Chapter 3
Evaluating and Enriching Online Knowledge Exchange: A Socio-epistemological Perspective

Paul Matthews [1] and Judith Simon [2]

[1] Department of Computer Science and Creative Technologies, Faculty of Environment and Technology, University of the West of England, Frenchay Campus, Bristol BS16 1QY, UK, paul2.matthews@uwe.ac.uk

[2] Department of Philosophy, University of Vienna, Austria & Institute for Technology Assessment and Systems Analysis, Karlsruhe Institute of Technology, Germany, judith.simon@univie.ac.at

Abstract We apply approaches from social epistemology and social psychology to the study of knowledge in online communities. We first provide a descriptive analysis of knowledge behaviours informed by social psychology and in particular Arie Kruglanski's lay epistemics. This is followed by normative considerations concerning epistemic, i.e. knowledge-related, practices based upon insights from the philosophical field of social epistemology. In particular, we refer to the works of Alvin Goldman, Helen Longino and Miranda Fricker. Outcomes of this dual analysis are norms and evaluation criteria as well as strategies to design online knowledge exchange through avoidance of bias, stimulation of diversity and warranted recognition of epistemic authority.

Keywords Knowledge • Social epistemology • Online communities • Diversity • Social psychology • Lay epistemics • Bias • Epistemic injustice • Values in design • Social Q&A • Knowledge exchange • QA

3.1 Introduction

The social web provides access to a wealth of minds willing and able to share knowledge, and when made available in online repositories, it forms a major source for learning amongst a wider public. Knowledge thus shared may be trivial, may contribute significantly to human welfare or may help us address the most pressing issues of our time. In all cases, we should acknowledge the new power of online community in social knowledge and seek to improve how it operates.

We believe that the rising philosophical field of social epistemology has much to offer as a foundation for the evaluation and design of online knowledge systems, providing a set of norms that may act both as evaluation criteria and design principles. Social epistemology is a branch of the philosophy of knowledge that foregrounds the

A.A. Lazakidou (ed.), *Virtual Communities, Social Networks and Collaboration*, Annals of Information Systems 15, DOI 10.1007/978-1-4614-3634-8_3,

social aspects of knowledge-related processes and often strives to demonstrate its applicability to real-world social systems. As a predominantly naturalist approach, it may find support and refinement in evidence from sociological, psychological or technology-related studies. Indeed, we argue that as it progresses from the abstract to the concrete, social epistemology needs to account for all of the situated and mediated reality of human interaction—particularly issues of trust, power and justice.

Through some well-known examples and a drawing on a range of interdisciplinary evidence, this chapter aims to understand social knowledge processes and demonstrate the value of social epistemology in guiding socio-technical design. In particular, we want to demonstrate how to fruitfully combine descriptive insights from social psychology with normative considerations from social epistemology and reflections on practical issues from the perspective of values in design (VID).

We therefore first offer some insights from social psychology concerning self-efficacy, sense of community and lay epistemics which are particularly important for knowledge sharing and acquisition. We then present the work of key thinkers in the field of social epistemology, focusing on epistemological and ethical guidelines for assessing and optimizing social knowledge practices. We then return to the specificities of knowledge in online communities and look at how these might benefit from the previous analyses. First, however, we would like to set the scene with some scoping, definitions and examples of knowledge-rich communities and how they have been shown to function.

3.2 Knowledge Online: Definitions and Examples

3.2.1 Definitions and Scope

Despite several problems and continuous discussion around the philosophical definition of knowledge as justified true belief, the broad definition nonetheless has its merits for apprehending epistemic practices in knowledge-related online communities: we desire members to acquire beliefs that are grounded in evidence, which are justified and ideally true.

In contrast to most approaches in epistemology which focus on individual knowers, social epistemology foregrounds different social aspects of knowledge creation, evaluation or sharing. Goldman [21] provides a useful taxonomy of topics in social epistemology. First, there are issues around social evidence, i.e. how individuals deal with evidence provided by other agents. Second, different issues come to the fore when we shift our focus to collective epistemic agents, such as groups or networks. Here, issues of judgement aggregation and decision-making loom large. Finally, there is the perspective of system-oriented social epistemology, where social or socio-technical systems ranging from national legal systems to specific web platforms such as Wikipedia may be assessed for their epistemological merits and limits.

We propose a socio-epistemological perspective for two reasons. First, we consider epistemology to deliver deeper insights that can enhance the understanding of the different phases of knowledge lifecycles as identified in knowledge management literature. Second, within epistemology, social epistemology appears most suited to apprehend online epistemic communities for its decided focus on different social aspects of knowledge.

Quite a number of knowledge process typologies were developed in knowledge management literature to describe phases in a (desirable) lifecycle. Several of these may be conceptually combined into the four key processes of creation, capture, transfer and utilization, with creation and exchange being key processes to assess in online community [50]. As we shall see, social epistemology and social psychology can offer insights for many of these areas, though we will largely focus on two facets of exchange: *sharing* and *acquisition*. Further qualitative typologies of knowledge are common, including the distinction between the tacit and the explicit [50]. Such distinction, however, may place too much emphasis on individuals and gives insufficient weight to where knowledge is situated in a community [57].

In terms of types of knowledge actors, we consider *lay* in addition to *expert* knowledge in our analysis, not least as the interplay between experts and novices is so central to the dynamics on online communities. Moreover, we will see that there are important similarities between lay and scientific knowledge, such that analytical considerations of the latter may be similarly applicable to the former.

We further define online communities as multi-user collectives where individuals interact (actively/ repeatedly) around a common purpose, governed by a set of implicit or explicit policies [30, 50] and take the term to encompass a range of models, ranging from closed, specialist communities to open, generalist and large-scale platforms.

While earlier conceptions of community perhaps placed more weight on long-term interaction and the development of personal relationships between community members [50], there is an increasing overlap between formal, purpose built community platforms and ad hoc interactions enabled through social networking and communication tools—in the sense that these tools may be used to delineate and structure temporary yet meaningful communities we include them in our consideration.

3.2.2 Illustrative Examples

To open our analysis, we will briefly introduce some better known, public web communities. We chose our examples because:

1. These sites have been quite extensively studied, and there is a foundation of evidence on how they work in practice.
2. Such larger sites have important structural similarities with many smaller, specialist or private communities, and some important aspects of the research are generalizable.

Table 3.1 Traffic summary for leading social knowledge sites

Site	Estimated daily users[a]	Global ranking[b]	Reach (% of global Internet users who visit)[b]	% Coming from search[b]
Wikipedia	25 M	7	15	50
eHow	2 M	128	1	45
Answers.com	3 M	166	0.7	47
Yahoo Answers	10 M	4 (all Yahoo)	24.8 to all Yahoo. Answers accounts for 5.6% of site traffic	–

Sources: [a]Google Trends; [b]Alexa Rankings, 25 July 2011

Table 3.1 shows visitor statistics and rankings for some of the most popular knowledge platforms on the web. What is clear from current trends is that the larger collaborative, social knowledge platforms are significant in terms of audience reach and in the number of users who reach them via web search. This is significant because these sites frequently featured in the top results returned for user queries—in other words, they serve an important epistemic (knowledge-related) function.

The way knowledge is constructed and represented in these platforms varies considerably, from the anonymously co-authored Wikipedia, through information aggregation sites (Answers.com) to collaborative question answering (Yahoo Answers). We will take Wikipedia and Yahoo Answers as examples for further exploration and draw out salient points of relevance in considering their role in knowledge exchange. Our areas of focus include scale, technical functions and features, and social characteristics of participating members.

Wikipedia currently has some 3.7 million content pages, averaging 19.39 edits per page and 142,000 active users (that part of over 15 million registered active within a 30-day period) [59]. While the wiki pages themselves represent the default public view of the content, much of the community work goes on in the "Discuss" pages, where edits are proposed and commented on—and there is some evidence that those pages with more discussions may be higher quality [43]. Over time, Wikipedia articles may arrive at a certain stability where they provide the balanced objectivity the site strives toward [58]—see also Fig. 3.1 for a visualisation of article evolution—though in contested cases, it may also be that one party wins the "edit war" [48]. In terms of quality, stable Wikipedia's articles have been blind tested to favourably compare to "traditional" encyclopaedias such as Britannica, though they perhaps contain more omissions than errors [39]. Those who point to bias in Wikipedia point to geographic bias—with more coverage of the global north—and contributor bias toward the technically literate with surplus time available to contribute [26]. Persistence of content on Wikipedia may be taken as a proxy for quality, and it has been noted that some 10% of editors contribute 86% of overall value in terms of word views [43]. Dubbing this minority of coordinating editors "cool farmers", Iba et al. show how they are responsible for most content edits and comments on the talk pages over time, though this group may be further divided into mediators and zealots—either seeking conciliation or the maintenance of a radical stance [27].

Yahoo Answers by contrast is a social question-answering site, where users may pose a question and other users contribute candidate answers. Other users may vote

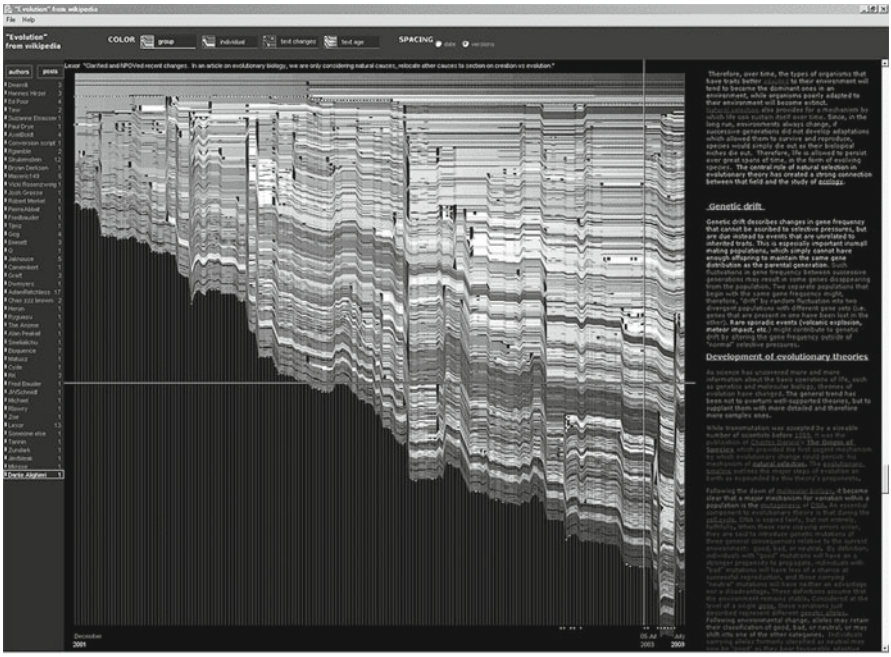

Fig. 3.1 Visualization of edits to a single Wikipedia page, showing evolution of article length (height of display) and increase in authors (each shown in a separate shade). Note persistence of contributions over time (Source: IBM History Flow Project, http://www.research.ibm.com/visual/projects/history_flow/index.htm)

the provided answers, and the original asker chooses a "best" answer amongst candidates. Content categories vary and include areas where specialized, more factual answers are common (e.g. mathematics) to broader general interest categories where answers are perhaps more a matter of opinion [1]. In exchange for supplying answers, users gain reputation points, providing several game-like interface aspects to the system. In 2009, the site was reported to have over 50 million resolved questions on its English site, with a further 5.3 million at the voting or open stage [63]. Studies of user typology on the site have revealed distinct types of answerer: those who respond from personal knowledge—"specialists"—and those who use other sources to construct answers, "synthesists". Notably, the latter tend to accumulate more reward points [16]. In terms of activity, users tend to be both contributors and consumers, and this may be taken as a sign of health of the community—as those communities with a dearth of answerers may be unsuccessful [52]. Characteristics of proffered answers may indicate likelihood of acceptance by the asker, particularly answer length and the previous number of accepted answers attributed to the user [1]. Figure 3.2 shows an example question and accepted answer pair on Yahoo Answers, with the accepted answer comment revealing the credibility cues used by the asker. In a detailed study of these acceptance criteria, Kim and Oh showed that

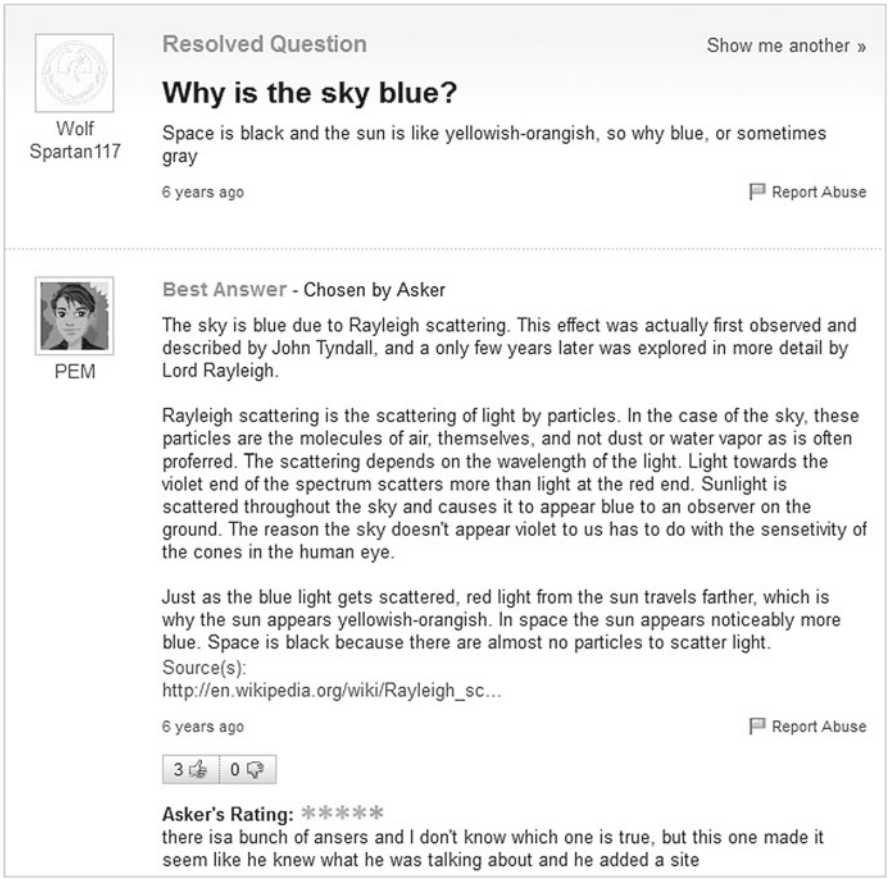

Fig. 3.2 Sample Yahoo answers question and accepted answer. Note the reasons for acceptance provided by the asker

content completeness, answer feasibility and attitude agreement/confirmation were important criteria for the asker, though emotional support, attitude and novelty of perspective or content were also significant [28].

3.3 Understanding Knowledge Exchange: Insights from Social Psychology and HCI

3.3.1 Why We Share: Motivational Aspects

A number of studies have investigated factors influencing participants' willingness to share knowledge in different types of online community. These have drawn on social psychology in addition to influential theories in sociology, communication

and HCI. In her review of recent studies, Kosonen grouped facilitating factors into individual motivations, personal characteristics, technical attributes and social capital. The majority of studies in the review used questionnaires informed by "established" constructs followed by statistical modelling, while perhaps fewer have looked at participant behaviour and interaction patterns [30].

In these studies, key individual motivators include the enjoyment of helping, knowledge-sharing self-efficacy, recognition and status. Personal characteristics include comfort with online communication, sportsmanship, "hobbyist" tendencies and being a "helpful" personality. The more social influencers have been found to be interaction ties, sense of community/structural embeddedness, trust and the norm of reciprocity. Technical aspects have much in common with information systems (IS) acceptance models and have shown system quality, usability, learnability and site management to be significant [30].

Other similar studies not covered in the Kosonen review highlight sociability and usability [44] and have underlined trust (above reciprocity) and self-efficacy [9]. A key idea of Bandura's self-efficacy is that experience of mastery heightens personal efficacy, which suggests a virtuous circle associated with successful sharing, which may be reinforced with feedback and attention [62]. A form of trust termed "swift trust" may be applicable in online communities, where suspension of doubt and risk taking guides behaviour in lieu of a history of trust-building interaction [50]. Ma takes the sense of community construct further and stresses "need to belong" as an overriding factor [38]. Blanchard argues that sense of community may be elaborated to sense of virtual community (SOVC) and found positive correlation between this and trust amongst members of the website babycenter.com, a forum for new parents [7].

Aside from social aspects, Schmitz-Justin introduced a "content quality" measure and found the perception of quality in existing content to be strongly related to knowledge contribution [50]. This study also noted a relatively strong effect of professional affiliation motivation, over and above social affiliation. While this may result from the community used in the study (a higher education forum), it perhaps captures an important subjective distinction we may try to make between our social and professional personas.

3.3.2 How We Acquire Knowledge: Insights from Lay Epistemics

The integrative work of Arie Kruglanski in social psychology marshals a range of evidence on human belief formation under different conditions, under the banner of "lay epistemics". While not exclusive to an online setting, the work provides an empirically sound platform for understanding the dynamics of knowledge acquisition online as well as identifying frailties in the way we seek and use social evidence in knowledge processes. Kruglanski introduces the interlocking paradigms of need for cognitive closure (NFCC), the unimodel theory of social judgement and the concept of epistemic authority [32]. NFCC may be specific or non-specific—desirous of a particular epistemic outcome or not—but the results of several studies

show that the higher the NFCC, the greater the tendency of human subjects to fixate on a conclusion and its associated beliefs earlier and to consult less detailed evidence in the process [31]. This may be further marked when there is increased accountability associated with the conclusion (e.g. where people are told they will have to have to justify the knowledge gained to a third party afterwards). High NFCC may be linked to more impersonal interaction styles, resistance to persuasion and intolerance of diversity [32]. NFCC also absorbs other cognitive theoretical movements such as dissonance and self-efficacy when applied to knowledge-seeking behaviour—we are motivated to seek greater predictability and avoid dissonance by confirming (or changing) our opinion, and we may only persevere in seeking when we have belief in our own effectiveness to do so [60].

The unimodel theory applies to the mechanism underlying hypothesis formation and validation. Here, Kruglanski posits an underlying syllogistic, if-then approach to the testing and assimilation of evidence, much of which comes from social sources, and it is here that an implicit hierarchy of epistemic authority applies, formed in childhood and developing in sophistication with age.

The foundational importance and interplay of the three components of lay epistemics become important in our consideration of knowledge exchange online. They reveal motivational aspects to knowledge seeking and precursors to confirmation bias and the avoidance of uncertainty, human tendencies that may be enhanced or mitigated through social and technical biases online. Notably, Kruglanski's conception of everyday belief formation is that it is fundamentally similar to the scientific method, in broad agreement with Popper—science as "common sense writ large" [32].

Studies of information-seeking behaviour tend to corroborate Kruglanski's claims, with NFCC doubtless being a key element in both formal and informal knowledge interactions. For instance, the speed, availability and convenience of online sources have been found to be important both to novice and expert seekers, to the extent of not considering options beyond those immediately available [11]. NFCC behaviours can also be linked to information satisficing, where a seeker will decide that "enough" evidence has been collected on a topic of research, in the absence of a full and rational analysis of options [46].

A note of caution is needed; however, while Kruglanski's approach provides broad brush explanations of socio-cognitive phenomena in knowledge acquisition, we should not neglect individual differences, both in experience level and cognitive style that may lead to significant differences in how individuals approach a knowledge task [12].

3.3.3 The Role of Norms: Evolution and Enaction

Community norms, overlaying and informing the development of the technical platform, provide the environment for knowledge exchange and are a crucial component to consider, as—overtly or otherwise—they represent the embodiment of the

community's underlying philosophy. Norms may be captured in explicit policies, enacted in the language used and the interactions between members or applied through moderation of content. In a relatively early study of news blog communities, for instance, Lackaff illustrates how community norms may be enacted through distributed or "heterarchical" moderation schemes, where power to edit, rate and sanction contributions is devolved to community members [35]. This has been in part a response to the onerous task of maintaining "top-down" moderation and also is illustrative of the value of granting autonomy to the communities' members. In this way, community norms emerge as a combination of the rules set by the community's creators and the way these rules are enforced by users. Importantly, applying qualitative categories to posts (e.g. "useful", "informative", "funny") is necessarily subjective, but group consensus through common ground and archetyping may emerge. Lackaff shows how different communities continually adjust moderation rules and influence changes to technical features.

Other evidence shows that core community members ("Wikipedians" in the Wikipedia community) are more likely to explicitly reference community policy in justifying modification of content and therefore act as "champions" for the community's values [43].

3.4 Assessing and Improving Knowledge Exchange: Insights from Social Epistemology

Social epistemology is the philosophical discipline exploring the ways and the extent to which knowledge and epistemic practices are social. Given the goals of this chapter, it appears to be a suitable theoretical framework to apprehend knowledge practices in which multiple agents are involved.

Social epistemology as a term is mostly used to refer to a field of discourse which emerged in the 1980s in a predominantly analytic tradition, although predecessors exist within and beyond philosophy [34].

In the following, we will briefly outline several approaches within social epistemology which appear suited to shed some light on different socio-epistemic processes in online communities.

First, we present Alvin Goldman's veritistic social epistemology and argue that his approach is well suited to understand less complex testimony-related processes of knowledge exchange, whereas it appears less suited to an understanding of knowledge creation.

We then turn to Helen Longino's critical contextual empiricism, which appears better suited to apprehend communities in which new knowledge gets created. Finally, we turn Miranda Fricker's notion of epistemic injustice to emphasize a crucial problem in designing systems that serve epistemic purpose: how to avoid bias and counter epistemic injustices. Two caveats appear necessary. First of all, it is obvious that given the scope of this chapter, we can only provide a much reduced portrayal of the approaches. Second, some may consider our approach "using" different

socio-epistemological approaches to shed light on difference aspects of online communities too eclectic, given that the approaches portrayed here differ profoundly in their philosophical premises. Given the scope and goals of this chapter, however, we rather wish to open up the debate and show the utility of social epistemology for online systems design and must refer to other publications for more detailed accounts of the differences between different approaches. The interested reader may confer Goldman and Whitcomb [22] for a collection of influential papers in analytical social epistemology. Overviews over the field are provided by Kusch [34] and Goldman [19]. For specific approaches within social epistemology, confer for instance Kusch [33], Longino [37] and Goldman [18].

3.4.1 Goldman's Normative Social Epistemology and Knowledge Exchange

Alvin Goldman is one of the most prominent social epistemologists, and two aspects make his work particularly relevant for this article. First of all, Goldman stressed the potential of social epistemology to provide guidance for systems design [21], and second, he explicitly addresses the role of information and communication technologies [18, 20]. However, as will also become obvious, ICT in his view are rather means of testimonial knowledge transfer conceptualized in analogy to simple sender-receiver models of information theory [53]. In the following, we will briefly outline central aspects of his veritistic social epistemology and show some of its pros and cons for systems design.

Epistemology quite generally for Goldman is "[…] a discipline that evaluates practices along truth-linked (veritistic) dimensions" and "[s]ocial epistemology evaluates specifically social practices along these dimensions" [18], and his veritistic social epistemology is specifically targeted at supporting truth and countering error and ignorance. The central concept is that of veritistic value.

Goldman argues that in everyday life as well as in science, certain value is placed on having true rather than false beliefs, and this value is what he calls veritistic value. For his assessment of epistemic social practices, it is essential that these (a) can have different veritistic outcomes, such as knowledge (positive veritistic value), error (negative veritistic value) or ignorance, and (b) that they can be evaluated according to these outcomes. Accordingly, an epistemic practice has a higher veritistic value if it yields more knowledge than another epistemic practice. While epistemic states, such as knowledge, error and ignorance, have fundamental veritistic values, epistemic practices, by contrast, have instrumental veritistic values for achieving these states [18]. Through different epistemic practices, the overall veritistic values in a society, i.e. the relationship between truth, error and ignorance, can be changed. And these changes in turn can then be used to assess the value of specific epistemic practices. Hence, the goal of epistemology in general is to evaluate practices along truth-linked dimension, and the goal of social epistemology consequently is to identify those social practices that have a comparably favourable effect on knowledge as opposed to error and ignorance.

Goldman's framework for the evaluation of social knowledge practices provides some important criteria and demonstrates his desire to progress from an egocentric to a workable social conception of knowledge. Specifically, Goldman acknowledges:

1. The core of the evaluative model, which should analyse the mechanics of the knowledge exchange practice itself: the inference practice of the "credal agent" (the believer), the speech practices of the speaker and the communication-control practice. This model is further developed in his analysis of system-oriented SE, where non-human agents may be considered part of this practice.
2. A role for personal interest as well as authority: "sensitivity to relative amounts of interest should play a modest role in assessing … epistemic credentials" [18].
3. "Need to know": Not everybody in a community needs to know everything; instead, knowledge should be available to those who need it most.
4. Accommodation of different types of question. For questions where no agreed "true" answer exists, it might be treated in a Popperian manner—you cannot know what is true, but you can reject specific hypotheses. For this to work, however, a full range of hypotheses need to be represented. As we have seen, this is in agreement with lay epistemics, which proposes an underlying, partly subconscious, logical mechanism to reasoning and belief formation.

3.4.1.1 Testimony and ICT

Testimony in Goldman's account refers to the process of communicating knowledge from one person to another. Stating that testimony is the most elementary and universal path to knowledge, Goldman focuses on the different roles and options for action that the provider and the recipient of testimony have. He conceptualizes testimony as a four-stage process. After the initial discovery of knowledge (stage 1), someone decides to transmit a message containing this knowledge to others (stage 2). This message has then to be received (stage 3) and accepted (stage 4) by others. Only if all four stages are mastered successfully, knowledge has been transmitted from one person to another. This model strongly resembles Shannon and Weaver's mathematical theory of information [53]. And in line with such a model of testimony as signal transfer, Goldman considers testimony to be responsible solely for the spread of knowledge which has been created before. The crucial question from a veritistic point of view then is under what conditions the report of testimony will produce the largest possible change of veritistic value [18].

This veritistic perspective on knowledge transfer from one person to another also informs Goldman's analyses of ICT with respect to their potential to increase the overall veritistic value in a society or community. This can be done in two ways: either more people get to know something or something new gets known, i.e. new knowledge gets discovered.

The epistemological task for ICT usage and design then consists in finding or designing systems so that they can help in spreading knowledge fast to many people. Besides the speed and breadth of spreading knowledge, it might be of additional

epistemic value that knowledge can be easily found and differentiated from non-knowledge. This aspect of quality assessment is increasingly relevant today where it is rather information overload than information scarcity which is problematic. Goldman analyses different tools, such as email, forums and the web, more generally with respect to their contribution to veritistic gain and shows how those media can either be used to raise the overall veritistic value in a society by distributing knowledge or to lower it, for example, by spreading error [18].

In a more recent paper, Goldman addresses role of blogging for democracy asking whether the web "[…] is better or worse in epistemic terms than the conventional media, in terms of public political knowledge generated by the respective communication structures" [20]. Emphasizing that conventional media often have the positive epistemic function of filtering information Goldman concludes, the veritistic impact of the blogosphere will depend on the users' motivations. Some users may use the abundance of unfiltered information to pick out only the information that corroborates their views, which are possibly biased or even utterly wrong. Others, however, might use the existing diversity to cross-check information from a variety of different sources in search for the truth in between.

Despite our general acknowledgment of Goldman's early identification of the relevance of ICT for social epistemology (and vice versa), his analyses of ICT suffer from a rather narrow perspective on ICT as being solely means of knowledge transmission. However, the web has enabled an abundance of socio-technical epistemic practices which are generative of new knowledge and not only a means of "signal transfer" under the conditions of noise. Furthermore, in neglecting issues of power, social status and trust, Goldman's social epistemology represents an "abstract social conception" [14]. There is a strong argument for philosophers to include such concerns to fully situate knowledge creation and exchange. To account for those practices, a more comprehensive and broader understanding of the socio-epistemic functions of ICT is needed, and Helen Longino's approach may be especially suited to provide guidance for the design of systems which enable knowledge creation.

3.4.2 Longino and Knowledge Creation in Communities

Helen Longino's Critical Contextual Empiricism as portrayed in her book *The fate of knowledge* also develops a social account of knowledge processes, which specifically aims at dissolving a dichotomous understanding of the social and the rational. While Longino does not explicitly address the epistemic functions of information and communication technologies, we nonetheless consider her approach particularly useful for the analysis and the design of ICT, especially for those tools that aim at knowledge creation rather than transfer of existing knowledge.

Longino acknowledges that the social dimensions of knowledge practices allow for biases and distortions in science and scientific knowledge [37]. However, sociality is not only a biasing factor—it can also deliver the means to counter such

biases. The role of a social epistemologist therefore is to show "[…] how the social dimensions of cognition have resources for the correction of those epistemically undermining possibilities" [37]. And these resources for correction lie within a combination of epistemic diversity with shared standards of evaluation.

To make sure that "[t]he social is not a corrupting but a validating element in knowledge" [36], Longino proposes four social norms for social knowledge:

1. *Venues*. There must be publicly recognized forums for the criticism of evidence, of methods, and of assumptions and reasoning. […]
2. *Uptake*. There must be uptake of criticism. The community must not merely tolerate dissent, but its beliefs and theories must change over time in response to the critical discourse taking place within it. […] Uptake is what makes criticism part of a constructive and justificatory practice. […]
3. *Public standards*. […] Participants in a dialogue must share some referring terms, some principles of inference, and some values or aims to be served by the shared activity of discursive interaction. […] A community's standards are themselves subordinated to its overall cognitive aims […] Finally, standards are not a static set but may themselves be criticized and transformed […] There is no particular act of adopting or establishing standards. […]
4. *Tempered equality*. Finally, communities must be characterized by equality of intellectual authority. [36]

These four norms are more than just ethical or political nice-to-have's. If knowledge involves justification and if knowledge as an epistemological status has to be attributed to a claim by a group of people, then knowledge fundamentally rests on communities and the social mechanisms and procedures they employ. As Longino states, "[w]hile intersubjective interaction is a necessary feature of scientific cognition, not just any form of interaction will do" [37]. Rather, interaction has to be effective to yield knowledge. And the four norms exemplify those criteria that have to be met for effective criticism, for checks and balances to take place. It is only through such effective criticism that the subjective can be transformed into the objective, "[…] not by canonizing one subjectivity over others, but by assuring that what is ratified as knowledge has survived criticism from multiple points of view" [37]. Thus, effective criticism is a prerequisite of knowledge and not just an addendum.

Longino considers her four norms as recommendations for the design and constitution of scientific communities. An ideal scientific community would fulfil all these four standards to a maximum. In such a community, all hypotheses and theories would be subject to the broadest range of valid criticism by fostering epistemic diversity, by encouraging dissenting voices instead of silencing them, and by listening and responding to criticism in ways that will further increase knowledge. The effect of such transformative criticism will be that the "[…] theories and hypotheses accepted in the community will not incorporate the idiosyncratic biases (heuristic or social) of an individual or subgroup" [37]. Given that these norms represent an ideal epistemic community, real epistemic communities will hardly ever fulfil all those norms completely; the quality of epistemic communities rather comes in degrees [37].

However, the fact that ideals might never be attained does not diminish their value as goals to strive for. After all, the possibility that ultimate truth may not be attained has never hindered people to strive for it.

The utility of Longino's norms for the design of online communities should be evident. If those norms can guide and amend *scientific* knowledge creation by giving concrete recommendations of how to design interaction in scientific communities, then the same norms may well be used to design and implement other socio-technical systems that serve epistemic purposes such as online epistemic communities.

3.4.3 Fricker's Epistemic Justice and the Problem of Bias

If one acknowledges that knowledge creation is a social process, one also has to deal with the possibility of social inequalities and epistemic injustices within such processes. Therefore, a normative social epistemology has to provide guidance as to how to avoid or counter such biases. Longino's normative standards for scientific communities provide guidance in how systems can be designed to foster the positive aspects of epistemic sociality while reducing the negative side effects such as the potentiality of biases and discrimination based upon social cues.

Another philosopher who has emphasized the political nature of epistemic practices is Miranda Fricker. In her book *Epistemic Injustice*, she links epistemology to ethics in order to shed a light on issues of power and prejudice in knowing. In particular, Fricker distinguishes between testimonial and hermeneutic injustices as two instances in which someone is wronged in his or her capacity as a knower based on his or her social position. According to her, "testimonial injustice occurs when prejudice causes a hearer to give a deflated level of credibility to a speaker's word", whereas hermeneutic injustice "[…] occurs at a prior stage, when a gap in collective interpretative resources puts someone at an unfair disadvantage when it comes to making sense of their social experience" [13].

Clearly, both forms of injustice are not only taking place frequently in our offline world, but also online. Quite generally, epistemic injustice occurs when people use invalid social indicators to make judgments about the epistemic merits of someone. To avoid such epistemic injustices, reputational cues and their epistemic usage have to be critically reflected upon and be kept open for constant scrutiny and revision.

Based upon these insights, we consider the critical analysis of trust and reputation metrics and of ranking and rating mechanisms in epistemic communities to be a particularly important topic in which social epistemology and feminist social epistemology in particular has a lot to offer for community design. In short, ratings and ranking as well as more sophisticated metrics by definition create biases because they support certain agents or their content in comparison to others. However, whether or not the specific ways in which those metrics are designed is epistemically beneficial and ethically acceptable should be analysed from case to case and by taking a look into the nitty-gritty of their algorithmic foundations (cf. [55] for a socio-epistemological analysis of trust metrics).

3.5 From Theory and Evidence to Design

If we were to design systems for knowledge creation and exchange that aim at learning from social psychology and at fulfilling the normative goals of social epistemology, how should we go about it? In other words, what are the characteristics of systems that support transformative criticism and which interface features might be useful? Of course, one should not assume that one can design socio-technical systems like machines that function according to plan. According to a more performative view, socio-technical systems evolve over time and change with practice. However, by designing systems just as much as by designing technical artefacts, one sets parameters that imply certain limits and affordances that make some forms of usage more likely than others.

Despite many euphoric claims regarding the democratic nature of the web and its potential to enable everyone to contribute and participate in public discourse or in knowledge creation, the web and the design of many applications make "differences that matter" [6]. First of all, there still exists a digital divide between those possessing the resources and skills needed to participate in the web, which needs to be addressed if social inequalities on the web should be reduced.

A second risk makes this vision of the web as a neutral place where everybody has the same rights even less realistic. Many tools on the web give different weights to different agents and their contributions; they make use of ratings and rankings or of implicit forms of endorsement (downloads, page views, etc.) to generate differences between different entities. Be it Slashdot's karma model, the hierarchies introduced in Wikipedia or even the weights of links in PageRank, the vast majority of tools on the web do not treat agents (users, institutions, websites) equally. Such weighting processes indeed have their epistemic merits because they help us to navigate in a sea of information. In some communities, for instance, moderation bias to automatically hide posts that are rated down may benefit the whole community in terms of information overload and removing noise from the system [35].

However, despite such benefits, such metrics also come with different problems, especially if they lack transparency. First of all, if the underlying algorithms which decide about what is ranked up or sorted out are not accessible or understandable to users of such systems, they lack epistemic responsibility, i.e. they have to blindly trust the epistemic criteria by which the systems rank entities [41, 55]. Think of the power of Google in structuring our knowledge universe.

Moreover and here again, Google's PageRank may be instructive; the web often reproduces social inequalities and has been said to support a Matthew effect in which the rich get richer, i.e. those with many hits and links get more hits and links. Such behaviour entails epistemic as well as ethical risks, especially if the link between ascribed epistemic authority and actual epistemic authority is weak or not even assessable.

As noted before, it has been primarily feminist epistemologists who have stressed the risk of epistemic injustices and biases that can occur when invalid reputational proxies are used to assess knowledge claims without critical reflection (cf. for

instance [4, 13, 49]). Feminist theoreticians have also convincingly argued that such biases and the unfair attribution of intellectual authority are not only ethical but also cognitive and epistemic failures, and have therefore requested a tempered equality of intellectual authority [37]. The crucial task in addition to raising awareness about these problems is to develop tools and strategies to counter different types of epistemic injustices. A common idea behind many feminist approaches to counter biases in socio-technical epistemic systems is to raise epistemic diversity (e.g. [4, 36, 37]). Crucially, this request for diversity is not an ethical nice-to-have, but rather an epistemic and cognitive necessity, which has not only been requested by feminist epistemologists but has also been demonstrated by computer simulations [42].

Taken together, our understanding of sharing and acquisition through the socio-epistemological lens would indicate that in design, special attention should be given to:

- Avoidance of detrimental bias, both biases crystallized in the system itself and biases inherent in the epistemic behaviour of users.
- Stimulating a diversity of members, voices and opinions, and actively recognising and incorporating dissent.
- Acknowledging authority only where warranted; criteria and mechanisms of reputation and authority attribution need to be made transparent and be bound to continuous evaluation, especially in systems where reputation attribution is automated.
- Underlining best practice along proven dimensions of usability, content quality, the empowerment of sharing and the sense of community.
- The development and support of community norms adhering to these principles.

3.5.1 Avoiding Biases and Epistemic Injustices

3.5.1.1 Pre-existent, Technical and Emergent Bias

Freedom from bias is a crucial goal for the development of any socio-technical system and much of the work in the field of science and technology studies (STS), and especially in feminist STS has focused on detecting and remedying bias in such systems. In their article on "Bias in Computer Systems", Friedman and Nissenbaum offer a taxonomy of biases that may be useful for understanding bias in online systems [15].

The two authors use bias "[…] to refer to computer systems that systematically and unfairly discriminate against certain individuals or groups of individuals in favour of others" [15]. Two notions are crucial: for bias to occur, it is not sufficient that people are treated unfairly; this treatment has to be systematic to be called bias. Moreover, the systematic distortion alone is also not considered to be bias, the distortion must also lead to unfair results. Friedman and Nissenbaum then identify three different categories of bias of relevance for computer systems, which we consider to be relevant also for our analyses: pre-existing bias, technical bias and emergent bias.

Pre-existing bias refers to "bias [which] has its roots in social institutions, practices, and attitudes" [15]. This is the type of bias that has stirred most analyses in critical science and technology studies. This form of bias is a classic example of all those societal injustices or personal prejudices that get inscribed into technology, be it intentionally or unintentionally.

Technical bias however is something different. This type of bias is not rooted in societal values, but rather arises within the process of technology design, when designers make technical decisions in certain ways and not in others, when they opt for one algorithm as opposed to another. The sources of technical bias that Friedman and Nissenbaum list are limitations of computer tools, decontextualized algorithms, methods of randomization and the biases that occur; then, human concepts have to be formalized to match the formats needed for computing.

Finally, the notion of emergent bias accounts for the fact that biases might occur later on through usage and appropriation of computer systems. Typically, such bias occurs when either the society in which the system is used changes, a process which Friedman and Nissenbaum describe as "new societal knowledge". The second reason for emergent bias has its roots in a mismatch between users and system design with respect to different expertise or values. This aspect has been empirically demonstrated by studies of Madeleine Akrich's in which technologies have been used in contexts other than the ones where they have been developed [3]. A further type of emergent bias is temporal in nature: in social question answering, for example, a delay in responding to questions may be detrimental where a ranking process is used, as rapidly provided yet relatively poor answers may rank more highly than higher quality answers provided later. This can also work at member level—long-standing community members may be conferred authority (e.g. high reputation)—though the actual quality of their contributed content does not warrant this [35].

Based on this taxonomy and the case studies they analyse, Friedman and Nissenbaum draw several conclusions concerning the necessity to avoid bias in computer systems and define tasks for designers involved in the creation of computer systems. First of all, to avoid pre-existing bias, designers have to be aware of different pre-existing biases and be sensitive to what they possibly inscribe into their systems. More precisely, they propose rapid prototyping and the inclusion of different user groups into the design process, formative evaluation and field testing as tasks that should minimize this type of bias. For technical and even more so for emergent bias, designers moreover have to envision the consequences and possible appropriations of their systems in different contexts. In order to achieve this, the authors propose the anticipation of probable contexts of usage, the articulation of constraints and the possibility to take corrective action once emergent bias becomes obvious.

3.5.1.2 Cognitive Bias

Very different types of bias are cognitive biases, such as confirmation bias and belief perseverance, which have been demonstrated in social and cognitive psychology

and are central to lay epistemic theory. While some of these biases may be harder to tackle than others, a first step is to make users more aware of such human tendencies to selectively use or ignore evidence. And this is a task that can be supported in online systems.

There is some reason to believe that heightened awareness of cognitive bias (introduced through training, induction or feedback) may serve to reduce it [8]. For the belief perseverance bias—where people may cling on to a belief in the face of overwhelming contradictory evidence—a further tack where there is some supporting evidence from cognitive science is to have community members rehearse alternative explanations. In an experimental setting, subjects asked to produce counter-explanations of observed data were more likely to depart from their initial belief [5].

The risks engendered by high need for cognitive closure (Kruglanski's NFCC) could be countered by interface features that make the key arguments quickly graspable but which also make it clear when "it's complicated", that time is needed to understand the debates around a topic. Recognition of complexity is arguably itself a form of interim, nonspecific closure. Rewarding or highlighting admissions of vacillation and evidence of sophisticated personal epistemology (awareness of the nature of knowledge) may be another tack, to some extent already practiced, that may serve to make an example of unbiased knowledge-sharing behaviour.

It might be argued that the more biased belief formation processes are essentially non-epistemic, in that they are too reliant on biases such as high NFCC, desire to conform or over-credulity. They do not involve the weighing of evidence or deeper investigation into a topic. In advocating the latter as a higher value approach, we need to stress the long-term pay-offs [45]. Poulter asserts a "delayed value of truth" principle:

"The benefits of valuing truth are (usually) long-term, higher-order benefits whereas the benefits of a bias to one particular opinion are (usually) short-term, lower order benefits" [45].

Again, this principle might be enacted through induction, feedback and community norms.

3.5.2 Encouraging Diversity

3.5.2.1 Diversity of Members

A major concern with the rise of social networking has been the tendency toward homophily that similar individuals gravitate toward each other, forming a relatively cliquey subnetwork that not only shares physical attributes—age, sex and ethnicity—but also shares opinions and attitudes. Network analyses have indeed shown that this can be the case for physical attributes [56], though analyses of shared opinion have found mixed results [17], indicating that inferred shared opinion is greater than actual shared opinion between online friends. Homophily nevertheless presents a major

challenge to a social epistemological ideal of diversity and may be an argument against the democratizing power of social technology [20].

The issue of homophily is tightly bound up with the design approach to identity and reputation. While taking on features from social networking tools—chiefly user profiles—may be a largely positive step in fostering an environment for knowledge sharing, it may also threaten the platform neutrality that enables dissenting voices to speak against the majority view.

A two-pronged approach of empowerment and recognition may be the key to fostering membership diversity. Examples such as Global Voices Online[1] show how contributors can be encouraged through local training initiatives, and their work may then be aggregated and amplified to become essential, widely read commentary on contemporary issues. A similar approach could work at smaller scale by targeting underrepresented groups and providing the tools and positive feedback needed to encourage sharing. Communities do have inherent scope for geographical variation that should be embraced. In specialist communities of practice, for instance, knowledge sharing is valued for enabling access to examples of practice in other locations [24].

3.5.2.2 Diversity of Views

Goldman's view on more qualitative or contested types of knowledge is that while one true answer is not achievable, we can do well by exploring and systematically ruling out possibilities [18]. The "insightful" tag in Slashdot came to serve this purpose quite well [35]. A more recent approach is shown in Opinion Space,[2] which highlights insightful comments that might be some distance from the user's own views. This needs to be weighted, however, in such a way to also indicate where consensus exists, thereby avoiding the "balance fallacy" where opposing though largely discredited views are given undue weight.

The example of Wikipedia is important in showing how the evolved idea is a balanced narrative, showing the range of opinion on both specialist and generalist topics. While this appears as a seamless whole, we have noted how the real range of views is made clear on the "discuss" pages [39]. Quora, a more recent social question-answering site,[3] combines a wiki-like topic "summary" above the individual posts and as such is perhaps more effective at making the connection between the diversity of opinion and the emergent discourse.

Where group discussion toward consensus building or decision-making is needed, genuine diversity of views militates against confirmation bias in new information seeking and assimilation [51]. In an unstructured communication environment, while

[1] http://globalvoicesonline.org/

[2] http://opinion.berkeley.edu/landing/

[3] http://www.quora.com

connections between individuals within a particular pole of a debate—homophily—may be stronger than between poles, the latter exchange can exist [64]. Importantly, metrics and increasingly network visualization can enable moderators to observe this process, forge new links between poles and recognize bridge builders in the community. Moreover, there recently has been a lot of research on diversity-based search and recommendation metrics, which also aims at providing information that not only conforms to the criteria of relevance and quality but adds the criterion of diversity of results [2, 23, 40].

3.5.3 Recognizing Authority

Online interaction provides additional barriers to epistemic authority assessment. The "calibration" of authority [29] may be more difficult as interaction may be fleeting. In the case of Wikipedia, we do not even know who has written particular passages or what their reputation is [61]. Instead, we may come to recognize the authority of a system as a whole that is "reliable enough" [54].

Authority in the online sphere is an area where different models are still being developed and tested, and as we have seen, where there may be bias and opportunities for gaming a system based on metrics which are not robust to system attacks. We can identify two possible errors in the attribution of epistemic authority:

1. Unjustly attributing authority where it is not warranted (e.g. Matthew effect), which has negative epistemic consequences
2. Unjustly denying epistemic authority, having both epistemological and ethical consequences (sensu Fricker)

We have noted how reputation metrics systems require special care since they are highly influential, often opaque, frequently bound to systematic bias and vulnerable to manipulation. While content cues alone can provide a significant amount of context regarding authority, supplementary cues may also be helpful. One innovation by the Quora question-answering platform is to allow different "topic biographies" for a single user (are they employed in the area, is it simply a hobby, etc.), thereby adding some evidence of a user's qualification to respond on a topic. When accompanied by links to other web presences, such biographies can provide a degree of verifiability (Fig. 3.3).

A specific problem occurs at the intersection of reputation and anonymity. While some users will appreciate the ability to disclose personal characteristics through their profile, anonymity (or at least minimal pseudonymity) should remain an option for others in order not to create (real or imagined) barriers to entry. Yet, irrespective of their qualities, those anonymous users may be down-ranked in systems employing reputation metrics based on the identifiability of contributors, which may be epistemically detrimental and inserts a systematic bias against anonymous users into the system.

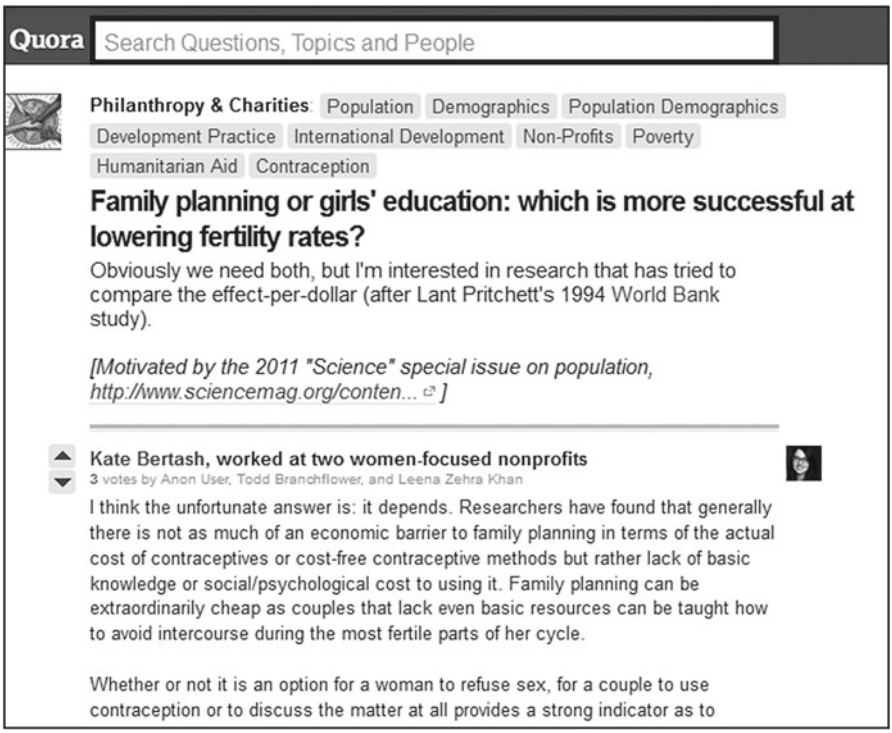

Fig. 3.3 The Quora interface showing authority cues for the highest rated answer

From a socio-epistemological perspective, several tasks for systems designers are considered crucial. First, it will be of utmost importance to design tools which not only help to assess and track authority and reputation of information sources and their providers, but which also make the criteria for authority transparent and assessable to users of varying competencies and expertise [55]. Second, the question of how to best combine traditional reputational cues (e.g. education-linked indicators) with new system-based cues (such as ratings) still needs to be fully addressed. Finally, systems need to be designed in ways to minimize the alpha and beta error of epistemic injustice: withholding epistemic authority where it would have been appropriate and attributions epistemic authority where it is not warranted.

3.5.4 Underlining Best Practice and Steering Norm Formation

Many of the principles and practices for developing a conducive environment for knowledge sharing are already well understood by community designers and managers. The motivational aspects we outline in Sect. 3.1 such as sense of community,

self-efficacy and other stimuli may be encouraged through the interaction between moderators and new or experienced members. Self-efficacy, for instance, may be encouraged through welcoming new members and providing reinforcement to initial interactions with the aim of encouraging system mastery. This type of activity may be devolved with time to experienced members [47].

It is well recognized that the ideal type of community governance is a balance between top-down moderation and the fostering of bottom-up empowerment and participation in the development and enaction of community policy, leading to norms that are accepted and shared between participants. Moderators are particularly important in providing an early exemplar, "performing the public drama of the foundation myth of the community" [47]. It would therefore fall not only to systems designers, but also to moderators and founding members to embody the socio-epistemological principles we have introduced.

While there is consensus that user censorship should be light touch and that guidelines may be preferable to hard and fast rules [25], it is also recognized that for the maintenance of a fair and civil environment, a degree of policing is needed. That said, the interpretation and enforcement of guidelines and policy by community managers is often improvised and constantly reinterpreted [10].

For larger communities, content "flagging" is becoming increasingly important—users are given the ability to draw moderator's attention to content that does not conform to community policy. These data can be aggregated and provide a list of priorities for moderator attention [25]. As we see in Wikipedia, lower quality or non-compliant content thus identified may be advertised as such, providing quality cues to users. Once more, we might envisage a kind of "epistemic flagging", where topics need more input of evidence, appear biased or seem to be ignoring valid yet dissenting views.

3.6 Conclusions

This chapter has reviewed theories from social psychology and social epistemology for their contribution to a deeper understanding of knowledge practices in communities as well as for their ability to inform design practice. We have shown how both descriptive and normative insights can be brought to bear on community structure through both interface aspects and collaborative policy formation. In the process, we have highlighted a number of areas where there remains controversy and where living experiments continue to be made. We hope to have shown that, while the ideals of social epistemology may be yet to be fully realized online, a number of examples and principles exist that may be drawn on to address clear epistemic failings and promote good epistemic practice. Importantly, these go beyond what is actually encoded in social and collaborative information building to consider behaviours of knowledge providers and consumers, a holistic and, we believe, more fruitful approach.

Although much research in social epistemology focuses on scientific knowledge, the insights seem appropriate to other forms of knowledge, including the types of

general knowledge exchanged on popular web platforms. It is therefore a valuable exercise to explore the work of key theorists in the field and draw lessons from them that apply as much in the diverse world of online communities as in scientific circles. Similarly, social psychology provides some quite sobering evidence of our tendencies to not behave according to good epistemic principles—even when we think we are. Our growing awareness of these biases and increased vigilance is a good first step toward their mitigation.

Acknowledgements This research was supported by the University of the West of England and the Austrian Science Fund (FWF): P 23770-G17.

References

1. Adamic, L., Zhang, J., Bakshy, E., & Ackerman, M. (2008) Knowledge sharing and Yahoo answers: Everyone knows something. In *17th international World Wide Web Conference (WWW2008), Beijing, China 2008.* New York: ACM.
2. Agrawal, R., Gollapudi, S., Halverson, A. & Leong, S. (2009). Diversifying search results. In R. E. A. Baeza-Yates (Ed.), *Proceedings of the Second ACM International Conference on Web Search and Data Mining* (pp. 5–14). New York: ACM Press.
3. Akrich, M. (1992). The de-scription of technical objects. In W. E. Bijker & J. Law (Eds.), *Shaping technology/building society: Studies in sociotechnical change*. Cambridge, MA: MIT Press.
4. Alcoff, L. M. (2001). On judging epistemic credibility: Is social identity relevant? In N. Tuana & S. Morgen (Eds.), *Engendering rationalities*. Albany, NY: SUNY Press.
5. Anderson, C. A., & Sechler, E. S. (1986). Effects of explanation and counterexplanation on the development and use of social theories. *Journal of Personality and Social Psychology, 50*, 24–34.
6. Barad, K. (2007). *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. Durham, NC: Duke University Press.
7. Blanchard, A. L. (2011). A model of online trust. *Information Communication Society, 14*, 76–106.
8. Burke, A. (2006). Improving prosecutorial decision making: Some lessons of cognitive science. *William & Mary Law Review, 47*, 1588–1632.
9. Chen, C., & Hung, S. (2010). To give or to receive? Factors influencing members' knowledge sharing and community promotion in professional virtual communities. *Information & Management, 47*, 226.
10. Colin, C. (2008). *Nasty as they wanna be? Policing Flickr.com – SFGate*. http://articles.sfgate.com/2008–09–29/business/17120099_1_flickr-policing-images. Accessed Sept 2011.
11. Connaway, L., Dickey, T., & Radford, M. (2011). "If it is too inconvenient, I'm not going after it": Convenience as a critical factor in information-seeking behaviors. *Library & Information Science Research, 33*, 179–190.
12. Ford, N., Wilson, T. D., Foster, A., Ellis, D., & Spink, A. (2002). Information seeking and mediated searching, Part 4: Cognitive styles in information seeking. *Journal of the American Society for Information Science and Technology, 53*, 728–735.
13. Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford, England: Oxford University Press.
14. Fricker, M. (2010). Scepticism and the genealogy of knowledge: Situating epistemology in time. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology*. Oxford, England: Oxford University Press.

15. Friedman, B., & Nissenbaum, H. (1997). Bias in computer systems. In B. Friedman (Ed.), *Human values and the design of computer technology*. Cambridge, England: Cambridge University Press.
16. Gazan, R. (2006). Specialists and synthesists in a question answering community. In Anon. (Ed.), *69th Annual Meeting of the American Society for Information Science and Technology (ASIST), Austin, US 3–8 November*. Silver Spring, MD: Richard B. Hill.
17. Goel, S., Mason, W., & Watts, D. (2010). Real and perceived attitude agreement in social networks. *Journal of Personality and Social Psychology, 99*(4), 611–621.
18. Goldman, A. I. (2003). *Knowledge in a social world*. Oxford, England: Clarendon Press.
19. Goldman, A. I. (2006). *Social epistemology*. http://plato.stanford.edu/entries/epistemology social/. Accessed Jan 2007.
20. Goldman, A. I. (2008). The social epistemology of blogging. In J. Weckert & J. V. D. Hoven (Eds.), *Information technology and moral philosophy*. New York: Cambridge University Press.
21. Goldman, A. I. (2010). Systems-oriented social epistemology. In J. Hawthorne & T. Gendler (Eds.), *Oxford studies in epistemology* (Vol. 3). Oxford, England: Oxford University Press.
22. Goldman, A. I., & Whitcomb, D. (2011). *Social epistemology: Essential readings*. New York/ Oxford: Oxford University Press.
23. Hadjieleftheriou, M., & Tsotras, V. (2009). Special issue on result diversity. *Bulletin of the Technical Committee on Data Engineering, 32*, 4.
24. Hara, N., & Khe, F. H. (2007). Knowledge-sharing in an online community of health-care professionals. *Information Technology & People, 20*, 235–261.
25. Haughey, M. (2007). *Some community tips for 2007*. http://fortuito.us/2007/05/some_com munity_tips_for_2007. Accessed Sept 2011.
26. Holloway, T., Bozicevic, M., & Borner, K. (2007). Analyzing and visualizing the semantic coverage of Wikipedia and its authors. *Complex, 12*, 30–40.
27. Iba, T., Nemoto, K., Peters, B., & Gloor, P. A. (2010). Analyzing the creative editing behavior of Wikipedia editors: Through dynamic social network analysis. *Procedia – Social and Behavioral Sciences, 2*, 6441.
28. 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*, 716–727.
29. Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. New York/Oxford: Oxford University Press.
30. Kosonen, M. (2009). Knowledge sharing in virtual communities a review of the empirical research. *International Journal of Web Based Communities, 5*, 144–163.
31. Kruglanski, A. W. (1990). Lay epistemic theory in social-cognitive psychology. *Psychological Inquiry, 1*, 181–197.
32. Kruglanski, A. W., Dechesne, M., Orehek, E., & Pierro, A. (2009). Three decades of lay epistemics: The why, how and who of knowledge formation. *European Review of Social Psychology, 20*, 146–191.
33. Kusch, M. (2002). *Knowledge by agreement: The programme of communitarian epistemology*. Oxford, England: Clarendon Press.
34. Kusch, M. (2010). Social epistemology. In S. Bernecker & D. Pritchard (Eds.), *The Routledge companion to epistemology*. London: Routledge.
35. Lackaff, D. (2005). *Norm maintenance in online communities: Analysis of heterarchical moderation regimes*. MA thesis, La Trobe University, Australia.
36. Longino, H. E. (1994). The fate of knowledge in social theories of science. In F. F. Schmitt (Ed.), *Social epistemology*. Lanham, MD: Rowman & Littlefield.
37. Longino, H. E. (2002). *The fate of knowledge*. Princeton, NJ: Princeton University Press.
38. Ma, W. W. K., & Yuen, A. H. K. (2011). Understanding online knowledge sharing: An interpersonal relationship perspective. *Computers in Education, 56*, 210.
39. Magnus, P. D. (2009). On trusting WIKIPEDIA. *Epistem, 6*, 74–90.
40. Minack, E., Demartini, G., & Nejdl, W. (2009, October). *Current approaches to search result diversification*. Proceedings of the 1st International Workshop on Living Web, Washington, DC.

41. Origgi, G., & Simon, J. (2010). On the epistemic value of reputation: The place of ratings and reputational tools in knowledge organization. In C. Gnoli & F. Mazzocchi (Eds.), *Advances in Knowledge Organization, Eleventh International ISKO Conference 2010: Paradigms and conceptual systems in knowledge organization* (pp. 417–423). Würzburg, Germany: Ergon.
42. Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton, NJ: Princeton University Press.
43. Panciera, K., Halfaker, A., & Terveen, L. (2009). Wikipedians are born, not made: A study of power editors on Wikipedia. In Anon. (Ed.), *Proceedings of the ACM 2009 International Conference on Supporting Group Work, Sanibel Island, Florida, USA* (pp. 51–60). New York: ACM.
44. Phang, C. W., Kankanhalli, A., & Sabherwal, R. (2009). Usability and sociability in online communities: A comparative study of knowledge seeking and contribution. *Journal of the Association for Information Systems, 10*, 721–747.
45. Poulter, M. (2003). *Value and belief*. Ph.D. thesis, University of Bristol, Bristol, UK.
46. Prabha, C., Lynn, S. C., Olszewski, L., & Jenkins, L. R. (2007). What is enough? Satisficing information needs. *Journal of Documentation, 63*, 74–89.
47. Rheingold, H. (1998). *The art of hosting good conversations online*. http://www.rheingold.com/texts/artonlinehost.html. Accessed Sept 2011.
48. Sanger, L. M. (2009). The fate of expertise after WIKIPEDIA. *Epistem, 6*, 52–73.
49. Scheman, N. (2001). Epistemology resuscitated: Objectivity as trustworthiness. In N. Tuana & S. Morgen (Eds.), *Engendering rationalities*. Albany, NY: SUNY Press.
50. Schmitz-Justin, F. J. (2006). *Knowledge factors: How to animate members of online communities to create knowledge-relevant content*. New York: Lang.
51. Schulz-Hardt, S., Jochims, M., & Frey, D. (2002). Productive conflict in group decision making: Genuine and contrived dissent as strategies to counteract biased information seeking. *Organizational Behavior and Human Decision Processes, 88*, 563–586.
52. Shah, C., Jung, S. O., & Oh, S. (2008). Exploring characteristics and effects of user participation in online social Q&A sites. http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/2182/2028. *First Monday*, *13*(9).
53. Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*, 379–423.
54. Shirky, C. (2009). *A speculative post on the idea of algorithmic authority*. http://www.shirky.com/weblog/2009/11/a-speculative-post-on-the-idea-of-algorithmic-authority/. Accessed Sept 2011.
55. Simon, J. (2010). The entanglement of trust and knowledge on the web. *Ethics and Information Technology, 12*, 343–355.
56. Thelwall, M. (2009). Homophily in MySpace. *Journal of the American Society for Information Science and Technology, 60*, 219–231.
57. Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge, UK: Cambridge University Press.
58. Wikipedia. (2011a). *Wikipedia – Five pillars*. http://en.wikipedia.org/wiki/Wikipedia:Five_pillars. Accessed 27 July 2011.
59. Wikipedia. (2011b). *Wikipedia: Statistics*. http://en.wikipedia.org/wiki/Special:Statistics. Accessed July 2011.
60. Wilson, T. D. (1997). Information behaviour: An interdisciplinary perspective. *Information Processing and Management, 33*, 551.
61. Wray, K. B. (2009). The epistemic cultures of science and WIKIPEDIA: A comparison. *Epistem, 6*, 38–51.
62. Wu, F., & Wilkinson, D. H. B. (2009). *Feedback loops of attention in peer production*. Ithaca, NY: ArXiv e-prints, Cornell University Library.
63. Yahzmin (2009). *Are there statistics about "yahoo answers"?* http://answers.yahoo.com/question/index?qid=20090505082755AAW8t2U. Accessed July 2011.
64. Yardi, S., & Boyd, D. (2010). Dynamic debates: An analysis of group polarization over time on Twitter. *Bulletin of Science Technology Society, 30*, 316–327.