INQUIRY · INQUIRY IN PRACTICE

Question-Answer Platforms and Collective Inquiry Online

Question-and-answer platforms turned asking and answering into a public, collective activity on a vast scale. This article reviews the research on how these platforms work, what makes answers good, why people answer strangers' questions, how crowds can be wise or misled, and what the platforms reveal about collective inquiry.

ModQu research article · 2,881 words · 27 references (APA 7th edition)

For most of history, asking a question meant asking someone you knew or someone whose job it was to answer: a teacher, a librarian, a doctor, a clerk. Online question-and-answer platforms changed that. On them, anyone can ask a question in public, anyone can answer it, and the question and its answers remain available for others with the same question to find years later. Millions of people have taken part, asking about computer code, health, law, cooking, relationships and almost everything else. The platforms form a large and unusual experiment in *collective inquiry*, the shared pursuit of answers by people who mostly do not know one another. Researchers have studied them closely, and their findings reveal both the power and the limits of asking the crowd.

From reference desks to open platforms

Online communities were answering questions long before dedicated platforms appeared. Discussion groups on early computer networks developed lists of frequently asked questions to save members from answering the same questions repeatedly. Within organisations, Mark Ackerman and Thomas Malone (1990) designed *Answer Garden*, a system that captured answers to common questions so that an organisation could build up a memory of what its experts knew, directing new questions to experts only when no stored answer existed.

Dedicated public platforms followed in the 2000s. Figure 1 shows some milestones. Some platforms paid researchers to answer; most relied on volunteers motivated by reputation points, badges and the satisfaction of helping. Others were built around particular communities, such as programmers.

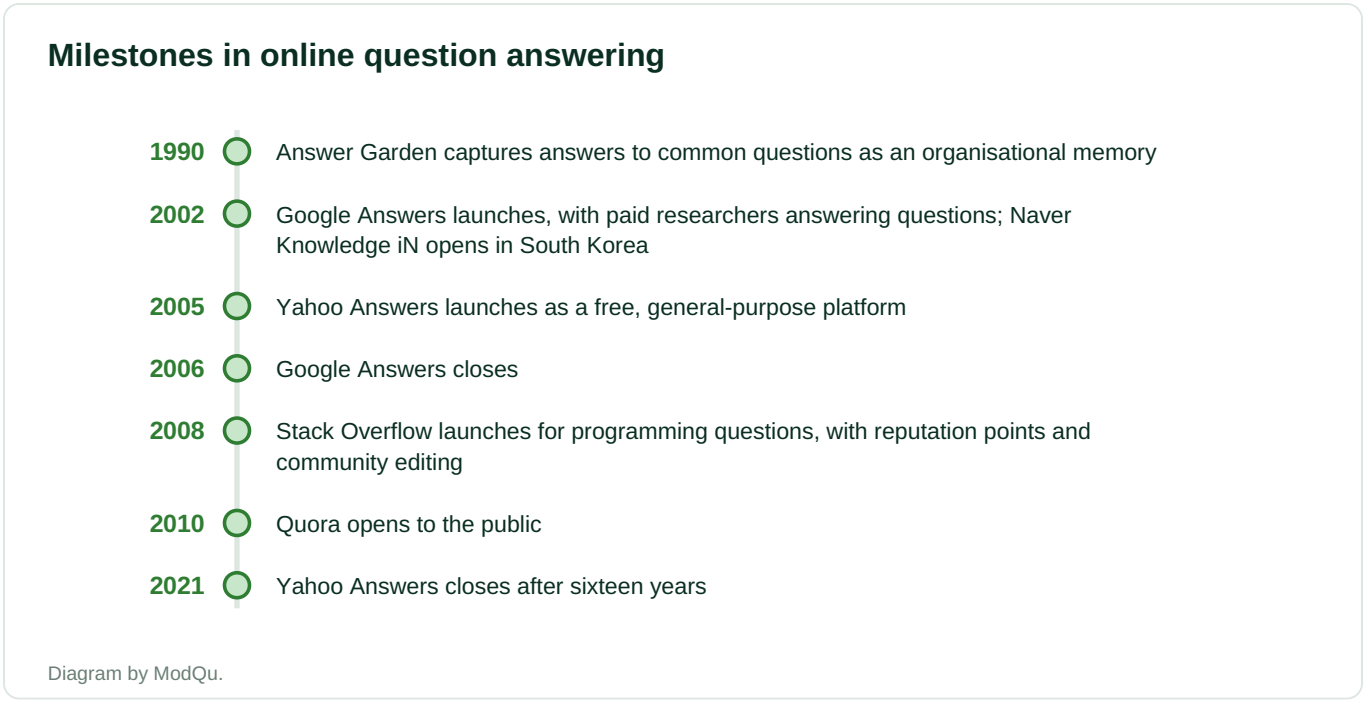


Figure 1. Milestones in online question answering. Diagram by ModQu.

Chirag Shah and colleagues (2009) defined *social Q&A* as a service in which people express information needs in natural language, other people answer, and the community can evaluate the answers. This differs from search, in which a person matches keywords against existing documents, in two ways. The asker can describe a specific, personal situation, and the answer is written for that situation by a person.

What people ask, and whom

People do not only use dedicated platforms. Meredith Ringel Morris, Jaime Teevan and Katrina Panovich (2010) surveyed people who had asked questions through status messages on social networking sites. Most questions sought recommendations or opinions, and people said they asked their networks rather than searching because they trusted their friends, wanted tailored answers, or felt that a search engine would not handle subjective questions well.

On public platforms, questions range from factual queries with a single correct answer to requests for advice, opinion and discussion. Lada Adamic and colleagues (2008) analysed a large general-purpose platform and found that categories differed sharply in character. Technical categories resembled focused exchanges, with many askers and a smaller group of knowledgeable answerers; discussion-oriented categories resembled forums, with longer threads and more interaction. They also found that users who focused their answers on a narrow range of topics tended to receive higher ratings for their answers than those who answered widely, consistent with the idea that expertise is specific.

How fast, and how good?

The speed of some platforms is remarkable. Mamykina et al. (2011) studied a leading platform for programming questions and found that most questions received an answer, and that the median time to a first answer was about eleven minutes. They attributed this success partly to the platform's design, including its reputation system and its focus on a single well-defined community, and partly to the active involvement of its founders in shaping community norms.

Answer quality varies considerably. Harper et al. (2008) compared answers to the same questions on several platforms, including a service on which askers paid for answers, and found that the fee-based service produced higher-quality answers than most free services, and that paying more produced better answers. However, a large free community also produced answers of good quality, apparently because many people could contribute. The number of people available to answer mattered as much as the incentive offered.

Anderson et al. (2012) studied the programming platform and found that the value of a question and its answers was not only to the asker. Many questions continued to attract visits long after they were answered, because other people with the same problem found them through search. The best answers emerged from an interaction among several contributors, in which early answers were improved, corrected or supplemented by later ones. A question with several good answers, discussed and refined over time, became a lasting resource.

Why strangers answer

Answering questions takes time and effort, and most answerers are not paid. Why do they do it? Peter Kollock (1999) analysed online cooperation as a form of gift exchange and public good, and identified motivations including anticipated reciprocity, the building of reputation, a sense of efficacy, and attachment to a group. Molly Wasko and Samer Faraj (2005) surveyed members of an online network of legal professionals and found that people contributed more when they expected to enhance their professional reputation and when they had experience to share, and less when they were driven only by a sense of obligation to the community.

Karim Lakhani and Eric von Hippel (2003) studied volunteer helpers who answered questions about open-source web server software. Most answers took only a few minutes, because helpers could draw on knowledge they already had, and helpers reported that they benefited from what they learned by reading others' questions and answers. The cost of answering was low and the benefit to the helper was real, which helped explain why so much help was given freely.

Platform design shapes these motivations. Robert Kraut and Paul Resnick (2012) reviewed research on online communities and set out design principles for encouraging contribution, such as making the need for contributions visible, asking specific people for help, giving feedback on the value of contributions and building a sense of identity with the community. Reputation points and badges, now standard on question-and-answer platforms, apply several of these principles.

Participation is unequal

In almost every online community, a small number of members contribute most of the content. Jakob Nielsen (2006) described a rule of thumb, sometimes called the 90-9-1 rule: about 90% of users read but never contribute, 9% contribute occasionally and 1% account for most contributions. Figure 2 shows this pattern. The proportions vary between communities, but the inequality is consistent.

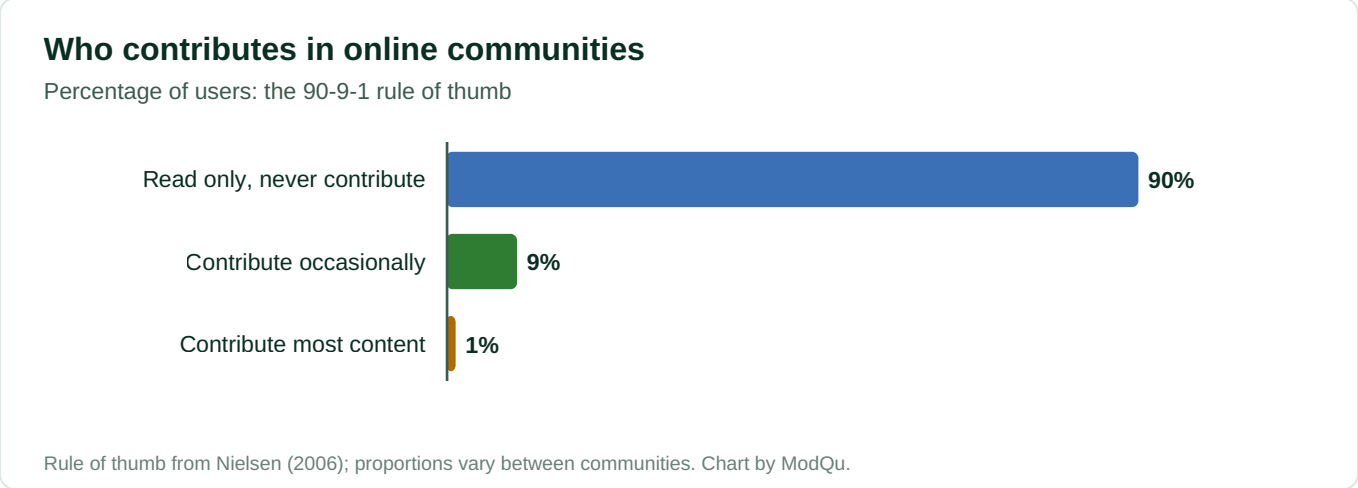


Figure 2. Participation inequality in online communities. A rule of thumb rather than a measurement of any particular site; based on Nielsen (2006). Chart by ModQu.

This inequality has consequences for collective inquiry. The answers available on a platform reflect the knowledge, interests and assumptions of its most active contributors, who may not represent the wider population of users or experts. Questions in areas where few active members have expertise may go unanswered or be answered poorly. Inequality also makes communities fragile. When a handful of prolific contributors lose interest or leave, the speed and quality of answers can fall sharply, and newcomers may find it hard to break into a community whose norms were set by a small established group. Platforms that welcome newcomers, explain their norms and recognise first contributions tend to sustain participation better than those that treat newcomers' mistakes harshly.

Can crowds be wise?

The idea that groups can know more than individuals has a long history. In 1906, Francis Galton (1907) examined the guesses submitted by visitors to a country fair who had paid to estimate the weight of an ox once slaughtered and dressed. The middle value of the roughly 800 guesses was within about 1% of the true weight. James Surowiecki (2004) argued that crowds can be wise when their members have diverse information, form their judgments independently and have their judgments combined in a suitable way.

Independence is the condition most easily lost online. Lorenz et al. (2011) asked people to estimate quantities and then showed some of them the estimates of others. Social influence reduced the diversity of the estimates without making them more accurate, and it made participants more confident. Salganik et al. (2006) created an artificial music market in

which participants could download songs. When participants could see how many times each song had been downloaded, success became both more unequal and more unpredictable: the same song could do well in one version of the market and poorly in another, depending on early downloads. Muchnik et al. (2013) ran a randomised experiment on a news discussion website, giving some comments a single positive vote when they were posted. Those comments were significantly more likely to receive further positive votes, and they ended with higher ratings on average. Early signals of approval, even arbitrary ones, shaped what the crowd later judged to be good.

Question-and-answer platforms are vulnerable to the same effects. The first plausible answer may attract votes that make it appear authoritative, and later, better answers may struggle to overtake it. Platform designs that sort answers by votes can therefore amplify early mistakes as well as early insights.

Errors and their correction

Open contribution brings errors as well as knowledge. Eric Raymond (1999), describing open-source software development, argued that "given enough eyeballs, all bugs are shallow": when many people examine work, errors are found and fixed quickly. The same principle can apply to answers. Jim Giles (2005) reported a comparison, organised by the journal *Nature*, in which experts reviewed science articles from an open, volunteer-written online encyclopaedia and from an established commercial encyclopaedia. The volunteer-written articles contained somewhat more errors on average, but the difference was smaller than many had expected. Aniket Kittur and Robert Kraut (2008) studied the same encyclopaedia and found that the quality of articles depended not simply on how many people edited them, but on how they coordinated. Articles improved most when a small number of editors did most of the work, coordinating implicitly, or when larger groups coordinated explicitly through discussion.

Correction is not guaranteed. Del Vicario et al. (2016) studied the spread of conspiracy theories and science news on a social networking site and found that users tended to cluster in communities that shared and reinforced the same kind of content. Within such communities, false claims could circulate without effective challenge. Vosoughi et al. (2018) found that false news stories spread further and faster than true ones on a major social media platform. Collective inquiry works when communities include people who know the subject and are willing to challenge errors; it fails when like-minded people reinforce one another.

Expertise in the crowd

Not all answerers are equally knowledgeable, and askers often cannot tell who knows what. Sanghee Oh (2012) studied people who answered health questions on a social Q&A platform and found that many drew on personal experience, either their own or that of family members, rather than professional training, and were motivated by altruism and a wish to help people in similar situations. Experiential knowledge can be valuable, offering practical advice and emotional support that professionals may not provide, but it is not a substitute for medical expertise.

Researchers have developed many methods for identifying good answers and expert answerers automatically, using signals such as past answer ratings, the answerer's topic focus and the structure of the answer. Ivan Srba and Mária Bieliková (2016) reviewed this body of research, which includes approaches to routing questions to people likely to answer them well, predicting answer quality and detecting duplicate questions.

Questions that go unanswered

Even on the most active platforms, some questions receive no answer, and the unanswered questions are informative in their own right. Some are poorly phrased, too vague or missing the details an answerer would need. Others are about topics where few knowledgeable members are active. Others again are simply hard: they concern problems that nobody in the community has solved. Platforms have developed responses to each case, such as prompting askers to add details, editing questions for clarity and allowing members to offer rewards for difficult questions. For anyone studying a field, a collection of persistent unanswered questions is a map of where knowledge runs out.

Question answering inside organisations

The same dynamics operate within organisations. Many companies run internal question-and-answer systems, discussion channels or expert directories in the hope of capturing what their staff know, much as Ackerman and Malone (1990) proposed. Their success depends on the same conditions as public platforms. Staff answer when it takes little time, when their contribution is recognised and when they see others doing the same. They stop when questions disappear into a void, when the same questions are asked repeatedly without anyone collecting the answers, or when answering is seen as a distraction from "real" work.

Internal systems also face participation inequality. A few knowledgeable people may answer most questions, and when they leave, much of the organisation's practical knowledge leaves with them unless their answers have been recorded and organised. Treating good answers as lasting assets, by curating them, keeping them current and making them easy to find, turns a stream of conversations into an organisational memory.

Automated answers

Automated systems that generate answers to questions in natural language now handle many questions that people once posted to platforms. They raise the same issues the research on platforms has identified: whether the answers are accurate, whether their sources can be checked, whether early or popular answers crowd out better ones, and whether errors can be corrected. The lesson of collective inquiry applies here too. An answer is only as trustworthy as the process that produced it and the means available to check it.

Collective inquiry and its conditions

The research suggests that question-and-answer platforms support good collective inquiry under certain conditions, summarised below.

- **A focused community.** Platforms built around a well-defined domain, with shared standards of what counts as a good answer, perform better than general ones.
- **Enough knowledgeable participants.** The quality of answers depends on the presence of people who know the subject, and on enough of them to check one another.
- **Independent judgment.** Systems that reveal popularity too early can amplify initial errors.
- **Visible correction.** Mechanisms for editing, commenting on and improving answers allow errors to be caught.

- **Lasting records.** Questions and answers that remain available create value for many future askers, not just the first.

These conditions echo older ideas about inquiry. Pierre Lévy (1997) described *collective intelligence* as intelligence distributed throughout a group and coordinated in real time. Yochai Benkler (2006) described *commons-based peer production*, in which large numbers of people cooperate to produce information goods without markets or managers. Etienne Wenger (1998) described *communities of practice*, groups bound by a shared domain and practice, in which members learn through participation. The most successful question-and-answer platforms combine elements of all three.

What platforms reveal about public inquiry

For researchers and organisations, the questions asked on these platforms are themselves evidence. They show what people want to know, in their own words, and what puzzles them enough to ask strangers. The answers show what people believe and what advice circulates among them. The votes show what a community considers helpful.

Consider a hypothetical example. A manufacturer of home espresso machines reads the questions posted about its products on a large cooking platform. It finds that the most frequent questions are not about features but about descaling, and that the highest-voted answers recommend a method that the company's engineers know can damage the machine. The company improves its instructions, publishes a clear guide, and commissions a survey of owners to learn how widespread the problem is and how owners currently care for their machines. The platform revealed the question; the survey measured its extent. The survey finds that about a third of owners have never descaled their machines at all, a group the platform had not revealed, because people who do not know a task is needed do not ask how to do it. The company adds a reminder to its machines' displays. Without the platform, the company would not have known what to ask; without the survey, it would not have found the larger group of owners who were not asking anything.

Such evidence must be read with care. Platform users are not a representative sample of the population, the most visible questions may come from the most active users, and popular answers are not necessarily correct. As with search data, the questions on these platforms are best treated as a guide to what deserves further investigation by methods that can establish who is asking, how many and why.

Conclusion

Question-and-answer platforms have made collective inquiry possible on a scale never seen before. At their best, they connect people with questions to people who know the answers within minutes, and build lasting archives of practical knowledge. Research shows that their success depends on focused communities, motivated and knowledgeable contributors, independent judgment and visible correction. Where those conditions fail, the same platforms can spread error and give early, popular answers more authority than they deserve. The lessons apply to any setting in which groups pursue answers together: diverse knowledge, independent judgment and a way of correcting mistakes are what make a crowd wise.

References

Ackerman, M. S., & Malone, T. W. (1990). Answer Garden: A tool for growing organizational memory. In *Proceedings of the ACM SIGOIS and IEEE CS TC-OA Conference on Office Information Systems* (pp. 31–39). Association for Computing Machinery. https://doi.org/10.1145/91474.91485

Adamic, L. A., Zhang, J., Bakshy, E., & Ackerman, M. S. (2008). Knowledge sharing and Yahoo Answers: Everyone knows something. In *Proceedings of the 17th International Conference on World Wide Web* (pp. 665–674). Association for Computing Machinery. https://doi.org/10.1145/1367497.1367587

Anderson, A., Huttenlocher, D., Kleinberg, J., & Leskovec, J. (2012). Discovering value from community activity on focused question answering sites: A case study of Stack Overflow. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 850–858). Association for Computing Machinery. https://doi.org/10.1145/2339530.2339665

Benkler, Y. (2006). *The wealth of networks: How social production transforms markets and freedom*. Yale University Press.

Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences, 113*(3), 554–559. https://doi.org/10.1073/pnas.1517441113

Galton, F. (1907). Vox populi. *Nature, 75*(1949), 450–451. https://doi.org/10.1038/075450a0

Giles, J. (2005). Internet encyclopaedias go head to head. *Nature*, *438*(7070), 900–901. https://doi.org/10.1038/438900a

Harper, F. M., Raban, D., Rafaeli, S., & Konstan, J. A. (2008). Predictors of answer quality in online Q&A sites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 865–874). Association for Computing Machinery. https://doi.org/10.1145/1357054.1357191

Kittur, A., & Kraut, R. E. (2008). Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work* (pp. 37–46). Association for Computing Machinery. https://doi.org/10.1145/1460563.1460572

Kollock, P. (1999). The economies of online cooperation: Gifts and public goods in cyberspace. In M. A. Smith & P. Kollock (Eds.), *Communities in cyberspace* (pp. 220–239). Routledge.

Kraut, R. E., & Resnick, P. (2012). *Building successful online communities: Evidence-based social design*. MIT Press.

Lakhani, K. R., & von Hippel, E. (2003). How open source software works: "Free" user-to-user assistance. *Research Policy, 32*(6), 923–943. https://doi.org/10.1016/S0048-7333(02)00095-1

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences, 108*(22), 9020–9025. https://doi.org/10.1073/pnas.1008636108

Lévy, P. (1997). *Collective intelligence: Mankind's emerging world in cyberspace* (R. Bononno, Trans.). Plenum.

Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G., & Hartmann, B. (2011). Design lessons from the fastest Q&A site in the west. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2857–2866). Association for Computing Machinery. https://doi.org/10.1145/1978942.1979366

Morris, M. R., Teevan, J., & Panovich, K. (2010). What do people ask their social networks, and why? A survey study of status message Q&A behavior. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1739–1748). Association for Computing Machinery. https://doi.org/10.1145/1753326.1753587

Muchnik, L., Aral, S., & Taylor, S. J. (2013). Social influence bias: A randomized experiment. *Science, 341*(6146), 647–651. https://doi.org/10.1126/science.1240466

Nielsen, J. (2006, October 9). *Participation inequality: Encouraging more users to contribute*. Nielsen Norman Group.

Oh, S. (2012). The characteristics and motivations of health answerers for sharing information, knowledge, and experiences in online environments. *Journal of the American Society for Information Science and Technology, 63*(3), 543–557. https://doi.org/10.1002/asi.21676

Raymond, E. S. (1999). *The cathedral and the bazaar: Musings on Linux and open source by an accidental revolutionary*. O'Reilly.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science, 311*(5762), 854–856. https://doi.org/10.1126/science.1121066

Shah, C., Oh, S., & Oh, J. S. (2009). Research agenda for social Q&A. *Library & Information Science Research, 31*(4), 205–209. https://doi.org/10.1016/j.lisr.2009.07.006

Srba, I., & Bieliková, M. (2016). A comprehensive survey and classification of approaches for community question answering. *ACM Transactions on the Web, 10*(3), Article 18. https://doi.org/10.1145/2934687

Surowiecki, J. (2004). *The wisdom of crowds*. Doubleday.

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science, 359*(6380), 1146–1151. https://doi.org/10.1126/science.aap9559

Wasko, M. M., & Faraj, S. (2005). Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. *MIS Quarterly, 29*(1), 35–57. https://doi.org/10.2307/25148667

Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.