

Social Epistemology and Online Knowledge Exchange

Paul Daniel Matthews

A statement submitted in partial fulfilment of the requirements of the University of the West of England, Bristol for the degree of DPhil

Faculty of Environment and Technology, University of the West of England, Bristol

October 2015

Abstract

This document summarises the submitted research, which has investigated online knowledge exchange and related it to the philosophical field of social epistemology. The broad aims have been: firstly to investigate what social epistemology theory can offer in the way of guidance and evaluative frameworks for the design of knowledge systems; and secondly, to determine what the empirical study of knowledge exchange platforms can tell us about knowledge as emerging from online practice. The submitted work consists of six papers that are a mixture of review/position papers and reports of empirical investigation. These have been published in information science journals and conference proceedings. However, following the established tradition of information science, the work is positioned as being cross-disciplinary in ambition.

After introducing the submitted papers and the inspiration for the research, the main theoretical positions of the research are outlined and justified. These were a naturalised social epistemological position, inspired by Alvin Goldman, but widened to a situated and systems-oriented view. The naturalised view of epistemology allows for consideration of evidence from psychology, and here some key theories in social and cognitive psychology are outlined. Finally, as the subject is human-computer-human interaction, the sociotechnical setting is established. Further, the main platforms of study in the empirical work — social question answering systems — are introduced and described.

The main methodology and research approaches followed are presented next. A mixed methods philosophy was deemed suitable for this area of research and — alongside the review work — the broad *web science* approach of combining network and data investigation with qualitative methods is justified. Review work included early collaborations with an information scientist and a philosopher, which helped to bring together and clarify epistemological, and sociotechnical themes.

The discussion section presents some of the main themes and conclusions of the submitted work, including: 1) The identification of knowledge patterns and practices online; 2) Criteria for online knowledge exchange distilled from the social epistemology literature; 3) Some *triangulations* where theory from philosophy and psychology seemed to corroborate and serve to explain online behaviour; 4) Socio-temporal aspects to online knowledge exchange that are perhaps under developed in philosophy but apparent in practice; 5) Credibility cues and bias, seen as crucial to a rounded study of user interaction with online sources; and finally 6) Interventions suggested by the research which would aim to raise the quality and effectiveness of social media knowledge systems.

Finally, conclusions and suggestions for further work are presented. These follow on from the submitted strands of research and present possibilities for how the work may be extended and improved upon. In common with the research, these combine philosophy, modelling, interaction design and qualitative methods. Such a combination is seen as essential to developing an enhanced understanding of how the web *serves* and *could serve* as a platform for human knowledge.

Contents

Introduction & Claim to Originality.....	3
Inspiration	5
Interdisciplinarity.....	6
Reciprocity of philosophy.....	7
Background and Position	8
Philosophical position.....	8
A working definition of knowledge	8
Social Epistemology.....	10
Philosophy of Information.....	12
Testimony	13
The case for normativity.....	14
Social and cognitive psychology	15
Psychological influences	15
Sociotechnical position	17
Socio technical views	17
Social epistemic software	18
Subjects & Methodologies Followed.....	19
Information and web science approach	19
Methodology.....	20
Methods used	21
Discussion and Main Contributions	22
Knowledge patterns and practices in online knowledge sharing environments	22
Normative SE for the evaluation and design of online systems	24
Triangulations	25
Spatiotemporal aspects in knowledge sharing	26
Credibility cues and bias in SQA.....	28
Identifying interventions.....	28
Conclusions	31
Core goal / contribution	31
Limitations and Further work.....	32
Limitations	32
Further work	33
Submitted papers	35
Administration Statements	36

Introduction & Claim to Originality

The central thesis of this work is that social epistemology (SE) can be used to inform the design and evaluation of online knowledge sharing systems. While many of the themes of the research have a relatively wide following in information science and psychology (e.g. credibility, knowledge sharing, information seeking behaviour), very few researchers have adopted a strong SE-oriented approach to the study of web systems — though there are some exceptions (e.g.:Fallis 2009, Simpson 2012, Chan, Li et al. 2014). This is therefore considered to be the main contribution of the submitted work, which is further differentiated by the candidate's primary grounding in web technologies and information science rather than philosophy.

The submitted work commenced with discussion and review papers that served to set the scene for the programme of research and then progressed into empirical work which picked up on some of the issues raised in the former papers and studied them in more detail in the social question-answering (SQA) environment.

I will henceforth refer to the submitted papers as [1] to [6], as presented in the first part of the references section.

The initial papers combined a descriptive look at aspects of social interaction in online knowledge environments with arguments for employing a normative SE to provide evaluation criteria and design principles for knowledge exchange. Papers [1] and especially [2] helped significantly in connecting SE theory with interface and social norm aspects in online communities and applications. Both papers also looked at relevant theory from social psychology as well as empirical research in online community dynamics to provide a rounded, naturalistic view of knowledge exchange. Paper [1] served to highlight and characterise epistemically positive and negative online interaction patterns (in the tradition of design patterns from the fields of architecture and software engineering). Paper [2] distilled SE theory into design principles for online knowledge communities and concluded that an underlying SE-inspired approach can help to correct some of the observed problems in these communities such as various forms of bias.

Article [3] was a contribution to an international workshop (CSCW 2013 workshop on Social Media Question Answering) focusing on research priorities in SQA and summarised some of the key points raised in [1] and [2] in addition to some new work. In workshop discussion, epistemic (particularly credibility) aspects of SQA were highlighted along with other research priorities, which included the impact of SQA on users and the need to look for generalisability of empirical findings across SQA applications.

In article [4], the study of specific interactions in an SQA community was used to study bias and temporal aspects of knowledge exchange, both at the single question-answer level and at the overall level of community dynamics. This paper raised issues relevant to the archival contribution of SQA resources and triangulated with earlier work on Wikipedia and work on other SQA platforms than the one spotlighted (Stack Overflow for computer programming knowledge exchange). The first mover advantage and other temporal bias effects were noted and visualised. Evidence that temporal effects contribute as much as the quality of answer to knowledge acceptance was put forward.

Article [5] took the empirical work further by studying the impact of credibility cues in SQA with users (and is the first known work to do so). Again, the Stack Exchange family of websites were used and the study extended to look not only at programming knowledge but also library and information science and philosophy-oriented knowledge exchanges, (thereby in part addressing the research priority identified in [3]). This work was felt to be important in establishing the use of peripheral cues to authority which has an important bearing on acceptance of testimony online. The bandwagon heuristic was shown to be relied on by users (both experienced and novice), though it was argued that this is often epistemically positive. A negative effect, however, is the Matthew effect which bypasses potentially useful sources to focus more on those already highly ranked. It was argued that this might be an epistemic injustice in SE terms.

Paper [6] reviews the role of knowledge intermediaries online, particularly that of the *synthesist* who collects and summarises evidence on behalf of an information seeker. The paper applied quality criteria for synthesis to social media synthesis and concludes that current content may approach the standard created by these criteria but largely fails to meet them.

The submitted articles attempted to derive normative conclusions and pragmatic recommendations from the evidence reviewed within. A key theme in [1] and [6] was the contribution of information professionals to popular knowledge exchange as a key intervention in raising their epistemic standards (these articles appeared in journals associated with a Library and Information Science readership). Other recommendations arising from the work involved the social and interface design of knowledge exchange as the artifactualisation of normative SE theory.

The main novel contributions of the submitted papers are presented more fully in the Discussion section, below. In the following sections, the overarching approach of the research is introduced, beginning with the main sources of inspiration.

Inspiration

“Many of the social institutions we rely on to judge trustworthiness and veracity are missing from our online information life. Being able to engineer the Web of the future requires not only understanding it as a computational structure but also how it interacts with and supports interaction among its users. “ (Hendler,Shadbolt et al. 2008 p67)

The research described here was inspired by the observation of the satisfying power of the web as a knowledge tool. Since the early years of the web, online forums have grown to provide a fantastic resource for answering questions or getting help with what at first seem to be intractably idiosyncratic problems. Clearly, the scale of knowledge sharing on the web — and how people have come to trust it as a source of *downloadable beliefs* (Schiltz,Truyen et al. 2007) — has implications for theories of knowledge and knowing. It seemed therefore that a priority was to investigate how existing theories of knowledge adequately catered for what was happening on this new collaborative and massively interconnected medium. Such work should have the twin aims of: 1) understanding which theories could adequately describe and predict what were now rather more visible, traceable and networked epistemic (knowledge-related) interactions and 2) providing a basis to evaluate and guide the design of online knowledge communities and their supporting software.

The emerging philosophical field of social epistemology seemed to be an excellent starting point for such a *cyber-epistemology*. A common aim of many of its proponents is to make sense of real-world knowledge practices and provide a basis for improving the design of social knowledge systems, with the objective of providing the associated societal benefits:

“I maintain that social practices can make both positive and negative contributions to knowledge. The task is to show just which social practices, under what conditions, will promote knowledge rather than subvert it.” (Goldman 1999 p7)

As we shall see, social epistemology is a broad church but several of its prominent authorities have ideas of relevance to web phenomena. They do, however, sometimes fall short on specifics, particularly in regard to those social and interface patterns that we find in large online communities of different types — collaborative filtering, voting, friend following and reputation proxies to name but a few. That said, theories of testimony and trust provide many of the necessary elements, which in many cases simply need extending or qualifying for the web domain.

There are still surprisingly few researchers who are taking such a grass roots knowledge-oriented view of the web. Philosophers on the whole will use web-based examples to illustrate particular theories or arguments, but are usually less interested in following this up into practice or application. The information systems and HCI communities are very adept at describing and documenting

interactions and connecting to relevant theory from communications, sociology and psychology - but look less often to philosophy. Perhaps the most directly involved in philosophical questions are the system designers and community organisers themselves, though here the philosophy is not explicit or used as a clear reference point. Of course, to this there are notable exceptions and these include philosophers and information scientists whose work is cited in the submitted papers and who may also be discussed below, as well as the author's collaborators on the papers for consideration.

Interdisciplinarity

As will be clear from the above, the setting for the submitted work is strongly cross-disciplinary. I will argue that this is both necessary and apt, fitting as it does within established traditions in information science, within which this work can be situated. Furthermore, it is strongly consonant with a naturalistic approach to philosophy, which requires that evidence from a range of disciplines should inform and shape philosophy. This latter position will be further explored in a subsequent section. For the moment, it will be helpful to discuss the additional benefits of cross-disciplinarity and to address possible objections.

While multi- cross- and interdisciplinarity are often used interchangeably, interdisciplinarity may be defined as a broader approach than multidisciplinary in that it does not imply staying within a single field (Alvargonzález 2011):

“Interdisciplinarity analyzes, synthesizes and harmonizes links between disciplines into a coordinated and coherent whole” Choi and Pak, quoted in (Alvargonzález 2011 p388)

It is to this type of integrative work that I aspire, with due regard to the risk of such holistic approaches being open to criticism for overreaching, misrepresenting, or even having a political agenda (Alvargonzález 2011). I would certainly agree with Steve Fuller's view of the need for philosophy itself to take a more central, interdisciplinary and integrative role (which he has dubbed “deviant interdisciplinarity” in contrast to a “normal” view where philosophy is seen as a discrete and somewhat separate layer) (Fuller 2013). This may be the best way for philosophy to achieve an impact on society.

Reciprocity of philosophy

To elaborate on the key motivating approach of the submitted work, it was to seek a reciprocal relationship between philosophy and sociotechnical design (design here meaning both design as pre-planned and design as accidental) in the area of knowledge exchange. This has a tradition in the naturalistic approach to epistemology — which will be further expanded below — and which finds such reciprocity in the relationship between philosophy and science:

“There is thus reciprocal containment, though containment in different senses: epistemology in natural science and natural science in epistemology” (Quine 1969 p83)

This approach has also been dubbed “cooperative naturalism” (Huaping and Xiaoming 2007). There is a need for an independent, analytical view of how knowledge works/should work in a social setting, and there is a need to test this view with real evidence “in the wild”. Furthermore, the design and practice found in sociotechnical knowledge systems provides an emergent, enacted philosophy which we may compare to foundational and analytic “ideals”:

“The process of having philosophical ideas take a concrete form via technology lends to them often radically new characteristics, transforming these very concepts in the process..”
(Monnin and Halpin 2014 p5)

Background and Position

Having established the broad motivating factors of the research, it will be beneficial to establish the theoretical position that I have taken in the submitted work. Clearly, philosophy is to the fore and the blend of social epistemology that has been influential should be briefly outlined. Furthermore, an overview of the naturalistic and normative nature of the philosophical approach taken should help to justify the use of theory from other areas of science, notably that which is sociotechnical and psychological in nature.

As a key desire has been to relate philosophy to sociotechnical systems for knowledge exchange found on the web, I will visit some sociotechnical models that have been found to be relevant to the work. I will go on to explore the precise subset of online systems that have informed the research - social question-answering services. These have featured to a greater or lesser extent in all of the submitted studies.

Finally, as I approve of a close relationship between philosophy and psychology, the theory used from the fields of social psychology and cognition will be outlined. Psychology is a useful intermediary that can provide insight into the usefulness or otherwise of the various philosophical approaches to knowledge.

Philosophical position

A working definition of knowledge

As the emphasis of the work has been knowledge exchange, it is helpful to draw on epistemology to try to clarify the kind of knowledge I have had in mind. Some preliminary caveats: knowledge has always been a rather fuzzy notion (Russell 1992) and different definitions are not necessarily mutually exclusive (Zagzebski 1998). Zagzebski usefully points to the need to first clarify the purpose of one's definition and that this may be either practical or theoretical. If one's aim is more practical, a definition can help us recognise knowledge and help us to get more of it. Zagzebski claims that practical aims can be satisfied by a contingent definition (Zagzebski 1998)

On the theoretical side, traditional analytical epistemology has been built around a tripartite theory of knowledge as *justified true belief*, with many ongoing debates as to the nature of all three of these pillars, together with the sufficiency of the three together. Some have concluded that this approach has not been successful due to its ongoing lack of agreement and the vulnerability of each proposed improvement to counter examples, usually involving luck. So it may in fact not be possible to specify

necessary and sufficient conditions for knowledge (Millar, in Pritchard, Millar et al. 2010) or definitions need an explicit anti-luck qualification (Pritchard, Millar et al. 2010).

Ultimately, an elaborated tripartite theory may place such stringent conditions on knowledge that it becomes hard to see how it can ever be achieved in reality, whereas in fact it can be hard to even distinguish belief from knowledge:

“There is no way of deciding how much uncertainty makes a belief unworthy to be called ‘knowledge’, any more than how much loss of hair makes a man bald.” (Russell 1992 p113)

So to put forward a more practical — yet fuzzy — definition of knowledge, it can be seen to be based on *strong belief*, with other conditions difficult to resolve (and anyway less important for the present purposes). Indeed, even in the tripartite theory, Plato originally intended *certainty* or *conviction* rather than *belief*, according with this idea of strength (Audi 1998). This in turn accords with a Bayesian view of knowledge and belief which allows for degrees of power to both (Goldman 1999).

We still of course need to distinguish *rightness* and *wrongness of belief*, given that we want to privilege *right* knowledge as an ultimate moral good (Zagzebski 1998). Here I feel it is helpful to look at outcomes as much as the nature and origin of the belief itself. We can in fact judge beliefs by their predictive and explanatory power (Nilsson 2014). As Russell has it, knowledge is preparation for action, and the success of this action distinguishes it from belief (Russell 1992). I would extend this success criterion to a view of *truthiness* as agreement with the most salient and prevalent evidence — including that received through the testimony of others.

I would further justify this *success* approach to knowledge through the content of a large amount of online knowledge exchange, which is often more about procedural knowledge. This *knowing how*, or *practical* knowledge, is seen by many to be distinct from propositional or abstract knowledge (Fantl 2014). People seek such knowledge on the web so that they can *move on*, whether this be with a practical or cognitive task. The success of that task’s completion will greatly determine how that rate the quality of the knowledge received, how much it will become part of their own belief system and how readily it will be conveyed to others.

In closing, it is worth noting the compelling thesis that our very concept of knowledge may have evolved from our ancestral need to detect reliable informants (Craig, referenced in Pritchard, Millar et al. 2010). That is, the very need to have a separate concept of knowledge may be a result of human sociality and communication.

Social Epistemology

The term “social epistemology” was first coined in a farsighted work by Egan and Shera (1952), librarians working with bibliographic systems and seeking a new foundational theory for the work of knowledge organisation and dissemination. The stated aim was to develop a *macrocosmic* method and to acknowledge and accommodate contextual information needs, communication and socio-cognitive factors in knowledge transmission. Many of the ideas in this work presaged the later appropriation of SE by philosophers as well as the field of information seeking behaviour in LIS. That said, the emphasis in later work by Shera was more heavily oriented toward the storage of knowledge than the dynamics of sharing (Shera 1973).

SE has been further promoted as a theoretical foundation for library and information science by Don Fallis (Fallis 2006). Fallis has argued convincingly that the — often unspoken — aim of information services is the maximisation of knowledge acquisition and that the setting is inherently social, so the work of social epistemologists should be consulted in order to inform the design of social information systems.

Alvin Goldman set out his vision of Social Epistemology (SE) in his 1999 book *Knowledge in a Social World (KISW)* (Goldman 1999). Goldman’s view comes from the tradition of analytic epistemology, but in KISW he admits the importance of the social sphere as a source of knowledge and outlines how epistemology might be extended to better accommodate social evidence and to evaluate social knowledge practices. Goldman makes a useful distinction between fundamental and instrumental values of knowledge patterns and practices and puts truth-value centre stage as the measure of success:

“What I am contemplating, then, is an evaluative structure of two parts. States like knowledge, error, and ignorance have fundamental veritistic value or disvalue. Practices have instrumental veritistic value insofar as they promote or impede the acquisition of fundamental veritistic value. The structure here is perfectly analogous to the structures of consequentialist schemes in moral theory. One type of state, such as happiness or utility, is taken to have fundamental or intrinsic moral value, and other items, such as actions, rules, or institutions, are taken to have instrumental value insofar as they tend to produce (token) states with fundamental value.” (Goldman 1999 p87)

Goldman’s approach, then, is deliberately normative and applied, as KISW first develops a framework for SE and then attempts to apply it in various real world setting such as democracy and law. It is also influenced by Bayesian approaches to belief and therefore steers usefully away from formalisms based around absolute propositional knowledge. Centred as his SE is on truth, Goldman is

rather scathing about relativist / postmodernist views of knowledge and truth as socially constructed, insisting on an approachable and shared truth about the external world.

KISW has been very influential, justifying Goldman's attempt to apply epistemology to real-world practices. But criticism has been levelled at his approach from three perspectives: the epistemological, the broader social epistemic and the ethical. Firstly, some have seen his "weak knowledge" approach — knowledge as merely true belief — as problematic in removing the condition of having a *reason* to believe and also more vulnerable to cases of epistemic luck (Brendel 2009)

At a broader level, Goldman's SE has also been criticised by those who take a more sociological or communitarian approach to knowledge:

“For all his criticism of the sociology of scientific knowledge, Goldman never addresses its most central contention about knowledge, that is, that the primary knower is the group rather than the individual. Goldman agrees that traditional epistemology is excessively individualistic. But he thinks this shortcoming can be remedied simply by adding a social wing to the old edifice.” (Kusch 2001 p188)

Similarly, the KISW focus on veritistic consequentialism sidelines important social epistemic goods such as self-efficacy, group communication and wellbeing (Vähämaa 2013).

Despite these concerns, KISW still provides one of the most thought-through models of SE and still has much that can be usefully adapted to other fields. Indeed, Goldman himself has acknowledged areas that needed updating since KISW and has worked to better accommodate sociological and systems approaches in addition to theoretical aspects of distributed cognition (Goldman 2010).

Where Goldman has probably been correctly criticised is in his failure to sufficiently acknowledge important influences on “knowledge in the wild” and to limit himself to an abstract account of SE (Fricker 2010). Those arguing for a more situated ecological view of knowledge that admits the role of reputation, power relations and advocacy provide compelling arguments for widening the field of vision to include *ethical* considerations (Fricker 2010, Code 2010, Origgi 2012). While clearly more rounded, these views add complexity to the analysis of social knowledge systems.

Philosophy of Information

Luciano Floridi has proposed the Philosophy of Information (PI) as a philosophical base to information science and computing, inspired by the mounting management challenges and ethical dilemmas posed by digital information (e.g. Floridi 2004, Floridi 2011). His approach is an info-centric one that deemphasises some core epistemological questions and instead focuses on the semantic value of information. According to Floridi, information carries truth by definition, with the entailment that there can be no such thing as false information (Floridi 2011 p96). This actually seems to put him close to Goldman in recognising the importance of fundamental veritistic value - though for Floridi, this truth is already encapsulated in information whereas Goldman it lies in how information is processed into knowledge.

Fundamentally, Floridi's PI and Goldman's SE seem to have different aims: PI is seeking an introversive normativity whereas SE purposefully seeks extroversive applied value that encompasses social practices. While PI may well move on to having more of an external face, it is not yet clear what shape this move will take.

For Floridi, PI is more suitable as a basis for Library and Information Science as LIS should be more about organising and managing access to information than about how it should be put to use:

“Both classic and revolutionary ESK [Epistemology of Social Knowledge] are prescriptive. Their ultimate aim is to establish, for example, not what one believes about the stars, but what one should, and is justified to, believe about them. Both share in principle the same scope of investigation, namely the social dynamics of (allegedly, for the revolutionary brand) epistemic phenomena. All this is at the same time too much and too little to provide a satisfactory foundation for LIS. It is too much in terms of aim, because LIS can be normative but it does not and should not be epistemologically prescriptive.” (Floridi 2002 p39)

The current work argues strongly for an interventionist role for LIS that has more in common with Goldman's instrumental veritistic evaluation than Floridi's assumption that SE should not exert evaluative power. Only through such an evaluative framework can we compare and contrast epistemic systems. Floridi's document-centric view of the LIS profession has disturbed some who have seen it to be neglecting the socially mediated function of knowledge transfer:

“LIS cannot be reduced to the techniques of document management and should not be confused with the job of running a library. The library job, in fact, relates closely to the sociology of social knowledge that Floridi rejects as a foundation for LIS.” (Cornelius 2004 p384)

Testimony

For social epistemologists, the testimony of others needs to regain an important place in our consideration of sources of/conditions for knowledge and thus has merited further attention in recent years. A key consideration has been whether testimonial justification alone is enough to lead to knowledge (a non-reductionist position) or whether testimony is only useful as a source of justification when the individual reduces its content to more foundational sources such as perception and deduction (the reductionist position). Many social epistemologists lean toward non-reductionism, particularly in view of the evolutionary importance of social knowledge to human survival (Fricker 2010). Indeed the non-reductionist position is one that seems consistent with the psychological evidence for how beliefs are formed (Gilbert 1991). But some theorists have argued convincingly for a false dichotomy here (Lackey 2008). While pure reductionism is clearly unfeasible, outright non-reductionism can be argued to lead to irrationality and the perpetuation of error. That said, I feel that the *principle of credulity* by which we trust to the sincerity of the bearer of testimony in the absence of contrary evidence, is strongly borne out by human behaviour and online interaction. A position on reductionism seems important when considering online knowledge sources and has been discussed in relation to Wikipedia (Tollefsen 2009) and the fact that in many cases the user of such sources has little prior scaffolding for assessing testimony through a reductionist process. Instead, as Tollefsen (op.cit) and others have argued, some form of non-reductionism may be in operation, with the help of some epistemic vigilance at the level of individual testimony or overall system credibility.

Many theories of testimony had the verbal form in mind when establishing conditions for trust and the transmission of knowledge. But social epistemologists have seen recorded testimony as merely an alternative kind of testimony. This has more recently been extended to online testimony in philosophical appraisals of resources such as Wikipedia (Magnus 2009). Such analyses have noted how much of the information available with personal testimony is stripped away in the online context — almost to the extent where we no longer have a yardstick for assessment. I would suggest that we need to be careful in considering online testimony to be merely a special case of the verbal — it has properties and characteristics that compensate for, or replace, the verbal cues and in some ways needs to be treated differently.

The case for normativity

As noted above, Goldman's approach in KISW was strongly normative, in trying to establish criteria for evaluation:

“social veritistic epistemology does not merely seek to describe social practices that are actually in place, nor to trace their historical development. It has the distinctive normative purpose of evaluating or appraising such practices on the veritistic dimension, that is, in terms of their respective knowledge consequences. Practices currently in place will be veritistically good or bad in varying degrees they will rarely be ideal. To investigate prospects for improvement, social epistemology must be prepared to transcend previously realized practices. It must be ready to consider the probable veritistic properties of practices that have not yet been, but might be, adopted. Thus, veritistic epistemology tackles the admittedly nontrivial task of assessing both actual and possible practices in terms of their foreseeable informational bounty” (Goldman 1999 pp6-7)

Establishing such normativity is essential to the submitted research, for without it philosophy would become a static, interpretive framework rather than a dynamic structuring tool.

Both naturalism and normativity in epistemology have come under criticism in that a scientifically-led approach to knowledge can be seen as purely descriptive, with a jump to normativity not having any rational grounds. This *is-ought* problem is addressed by Goldman through his use of “veritistic value” as an outcome measure with an an a priori positive valence (Wrenn 2004).

Of equal importance is an appeal to both ethics and to taking an evolutionary approach to knowledge. We have already seen there is close connection between social epistemology and consequentialist ethics and to this Fricker has added justice (Fricker 2007). Yes, we do need to make a value-based step from description to normativity, but have a number of ethical guides in doing so. Also, we need to give up on an idea of perfectionism, but can use such guides in evolving systems toward an improved state. Such a state is itself only a point in time and with increased evaluative evidence we will probably still have a large scope for improvement. Importantly also, this is where the feedback loop from practice to philosophy has its place, since idealised systems may have perverse consequences when put into practice.

Social and cognitive psychology

Psychological influences

I hope to have established that a naturalised, social epistemological approach to knowledge not only justifies, but necessitates the consideration of evidence from psychology and cognitive science. And these disciplines offer several lines of research that greatly aid our understanding of knowledge exchange.

Table: Relevant psychological theories for knowledge exchange

Aspect	Theory/Model and Proponents
Knowledge Sharing	Self-efficacy - (Bandura 1994) Social exchange theory (esp. Reciprocity) - (Homans 1958)
Knowledge Seeking / Knowledge Acceptance	Lay epistemics - (Kruglanski, Dechesne et al. 2009) Cognitive bias - (various, e.g. Shermer 2012) Dual process / Elaboration likelihood - (e.g. Kahneman 2011)
Overall Knowledge Systems	Distributed cognition (e.g. Dror and Harnad 2008)

A large body of studies investigating motivations for knowledge sharing have found evidence for self-efficacy and reciprocity as key intrinsic motivators (Kosonen 2009, Ma and Agarwal 2007, Hsu, Ju et al. 2007). Self-efficacy may be defined as “a judgement of one’s ability to organize and execute given types of performances” (Bandura, quoted in Kosonen 2009). It may be achieved both through an individual’s own practice and through the observation of others. Similarly, many of the norms for social exchange, originally developed through studying interactions in face-to-face groups, have been shown to hold online (Kosonen 2009). Among these, the norm of reciprocity — the expectation that if you share your own knowledge, you will also be able to benefit from the knowledge of others — is especially important. Together, self-efficacy and social exchange norms help to explain the

emergence of amateur yet enthusiastic and prolific users of knowledge sharing platforms.

Chief among the psychological theories relating to knowledge seeking has been the lay epistemic system developed by Kruglanski (Kruglanski, Dechesne et al. 2009, Kruglanski 1990). Kruglanski demonstrated that *need for cognitive closure* greatly influences the way that people seek and use information. His unimodal theory tackles hypothesis formation and evidence gathering, while he also describes how the phenomenon of epistemic authority has been empirically investigated. Kruglanski provides a strong theoretical framework to understand how users interact with information and this seems to translate very well to the kind of online systems studied in the submitted research.

Importantly, the theory applies to general, everyday knowledge formation, whereas much other work on knowledge gathering and use of evidence has been focused specifically on scientific knowledge.

A host of psychological studies have identified and profiled various cognitive biases in information seeking and usage and these echo the concerns over the detrimental influence of bias voiced by social epistemologists (Goldman 1999). Bias threatens many of the social epistemic ideals developed by epistemologists. Among the significant biases are those of confirmation — confining of knowledge seeking and acceptance to that which confirms your existing beliefs authority bias — the tendency to value the opinions of an authority, especially when we have low levels of knowledge on a topic the bandwagon effect — the tendency to defer to the opinions of a group and the availability heuristic — the tendency to make decisions based on recent and prominent examples (Shermer 2012). The pervasive nature of such biases and an understanding of their modification was seen to be essential, hence their significant treatment in the submitted work.

A further set of literature relating to dual process theory in cognition has been influential to information science — particularly studies of credibility of online information. Evidence from a range of studies points to a fast, heuristically-guided processing system (System 1) that is used for many familiar tasks but which — while often giving good results — is prone to error. A slower, deliberative system is engaged for deeper processing, but this is effortful and serial in nature (System 2). Kahneman is one of several authors who point to belief formation being often handled by System 1, with refutation of beliefs being a System 2 function, echoing *Spinoza's conjecture* (Shermer 2012) that (rational, rather than reactionary) scepticism is more difficult than credulity (e.g. Kahneman 2011)

A final theoretical movement which dovetails well with much thinking in philosophy, is that of distributed cognition (e.g. Dror and Harnad 2008). This raises the level of the study of cognition to human-human and human-computer systems, and is pragmatic in not expecting any one individual to fully understand a process or system of knowledge. As noted above, the reconciliation of the *single knower* of traditional epistemology with a wider knowledge group or system is not always easily

made, and does present a problem with some accounts of knowledge. That said, the importance of testimony and trust to SE helps to provide a bridge to the more group-oriented view.

Sociotechnical position

Socio technical views

As has been mentioned, Goldman's *systems-oriented* SE is a more recent development in his thinking and was partly in response to sociotechnical theory that admits the importance of external actors and artifacts in cognition. Indeed, the potential for computers to be thought of as epistemic agents has been proposed within philosophy (Humphreys 2009) and is closely related to the *extended mind* hypothesis, that the mind can be considered to be wider than the individual and can include objects and tools used to both manipulate and record knowledge (Clark and Chalmers 1998).

If we accept the importance of epistemic technology to the process of knowledge acquisition and exchange, then design and evaluation questions become paramount, as we see in technical systems the embodiment of a range of potentially competing philosophical positions:

“The material a priori of technical systems such as the Web is brought about by what we call ‘artifactualization’, a process where concepts become ‘embodied’ in materiality — with lasting consequences” (Monnin and Halpin 2014 p5)

Monnin's artifactualization can be cast as a socio-technical version of Quine's naturalization, a process which will inevitably occur with or without conscious philosophical scrutiny (Monnin and Halpin 2014)

Perhaps crucially for the research described herein, a socio-technical approach to evaluation should evaluate both technical designs (GUIs, moderation systems, reward systems, voting capabilities) and social designs (community norms, motivations, authority, moderation regimes — the way that the technical capabilities are implemented). This motivates a holistic and mixed-methods approach to the investigation of online systems. This is very much in line with the *web science* foundational principle, that the design of the web is bipartite and includes both social and technological protocols (Berners-Lee, Hall et al. 2006a)

Web science has emerged as an interdisciplinary approach to studying the web and its impact on society (Berners-Lee, Hall et al. 2006b). A key aim is to develop design improvements on the web that fix known issues such as privacy and intellectual property. Besides understanding technical protocols, web science needs to extend itself to understanding how human interactions contribute to successful outcomes on the web — it needs to understand the *social machine* — the technology

coupled with the underlying “rules, policies and organizational structures used to manage the technology” (Hendler, Shadbolt et al. 2008)

Social epistemic software

As web-based knowledge exchange is supported in a range of systems, from the generalist to the highly specialist it was necessary to focus on a subset of systems for the evaluation, examples, and empirical work used in the research. As the work was done concurrently to the rise of social, *web 2.0* platforms, these were a natural candidate for investigation. These social platforms are typified by large-scale collaboration, a read-write approach and the *de facto* delegation of content-generation, authority and often moderation to users not directly affiliated with the tool provider.

While including evidence from activity on wikis, blogs, social networks and online forums, the empirical work and the majority of evidence in the submitted work has come from research into *social question-answering platforms*, henceforth SQA. These platforms have been in existence since the early 2000s, but saw a significant rise in popularity and usage in the middle and later parts of that decade. SQA was seen to be a natural fit to the investigation of SE — systems that are by their nature social and which focus on exchange of knowledge through the structured collection and comparison of written testimony.

The typical SQA design allows for user-contributed questions, to which other users contribute answers. The answers are voted up or down, both by the original questioner and the user community in general. This gives a ranking order to user contributions. There are often points awarded for contributing answers that receive up-votes and for asking popular questions.

While SQA emerged from the more technically-oriented, computing communities on the web, it was soon being used for asking questions on numerous topics — both academic and social, specialist and generalist. With their high visibility on search engines, SQA question-answer resources are widely used by the web community at large, not simply the users of particular platforms. For this reason, the knowledge shared on these sites may have high impact and reach.

In the late 2000’s, influential research on SQA began to emerge within information science which took both a network-oriented and content-analytic approach to understanding how people ask and answer questions and their criteria for ranking answers (e.g. Adamic, Zhang et al. 2008, Kim and Oh 2009). This work was highly influential in the current research, both methodologically and in terms of what it said about online knowledge exchange dynamics.

Subjects & Methodologies Followed

In the preceding I have tried to establish the overall theoretical framework and influences for the research, with a naturalised SE as the overarching theme. I have noted psychological influences that work well within this framework and translate well to the online context. I have also introduced relevant sociotechnical positions — that motivate a decidedly systems-oriented SE — and the particular type of knowledge exchange platforms that have been referenced in the submitted research.

This section will outline in more detail the range of methods used in the submitted papers, along with further justification of the methodology and what it could potentially add to our picture of web-based knowledge exchange. I shall begin with the grounding of the methodology within information science and the new subfield of web science.

Information and web science approach

Information science regards itself as a multidisciplinary, integrating science (Bawden 2009). Given the range of viewpoints and practice in the discipline it can also be viewed as a number of disciplines working at different integrative levels. That said, its repertoire of primary methods (regardless of positivist or interpretive viewpoint) may be viewed as all stemming from observation (Wilson 2002)

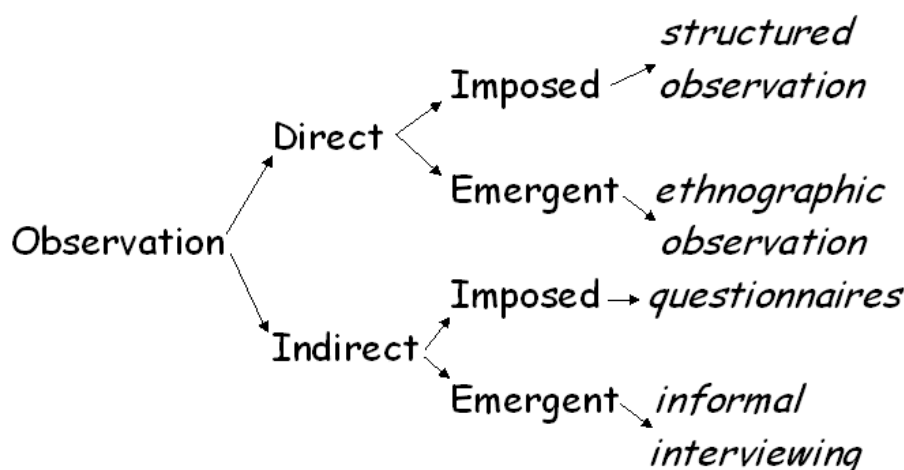


Figure: Information science methods typology (Wilson 2002). Used under a creative commons license (<http://creativecommons.org/licenses/by-nc-nd/3.0/>)

As an emerging subdiscipline of information science, web science has a strong background in network science and graph theory but admits the need to go wider than this approach in the

development of new models (Hendler, Shadbolt et al. 2008). In web science, a strong interdisciplinarity motivates a mixed methods approach in which “a range of epistemologies and methodologies must be harnessed to understand the Web at both micro and macro level” (Tinati, Halford et al. 2012, based on Halford et al’s Manifesto for Web Science)

Methodology

“Ontologically mixed method research accepts a mental world as well as a physical reality within which the communication process is embedded. It acknowledges that social science aims to understand the modes of practice of human beings within their lived environments and as such recognises many different aspects of human action, including qualitative intentions, attitudes, beliefs, experience and culture as well as more quantitative reductive and more easily measurable phenomena” (Hodgson and Cofta 2009)

The methodological approach echoed the veritistic SE view of acknowledging external truths, but the importance of individual internal beliefs to their approximation. In information science, such a positivist stance can coexist productively with an interpretive view and lead to a fuller overall picture of the phenomena under study (Tinati, Halford et al. 2012)

While social media activity data is hugely informative in understanding the “what” of knowledge exchange processes, it was also considered necessary to understand the “why” (Hodgson and Cofta 2009) through interpretive content analysis and more structured empirical experimentation with knowledge exchange interfaces. So the research methodology was one of *paradigm relativism* — selecting a methodology that was appropriate and enabling *concurrent triangulation* between methods (Terrell and Steven R Terrell 2012)

Open social datasets provide a great opportunity to study social interaction patterns (Halavais 2011) and provide a source of *natural* data (Speer 2008). The availability of user comments and discussion is also such a natural source of evidence. However, in order to understand impact and information usage, the later work involved a more *artificial* experimental setting.

Chronologically, the work began by taking a review and positioning approach [1, 2, 3], followed in the later papers by a more primarily empirical approach [4, 5]. Review work was needed to integrate evidence from philosophy, psychology and information science/computing and to establish and explore links between philosophy and sociotechnical design. Further, work was done to define and test evaluative frameworks for online knowledge exchange that were grounded in SE theory.

Methods used

The information, web and data science footing, then, indicated a mixed methods concurrent triangulation approach. While a direct, emergence-oriented observation (following Wilson's typology, above) was possible for web interaction data — facilitated by open data for the sites studied — there was a clear need to pair this with more structured approaches to investigation of user interaction with knowledge-exchange web applications. The latter would need to be both direct and indirect. Direct, in the sense that an imposed experimental design could help to model interaction processes relating to the rating and credibility of online testimony, and indirect, in that it was necessary to better understand the impact of online testimony on the user through accompanying survey and content analysis.

For the structural study of SQA [4], a pipeline approach was used which took the public domain dataset from the Stack Overflow SQA service, uploaded this to a database store, then used the *R statistics* platform to sample, analyse and visualise question and answer patterns and moderation activity. The dataset consisted of just under 6.5 million posts (questions and answers) from 800,000 users. For many of the analyses, pre-processing of the data was required and was achieved using database procedures before being imported into R. As the dataset was so large, random sampling was used for several of the analyses.

To accompany the structural work, additional content analysis was conducted on the discussion pages populated by members of the community, where “meta” conversations take place about the functioning of the site and norm and moderation issues. The researcher was additionally able to gain extra insight into the community through participating himself in questioning and answering activity thereby adding a small participant-observer dimension to the method.

For the study of SQA users [5], level 2, 3 and postgraduate students at UWE were recruited as subjects. An experimental interface was created, using questions and answers from the StackOverflow programming SQA service (and also its sister sites for Library Services and Philosophy). Students were presented with question and answers and asked to rate their own knowledge on the topic and the credibility of the proffered answers. Two conditions were used, one in which extra credibility cues were visible and another where they were hidden, in order to determine their impact on credibility ratings. Additionally, qualitative comments were collected to complement quantitative ratings. These provided additional insight into credibility judgements.

Data generated by [4] and [5] was analysed in R and subsequently anonymised and made available on an online repository. This was consistent with the licensing of the Stack Overflow dataset and with the subjects' agreement as obtained through participation consent forms.

Discussion and Main Contributions

This section will highlight some of the main theses and conclusions of papers [1-6]. I have grouped these into: *Knowledge Patterns and Practices*, which highlights the main contribution of the synthesis work in papers [1] and [6] *Normative SE*, which looks at the interpretation of SE for the evaluation of knowledge exchange focusing mostly on [2] *Naturalistic Triangulations*, looking at links between SE and psychological/cognitive theories and based on work in [1,2,6] *Spatio-temporal Aspects*, focusing on [4] *Credibility Cues and Bias*, referencing [5] and finally, *Identifying Interventions*, which references the full body of work and highlights conclusions and recommendations for practice.

In terms of their bearing on SE, the work consists of both descriptive and normative enquiry. In all papers, primary or secondary descriptive evidence – which considers how knowledge exchange operates at the participant level – is accompanied by normative analysis of knowledge sharing efficacy and recommendations for improving systems to become more epistemically valuable. Of the papers submitted [2] and [3] are perhaps the most explicitly and intentionally normative, while [1] and [6] seek normative recommendations based on a range of empirical evidence. The primary empirical work in papers [4] and [5] could be considered simultaneously descriptive and evaluative in that those paper attempt to ground assessments of system suitability against evidence of the operational dynamics of SQA and the impact of SQA systems on users.

Knowledge patterns and practices in online knowledge sharing environments

The first parts of the submitted work sought to bring together evidence on knowledge exchange from information science, psychology and philosophy and derive some overarching principles and patterns. One such principle foregrounded in [1] was the research evidence pointing to the important role of socio-emotional support to knowledge sharing online:

“The focus on sociability has shown that in presenting evidence and information, the manner, language and structures used are as important as the content in impacting on pervasive belief.” [1, p550]

This was more recently expanded in [6] which links such support to the characteristics of effective online knowledge brokerage. These aspects are notably lacking from much of SE theory and information science work on knowledge exchange, which tends to focus on content rather than the *caring* delivery of knowledge in testimonial exchanges.

What is perhaps distinct here is that in the online environment these socio-emotional effects need to

be *more* pronounced due to the absence of other interpersonal cues. This is coupled with the need to also provide a more transparent warrant to compensate for relative anonymity or pseudonymity:

“if anything, the message with online is to provide even more evidence of your bedside manner as the nonverbal cues are lacking, and (to switch from medical to mathematical metaphors) to show even more of your working.” [1, p545]

A further key contribution of [1] was to propose *knowledge patterns*, construing patterns in much the same way that they have been thought of in the architecture and software domains, being named solutions to recurring design problems and challenges. The following patterns were recognised in [1] as epistemically significant, and presented as patterns and *anti-patterns* — those with harmful consequences.

Mostly positive:

- Reputation indicators as credibility cues
- Threading / Braiding - objective synthesis of multiple viewpoints into a single narrative
- Conversation and argumentation
- Narrative and personal testimony
- Collaborative filtering (voting and ranking)
- Repackaging - essentially the same information but reposted or repurposed to appear in a different context.

Often negative:

- Echo chambers / filter bubble
- Persuasion and influence
- Sophistry - trolling and provocation
- Homophily - the tendency of online social networks to coalesce around people of similar race, gender and socioeconomic status.

Paper [1] supplies caveats for positive patterns in that certain conditions must be met for them to result in positive outcomes from online exchanges. In the case of argumentation, for instance, the ideal conditions for veristically positive argumentation are laid down by Goldman (1999). These

include that the speaker must believe premises and conclusions, and must be justified in these beliefs, but also that there is potential for the audience to understand and properly evaluate arguments and be swayed by them. Goldman cites Grice's conversational maxims and argues that, like conversation, there are tacit "folk" rules that reinforce positive cooperative endeavours. Clearly though, where such a sense of community is more fragile, or other motivations than cooperation take precedence, then argumentation may more resemble sophistry and become a negative pattern - even though it uses the same devices as sound argumentation.

The patterns described above were revisited in subsequent papers, notably [5] on reputational indicators and voting, and [6] on threading and synthesis. Some tally well with existing epistemological analyses of online systems, though others have not necessarily received significant treatment from an SE perspective.

Normative SE for the evaluation and design of online systems

Perhaps one of the most significant contributions to the literature relating SE theory to online knowledge was paper [2]. This work was a collaboration with Judith Simon, a philosopher at the University of Vienna whose PhD — completed in 2010 — was on SE and socio-technical systems (Simon 2010) and was one of the few authors publishing on this topic. This paper focused on online communities and sought to first review evidence of knowledge-related community dynamics, then to draw out recommendations for community design and interaction based on normative SE theory from Alvin Goldman, Helen Longino and Miranda Fricker [2]. These various influences helped to shape a useful view of SE that had individualistic and communitarian aspects and also combined analytic with naturalistic approaches.

A notable assumption in this work was the similarity of scientific knowledge to general *folk* knowledge. While some SE theorists have focused particularly on knowledge transmission in scientific communities, we wanted this to apply more widely to different types of knowledge exchanged in online communities. There were quite compelling reasons for us to make this step — including the evidence from cognitive psychology for the similarity between personal hypothesis formation and testing and the scientific method (of which, more below).

A further point we made in [2], also alluded to in [1], was to outline the limitations of Goldman's view of information, which we felt limited the testimonial exchange to one of message transmission, missing out important conditions relating to the relative status and relations of the interlocutors. This shortcoming we felt was met by the work of Longino and Fricker, focusing as it does more fully on

questions of authority and bias.

Based then on our reading of the SE literature the following guidelines for online knowledge exchange were developed in [2]:

- Avoidance of detrimental bias, both biases crystallised 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.

We concluded that this SE evaluation is essential to the productive functioning and epistemic design of online community:

“Systems need to be designed in ways to minimise the alpha and beta error of epistemic injustice: withholding epistemic authority where it would have been appropriate and attributing epistemic authority where it is not warranted.” [2, p55]

We further proposed that the best approach to realising some of these SE-derived desiderata was through *values in design*, the overt effort to encapsulate a philosophical standpoint in a community design process.

Triangulations

It is worth highlighting two significant cases of *naturalistic triangulation* — the novel connection of theory from different domains — notable in the submitted work.

Firstly, the connection of *personal epistemology* theory with Kruglanski’s *need for cognitive closure (NFCC)* highlights the tension between searching for quick answers and a deepened understanding of the complexity and contingent nature of much human knowledge:

“Following Kruglanski’s (1990) analysis, we also need to become adept at recognising and diffusing frustration associated with the relative novice discovering complexity and contradiction when being driven by the urgency for epistemic closure, a pressure only further exacerbated by the immediacy of the online medium. In this way, knowledge seekers can be “managed up” the chain in their own personal epistemology” [1, p 545]

According to various stage models in personal epistemology, learners develop from a simplistic, dualist approach to knowledge to a more constructivist position appreciative of argument dialectics (Hofer and Pintrich 1997). An interesting question is whether the various models of online knowledge exchange encourage or hinder such development. This idea is picked up in [6] and is worthy of further research, as online measures that help to mitigate the biases found in high NFCC epistemic behaviour would be particularly significant.

A second case was the coherence of hypothesis formulation and testing in knowledge building that is part of both Goldman’s SE and Kruglanski’s lay epistemics. Both theoreticians take a Popperian view that all rational (lay) agents exhibit a form of scientific thinking in relation to belief formation and that falsification is an essential part of this:

“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” [2, p42]

This seems coherent also with Goldman’s Bayesian approach to belief and with my own experience of the reported influence of new online testimony on our personal belief systems. The patchy nature of evidence and the need to trust others on many matters helps to explain why such *abductive* reasoning may be much more common than often thought, both in everyday life and the scientific method (Douven 2011).

Spatiotemporal aspects in knowledge sharing

As the work focused in on SQA as an exemplar of online knowledge exchange, it became increasingly clear that both spatial and temporal factors were essential in accounting for the patterns of seeking and sharing that were observed. These seemed to be areas where SE had had relatively little to say, but had been highlighted as central in the knowledge management literature (e.g. Nonaka 2005). There seemed to be important connections, too, between the types of bias and exclusion mentioned in the SE evaluation above, and the role played by the constraints of the SQA interface in shaping and aggregating knowledge. A snapshot of an SQA resource, as we might see it at any

moment on the web, does not provide a full picture of the dynamics behind it and the way it has developed over time. This is why the work outlined in [4] was undertaken, to better understand and represent the temporal patterns inside an SQA platform.

Paper [4] helped to illustrate that the *first-mover* advantage, identified in other collaborative knowledge sharing web sites, was also present in SQA sites and that this had a lasting effect on resource quality and the way in which the testimony was represented to subsequent consumers. That is, the effect of the SQA exchange is to *fix* knowledge in time, after which there is some gradual, additional enhancement, though this rarely is enough to supplant the initial content. The drivers for this are the community reward structures, in addition to the ephemeral attention received by *active* question threads. In Goldman's veritistic terms, then, there seems to be a systemic constraint to the ability for the aggregate V-value (truth value) of the system to increase far beyond its starting value.

Goldman does consider two dimensions of knowledge systems of relevance to this analysis, being that of *speed* of knowledge acquisition and of the *interest level* of the individual members in a knowledge community. In these regards, the work in [4] points to the suitability of SQA. The speed of responses — though it entails the fixation issue — does make for efficient knowledge transfer in many cases. Furthermore, the use of semantic tagging to flag question topics gives users the ability to hone in on questions and answers of interest, thereby addressing Goldman's interest requirement. The thematic partitioning of the SQA resource is indeed a very powerful feature allowing for the distribution of knowledge and for the better targeting of queries.

In addition to the timing of question and answer threads, paper [4] also considered the voting patterns and cumulative votes received for a range of types of question in the SQA system under study. It was notable that voting behaviour was *bursty* (periods of inactivity followed by bursts of voting from different users), partly due to the varying visibility of the question on the web site, but probably also due to other factors that it was not possible to investigate in this particular study (e.g. that there was some off-site sharing of answers that served to garner votes). What became clear was that there were a considerable number of questions attracting votes from the general community over and above the few individuals providing the question and answers. Together with the reputation scores of community members, the voting totals became the focus of the subsequent work. Voting and reputation systems in the real world had been treated to epistemological analysis and it was therefore useful to see how this translated to the online medium.

Credibility cues and bias in SQA

Given prior work on the credibility of web pages from various sources, an interesting research question was how far these findings generalised to social and collaborative web systems, including SQA. This question also had a direct bearing on the nature of testimony in online systems and what features of another's testimony are attended to online. Some evidence from other work had already highlighted the potential importance of peripheral cues, particularly in conditions of cognitive load or distraction. But [5] was the first known work to take an experimental approach to credibility in SQA.

The empirical work reported in [5], then, looked at the effect of credibility cues — namely voting score and user reputation — on the subjective impact of answers posted to questions in SQA. Again, the particular systems studied were from the StackExchange system of SQA. In this research, students were given question and answer pairings, with the credibility cues visible or invisible in order to discern the impact that the cues had on overall ratings of credibility.

Findings showed that the credibility cues were significant in promoting a *bandwagon effect* making users more likely to up-vote what were already highly scoring answers. It was shown, however, that this effect was largely positive and useful in promoting what tended to be rated independently as higher quality answers. That said, there was clearly a polarising effect brought about by the voting behaviour, as without the explicit cues being present, quality judgements would have been more evenly spread. We can therefore see this type of SQA as open to the *filter bubble* criticism of collaborative filtering (Pariser 2011), whereby users tend to be drawn to the same sources and to ignore other valid yet poorly signposted sources. The filter bubble has undesirable epistemic consequences which are counter to norms of testimony (that each voice should have an equal chance of being heard) — an example of Fricker's testimonial injustice (Fricker 2007)

Identifying interventions

We have seen that the approach to SE taken here was strongly normative and evaluative. Through the description and critique of knowledge exchange systems, it was desirable to arrive at conclusions and recommendations for their improvement. All of the submitted papers attempted to do this, with the type and level of intervention changing as the analyses progressed from general to more specific aspects on knowledge exchange and with the emerging focus on SQA. Interventions discussed in the submitted work thus included enhanced *institutional* contributions, *community* level norms, *interface/interaction* mechanics and *individual cognitive development*.

With their publication in library and information science journals, [1] and [6] especially spoke to the

information profession, and called for more proactive involvement in online communities, given the skills that the professional can bring to the summary of topics and the provision and signposting of evidence. Indeed, sources quoted in [6] showed where such activities are already happening — albeit to a limited extent and not always sanctioned by employees’ institutions. There remains an enormous amount that can be done here in bridging the gap between online knowledge seekers and institutional repositories, requiring something of a step change and a recognition of the shift in media reliance among the majority of consumers.

It has been noted that social design aspects can be even more important than technical design for the successful functioning of online communities and social design is the stuff of SE actuation. Paper [2] therefore concluded that encouraging the development of norms to enable unbiased, fair and diverse exchange of knowledge is perhaps the most important intervention that can be made in the online sphere. That paper pointed to some examples and [6] noted how the demonstration of norms such as balance, objectivity and verifiability might be better rewarded through interface features.

Studies of SQA have noted the impact of the *gamification* of the interface on community dynamics and resource quality. Gamification is largely used — mostly to good effect — to encourage engagement and an ongoing commitment to an epistemic community. As there can clearly be some perverse consequences of this gamification, it is also likely that different kinds of gamification can be used to stimulate positive epistemic behaviour:

“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” [2, p56]

This is a ripe research area but one where experimentation has really been confined to the community managers / designers, who are motivated chiefly by engagement, although quality and positive epistemic traits certainly feature to varying degrees. Given the importance of bias mitigation noted in [2], a great deal of new research on this from psychology has relevance to interaction design in epistemic systems. Some small examples have been highlighted in [5]. Given that many web-based interaction experiments are driven toward persuasion and a *conversion* or sale, we might use the same techniques to maximise positive epistemic qualities.

A further strong trend that is in need of countering and noted above, is that toward absolutism in personal epistemology and knowledge seeking online. Interface patterns (and community norms) might provide a means to push back against certainty:

“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 practised, that may serve to make an example of unbiased knowledge-sharing

behaviour.” [2, p52]

This needs to be connected also to affective factors associated with different stages of the information seeking process and which have been identified in established models, but little researched (Savolainen 2014). Here, the crucial stage for intervention is when the seeker is the transition from *pre-focus* (becoming informed) to *focus* (predicting an outcome) in topic exploration.

Finally, the individual’s perception of the online knowledge resource and its relative impact is certainly one of the most interesting dimensions to the research, albeit one that was only developed in a tentative way chiefly in [5]. Here, an awareness of how an epistemic system works — what I have since been terming *algorithmic metacognition* — was proposed as an important counteraction to the kind of algorithmic, filter-bubble bias that has been frequently observed. I have noted this in myself, that since better understanding SQA dynamics and bias, I am a more informed and conscientious user of these systems. This seems to be wholly consonant with an extended / distributed cognitive environment and extends the self-monitoring aspect also beyond the individual. Again, here is a research area that is under-emphasised but hugely important as we offload more and more cognitive responsibilities onto the web.

Conclusions

The research submitted herewith is an attempt to relate social epistemological theory with the practice of internet-mediated knowledge sharing, taking into account relevant evidence from information science, psychology and human-computer interaction. The research has focused specifically on knowledge sharing as conducted in social question-answering services. Both through the review and integration of research evidence, and later some empirical investigation, the work has attempted to characterise online knowledge behaviour and to assess the extent to which it meets criteria developed within SE for ensuring that systems have optimal designs for creating, storing and transmitting true belief.

We have seen that the work painted a mixed picture of the reality of knowledge exchange on systems such as SQA. While the mechanisms that have evolved to stimulate engagement, contribution and collaboration in many ways serve their purposes well, a variety of biases and sub-optimality — in knowledge terms — remain. These in many ways illustrate the tensions between needs for immediacy and *satisficing* and the *higher* ideals of epistemic quality. At the same time I believe that I have also illustrated the importance of an ongoing programme of research that highlights epistemic quality over and above informational transactionalism.

Core goal / contribution

My experience of taking this research to various research and design communities has shown that information scientists and designers *appreciate* philosophy and attempts to apply its theory in the way I have done. I believe this is because such analysis goes beyond social and behavioural to core human goals and values. I therefore believe that this research has made a contribution through drawing attention to thought in social epistemology and through this, strengthening the link between ideal epistemic design and socio technical practice.

Limitations and Further work

Limitations

The submitted research has admitted limitations relating to the cross-disciplinary nature of the research and some of the specific methods of the submitted papers.

A cross-disciplinary research project such as this will suffer inevitably from an increased risk of missing significant areas of the literature of potential use and relevance. The most heavily monitored literature during the research was that of *information science*, which features much of the current SQA work as well as work on online communities and knowledge sharing more generally. As a result of this information science bias, it may well be that work from psychology and philosophy that would have proven useful was under-exploited.

A significant area that was under-exploited was that of formal social epistemology. Here, testimonial exchange and degrees of belief under different conditions can be modelled and compared with real-world behaviour and phenomena such as opinion spread and judgement aggregation. This can provide an arguably more thorough and testable approach to epistemology that might be more directly relatable to systems design. This approach will be discussed further in the next section on future work.

Other limitations relate to specific studies. With the empirical work in [5], subjects were drawn from a student population and, while the sample size was quite good, greater numbers of postgraduates and students from subjects besides computing would have helped to balance the data gathered. While the use of student populations is often criticised as affecting validity, in this case the participants represent a key audience for the SQA platforms — as learners who frequently use them to solve problems and develop their knowledge.

With the data analysis work undertaken in [4], issues were encountered in relation to pre-treatment of the data prior to public release. This meant that some of the temporal aspects were hard to model, due to deliberately introduced obfuscation. Without primary ownership of the data, however, this would be hard to improve upon. Moreover, as the accuracy and completeness of the data increases, so do concerns about privacy and ethical use of the data.

In terms of papers [1],[2] and [6], in taking a *position* rather than *review* paper approach they could be seen to be less structured and thorough than systematic reviews of available evidence and that they were driven by *a priori* conclusions. While it is certainly true that they took a less precisely prescribed approach to the choice of material, I would argue that no encountered evidence has been knowingly

ignored in these papers and that their strength lies in representing an interpretation formed by a reading of a broad cross-section of relevant literature from a number of disciplines.

Further work

A range of possibilities exist for taking the submitted research further. The work presented in [5], on the credibility impacts of the social interface, can be explored in a number of ways. The interaction of content and presentational aspects of the interface can be investigated through further controlled experiments — and this might be extended to incorporate not only existing credibility cues but suggested additional cues. Further longitudinal effects might also be tested, such as the test-retest reliability of credibility ratings and the longer term impact of SQA usage on domain knowledge. Furthermore, the limitations in the user population mentioned above can be addressed by using larger samples drawn from a more diverse demographic.

The development of personal, social and systemic metacognition, mentioned above, is a further interesting avenue for exploration and one where other researchers have also pointed to a research priority (Chan, Li et al. 2014). While questionnaire instruments exist for exploring personal epistemologies, these have been confined to a view of knowledge gained through individual learning and less so through social interaction online. Enhancement of these instruments to incorporate knowledge of community dynamics, sources of online bias and ranking and filtering algorithms would provide a way of benchmarking digital literacy and determining the impact of education and familiarisation interventions.

A potentially fruitful source of evidence is in the agent-based simulation of epistemic communities. Here, community dynamics, moderation and social interaction can be modelled and outcomes compared, both to theoretical outcomes and to the kind of empirical findings discussed above. The application of simulation can reveal both intended and unintended consequences of the application of particular combinations of norms and interaction constraints and has already shown its value in areas such as the epistemology of scientific communities (Mayo-Wilson 2014 e.g.). This kind of analysis might serve to improve the connection between analytic, formal epistemology and prescriptive system design.

A related kind of analysis would take advantage of the game-like qualities of knowledge exchange. Given the importance of gamification features in online systems such as SQA, game-theoretical analyses of online communities such as (Jain, Chen et al. 2009) can help to shed light on optimal incentives for achieving the communities goals by way of reward systems and feedback mechanisms.

As established in [2], the formation of community norms is an important dimension to the establishment of epistemic communities. An interesting research question here is how norms vary

given a relatively stable or uniform interface and platform. With the Stack Exchange SQA service studied in [4] and [5], it will be possible to answer this by studying variations in the way the platform is implemented between communities of interest. I do not know of any prior work on this but can see it providing a useful comparison of the kind identified as a research priority in the workshop discussions of [3].

Finally, improving the involvement and abilities of information professionals in online knowledge exchange is a strong theme of the research that deserves to be taken further. While the work has helped to identify opportunities and interesting existing pilot projects, more can be done to investigate how to measure the impact of these projects and how skill sets need to be adapted or supplemented to enable high quality contribution to online services, along with moderation and community building. There are a number of challenges here, notably how to justify the contribution of particular national institutions to a more global, online sphere; how to identify the platforms, communities and topics that deserve strategic targeting; and how to assess the impact of interventions. What is clear is that online intervention provides a clear opportunity for a profession making the transition from the physical and situated to the digital and distributed.

Submitted papers

1. Matthews, P. and Stephens, R. (2010) Sociable knowledge sharing online: philosophy, patterns and intervention. *Aslib Proceedings*, 62 (6). pp. 539-553. ISSN 0001-253X
2. Matthews, P. and Simon, J. (2012) Evaluating and enriching online knowledge exchange: A socio-epistemological perspective. In: Lazakidou, A., ed. (2012) *Virtual Communities, Social Networks and Collaboration*. Springer, pp. 35-59. ISBN 9781461436348
3. Matthews, P. (2013) Epistemological Dimensions of Q&A. CSCW 2013 Workshop on Social Media Question Asking

<http://research.microsoft.com/en-us/events/cscw2013smqaworkshop/papers.aspx>
4. Matthews, P. (2014) Knowledge fixation and accretion: Longitudinal analysis of a social question-answering site. *Journal of Documentation*, 70 (5). pp. 711-733. ISSN 0022-0418
5. Matthews, P. (in press) Use of credibility heuristics in a social question-answering service. Presented at Information Seeking in Context conference - ISIC 2014 and forthcoming in *Information Research*
6. Matthews, P. (in press) Search delegation, synthesists and expertise on social media. Presented at workshop on Social Media Information Seeking, ISIC 2014 and forthcoming in LIBRES, Library and Information Science Research.

Administration Statements

Joint Authorship

[1] Matthews and Stephens, is approximately 80% the work of the first author

[2] Matthews and Simon, is approximately 60% the work of the first author

Training Requirement Statement

Year	Module	Outcome	Credits
2011	UFMEL9-20-M Statistics for Research	70 Pass	20
2008	UFPEL6-15-M Research Writing	70 Pass	15
2012	UFCE3J-20-3 Agent- based Computing	67 Pass	20
2013	UFMEL3-15-M Research Portfolio	75 Pass	15
		TOTAL	70

Online Courses

Title	Provider	Year	Mark
Introduction to	University of	2013	69.3

Philosophy	Edinburgh		
Introduction to Mathematical Philosophy	Ludwig-Maximilians-Universität München (LMU)	2013	83.3
Introduction to Data Science	University of Washington, USA	2013	100
Web Science	University of Southampton	2013	N/A

References

- Adamic, L., Zhang, J., et al. (2008) Knowledge Sharing and Yahoo Answers: Everyone Knows Something. In: Anon. *17th International World Wide Web Conference (WWW2008), Beijing, China 2008*. ACM, Available from: <http://www2008.org/papers/fp840.html> [Accessed June 2009].
- Alvargonzález, D. (2011) Multidisciplinarity, Interdisciplinarity, Transdisciplinarity, and the Sciences. *International Studies in the Philosophy of Science*. 25(4), 387-403.
- Audi, R. (1998) *Epistemology: a contemporary introduction to the theory of knowledge*. London New York: Routledge.
- Bandura, A. (1994) Self-efficacy. In: V.S. Ramachandran, ed. *Encyclopedia of human behavior*. Vol. 4 ed. New York: Academic Press.
- Bawden, D. (2009) Smoother pebbles and the shoulders of giants: The developing foundations of information science. In: A. Gilchrist, ed. *Information science in transition*. London: Facet.
- Berners-Lee, T., Hall, W., et al. (2006a) Creating a Science of the Web.
- Berners-Lee, T., Hall, W., et al. (2006b) A Framework for Web Science. *Foundations and Trends in Web Science*. 1(1), 1-130.
- Brendel, E. (2009) Truth and Weak Knowledge in Goldman's Veritistic Social Epistemology. *Grazer Philosophische Studien*. 79(1), 3-17.
- Chan, R.Y., Li, S., et al. (2014) Social Epistemic Cognition in Online Interactions. In: Anon. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Toronto, Ontario, Canada* New York, NY, USA: ACM, pp 3289-3298. Available from: <http://doi.acm.org/10.1145/2556288.2556977>.
- Clark, A., and Chalmers, D. (1998) The Extended Mind. *Analysis*. 587-19.
- Code, L. (2010) Testimony, advocacy, ignorance: Thinking ecologically about social knowledge. In: A. Haddock, A. Millar and D. Pritchard, eds. *Social Epistemology*. UK: Oxford.
- Cornelius, I. (2004) Information and Its Philosophy. *Library Trends*. 52(3), 377-386.
- Douven, I. (2011) Abduction. In: Edward N. Zalta, ed. *The Stanford Encyclopedia of Philosophy*. [Online]. Spring 2011 ed.
- Dror, I. and Harnad, S. (2008) Offloading cognition onto cognitive technology. In: I. Dror and S. Harnad, eds. *Cognition distributed: how cognitive technology extends our minds*. Amsterdam Philadelphia: John Benjamins Pub.
- Egan, M.A., and Shera, J.H. (1952) Foundations of a Theory of Bibliography. *Library Quarterly*. 22(2), 125-137.
- Fallis, D. (2009) Introduction: The Epistemology of Mass Collaboration. *Episteme*. 6(1), 1-7.
- Fallis, D. (2006) Social epistemology and information science. *Annual Review of Information Science and Technology*. 40(1), 475-519.

- Fantl, J. (2014) Knowledge How. *In*: Edward N. Zalta, ed. *The Stanford Encyclopedia of Philosophy*. [Online]. Fall 2014 ed.
- Floridi, L. (2011) *The philosophy of information*. Oxford New York: Oxford University Press.
- Floridi, L. (2004) Open Problems in the Philosophy of Information. *Metaphilosophy*. 35(4), 554-582.
- Floridi, L. (2002) On defining library and information science as applied philosophy of information. *Social Epistemology*. 16(1), 37-49.
- Fricker, M. (2010) Scepticism and the genealogy of knowledge: Situating epistemology in time. *In*: A. Haddock, A. Millar and D. Pritchard, eds. *Social Epistemology*. UK: Oxford.
- Fricker, M. (2007) *Epistemic injustice. Power and the ethics of knowing*. Oxford: Oxford University Press.
- Fuller, S. (2013) Deviant interdisciplinarity as philosophical practice: prolegomena to deep intellectual history. *Synthese* 190(11), 1899-1916.
- Gilbert, D. (1991) How Mental Systems Believe. *American Psychologist*. 46(3), 107-119.
- Goldman, A.I. (2010) Systems-oriented social epistemology. *In*: T.S. Gendler and J. Hawthorne, eds. *Oxford Studies in Epistemology, Volume 3*. Oxford: Oxford University Press.
- Goldman, A.I. (1999) *Knowledge in a Social World*. Oxford.
- Halavais, A. (2011) Social science: Open up online research. *Nature*. 480(7376), 174-175.
- Hendler, J., Shadbolt, N., et al. (2008) Web Science: An Interdisciplinary Approach to Understanding the Web. *Commun.ACM*. 51(7), 60-69.
- Hodgson, P. and Cofta, P. (2009) Towards a Methodology for Research on Trust. *In*: Anon. *Proceedings of the WebSci'09: Society On-Line, Athens, Greece 2009*. Available from: <http://journal.webscience.org/164/> [Accessed 10/24/2014].
- Hofer, B.K., and Pintrich, P.R. (1997) The Development of Epistemological Theories: Beliefs about Knowledge and Knowing and Their Relation to Learning. *Review of Educational Research*. 67(1), 88-140.
- Homans, G.C. (1958) Social Behavior as Exchange. *American Journal of Sociology*. 63(6, Emile Durkheim-Georg Simmel, 1858-1958), 597-606.
- Hsu, M., Ju, T.L., et al. (2007) Knowledge sharing behavior in virtual communities: The relationship between trust, self-efficacy, and outcome expectations. *International Journal of Human-Computer Studies*. 65(2), 153.
- Huaping, W., and Xiaoming, S. (2007) Cooperative Naturalism. *Frontiers of Philosophy in China*. 2(4), 601-613.
- Humphreys, P. (2009) Network Epistemology. *Episteme*. 6(2), 221-229.
- Jain, S., Chen, Y., et al. (2009) Designing Incentives for Online Question and Answer Forums. *In*: Anon. *Proceedings of the 10th ACM conference on Electronic commerce, Stanford, California, USA* New York, NY, USA: ACM, pp 129-138. Available from: <http://doi.acm.org/10.1145/1566374.1566393> .

- Kahneman, D. (2011) *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Kim, S., and 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), 716-727.
- Kosonen, M. (2009) Knowledge sharing in virtual communities a review of the empirical research. *International Journal of Web Based Communities*. 5144-163.
- Kruglanski, A.W. (1990) Lay Epistemic Theory in Social-Cognitive Psychology. *Psychological Inquiry*. 1(3), 181-197.
- Kruglanski, A.W., Dechesne, M., et al. (2009) Three decades of lay epistemics: The why, how and who of knowledge formation. *European Review of Social Psychology*. 20146-191.
- Kusch, M. (2001) 'A general theory of societal knowledge?': Aspirations and shortcomings of Alvin Goldman's social epistemology. *Studies in History and Philosophy of Science Part A*. 32(1), 183-192.
- Lackey, J. (2008) *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press, USA.
- Ma, M., and Agarwal, R. (2007) Through a Glass Darkly: Information Technology Design, Identity Verification, and Knowledge Contribution in Online Communities. *Information Systems Research*. 18(1), 42-67.
- Magnus, P.D. (2009) On Trusting Wikipedia. *Episteme*. 6(1), 74-90.
- Mayo-Wilson, C. (2014) The Reliability of Testimonial Norms in Scientific Communities. *Synthese*. 191(1), 55-78.
- Monnin, A. and Halpin, H. (2014) Toward a philosophy of the web: foundations and open problems. In: H. Halpin and A. Monnin, eds. *Philosophical Engineering: Toward a Philosophy of the Web*. Wiley Blackwell.
- Nilsson, N.J. (2014) *Understanding beliefs*. Cambridge, Massachusetts: MIT Press.
- Nonaka, I. (2005) *Knowledge management: critical perspectives on business and management*. [Online]. New York: Routledge.
- Origg, G. (2012) A Social Epistemology of Reputation. *Social Epistemology*. 26(3-4), 399-418.
- Pariser, E. (2011) *The filter bubble: what the Internet is hiding from you*. New York: Penguin Press.
- Pritchard, D., Millar, A., et al. (2010) *The nature and value of knowledge: three investigations*. [Online]. Oxford New York: Oxford University Press.
- Quine, W.V. (1969) *Ontological relativity, and other essays, by W. V. Quine*. [Online]. New York, Columbia University Press.
- Russell, B. (1992) *Human knowledge: its scope and limits*. London: Routledge.
- Savolainen, R. (2014) Approaching the Affective Factors of Information Seeking: The Viewpoint of the Information Search Process Model. In: Anon. *Proceedings of ISIC, the Information Behaviour conference*, Leeds: pp 444-460.
- Schiltz, M., Truyen, F., et al. (2007) Cutting the Trees of Knowledge: Social Software, Information Architecture and Their Epistemic Consequences. *Thesis Eleven*. 89(1), 94-114.

- Shera, J. (1973) Toward a theory of librarianship and information science. In: J. Shera, ed. *Knowing Books and Men Knowing Computers, Too*. Libraries Unlimited, Inc.
- Shermer, M. (2012) *The believing brain : from spiritual faiths to political convictions : how we construct beliefs and reinforce them as truths*. London: Robinson.
- Simon, J. (2010) *Knowing Together: A Social Epistemology for Socio-Technical Epistemic Systems*. Wikipedia. PhD, University of Vienna. Available from: <http://othes.univie.ac.at/10285/>.
- Simpson, T.W. (2012) Evaluating Google as an Epistemic Tool. *Metaphilosophy*. 43(4), 426-445.
- Speer, S. (2008) Natural and Contrived Data. In: P. Alasuutari, L. Bickman and J. Brannen, eds. *The SAGE Handbook of Social Research Methods*. [Online]. GB: Sage Publications Ltd.
- Terrell, S., and Steven R Terrell. (2012) Mixed-Methods Research Methodologies. *Qualitative report*. 17(1), Oct 2014-280. Available from: <http://www.nova.edu/ssss/QR/QR17-1/terrell.pdf>
<http://www.nova.edu/ssss/QR/QR17-1/terrell.pdf>.
- Tinati, R., Halford, S., et al. (2012) Mixing Methods and Theory to Explore Web Activity. In: Anon. *Third Annual Web Science Conference (WebSci2012), June*. Available from: <http://eprints.soton.ac.uk/344805/>.
- Tollefsen, D.P. (2009) Wikipedia and the Epistemology of Testimony. *Episteme*. 6(1), 8-24.
- Vähämaa, M. (2013) Groups as Epistemic Communities: Social Forces and Affect as Antecedents to Knowledge. *Social Epistemology*. 27(1), 3-20.
- Wilson, T. (2002) *'Information science' and research methods* Wikipedia Available from: <http://www.informationr.net/tdw/publ/papers/slovak02.html>[Accessed 10/10/2014 2014].
- Wrenn, C.B. (2004) *Naturalistic Epistemology* Wikipedia Available from: <http://www.iep.utm.edu/nat-epis/>[Accessed 10/2/2014 2014].
- Zagzebski, L. (1998) What is Knowledge. In: J. Greco and E. Sosa, eds. *The Blackwell Guide to Epistemology*. Blackwell.