-.250(**)	-.250(**)
	Sig. (2-tailed)	.000	.000	.000	.000
AveIndegree	Pearson Correlation	-.112(**)	-.240(**)	.061(*)	-.115(**)
	Sig. (2-tailed)	.000	.000	.040	.000

* non-significant correlations omitted

We find that all measures of task centrality (total degree, outdegree, closeness and betweenness) are negatively correlated with the average indegree and PageRank of the users submitting to the task (except for task closeness and average indegree). These correlations, shown in Table 4, are explained in part by the way in which the user and task networks are constructed. A task with high outdegree has a winner who has lost in several other tasks. Since we saw earlier that pairwise outcomes between two players are likely to be repeated, the other participants are likely to have even lower prestige than the winner. This then explains the strong negative correlation between the average participant PageRank and task outdegree. Another interesting observation is that closeness and betweenness also have a negative, albeit smaller, correlation with the average PageRank of the participants. This hints that there is a dense core of tasks where many users, both expert and not, participate. Central tasks have more mutual participants while the tasks on the periphery of the network have their own participants who do not participate much in other tasks. The result shows that those tasks have on average more prestigious participants than the central tasks. This could be because the more central tasks may be more accessible (easier) for a wider range of participants, while the peripheral tasks require higher or more specific levels of expertise. We intend to examine this more closely in future work.

The Winner's Expertise Level

The average expertise measure of all submitters of the task makes the task competitive and increases the chance to obtain good answers. However, another measure of task success, one that matters to the requester selecting the best solution, is the prestige of the winning user. Consistent with the above results, we observe that not only did monetary reward not attract a greater number of submissions, but it also did not attract a winning submission from a significantly more prestigious user.

Table 4 Correlation between winner and task.

Winner	Submit	Award	In Degree	Out Degree	Total Degree	Page Rank	Closeness	Betweenness
PageRank ρ	.281**	.039	.677**	.288**	.662**	.364**	.275**	.384**
Sig.	.000	.196	.000	.000	.000	.000	.000	.000
Closeness ρ	.129**	.090**	.445**	.482**	.599**	.210**	.791**	.306**
Sig.	.000	.003	.000	.000	.000	.000	.000	.000

The remaining correlations are in part influenced by the way the user and task networks are constructed. There is a correlation between the winner's expertise measures and the number of submitters. More directly, the task indegree reflects the participation of the winners of other tasks, while these winners have not won in the current task. This event will also boost the winner's PageRank and other prestige measures, by sharing some of those users' prestige with the winner. Since the event of a user winning a task influences both the task and user centrality simultaneously, we next perform a regression that excludes the user's performance on the given task and predicts the likelihood of winning the task.

Winning Probability of a User in a Task

We hypothesize that the number of submitters of the task negatively affects the chance any particular user wins, while the user's prestige (here represented by indegree) and history of past wins strengthens the probability. The first model includes only the number of submitters to the task and users' prestige.

$$\text{Logit } P_{\text{win}=1, \text{lose}=0} = \beta_1 X_{\text{number of submitters}} + \beta_2 X_{\text{userIndegree}} + \epsilon$$

To do the prediction, we divided all 1130 tasks into two groups: 55% older tasks for generating predictor parameters and 45% newer tasks for regression. We then ran logistic regression on all task-to-user interactions in the later 45% tasks (when a user participated in a task and either won or lost).

Table 6 Model summary.

Step	-2 Log likelihood	Cox & Snell R Square	Nagelkerke R Square
1	2300.046(a)	.681	.908
2	2287.073(b)	.681	.909

Table 7 Variables in the equation.

		B	S.E.	Wald	df	Sig.	Exp(B)
Step 1(a)	logSubmit	-.818	.015	2870.021	1	.000	.441
Step 2(b)	logSubmit	-.833	.016	2653.267	1	.000	.435
	logUserIndegree	.191	.048	15.556	1	.000	1.211

The test shows that both predictor factors have a significant effect on the prediction (Cox & Snell R Square = .681, Nagelkerke R Square = .909). In addition, each estimated parameter has the expected sign: the number of submitters negatively influences winning chance, while the user's indegree enhances the winning chance.⁷ However, the number of submitters has much more predictive power than the user's indegree. In other words, when two users participate in the same task, the expertise level could slightly affect their chance of winning, but if they have

⁷ The magnitude of the estimated parameters (B in the table or Exp (B)) is sensitive to the measure units. We mainly look at the change in the -2 Log Likelihood for the effect size of each predictor factor.

different strategies to participate in different tasks that have very different numbers of submitters, then the one who attempts a less popular task would have a higher probability of winning. This result further suggests that users' participation strategy is quite important in leading to their better performance on the website.

The difficulty with the above approach is that the activity level of users is highly skewed, meaning that for most users, we did not have participation in both the test and training data. We tried a second, non-network based approach that takes into account a user's performance on all tasks excluding the task being predicted.

$$\text{Logit } P_{\text{win}} = \beta_1 X_{\text{number of submitters}} + \beta_2 X_{\text{userWins}} + \beta_3 X_{\text{userAttempts}} + \epsilon$$

We formed a balanced set of outcomes (50% wins and 50% losses), and used ten-fold cross validation to predict the outcome.

Table 8. Predicting user wins.

Variable	β
Log(# submissions)	-0.151
Log(# other attempts by user)	-0.135
# other wins by user	0.029
$R^2 = 0.708$	

As before, the number of submissions negatively impacts the probability that any single user wins the task, but the user's history of previous wins is fairly predictive of the outcome of the particular task ($R^2 = 0.57$). In short, we find that the amount of competition and the track record of the user are predictive of her probability of winning.

Conclusion and Future Work

This paper presents some general participation patterns for Taskcn.com, a knowledge sharing website where people pay other users for solutions to a variety of tasks. First, we observed that both users and tasks present scale-free characteristics, indicating users' uneven participation in tasks. The majority of users participate in a few tasks and win even fewer, while only a handful of users actually win the monetary awards. Most of the tasks have dozens of submissions while there are some "hot" tasks that attract even thousands of participants. We also noted some differences corresponding to task category that need to be further explored in subsequent work.

Adapting the idea of structural prestige in social network analysis, we were able to construct prestige networks for both users and tasks. We employed graph-based algorithms to measure users' expertise and tasks' prestige. These networks, in turn, allowed us to evaluate the interaction dynamics in Taskcn.

One surprising, and perhaps counterintuitive result is that monetary award is not a significant incentive for people to participate in a task on this website. Our results also show that a task's prestige (i.e., perceived expertise requirement) can slightly hinder people's participation. This means that users do consider the probability of winning when participating in Taskcn activities. However, there are likely to be many more factors that attract people to participate than just perceived difficulty.

In addition, we used prestige to evaluate the quality of all participants of a task. The result shows that a task's centrality in the network is correlated with a significantly lower average expertise level for its participants, while peripheral tasks attract more prestigious participants on average. As noted above, this could be because the more central tasks may be easier, and therefore more accessible, for a wider range of participants, but the peripheral tasks may require higher or more specific levels of expertise. This is to be explored in future work also.

These results suggest that design factors for these sites are likely to be tricky. For this particular site, monetary awards appear to be superfluous, and they have failed as an incentive mechanism to attract considerably better answerers and better solutions. This indicates that the efficiency of similar monetary incentives on websites is open to question.

We find that both the number of submitters and the user's prestige are significant factors in predicting the user's chance of winning a task. Relative to the user's prestige (indegree or number of wins), the number of competing submitters accounts for much more of the variance in a user's probability of winning. This means that when two users participate in the same task, their relative expertise level could affect who wins; but, for the individual user, even more important is the selection of tasks with fewer competitors. This again brings us back to efficiency, but from the perspective of the users. In future work, we would like to examine users' strategies over time in distributing their effort with respect to competitive tasks and the potential rewards. We especially want to examine how similar sites can provide the optimal amount of participation by both high prestige *and* naïve users.

In conclusion, in spite of their success in gathering many participants, Witkey websites face the challenges of efficiently allocating people's various levels of expertise and efforts onto different needs sources. And hence, our future work will be to further examine how to build these structural prestige measures into designing the incentive mechanism, thereby improving the system efficiency in knowledge exchange.

Acknowledgements

We would like to thank George Furnas, Jun Zhang, and Xiao Wei for insightful discussions. This work was funded in part by Intel Research, the Army Research Institute, and the National Science Foundation (grant IRI-9702904).

References

Ackerman, M.; Pipek, V.; and Wulf, V. 2003. *Sharing Expertise: Beyond Knowledge Management*. MIT Press.

Adamic, L.; Zhang, J.; Bakshy, E.; and Ackerman, M. S. 2008. Knowledge sharing and Yahoo Answers: Everyone knows something. *Proceedings of WWW'08*.

Barabasi, A.-L.; and Albert, R. 1999. Emergence of Scaling in Random Networks. *Science*, 286(5439): 509-512.

Berkhin, P. 2005. A Survey on PageRank Computing. *Internet Math*, 2: 48.

Broder, A.; Kumar, R.; Maghoul, F.; Raghavan, P.; Rajagopalan, S.; Stata, R.; Tomkins, A.; and Wiener, J. 2000. Graph structure in the Web. *Computer Networks*, 33(1): 309-320.

Campbell, C. S.; Maglio, P. P.; Cozzi, A.; and Dom, B. 2003. Expertise identification using email communications. *Proceedings of CIKM'03*: 528-531.

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

Dom, B.; Eiron, I.; Cozzi, A.; and Zhang, Y. 2003. Graph-based ranking algorithms for e-mail expertise analysis. *Proceedings of SIGMOD workshop on research issues in data mining and knowledge discovery*: 42-48.

Fisher, D.; Smith, M.; and Welsler, H. T. 2006. You Are Who You Talk To: Detecting Roles in Usenet Newsgroups. *Hawaii International Conference on System Sciences*.

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

Holloway, T.; Bozicevic, M.; and Börner, K. 2007. Analyzing and visualizing the semantic coverage of Wikipedia and its authors. *Complexity*, 12(3): 30-40.

Holme, P.; Edling, C. R.; and Liljeros, F. 2004. Structure and time evolution of an Internet dating community. *Social Networks*, 26(2): 155-174.

Kautz, H.; Selman, B.; and Shah, M. 1997. Referral Web: combining social networks and collaborative filtering. *Commun. ACM*, 40(3): 63-65.

Kou, Z., and Zhang, C. 2003. Reply networks on a bulletin board system. *Physical Review E*, 67(3): 036117.

Malone, T. W. 2004. *The Future of Work*. Boston, MA: Harvard Business School Press.

Milo, R.; Itzkovitz, S.; Kashtan, N.; Levitt, R.; Shen-Orr, S.; Ayzenshtat, I.; Sheffer, M.; and Alon, U. 2004. Superfamilies of evolved and designed networks. *Science*, 303(5663): 1538-1542.

Milo, R.; Shen-Orr, S.; Itzkovitz, S.; Kashtan, N.; Chklovskii, D.; and Alon, U. 2002. Network motifs. *Science*, 298(5594): 824-827.

Page, L.; Brin, S.; Motwani, R.; and Winograd, T. 1998. The pagerank citation ranking: Bringing order to the web: Stanford Digital Library Technologies Project.

Turner, T. C.; Smith, M. A.; Fisher, D.; and Welsler, H. T. 2005. Picturing Usenet: Mapping Computer-Mediated Collective Action. *Journal of Computer-Mediated Communication*, 10(4).

Wasserman, S., and Faust, K. 1994. *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.

Wenger, E. 1998. *Communities of Practice*. Cambridge: Cambridge University Press.

Wernicke, S., and Rasche, F. 2006. FANMOD. *Bioinformatics*, 22(9): 1152-1153.

Zhang, J.; Ackerman, M. S.; and Adamic, L. 2007. Expertise networks in online communities. *Proceedings of the 16th international conference on World Wide Web*.