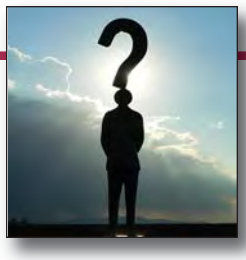

Don't Ask Me
Designing Social Q&A Services

By Gary Hsieh

Searching for information online has become an integral part of our everyday lives. However, sometimes we don't know the specific search terms to use, while other times, the specific information we're seeking hasn't been recorded online yet.

What we often resort to, after a few minutes of searching, is asking someone else.

This type of information-seeking behavior is one of the primary reasons social question and answer (Q&A) sites have become more popular. In fact, these sites—Yahoo! Answers is one example—enable tens of thousands of user-asked questions to be answered daily by other users.

The basic premise of these sites is that anyone who has a question can post it, and others in the community can respond and share their expertise or knowledge.

However, if you've ever used one of these social Q&A services, you'll know that they are not designed to accommodate differences in individuals' needs and constraints, which can result in inefficiency. People with questions who are in dire need of answers typically have no way to indicate their urgency and may not get an answers in time, while people who answer questions may feel overwhelmed by all the questions, especially if they're directed by potentially disruptive communication channels, such as instant messaging.

Quickly browse through Yahoo! Answers and you'll notice that seemingly important questions are presented alongside a substantial number of frivolous non-questions. For example, right below one sincere question, "Where can I go in Palm Bay, Florida, to get assistance with deposit money for an apartment?" there is a more frivolous one, "Is it just me or do you think Jeff from Big Brother 11 usa looks alot like the actor Jason Bateman?"

How can we better design social Q&A sites so that they are more sensitive to users' needs and constraints? In this article, I will discuss the strengths and weaknesses of three types of solutions.

Context-Sharing

The most direct solution is to allow question askers and answerers to share contextual information. If these communicating parties can be more informed about each others' needs and constraints, they may be able to make better decisions. These solutions have been explored in general communication domains.

Much research on media spaces in the 1990s explored how to use technology to improve awareness of remote collaborators. Colleagues can glance into another's workspace and engage opportunistically [1, 2], even though they may be working thousands of miles apart.

More recently, different types of status update mechanisms have been incorporated into everyday communication. These updates allow users to share location and activity information, as well as business-related project updates. With this meta-level information, askers can potentially target answerers who are available, and answerers can better infer the needs of the askers.

In my own work, I have explored the use of instant messaging tags to provide communication initiators with a way to signal their information needs [3]. People often use text tags in email subject lines to denote the type of email they are sending. By having a programmable set of tags for communication, additional services can be automatically triggered. For example, using the tag [15m] in a message indicates a level of urgency—a response within 15 minutes is desired. An automated reminder can be triggered when the time is up.

This type of support can be easily extended to online Q&A services. Question askers can provide additional information on how urgently they need the information, and answerers can then respond accordingly.

The main advantage of this solution is that it is straightforward and intuitive to the users; people are accustomed to using available contextual information such as gestures, body positioning and verbal statements to handle face-to-face requests. However, there are two major problems with this type of solution. First, sharing contextual information is only beneficial if askers and answerers have an incentive to respect each others' constraints and needs. We may be able to expect this from communicating partners who have existing social relationships, but we cannot expect this when communicating with strangers. Consider the case of the spammer, who may send spam regardless of how busy the answerers are. Second, full information disclosure has potential privacy problems. Not everyone is willing to offer full-disclosure, especially to strangers.

Intelligent Mediation

Instead of disclosing relevant information and relying on askers and answerers to make the proper decision on how to handle question requests, the second type of solution uses computing mechanisms to mediate social Q&A. This includes using collaborative filters to reduce spam and minimize unwanted requests, as well as utilizing social networks to target the questions to more appropriate answerers, as is done with Aardvark and Answer Garden.

But "intelligent" solutions can also leverage machine learning models to help determine when and how to target answerers for impending information requests. Using sensors placed in the environment, machine learning models can predict the answerers' interruptibility (see, for example, Fogarty [4]), which can then be used to prevent interruptions at inopportune moments. Additional models have been built to examine the cost of deferring communication, which can be used to improve the mediated decision of when to interrupt answerers (see, for example, Horvitz, Jacobs, and Hovel [5]).

Intelligent mediation designs have the potential to reduce request overload and minimize interruption costs for real-time Q&A services. Furthermore, individuals' privacy may not be violated as contextual

information is not directly disclosed to other users. However, these solutions work by blocking unwanted exchanges, where unwanted is usually defined by how costly the exchange is to the answerers. Even if the models are 100 percent accurate, mediation may not be ideal.

Consider a social Q&A service designed to be used by askers and answerers with some sort of higher-level shared goals, such as employees within a company. There may be scenarios where answerers should be interrupted to incur an immediate cost in order to help another who is working on more important and urgent tasks.

Market Pricing

Market pricing offers a third approach to improve existing communication technologies. The basic idea is that the askers need to pay or commit to pay the answerers for the answers. Frivolous requests may be reduced because the incremental cost of sending a question forces initiators to be selective, sending requests only if they believe the value of the answer is more than or equal to the price [6].

> **"How can we better design social Q&A sites so that they are more sensitive to users' needs and constraints?"**

Also, the use of pricing can allow initiators to signal the value of communication abstractly, without disclosing personal information. Additionally, financial rewards may attract more answerers. This idea has been explored by Google in Google Answers and is currently incorporated into various pay-for-answer services, such as Mahalo Answers, Just Answers, UClue, and AskBright.

I've also demonstrated a pay-for-answer service with a real-time communication channel: mimir [7]. Question askers can broadcast their questions to answerers directly, while offering mims, the virtual currency. On the other side, answerers can filter incoming requests based on the mims. When they do choose to answer, and if their answers are chosen as the best answers, they are rewarded with mims. Figure 1 shows the mimir interface.

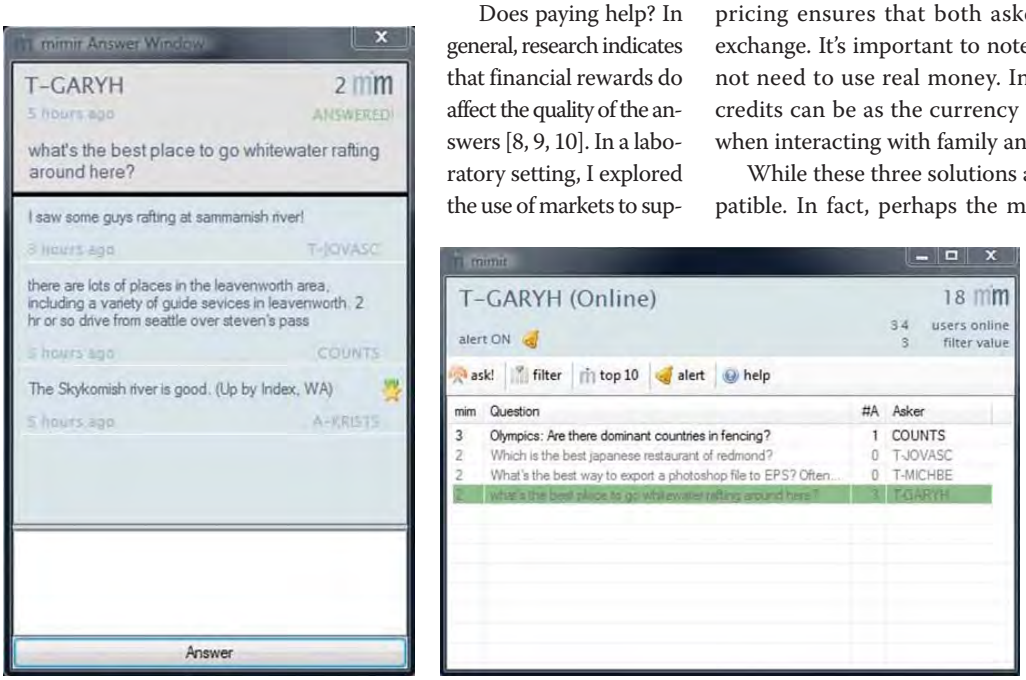

Figure 1: Interface for mimir, a market-based real-time social Q&A service.

Does paying help? In general, research indicates that financial rewards do affect the quality of the answers [8, 9, 10]. In a laboratory setting, I explored the use of markets to support help-seeking and help-giving. I found that allowing people to pay for help can indeed increase overall welfare for parties involved, although there may be overhead costs in deciding the value of getting help and the cost of giving it [11]. If users are expected to make these types of decisions many times a day, a complicated market may actually incur too much overhead cost and reduce potential gains. Related studies on pay-for-answer Q&A services also show that paying can improve overall answer quality [7, 8], and the length of the answers [9], but not necessarily the quality of the single best or chosen answer [9].

One area of research that needs further exploration is how financial rewards impact who uses these systems and the social relationships between users. Are these users more financially motivated than users on free social Q&A services? Related research in behavioral economics suggests that financial rewards can change the framing of the interaction between people, from social to transactional [11, 12]. Using financial incentives for Q&A may reduce the amount of social interactions that are vital in sustaining an online community. Further research is needed to determine if this also applies to social Q&A sites.

Applying the Solutions

Of the three types of solutions to improve social Q&A—context-sharing, intelligent mediation and market pricing—context-sharing may be most appropriate for Q&A services to be used with close family and friends who would use the contextual information in a positive way. Intelligent mediation can help Q&A between strangers as it reduces unwanted communication requests.

However, users may prefer to handle certain requests themselves when communicating with others with shared goals, as there may be cases where partners' gains outweighs one's own costs. Finally, market pricing ensures that both askers and answerers benefit from the exchange. It's important to note that the market pricing system does not need to use real money. In a corporate setting, company store credits can be as the currency of exchange, and tokens can be used when interacting with family and friends.

While these three solutions are fairly distinct, they are not incompatible. In fact, perhaps the most optimal solution is to leverage a hybrid system. For example, one could utilize both intelligent mediation and market pricing for social Q&A between strangers or other weak social-ties. As the relationship between users strengthens, the system could be gradually transitioned to use context-sharing and simple pricing rules with tokens.

This design will protect users from spam when interacting with strangers but offers flexibility when interacting with closer and more frequent communication partners. While more research is needed to explore hybrid strategies, it is cer-

tain that improvements to existing social Q&A services are necessary to allow people to fully leverage their knowledge and gain expertise from other users. The three solutions presented all are capable of resulting in social Q&A services that are more sensitive to our individual needs and constraints.

Biography

Gary Hsieh is a PhD candidate in the Human-Computer Interaction Institute at Carnegie Mellon University. His broad interests are in Human-Computer Interaction, Computer Mediated Communication and Behavioral Economics. Specifically, he is interested in studying, designing and developing technologies to enable people to communicate and interact in ways that are efficient and socially optimal.

Acknowledgements

The author thanks his collaborators on these projects: Scott Hudson, Robert Kraut, Scott Counts, and Jennifer Lai, and Karen Tang for offering insightful feedback on how to frame this article.

References

1. Gaver, W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K., and Buxton, W. 1992. Realizing a video environment: EuroPARC's RAVE system. In *Proceedings of CHI'92*. 27-35.
2. Fish, R. S., Kraut, R. E., Root, R. W., and Rice, R. E. 1993. Video as a technology for informal communication. *Comm. ACM 36*, 1. 48-61.
3. Hsieh, G., Lai, J., Hudson, S. E., and Kraut, R. E. 2008. Using tags to assist near-synchronous communication. In *Proceedings of CHI'08*. 223-226.
4. Fogarty, J. and Hudson, S. E. 2007. Toolkit support for developing and deploying sensor-based statistical models of human situations. In *Proceedings of CHI'07*. 135-144.
5. Horvitz, E., Jacobs, A., and Hovel. D. 1999. Attention-sensitive alerting. In *Proceedings of the Conference on Uncertainty in Artifical Intelligence* (*UAI'99*). 305-313.
6. Kraut, R. E., Sunder, S., Telang, R., and Morris, J. 2005. Pricing electronic mail to solve the problem of spam. *Hum. Comput. Interact. 20*. 195-223.
7. Hsieh, G. and Counts, S. 2009. mimir: A Market-Based Real-Time Question and Answer Service. In *Proceedings of CHI'09*. 769-778.
8. Harper, F. M., Raban, D., Rafaeli, S., and Konstan, J. A. 2008. Predictors of answer quality in online Q&A sites. In *Proceedings of CHI'08*. 865-874.
9. Chen, Y., Ho, T., and Kim, Y. 2009. Knowledge Market design: A field experiment on Google answers. *J. Public Economic Theory*. To appear.
10. Hsieh, G., Kraut, R., Hudson, S. E., and Weber, R. 2008. Can markets help?: Applying market mechanisms to improve synchronous communication. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (*CSCW'08*). 535-544.
11. Gneezy, U. and Rustichini, A. 2000. A fine is a price. *J. Legal Studies XXIX*, 1, part 1. 1-18.
12. Heyman, J. and Ariely, D. 2004. Effort for payment: A tale of two markets. *Psychol. Science 15*, 11. 787-793.