THE GENERATION OF KNOWLEDGE · EPISTEMOLOGY

Social Epistemology: Knowledge as a Collective Achievement

Most knowledge is produced, checked and held by groups, not by isolated individuals. This article explains social epistemology, from trust and the division of cognitive labour to group belief, diversity and communication networks, and what it means for organisations that generate knowledge.

ModQu research article · 2,885 words · 27 references (APA 7th edition)

For much of its history, the theory of knowledge pictured a single mind facing the world. Descartes, alone by his stove, doubted everything and tried to rebuild knowledge from his own resources. The picture is appealing, but it describes almost nobody's actual situation. A modern scientist relies on instruments built by others, data collected by colleagues, results published by strangers and methods developed over centuries. A business decision draws on market research, accounting systems, expert advice and the experience of many employees. Knowledge, in practice, is a collective achievement.

Social epistemology is the study of the social dimensions of knowledge: how people acquire knowledge from one another, how groups come to know things, and how social structures and institutions help or hinder the pursuit of truth. Goldman (1999) defined it as the study of the social practices that promote or impede true belief. This article introduces its main themes and draws out their implications for organisations whose business is the generation of knowledge.

From individual to social

Dependence on others

The starting point of social epistemology is the recognition that individuals depend on others for most of what they know. Hardwig (1985) put the point sharply. A physicist who relies on a colleague's experimental result, which they lack the specialist skill to check, nonetheless has good reason to believe it. A layperson who believes what the expert says, without understanding why, can be acting rationally. Hardwig called this **epistemic dependence** and argued that it is not a regrettable weakness but a basic feature of knowledge in a complex society. No one can check everything; the rational course is often to trust those who are in a better position to know.

Epistemic dependence raises the question of whom to trust. Goldman (2001) set out the evidence a layperson can use to choose between experts who disagree: the arguments they offer, the agreement of other experts, independent evidence of expertise, signs of bias or interest, and past track records. Trust in others is rational when it is calibrated to such evidence. It becomes a problem when it is given to the wrong people or withheld from the right ones.

The social history of knowledge

Historians and sociologists have shown how deeply social factors shape what counts as knowledge. Fleck (1979), in a study first published in 1935, described how the medical concept of syphilis developed within a "thought collective" whose shared assumptions shaped what its members could see. Kuhn (1962) argued that normal science takes place within paradigms, shared frameworks of theory and practice that are passed on through training. Shapin (1994) showed that in seventeenth-century England the credibility of a scientific report depended partly on the social standing of the person who made it; a gentleman's word carried weight that a servant's did not.

Some sociologists went further. Bloor (1976) proposed a "strong programme" in the sociology of knowledge, arguing that true and false beliefs alike should be explained by social causes. This approach was controversial, because it seemed to reduce knowledge to social agreement. Most social epistemologists take a middle position: social factors shape knowledge profoundly, but some social practices are better than others at producing true beliefs, and it is possible to say which.

Science as a social enterprise

The norms of science

The
sociologist Robert Merton described the institutional norms that, in his view, made science effective (Merton, 1973). **Communalism** requires that findings be shared openly. **Universalism** requires that claims be judged by impersonal criteria, not by the identity of the person making them. **Disinterestedness** requires that scientists act for the benefit of the shared enterprise rather than personal gain. **Organised scepticism** requires that all claims be subjected to critical scrutiny before they are accepted. These norms describe a social system designed to produce reliable knowledge from the work of fallible and self-interested individuals. Merton did not claim that scientists always live up to them. His point was that the institutions of science, such as publication, priority claims, citation and peer review, reward behaviour that serves the shared goal. A scientist who wants recognition must publish results so others can use them, and must expose those results to criticism from rivals with every incentive to find errors. Individual ambition is harnessed to collective knowledge.

Objectivity as a social achievement

Longino (1990) developed this idea into an account of scientific objectivity. Individual scientists, she argued, cannot escape their own assumptions and values. Objectivity is achieved not by individuals freeing themselves of bias but by communities that subject each other's work to criticism. Longino identified four conditions for such a community: recognised venues for criticism, such as journals and conferences; real uptake of criticism, so that views change in response to it; shared standards by which claims are judged; and **tempered equality** of intellectual authority, so that criticism from any qualified member is taken seriously. On this account, objectivity is a property of how a community is organised, not of individual minds.

Feminist philosophers of science have pressed related points. Haraway (1988) argued that all knowledge is produced from particular positions, and that objectivity is better served by acknowledging those positions than by pretending to a "view from nowhere". Harding (1991) argued that research is more objective when it begins from the perspectives of groups whose experiences have been neglected, because those perspectives reveal assumptions that dominant groups take for granted.

The division of cognitive labour

Science is also organised as a division of labour. Kitcher (1990) asked how a scientific community should distribute its efforts among competing research programmes. If every scientist pursued the approach that currently looks most promising, the community would put all its eggs in one basket and might never discover that a less promising approach was right. Kitcher showed that a community can benefit when individuals are motivated partly by credit and reputation, because some will then choose to work on neglected approaches where they have a better chance of making a distinctive contribution. The pursuit of individual recognition can, under the right conditions, produce a healthy distribution of effort for the community as a whole (Kitcher, 1993).

The shift to collaborative work makes this social dimension ever more visible. Wuchty et al. (2007) analysed millions of research papers and patents and found that teams increasingly dominate the production of knowledge in almost every field, and that team-authored papers are more highly cited on average than those by single authors. Wray (2002) argued that collaborative research is epistemically significant in its own right, because it allows research to be done that no individual could do alone, not merely because it pools labour.

Group knowledge

Can groups know things?

We often speak as though groups have beliefs and knowledge. "The committee believes the proposal is too expensive." "The company knew about the defect." Are such statements just shorthand for what individual members believe, or can a group believe something that none of its members does?

Gilbert (1987) argued that collective belief is a real phenomenon. A group believes something when its members have jointly committed to accepting it as the group's view, even if some members privately doubt it. A jury that returns a verdict, or a committee that adopts a finding, holds a belief in this sense. List and Pettit (2011) showed that groups can be agents in their own right, with attitudes that do not simply mirror those of their members. A famous result, the **discursive dilemma**, shows that a group that decides each premise of an argument by majority vote can end up accepting premises that imply a conclusion the majority rejects. To be consistent, groups need procedures for forming their views that go beyond simple aggregation of individual opinions.

Lackey (2021) offered a detailed theory of group belief and knowledge, arguing that a group's knowledge depends both on the beliefs of its members and on the evidence and processes that the group as a whole has access to and uses. On her account, a company can be responsible for knowing about a defect if the information was available within it and its processes should have brought it to the relevant decision-makers.

When groups are wiser than individuals

Under the right conditions, groups can know more than their members. The classic demonstration, retold by Surowiecki (2004), is Francis Galton's ox-weighing competition, in which the median of hundreds of guesses was almost exactly right. Surowiecki argued that such "wisdom of crowds" depends on diversity of opinion, independence of judgement and a good method of aggregation.

Diversity matters in a second way, through the range of approaches a group brings to a problem. Hong and Page (2004) used computational models to show that a group of diverse problem-solvers can outperform a group of individually more able but similar problem-solvers, because the diverse group explores more of the possible solutions. Woolley et al. (2010) found evidence for a general **collective intelligence** factor that predicts how well a group performs across a wide range of tasks. It was not strongly related to the average or maximum intelligence of members, but it was related to the social sensitivity of members and the equality of their contributions to conversation.

When groups are worse than individuals

Groups can also do much worse than their members. Janis (1972) described **groupthink**, in which cohesive groups under pressure suppress dissent and converge on poor decisions, as in several American foreign policy failures. Stasser and Titus (1985) found that when information is distributed among group members, discussions tend to focus on information that everyone already shares, and information held by only one member is often never mentioned. The group as a whole then decides as if the unshared information did not exist, even when it would change the decision. Sunstein (2006) reviewed evidence that deliberation can amplify errors, polarise opinion and silence minority views, and argued that methods such as prediction markets and anonymous aggregation can sometimes extract knowledge from groups more reliably than discussion.

Networks and communication

How information flows within a community also affects what it comes to know. Zollman (2007) used computer simulations of scientists choosing between two theories and sharing their experimental results. Surprisingly, communities in which everyone shared results with everyone else were **less** likely to settle on the correct theory than communities with sparser communication. In a fully connected network, a few early misleading results could quickly persuade everyone to abandon the better theory before it had been properly tested. In a sparser network, different parts of the community kept exploring different options for longer, so the community was more likely to find the truth, though it took longer to agree.

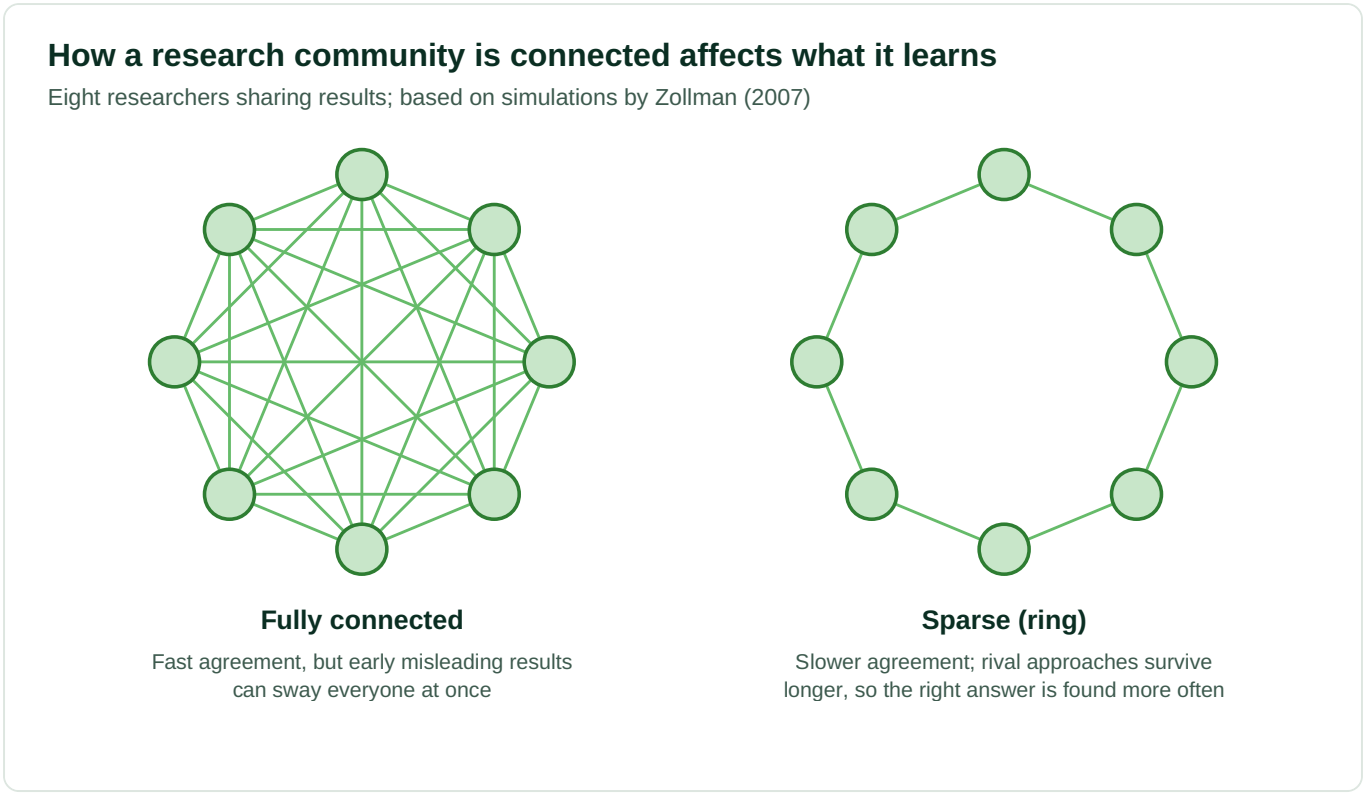


Figure 1. Communication structure and collective learning. In simulations by Zollman (2007), research communities with sparser communication were more likely to reach the correct answer than fully connected communities, because they preserved diversity of approach for longer. Diagram by ModQu.

The finding has a counterpart in modern information environments. Nguyen (2020) distinguished between **epistemic bubbles**, in which people simply do not encounter opposing views, and **echo chambers**, in which people are actively taught to distrust outside sources. Bubbles can be burst by exposure to new information; echo chambers are harder to escape, because members have been prepared to dismiss that information. Both show that the social structure of communication shapes what groups come to believe, for good or ill.

Fairness and knowledge

Social epistemology also has an ethical dimension. Fricker (2007) identified two forms of **epistemic injustice**. **Testimonial injustice** occurs when a speaker is given less credibility than they deserve because of prejudice, as when a woman's report is discounted because of her gender. **Hermeneutical injustice** occurs when a group lacks the shared concepts needed to make sense of its own experience, as before the term "sexual harassment" was coined. Both are wrongs to the people affected, and both are failures of knowledge: a community that discounts some voices or lacks the concepts to understand some experiences knows less than it could.

Surveys as social epistemology in action

Survey research is itself a social epistemic practice, and a useful illustration of the themes above. A survey collects testimony from many people and combines it into a collective picture that no individual could provide. Its reliability depends on exactly the factors social epistemologists study.

First, it depends on **trust calibrated to evidence**. Analysts must decide how far to accept what respondents report, using knowledge of where self-reports are accurate and where they are not. That is Goldman's (2001) problem of evaluating sources, applied to thousands of sources at once.

Second, it depends on **independence and aggregation**. A survey gains its power from combining many independent answers. If respondents influence one another, for example by discussing the questions before answering, or if the same people answer many times, the effective amount of independent information shrinks. The arguments of Surowiecki (2004) and Sunstein (2006) about when aggregation works apply directly.

Third, it depends on **inclusion**. A survey that reaches only some groups, or that uses language and concepts that fit only some people's experience, produces a collective picture that is partial in both senses: incomplete and biased. Fricker's (2007) analysis of hermeneutical injustice has a practical counterpart in questionnaire design. If the response options available do not describe a group's situation, that group's experience is lost from the data, however carefully the rest of the survey is done.

Finally, the knowledge produced by a survey is itself passed on as testimony, to clients, to the public and to other researchers. Its users depend on the research organisation in the way Hardwig (1985) described, usually without being able to check the work. That dependence places a responsibility on the organisation to be transparent about its methods, so that trust in its findings can be calibrated to evidence rather than given blindly.

Implications for knowledge-generating organisations

Social epistemology shows that the reliability of knowledge depends not only on the skill of individuals but on how they are organised. Several practical lessons follow for organisations that produce research, analysis or advice.

Build in organised scepticism. Merton's norm and Longino's account of objectivity both point to the same practice: findings should be exposed to criticism before they are accepted. In a research organisation this means internal review of questionnaires, analyses and reports by people who did not produce them, and a culture in which challenging a finding is welcomed rather than resented.

Keep criticism open to everyone qualified. Longino's condition of tempered equality is easily violated in hierarchical organisations. If a junior analyst's objection carries less weight than a senior manager's opinion simply because of rank, the organisation loses knowledge. The findings of Woolley et al. (2010) suggest that groups in which members contribute more equally tend to perform better.

Surface unshared information. The hidden-profile findings of Stasser and Titus (1985) warn that meetings tend to discuss what everyone already knows. Asking each person to state relevant information before discussion begins, or collecting views independently in writing, helps bring unshared knowledge into the open.

Protect independence before aggregating. When an organisation wants a collective judgement, such as a forecast or an assessment of risk, it should collect individual judgements independently before any discussion, so that early opinions do not anchor everyone else. This preserves the diversity on which the wisdom of crowds depends.

Preserve diversity of approach. Zollman's (2007) results and Kitcher's (1990) analysis both suggest that it can be valuable to keep pursuing alternative approaches even after one appears to be ahead. In research practice this might mean testing several explanations for a finding, or commissioning a second method to check the first, rather than converging quickly on a single story.

Seek out neglected perspectives. Harding's (1991) and Fricker's (2007) arguments have a practical side. A survey designed only from the perspective of the majority may miss questions that matter to minorities, or use concepts that do not fit their experience. Involving diverse people in designing research, and testing instruments with different groups, produces knowledge that is both fairer and more accurate.

Be explicit about how the organisation forms its views. The discursive dilemma shows that the method a group uses to reach conclusions affects what conclusions it reaches. A research organisation should be clear about how its published findings are decided: who signs off, on what criteria, and how disagreements are resolved.

Conclusion

Social epistemology replaces the picture of the solitary knower with a more realistic one. Knowledge is produced by people who depend on one another, who divide their cognitive labour, who check each other's work and who form collective views. The quality of that knowledge depends on whether trust is well placed, whether criticism flows freely, whether diverse perspectives are included and whether information is shared and combined well. These are features of organisations and communities, and they can be designed. They can also be neglected. An organisation can hire excellent individuals and still produce unreliable knowledge if those individuals never challenge one another, if dissent is punished, if information stays in separate departments, or if the same narrow group designs every study. Equally, a team of ordinary ability can produce excellent knowledge if its procedures bring diverse evidence together, expose every claim to criticism and keep a clear record of what has been learned. For any enterprise devoted to generating knowledge, getting the social structure right is as important as getting the methods right.

References

Bloor, D. (1976). *Knowledge and social imagery*. Routledge & Kegan Paul.

Fleck, L. (1979). *Genesis and development of a scientific fact* (T. J. Trenn & R. K. Merton, Eds.; F. Bradley & T. J. Trenn, Trans.). University of Chicago Press. (Original work published 1935)

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198237907.001.0001

Gilbert, M. (1987). Modelling collective belief. *Synthese, 73*(1), 185–204. https://doi.org/10.1007/BF00485446

Goldman, A. I. (1999). *Knowledge in a social world*. Oxford University Press. https://doi.org/10.1093/0198238207.001.0001

Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research, 63*(1), 85–110. https://doi.org/10.1111/j.1933-1592.2001.tb00093.x

Haraway, D. (1988). Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies, 14*(3), 575–599. https://doi.org/10.2307/3178066

Harding, S. (1991). *Whose science? Whose knowledge? Thinking from women's lives*. Cornell University Press.

Hardwig, J. (1985). Epistemic dependence. *The Journal of Philosophy, 82*(7), 335–349. https://doi.org/10.2307/2026523

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences, 101*(46), 16385–16389. https://doi.org/10.1073/pnas.0403723101

Janis, I. L. (1972). *Victims of groupthink: A psychological study of foreign-policy decisions and fiascoes*. Houghton Mifflin.

Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy, 87*(1), 5–22. https://doi.org/10.2307/2026796

Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford University Press.

Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago Press.

Lackey, J. (2021). *The epistemology of groups*. Oxford University Press. https://doi.org/10.1093/oso/9780199656608.001.0001

List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199591565.001.0001

Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton University Press.

Merton, R. K. (1973). *The sociology of science: Theoretical and empirical investigations* (N. W. Storer, Ed.). University of Chicago Press.

Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme, 17*(2), 141–161. https://doi.org/10.1017/epi.2018.32

Shapin, S. (1994). *A social history of truth: Civility and science in seventeenth-century England*. University of Chicago Press.

Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology, 48*(6), 1467–1478. https://doi.org/10.1037/0022-3514.48.6.1467

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

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

Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science, 330*(6004), 686–688. https://doi.org/10.1126/science.1193147

Wray, K. B. (2002). The epistemic significance of collaborative research. *Philosophy of Science, 69*(1), 150–168. https://doi.org/10.1086/338946

Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science, 316*(5827), 1036–1039. https://doi.org/10.1126/science.1136099

Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science, 74*(5), 574–587. https://doi.org/10.1086/525605