

Use of credibility heuristics in a social question-answering service

Paul Matthews

**University of the West of England, Bristol, Coldharbour Lane, Bristol
BS16 1QY**

Abstract

Introduction. This study looked at the effect of community peripheral cues (specifically voting score and answerer's reputation) on the user's credibility rating of answers.

Method. Students in technology and philosophy were asked to assess the credibility of answers to questions posted on a social question-answering platform. Through the use of a test interface, the availability of cues as to community rating and the reputation of the answer providers were manipulated.

Analysis. The main analysis concerned the correlations between the students' credibility rating and the community votes the answers had received.

Results. Findings showed that across subjects and knowledge levels, the availability of cues increased the level of agreement between the test subjects and the community aggregate vote. Agreement with answerer reputation was weaker but still significant. Furthermore, at low knowledge levels cue visibility increased confidence regarding credibility judgements.

Conclusions. The findings are discussed within a framework of social epistemology, which provides a theoretical foundation for the assessment of knowledge systems. Some suggestions are provided for future work aimed at correcting overreliance on heuristics in credibility decisions.

Introduction

Knowledge is of two kinds. We know a subject ourselves, or we know where we can find information upon it - Samuel Johnson

To the aggregate of individuals we need to add the morally textured relations between them, notions like authority and trust and the socially situated norms which identify who is to be trusted, and at what price trust is to be withheld - Steven Shapin

Over the years, communities of interest on the web have developed extremely useful knowledge exchange resources on a variety of subjects. Such platforms are widely used by learners and professionals in solving technical and theoretical problems in their field. Answer providers, motivated by community spirit and the willingness to help, often devote hours of their time to monitoring expertise areas and answering questions from the community.

A range of social and cognitive factors contribute to user's assessment of community-

contributed answers. Factors such as social agreement, expertise and authority have transferred to the online world and are used, sometimes without question, as credibility cues. Seen from the viewpoint of evaluation of online platforms as ideal knowledge exchange environments, there are potential problems with this use of broad-brush cues: social voting mechanisms may unfairly weight particular types of answer, or involve self-reinforcing mechanisms; online reputation may mean something very different to real-world reputation and the sheer speed on online interaction may bypass cognitive faculties designed for the critical assessment of evidence.

Because online credibility cues are so important yet relatively poorly understood, research is needed to assess their impact on belief formation and to better understand their benefits and risks. Hugely multiuser, collaboratively edited information resources have been identified as a particular priority ([Rieh and Danielson, 2007](#)). This study was therefore designed to investigate the power of social interface heuristics in credibility assessments amongst students. Before presenting the experimental approach, I will present the theoretical connection between online knowledge exchange practice and social epistemology, for what it tells us about how learners may accommodate new evidence and how we can design and evaluate epistemic (knowledge-oriented) systems. There will follow some discussion of the platforms under study and their emergent dynamics. Prior work on online credibility and with the experimental manipulation and evaluation of interfaces will then be discussed.

Philosophical dimensions of knowledge exchange

Social epistemology is the branch of philosophy concerned with studying knowledge and belief in a social setting. Amongst other things social epistemology is concerned with '*choices among alternative institutions, arrangements or characteristics of social systems that influence epistemic outcomes*' ([Goldman, 2010](#)). For online sociotechnical systems designed for knowledge exchange, social epistemology presents a framework for evaluation and a broad set of values that may guide design ([Matthews and Simon, 2012](#))

To many social epistemologists, personal testimony is a *first class* source of knowledge (the non-reductionist position). There are cogent and convincing arguments as to the importance of others' testimony to our knowledge of the world beyond our immediate experience. A strong argument in social epistemology is that acceptance of testimony is neither purely the role of the giver or the receiver but is achieved through the interaction of both with a commitment to reliability on both sides. ([Lackey, 2008](#))

The theory of credibility in testimony has largely been construed in face-to-face situations and thought experiments, though more recently epistemologists have considered online knowledge sources. Anonymity and pseudonymity in many of these has been raised as a significant issue in forming credibility judgements (e.g. ([Magnus, 2009](#))). At the same time, any available cues as to authority and credibility, or *indicator properties*, stand to be challenged ([Goldman, 2010](#)). Indeed, the idea of authority is multifaceted and evolving, with accessible, virtue-based systems winning over more entrenched, inherited approaches in some domains ([Origgi, 2012](#))

In terms of judgement aggregation, the Condorcet principle has been an influential part of social epistemology theory. Here, aggregated judgements may be more accurate than those of any individual or sub-group given certain conditions. Chief among these are that groups need to be large, independent and diverse ([Goldman, 2010](#), [Sunstein, 2006](#)). Amongst epistemologists, there has been some debate over whether weight of numbers toward a

particular position should confer stronger epistemological status, with one position being that this is the case only when each individual brings to bear some individual knowledge on the issue and is not simply parroting another's position ([Coady, 2012](#)). As related to online collaborative filtering, this is perhaps a caution against accepting weight of agreement as proportional to resource quality, particularly when many voters will not be experts.

Social question answering

[StackOverflow](#) was the first in what is now a family of topic-specific question answer sites, the StackExchange network. Members post questions on a labelled topic (the StackOverflow site being concerned with aspects of computer programming) and receive answers from the community. Other community members may up- or down- vote question and answers and the original poster can accept a particular answer as the most suitable. Up- votes and accepted answers earn reputation points for the site member. Increases in reputation points above particular thresholds are associated with elevated member privileges, such as the ability to down-vote or move for questions to be closed.

The most common type of question on StackOverflow is a *how-to*, instructional type, with queries regarding errors and discrepancies quite common, along with requests for help with design or technology decisions ([Treude, Barzilay and Storey, 2011](#))

Social question-answering sites such as StackOverflow, which expose their content to web indexing and search, attract a far higher number of passive viewers over and above those directly involved in the question and answering process. On StackOverflow it is not uncommon for popular threads to gain half a million views over a five year period but have twenty or fewer member contributions. External viewers cannot contribute (in the form of voting, editing or providing answers) without registering with the service. In many cases, then, the external viewer is interacting in a rather static way with the resource and basing judgements on the activities of the community members who created it.

Social question-answering sites such as StackOverflow employ a *gamified* interface in rewarding contributors with reputation points. This may lead to the *fastest gun* effect, where users compete to be the first to answer a question, thereby benefitting from the user *accepting* their answer as the best available ([Mamykina, 2011](#)). Analysis suggests that the first answer preferred has a 60% likelihood of being accepted, dropping to 35% by the third. Temporal features are therefore heavily correlated with the likelihood of user acceptance and high ranking in the answer list ([Matthews, 2014](#))

The later emergence of new, higher voted answers is a rarity, though it can happen in a small percentage of cases. Even in a short time window, however, the rate at which the later answers may accumulate votes, especially given a large answer pool, may be slow. A snapshot of a thread at any particular time may therefore give a misleading view of the relative answer qualities as perceived by the community.

Online credibility models and social question-answering

Studies of the credibility of web-based information have often divided the concept into characteristics of the source, the message/medium and the receiver ([Hilligoss and Rieh 2008](#), [Metzger, 2007](#), [Wathen and Burkell, 2002](#)). Source characteristics include expertise, honesty and reputation together with institutional affiliation and motive. On [Quora](#), a general

social question-answering site, a *real names* policy helped to ensure that user profiles were more genuine, improving the site's credibility and user satisfaction. On this network, the ability to follow particular users and view users answer history are ways used to assess respondent quality. ([Paul, Hong and Chi, 2012](#)). In QandA, reputation may be in algorithmic form, representing user contribution history, or derived from a user's real world identity. In MathOverflow, one of the StackExchange network sites, both online and offline reputation were found to contribute separately to the prediction of answer quality ([Tausczik and Pennebaker, 2011](#)).

Relevant message characteristics include objectivity, clarity, consistency, coverage and timeliness. Usage of these features may vary with knowledge level. For collaboratively edited articles non-expert users have been found to rely more on surface and source features of the content than semantic features ([Lucassen and Schraagen 2011](#)). Notably, credibility studies have traditionally tended to consider the medium and message as one and the same ([Wathen and Burkell 2002](#)), or only looked at general aspects of the interface such as usability and usefulness. When studies have looked at the effect of having the same information presented on different types of platform, those media with a collective, multi-user aspect or with more clearly visible gatekeeping may be perceived as more credible ([Kim and Sundar 2011](#)).

At the side of the receiver, prior conceptions and stereotypes may be significant, as are the general motivational state of the information seeker (strength of the information need, available time and personal stake). These last conditions may determine the extent to which the user engages with central aspects of the message or relies on more peripheral cues ([Hilligoss and Rieh, 2008](#)). The selection and application of heuristics in perception, judgement and decision-making has been studied by social psychologists in a range of settings. The dual process theory suggests a fast, automatic cognitive system (system one), attuned to pattern recognition and parallel processing, and a slower analytical system (system two) that is brought to bear on more difficult and complex tasks. Significantly, system one may be that involved in the formation of beliefs and system two with their later ratification or rejection ([Kahneman, 2011](#)).

Interface heuristics and credibility

Credibility heuristics are relatively simple to manipulate experimentally and evidence of their influence on credibility judgements exists. Kim and Sundar found '*bandwagon cues*' in the form of thread views, replies and share count to be more influential than authority cues in the assessment of the credibility of community-provided health information and the intention to share it ([Kim and Sundar, 2011](#)). In a similar study for product reviews, unequivocal bandwagon cues were found to be more powerful than seals of approval in determining purchasing intention ([Sundar, Xu, and Oeldorf-Hirsch, 2009](#))

For online news, source credibility may be a key heuristic, though *freshness* and the number of related stories may take over if the source credibility is not immediately recognisable ([Sundar, Knobloch-Westerwick and Hastall, 2007](#))

Popularity signals for community users may also function as a heuristic when present. On Quora, users also reported the tendency for people with large numbers of followers to receive many up-votes for their answers, regardless of the answer quality. This made it more difficult for other answers to receive attention ([Paul et al., 2012](#)).

Rationale for credibility judgements

In a large study of the answer selection criteria employed in in Yahoo Answers, Kim and Oh divided criteria into overall classes by content (accuracy, specificity, clarity, rationality, completeness and style), cognitive (novelty, understandability), utility (effectiveness, suitability), information source (use of external sources, answerer expertise), extrinsic (external verification, available alternatives, speed) and socio-emotional (agreement, emotional support, attitude, effort, taste and humour). Socio-emotional, content and utility classes dominated (Kim and Oh, 2009). The authors noted that the importance of the socio-emotional type was connected to the types of thread analysed, which included a number of discussion, opinion-based questions.

Method

The study involved forty-two undergraduates and eleven postgraduates on a range of courses: undergraduates in computing and information science, undergraduates in philosophy and a postgraduate programme in information management. Students were selected as attendees from the author's own teaching programmes except for philosophy, which was arranged separately. Students took part in the investigation voluntarily in class time, though the questions chosen for each group were of relevance to their programme of study, so that participation was made to feel a little like an in-class test. Students were informed that the study was to investigate how students interact with online question-answering systems

	Undergraduate computing and information science	Postgraduate information science	Undergraduate philosophy
Subjects	29	11	13
Questions presented	16	5	5
Rated answers (cues visible)	294	63	75
Rated answers (cues not visible)	243	66	114

Table 1: Summary of participants and conditions

An experimental interface for the study was created, employing data from the Stack Exchange family of Websites. Questions were preselected and question and answer text and metadata was downloaded to the interface on demand using the Stack Exchange API. Questions were chosen to be within the curricula of the students' courses, though not always covered in explicit detail or by the time of testing. For the undergraduate students, data was drawn from [Stack Overflow](#) or [Philosophy Stack Exchange](#), whereas for the postgraduate students it came from the [Library Stack Exchange Beta](#). Questions were chosen that had at least three viable answers and when more than three answers were available, a high, medium and low-ranking answer were selected from those available. Question order was randomised and in half of cases (by subject, chosen randomly) the credibility cues of votes, answerer reputation and user acceptance were hidden. In other cases this information was displayed prominently before the

text of the users' answers.

Question:

Difference between inner and outer join

What is the difference between **INNER JOIN** and **OUTER JOIN** ?

How much do you feel you already know about this?

Little or nothing 1 2 3 4 5 Quite a lot

Answer 1:

I don't see much details about performance and optimizer in the other answers.

Sometimes it is good to know that only inner join is associative which means the optimizer has the most option to play with it. It can reorder the join order to make it faster keeping the same result. The optimizer can use the most join modes.

Generally it is a good practice to try to use inner joins instead of the different kind of outers.

How convincing do you find this answer?

Not at all convincing 1 2 3 4 5 Very convincing

Please say why you chose that score:

Figure 1: Experimental interface showing question and first answer

Question:

Difference between inner and outer join

What is the difference between **INNER JOIN** and **OUTER JOIN** ?

How much do you feel you already know about this?

Little or nothing 1 2 3 4 5 Quite a lot

Answer 1:

User Reputation: 64402, Total Votes: 1870, Accepted as Best Answer by User?: yes

Assuming you're joining on columns with no duplicates, which is by far the most common case:

- An inner join of A and B gives the result of A intersect B, i.e. the inner part of a venn diagram intersection.
- An outer join of A and B gives the results of A union B, i.e. the outer parts of a venn diagram union.

Examples

Suppose you have two Tables, with a single column each, and data as follows:

Figure, 2: Experimental interface with cues shown (highlighting for illustration only)

Subjects were asked to rate the question in terms of the extent of their prior awareness of the topic, and to rate each answer in terms of how convincing they believed them to be (selected as proxies for prior knowledge and credibility respectively and validated with a pilot test). For each answer subjects were encouraged to write a few words to justify their allocated score. Subjects were given unlimited time and asked to review three to five questions presented in random order.

Accompanying comments were post-coded inductively via a two-pass method. Categories were created on a first pass, then were merged and further edited on a second pass. A number of comments were removed as they failed to add useful information as to the reasons for the credibility assignment.

For data analysis, prior knowledge scores were aggregated into three levels, with scores of 1 and 2 assigned to low, 3 to medium, and 4 and 5 to high. The same grouping was used for the credibility scores when it came to presentation of the qualitative comments.

Results

855 answer assessments were received, together with 532 usable comments relating to the assigned credibility rating (Table 1).

Effect of credibility cue availability

Table 2 provides the correlation between the subjects' credibility ratings of answers and the community score (the total up-votes received minus the down-votes during the lifespan of the question). Kendall's Tau was used as a correlation coefficient due to the non-normal distribution of answer and user reputation scores and the need to derive a probability estimate in the presence of ties. Overall the relationship was shown to be significant, though when the cues were visible the strength of the correlation more than doubled. This finding was consistent across the experimental groups, although was less marked for the postgraduates. A similar relationship was observed between the answer provider's reputation score (Table 3) and the user's credibility rating, though in this case the relationship appeared to only be significant when this cue was available.

	Cues not shown	Cues shown	Difference
Undergraduate computing and information science	0.17 (0.001)*	0.273 (0.000)*	0.103
Postgraduate information science	0.196 (0.054)	0.218 (0.052)	0.022
Undergraduate philosophy	0.132 (0.078)	0.294 (0.002)*	0.162
Overall	0.137 (0.000)*	0.284 (0.000)*	0.147

Table 2: Kendall's tau correlations between community up-votes and user credibility rating, by experimental condition and subject group. Probabilities of zero correlations are shown in

brackets.

When disaggregated instead by prior knowledge level (a measure taken at the level of the question), it was apparent that the low knowledge condition was associated with a heavy reliance on the cues and that in their absence the low knowledge users failed to agree with the community as to the quality of the questions (Table 3, Table 4).

	Cues not shown	Cues shown	Difference
Low	0.087 (0.115)	0.28 (0.000)*	0.193
Med	0.198 (0.003)*	0.31 (0.000)*	0.112
High	0.248 (0.005)*	0.3 (0.000)*	0.052

Table 3: Kendall's tau correlations between community votes and subject's credibility rating, by condition and prior knowledge level. Probabilities of zero correlations are shown in brackets.

The overall correlation between answerer's reputation and community votes was 0.498 ($p < 0.000$) for all answers. All coefficients in the study are low, because of the importance of additional factors besides aggregate up-votes in answer ratings. This is also evidence that highly voted answers are not necessarily the highest objective quality (see *fastest gun* effect above). The difference between correlations with and without cue visibility, however, is felt to be indicative of its use as a credibility heuristic.

	Cues not shown	Cues shown	Difference
Low	0.076 (0.162)	0.182 (0.001)*	0.106
Medium	-0.065 (0.309)	0.208 (0.003)*	0.273
High	0.204 (0.017)*	0.128 (0.057)	-0.076

Table 4: Kendall's tau correlations between answerer's reputation and subject's credibility rating, by condition and prior knowledge level. Probabilities of zero correlations are shown in brackets.

Further analysis of the behaviour of subjects under different knowledge conditions revealed that the availability of cues led to a significantly more confident credibility score being assigned on average only for the low knowledge condition. In the medium and low knowledge conditions, the credibility ratings did not appear to be significantly affected by cue availability.




Figure 3: Mean credibility ratings according to prior knowledge level in the two experimental conditions (light grey = cues are shown)

Supporting rationale

The coding that was assigned to the qualitative comments is summarised in Tables 5, 6 and 7. These broadly related well to Kim and Oh's classification ([Kim & Oh, 2009](#)), though those authors only looked at reasons for choosing the best answer and not on reasons for low credibility. Nevertheless, codes tended to be bivalent, with the negation of the code applicable to low credibility answers (e.g., clear/unclear).

The most common type of comment was where the subject engaged with the subject of the answer and provided their own argument for or against it. As might be expected, these types of justifications featured more heavily in the Philosophy students than in the others (Table 8). The clarity, accuracy and completeness of answers accounted for a large number of additional comments. Where students had prior experience of a question topic, they tended to cite this as a reason for finding an answer convincing or otherwise.

Notably, explicit reference to the community score or answerer's reputation score was only noted in three cases, indicating that users were not consciously using this as a decisive criterion (or were reluctant to record as much!).

Simplicity above over-complexity seemed to be valued most by the computing and IS

undergraduates (Table 5), illustrative of the technical nature of the subject matter and the range of possible approaches to answering questions. While providing alternative explanations and clear code examples seemed to be valued, longer answers were sometimes criticised for undue complexity.

Low	Medium	High
argument against answer (17)	unable to say why! (8)	argument for answer (13)
incomplete explanation or lacking detail (11)	incomplete explanation or lacking detail (7)	coherence (10)
lack of clarity (7)	clarity (5)	accuracy or completeness (7)
unable to say why! (7)	lack of detail (5)	multiple explanations (7)
over-complexity (6)	argument against answer (4)	clarity (6)
lack of knowledge of subject (5)	coherence (3)	examples offered (6)
tangential/does not answer question (5)	external references provided (3)	simplicity (5)
Doesn't fit preconception or prior experience (4)	undecided (3)	comparative (4)
lack of detail (4)	argument for answer (2)	fits preconception or prior experience (4)
lack of examples (3)	comparative (2)	succinct (4)

Table 5: Supporting comments by credibility rating (Undergraduate computing and information science)

Low	Medium	High
argument against answer (7)	argument against answer (4)	argument for answer (8)
opinion based (4)	argument for answer (4)	coherence (5)
incomplete explanation / lacking detail (2)	lack of knowledge of subject (3)	experience of answerer (4)
accuracy or completeness (1)	coherence (2)	practicality (4)
lack of detail (1)	incomplete explanation / lacking detail (2)	fits preconception or prior experience (3)
succinct (1)	tangential/does not answer question (2)	accuracy or completeness (2)
tangential or does not answer question (1)	accuracy or completeness (1)	examples offered (2)
	clarity (1)	incomplete explanation or lacking

	lack of confidence (1) lack of detail (1)	detail (2) sympathise (2) clarity (1)
--	--	---

Table 6: Supporting comments by credibility rating (Postgraduate)

Low	Medium	High
argument against answer (28) tangential/does not answer question (10) lack of clarity (3) lack of detail (3) incomplete explanation or lacking detail (2) lack of answerer understanding (1)	argument against answer (8) incomplete explanation or lacking detail (4) argument for answer (2) lack of detail (2) coherence (1) examples offered (1)	argument for answer (32) accuracy or completeness (4) succinct (4) coherence (2) answer presentation (1) clarity (1)

Table 7: Supporting comments by credibility rating (Undergraduate philosophy)

Discussion

The experimental data reveals the tendency for the availability of credibility cues to cause closer agreement between the users of the resource and user ratings resulting from prior community activity. While the effect was most marked in the condition of low knowledge it existed across the spectrum. Furthermore, in conditions of low knowledge, users became more confident in their judgements when cues were present. Additionally, there is provisional evidence that these effects exist across heterogeneous subject domains. The findings are in general agreement with Kim and Sundar (2011) and Sundar, Xu(2009), that social rating (bandwagon) cues seem to be more powerful than authority cues, though both seem to exert an influence.

We see here the double-edged nature of these credibility heuristics: in the absence of cues even those who know something about the topic will be more likely to disagree on the best answer. While much of the time, especially for low knowledge users, the cues can provide a useful guide, they may also draw the eye away from answers that are worthy of greater consideration. It is effectively creating a *casca*d which otherwise adapts well to a mix of expertise, akin to a Condorcet system but with bias introduced through prior deliberation (Sunstein, 2006). For some information tasks, it has been shown that even small amounts of social influence may act to reduce the diversity of the group without improving its accuracy. At the same time subjective confidence may (mistakenly) increase with this convergence (Lorenz, Rauhut, Schweitzer and Helbing, 2011)

A pragmatic approach to counteracting the undesirable influence of credibility heuristics might

rely on a three-pronged approach based on user education, additional social/content cues and interface adaptation.

While we may rely on heuristics at early phases in learning, we soon gain more autonomy and are able to apply our own experience and that of trusted others to challenge and revise heuristic evidence ([Origi, 2008](#)). At the same time, heuristics are not all-pervasive, but maybe counteracted with some small mental effort ([De Cruz, Boudry, De Smedt and Blanke, 2011](#)). It therefore seems reasonable that a meta-knowledge of heuristics may reduce the cascade effect and guide the resource user to a closer interaction with the important semantic aspects of the content.

Additionally, we may supplement existing cues with those which seek to nudge users in the right direction. In a deliberative task, small interface changes can encourage users to consider more alternative hypotheses and gather more evidence before making decisions ([Jianu and Laidlaw, 2012](#)). Given the tendency for question-answering sites to suffer from temporal and first-mover bias, greater emphasis on flagging the merits of newer answers is needed, be they complementary or competing (and a mixture of both is to be encouraged).

Conclusion

We have seen that both source and receiver characteristics are crucial in the exchange of testimonial evidence. Given that most community sites are able to track user profiles, it should be possible to adapt the experience and the prominence of cues to the use case. Adapting the interface to the user's stake in the information (reducing heuristics for a higher level of involvement) has been suggested, for instance. ([Sundar et al., 2009](#)). More experimental work on this and on the various ways of presenting the user's interaction history is needed to assess the potential for an improved system.

Social knowledge exchange and the question-answer format present an immensely powerful opportunity for the evolution of truly inclusive, blended knowledge resources. We still need to understand the power and also the risks of community filtering and weighting and how it influences the knowledge seeker. The harnessing of heuristics and social cues provide a layer with scope for small adjustments that can have a large impact on how testimony is weighted and considered.

Acknowledgements

This research was supported by the Faculty of Environment and Technology at the University of the West of England. The author would also like to thank the participating students and the ISIC reviewers for their helpful comments.

About the author

Paul Matthews is a Senior Lecturer in Information Science at the University of the West of England, Bristol, UK. He can be contacted at: paul2.matthews@uwe.ac.uk

References

- Coady, D. (2012). *What to believe now: Applying epistemology to contemporary issues*. Chichester, UK; Malden, MA: Wiley-Blackwell.
- De Cruz, H., Boudry, M., De Smedt, J. & Blanke, S. (2011). Evolutionary approaches to

- epistemic justification. *Dialectica*, 65(4), 517-535.
- Goldman, A. I. (2010). Systems-oriented social epistemology. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology*, volume 3. (pp. 189-214). Oxford: Oxford University Press.
- Hilligoss, B. & Rieh, S. Y. (2008). Developing a unifying framework of credibility assessment: Construct, heuristics, and interaction in context. *Information Processing & Management*, 44 (4), 1467-1484.
- Jianu, R. & Laidlaw, D. (2012). *An evaluation of how small user interface changes can improve scientists' analytic strategies*. Paper presented at the Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Austin, Texas, USA. 2953-2962.
- Kahneman, D. (2011). *Thinking, fast and slow* (1st ed.). New York, NY: Farrar, Straus and Giroux.
- Kim, H. & Sundar, S. S. (2011). *Using interface cues in online health community boards to change impressions and encourage user contribution*. Paper presented at the Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Vancouver, BC, Canada. 599-608.
- Kim, S. & Oh, S. (2009). Users' relevance criteria for evaluating answers in a social Q&A site. *Journal of the American Society for Information Science and Technology* 60(4),
- Lackey, J. (2008). *Learning from words: testimony as a source of knowledge* Oxford University Press, USA.
- 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*,
- Lucassen, T. & Schraagen, J. M. (2011). Factual accuracy and trust in information: The role of expertise. *Journal of the American Society for Information Science and Technology*, 62(7), 1232-1242.
- Magnus, P. D. (2009). On trusting Wikipedia. *Episteme*, 6(1), 74-90.
- Mamykina, L., Manoim, B., Mittal, M., Hripcsak, G. & Hartmann, B. (2011). Design lessons from the fastest Q&A site in the west. Paper presented at the *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems*, Vancouver, BC, Canada. 2857-2866
- Matthews, P. (2014). Knowledge fixation and accretion: Longitudinal analysis of a social question-answering site. *Journal of Documentation*, 70(5)
- Matthews, P. & Simon, J. (2012). Evaluating and enriching online knowledge exchange: A socio-epistemological perspective. In A. A. Lazakidou, R. Sharda & S. Voay (Eds.), *Virtual communities, social networks and collaboration* (pp. 35-59) Springer New York.
- Metzger, M. J. (2007). Making sense of credibility on the web: Models for evaluating online information and recommendations for future research. *Journal of the American Society for Information Science and Technology*, 58(13), 2078-2091.
- Origi, G. (2008). Wine epistemology. the role of reputation and rating systems in the world of wine. In Smith, B. (ed.), *Questions of Taste: The Philosophy of Wine*. Oxford University Press. 236-253
- Origi, G. (2012). A social epistemology of reputation. *Social Epistemology*, 26(3-4), 399-418.
- Paul, S. A., Hong, L. & Chi, E. H. (2012). [Who is authoritative? understanding reputation mechanisms in quora](http://arxiv.org/abs/1204.3724v1). *Arxiv.Org*, Retrieved from <http://arxiv.org/abs/1204.3724v1>
- Rieh, S. Y. & Danielson, D. R. (2007). Credibility: A multidisciplinary framework. *Annual Review of Information Science and Technology*, 41(1), 307-364.
- Sundar, S. S., Knobloch-Westerwick, S. & Hastall, M. R. (2007). News cues: Information

scent and cognitive heuristics. *Journal of the American Society for Information Science and Technology*, 58(3), 366-378.

Sundar, S. S., Xu, Q. & Oeldorf-Hirsch, A. (2009). Authority vs. peer: How interface cues influence users. Paper presented at the CHI, Boston, MA, USA. 4231-4236.

Sunstein, C. R. (2006). *Infotopia: How many minds produce knowledge*. New York: Oxford.

Tausczik, Y. R. & Pennebaker, J. W. (2011). *Predicting the perceived quality of online mathematical contributions from user's reputations*. Paper presented at the Proceedings of CHI 2011, 1885-1888.

Treude, C., Barzilay, O. & Storey, M. -. (2011). *How do programmers ask and answer questions on the web?: NIER track*. Paper presented at the Software Engineering (ICSE), 2011 33rd International Conference On, 804-807.

Wathen, C. N. & Burkell, J. (2002). Believe it or not: Factors influencing credibility on the web. *Journal of the American Society for Information Science and Technology*, 53(2), 134-144.