

Chapter 16

Knowledge in a Scientific Community

Abstract In earlier chapters various social aspects of scientific knowledge have been explored. These have been aspects which allow for social evidence to provide scientific knowledge to an individual. The focus in this chapter, however, moves beyond the study of individualistic characteristics of scientific knowledge by looking at science itself as an epistemic system. The thoroughgoing social nature of science leads to some characteristics which make it an epistemic system particularly well suited for adding to the store of scientific knowledge. In particular, the social nature of science leads to a division of cognitive labor. This division of cognitive labor both makes it so that trust plays an integral role in the generation of scientific knowledge and so that scientific progress is enhanced by the scientific community hedging its bets by scientists pursuing a wide variety of research projects utilizing a variety of methods. Although the individual scientists who make up the scientific community are not perfect, various social institutions in science help to make good use of their baser motivations.

Scientists do not operate in a vacuum. As we noted in the previous chapter, science is a thoroughgoing social activity (Goldman 1999; Longino 2013). After all, scientists are not isolated from society-at-large, nor do they work in isolation from one another. They “read journals, go to conferences, establish collaborations, seek out grant money, win acclaim and prestige, and operate within a labor market” (Muldoon 2013, p. 117). In fact, when we talk of “science” we are not speaking of some abstract object that exists in an unchanging, eternal form. We are talking about a social activity in which many people are actively engaged, one from which many people beyond those actively engaged enjoy benefits. As Philip Kitcher (1993, p. 179) explains, “The science of a time is constituted by the collection of individuals engaged in doing science, their relations to one another and to the broader community, their cognitive propensities and their individual practices, the consensus practices of the various fields and subfields.”

One of the most important facets of the social structure of science is the division of cognitive labor which it encourages. The fact that science incorporates a division of labor is relatively unsurprising. As any culture or society develops the members of the culture “divide up chores in ways that create different areas of expertise” (Keil 2006, p. 144). The idea that a division of labor is necessary for a society to be

sufficiently productive is not at all new—economists, philosophers, and sociologists have argued for this claim for centuries.¹ Importantly, “in most human cases, divisions of physical labor carry with them implications for divisions of cognitive labor” (Keil 2006, p. 144). Such divisions of cognitive labor seem to occur in just about every community (Lutz and Keil 2002). It seems to be clearly true that “the division of cognitive labor is an essential infrastructure that allows us to transcend the very limited understandings that exist in the mind of any one individual” (Keil 2006, p. 164). Even a mind as great as that of Isaac Newton is very limited in the grand scheme of things. As he himself is reported to have said when reflecting on his achievements as a scientist, “I don’t know what I may seem to the world, but, as to myself, I seem to have been only like a boy playing on the sea-shore, and diverting myself in now and then finding a smoother pebble or a prettier shell than ordinary, whilst the ocean of truth lay all undiscovered before me” (Gleick 2004, p. 4). A division of cognitive labor is necessary if science is to move beyond the amount of knowledge that a single person can have at a particular time. Thus, it is no surprise that science is structured in a way that involves various divisions of labor.

One obvious way that science involves division of cognitive labor is the various disciplines and sub-disciplines into which scientists tend to group themselves (Muldoon 2013). Physicists tend to work on different issues than chemists, who in turn work on different issues than biologists, and so on. Even within broader disciplinary divides there are numerous divisions into sub-disciplines—astrophysicists tend to work on different problems than quantum physicists, for example.

Although the relatively macro-level division of cognitive labor into scientific disciplines and sub-disciplines is interesting, there is another more micro-level division of cognitive labor which is perhaps more central to better understanding the social character of NOS. There are two primary aspects of this micro-level division of cognitive labor: how scientists rely on the work of others in forming their own beliefs about which scientific theories and claims are true and how scientists go about making choices as to which research projects to pursue. The first issue concerns the role of trust in science—something which is unavoidable given the division of cognitive labor that is present, and something necessary for effectively making scientific progress. The second issue concerns how labor is divided when it comes to pursuing research projects in science. It is these two vital aspects of the epistemic system that is science which will be our focus in this chapter.

16.1 The Role of Trust in Science

In addition to displaying a healthy dose of intellectual humility, the quote from Newton above eloquently expresses a simple truth which holds for science as a whole—there is much more to the universe than what science has revealed so far.

¹See Durkheim (1893/1997), Hume (1739–1740/1978), and Smith (1776/1904).

There is much scientific work to be done. Despite the fact that there are many truths yet to be uncovered by science, the extent of scientific knowledge we have already produced is daunting! In fact, there is so much scientific knowledge that no individual can reasonably claim to possess a significant share of it. As Michael Strevens (2010, p. 295) puts the point, “The expertise of even a professional scientist constitutes only a speck in the vast constellation of scientific knowledge . . . Our acceptance of the scientific image of the world is, then, based on trust in what scientists say, whether we are civilians or scientists ourselves.”²

The trust that we must place in others when it comes to science is a result of our limitations as finite individuals. We simply do not have the resources to run every experiment for ourselves or to independently confirm every theory which is widely accepted by experts. It is because of our limitations that “only some results are so checked and many are simply accepted on trust” (Longino 2013, p. 7). We do not seem to have any other choice if we are to make much progress in science at all. This is especially clear when one considers the various divisions of labor which exist in science. No one can be an expert in all aspects of science, and no one has the time to independently verify every important scientific result in even a single scientific discipline let alone several of them. If scientists constantly have to check each and every result of other scientists, they will never get around to building on the scientific knowledge which already exists because there will be no time left for new experiments and the development of new theories. There is no getting around it—trust plays a critical role in the growth of scientific knowledge.

Given the vast importance of trust in science one might be inclined to worry that this puts science, and the knowledge which scientists seeks to generate, in a precarious position. After all, this great need for trust in science means that scientists and laypersons alike are largely dependent upon others for their scientific knowledge. Any time one is dependent on other people there is a chance that those other people will fail to live up to expectations. Sometimes people will be dishonest or incompetent or fall short of being a reliable source of information for some other reason(s). Unfortunately, “scientists cannot spend very much time checking the work of others if they are to make contributions” of their own (Hull 1988, p. 394). So, one might worry that this fact gives us reason to question how solid the knowledge produced by science is—or to question if our foundation for scientific beliefs is really sufficiently firm for knowledge at all.

Although there is always a chance that we will be led astray when we rely on others, as we have seen in the previous chapter, there are good reasons for thinking that we can, for the most part, trust the information we gain from others. This is particularly true in science where there are social structures set up to reward and punish in such a way that scientists have strong incentive to report honestly and competently (Hull 1988). Scientists, like people in general, are often greatly affected by the quality of their reputations. Hence, the potential risk to one’s reputation by

²See Goldman (1999, 2001), Hardwig (1985), Hull (1988), Kitcher (1993), Longino (2013), and Shapin (1994) for considerations in support of thinking that trust is integral to science.

reporting information dishonestly or incompetently is a strong constraint on the quality of information that is shared among scientists and conveyed to laypeople (Adler 2014; Hull 1988). Additionally, things like the peer review process, the public availability of reports and findings, and the possibility of attempts to replicate experimental results help to ensure that the information circulated in the scientific community is trustworthy (Shatz 2004). Thus, it seems that we have good reason to trust that the foundation of knowledge arising from the scientific community is indeed solid.

Despite the constraints in place in science and our good reasons for trusting the information we receive from scientific experts, one might still worry about the reliance on trust which we find in science. When someone is a layperson, or even a scientist who is simply not an expert in the area of science in which she is relying on information from other scientists, one might worry that her reliance on the experts is “blind” (Hardwig 1985). That is to say, one might worry that the person, layperson or scientist, trusting an expert is in no position to evaluate the information she receives. In light of this, one might think that her position is the unenviable one of having to simply rely on information from an expert with no way of determining how reliable that expert really is. This is particularly disturbing since as we noted in the previous chapter, we should “always engage in some assessment of the speaker for trustworthiness” because “to believe what is asserted without doing so is to believe blindly, uncritically” (Fricker 1994, p. 145). If believing blindly is tantamount to believing uncritically, believing blindly does not seem to be very conducive to attaining knowledge. As Alvin Goldman (2001, p. 86) explains, if we really are “blind” in the way that some suggest we are when it comes to science, it “seems to imply that a layperson (or a scientist in a different field) cannot be *rationally justified* in trusting an expert”. Of course, if we cannot be rationally justified in trusting experts, it seems likely that we cannot come to have knowledge on the basis of what they tell us. Thus, if the trust which is a necessary component of science is truly blind, science’s ability to generate knowledge is dubious.

Fortunately, our trust in science (and in experts in general) is far from blind. We tend to be very good at determining who is an expert in what areas, and so, likely to provide us with reliable information on a given topic. “There are several distinct heuristics that can be used to figure out who knows what”, and even “very young children are sensitive to many of these heuristics” (Keil 2006, p. 164). Given this sensitivity to heuristics for discerning expertise, it is unsurprising that it is widely assumed in educational settings that children are able to evaluate the expertise of others—this is a foundational point for the peer tutoring movement in educational practice.³ In summarizing the results of their empirical research concerning the evaluation of expertise Frank Keil and his colleagues (2008, p. 298) claim:

One of the most powerful ways of thinking about the organization of knowledge is based on the following idea: There are domain-specific patterns in the world that experts know and use to understand a wide range of phenomena that arise from those patterns. If one

³See Cohen et al. (1982) and Fuchs et al. (1997).

can also grasp those patterns in some coarse yet reliable manner, one knows which sort of expert to approach for further understanding. We have shown here that this appreciation of discipline-based ways of understanding shows its first signs quite early in childhood and develops substantially during the elementary school years.

It seems that we can determine who is an expert in which area from a very young age. This ability casts doubt on the idea that the trust so integral to science is truly blind.

In addition to our having the ability to determine who is an expert in a particular field, there is good reason to think that we have ways of determining which experts are more reliable than others. For one thing we have ways of “calibrating” experts (Kitcher 1993). A scientist who is an expert in a particular field can directly calibrate the reliability of other experts in that same field by using her own information about that field. A scientist can indirectly calibrate a particular expert’s level of expertise by taking account of the opinions other scientists, who she has directly calibrated, have of the expert in question. As we noted above, reputation plays a significant role in science. One important function of reputation is to help with the calibration of expertise.

Finally, there are numerous additional sources of evidence which we (either as laypeople or as scientists) can consult when evaluating a particular expert. Alvin Goldman (2001, p. 93) lists five common forms of evidence we have when evaluating experts:

- (A) Arguments presented by the contending experts to support their own views and critique their rivals’ views.
- (B) Agreement from additional putative experts on one side or other of the subject in question.
- (C) Appraisals by “meta-experts” of the experts’ expertise (including appraisals reflected in formal credentials earned by the experts).
- (D) Evidence of the experts’ interests and biases vis-a-vis the question at issue.
- (E) Evidence of the experts’ past “track-records”

In the majority of cases at least some of these sources of evidence will be readily available. Thus, it seems that our trust in experts in science is far from blind.

While trust plays a key role in science, and trust always carries the possibility of disappointment, it seems that this should not give us pause when considering the ability of science to generate knowledge. Understanding the importance of trust in science, and why this trust is not simply blind faith in experts, helps to illuminate the social nature of science.

16.2 The Division of Cognitive Labor

We have noted that trust is vital for scientific progress. A key reason for the necessity of trust is the division of cognitive labor which we find in the sciences. Such a division of cognitive labor is necessary for science to progress beyond the

limitations of what individuals can know on their own. Although the division of cognitive labor plays a critical role in the discovery and distribution of scientific knowledge, it presents us with important questions. One of the most important questions, which will be our focus here, is how the division of cognitive labor leads scientists to choose various research projects to pursue.

In an ideal situation we would have an infinite amount of resources and a community of scientists who are each perfectly rational, purely altruistic, completely aware of what projects other scientists are pursuing, and motivated solely by the desire to make discoveries which increase the amount of scientific knowledge in the scientific community as a whole. The situation that we actually find ourselves in is far from this ideal though. Obviously, we do not have an infinite amount of resources. If we did, the need for a division of cognitive labor would be greatly diminished, if not entirely non-existent. Additionally, scientists, while typically rational, are like any other group of humans—they fall short of perfect rationality. They are also limited in their communication with one another, and so limited in their awareness of the projects that other scientists are pursuing. Despite being generally good people with a motivation to increase the total stock of scientific knowledge, scientists are not typically altruists, nor do they tend to be motivated solely by the desire to increase the amount of scientific knowledge possessed by the community. “Realistically . . . scientific research is driven by twin motives” the desire for scientific knowledge and for credit—“some amount of scientific credit, after all, is normally a precondition for job maintenance, and larger amounts of credit often conduce to a higher salary, a better job, and/or greater personal satisfaction” (Goldman 1999, p. 260). So, while scientists are probably, by and large, rational, good people with good intentions, like everyone else they are far from perfect.

It is because scientists are people too and our collective situation falls short of the ideal in the sorts of ways mentioned above that questions about the division of labor and how different research projects are chosen are worth considering. Our limitations are the source of tensions which can arise in science between what is rational for individual scientists and what is rational for the scientific community as a whole (Kuhn 1977). Interestingly, this sort of tension can even arise in cases where scientists are ideally motivated by a desire to attain scientific knowledge, and they are each behaving rationally as individuals. The reason for this is that it seems that what is rational for a group (such as the scientific community) is not merely a matter of each member of that group behaving in a way which is individually rational (Sarkar 1983).⁴

To see how group rationality can diverge from individual rationality, consider a simple case from Goldman (2011, p. 17). Assume that we have a group of three

⁴Thagard (1993) claims that it should be somewhat unsurprising that group rationality and individual rationality diverge in science. He points out that such divergences have long been recognized in decision theory in situations such as the “prisoners’ dilemma” and the “tragedy of the commons”.

individuals who are operating under the seemingly rational rule of majority vote. In other words, when at least two of the members agree on a judgment the group accepts that judgment. These three individuals are each evaluating the truth of three propositions: p , q , and $\langle \text{if } p, \text{ then } q \rangle$. Person A accepts all three propositions as true. Person B accepts p , but denies the other two propositions. Person C accepts $\langle \text{if } p, \text{ then } q \rangle$, but denies the other two propositions. Each individual accepts a set of propositions that is perfectly consistent, i.e. each individual makes a judgment which is plausibly rational. The group, however, does not end up with a rational set of judgments. The group ends up accepting p (A and B agree about it) and $\langle \text{if } p, \text{ then } q \rangle$ (A and C agree about it) while denying q (B and C both deny it). Each individual made a consistent judgment about the three propositions and the group employed the reasonable rule of going with the majority opinion on each proposition in forming the group's judgment. And yet, the group ends up with an inconsistent set of commitments: p and $\langle \text{if } p, \text{ then } q \rangle$ (which the group accepts) entail that q must be true, but the group denies q . Hence, just because each member of a group is acting rationally it does not follow that the group itself will act rationally.

This sort of divergence between individual and group rationality can occur in a scientific community. It can even occur in a scientific community composed entirely of "epistemically pure" scientists, who are only motivated by the desire to discover scientific truths (Kitcher 1993). For simplicity, assume that the scientific community is seeking to uncover the truth of whether a particular molecule exists, and there are only two possible research projects to pursue in order to find out the truth of the matter. Project 1 has a much better chance of success than Project 2. However, since science is not perfect, Project 2 also has some chance of success because the odds that Project 1 succeeds fall short of 100%. In this case the rational thing for an individual motivated solely by the desire to discover truth to do is to pursue Project 1 because doing so has the best odds of leading to a successful discovery concerning the molecule. Nevertheless, if each scientist is motivated only by the desire to discover truth and does what is individually rational, then each scientist will pursue Project 1. This, however, means that the scientific community as a whole is not pursuing projects in the most rational fashion. The rational thing for the community to do is to expend at least some resources in pursuing Project 2. After all, the odds of discovering the truth concerning the molecule are much higher if both projects are pursued rather than just one. As Philip Kitcher (1993, p. 344) aptly notes, "a community that is prepared to hedge its bets when the situation is unclear is likely to do better than a community that moves quickly to a state of uniform opinion" about what projects to pursue.⁵ Consequently, even a community

⁵It is worth noting that this situation does assume that scientists do not have perfect knowledge of what every other scientist is doing. If each individual scientist had this knowledge and was altruistic, then some might rationally accept the lower chances of their making the discovery by pursuing Project 2 for the greater good of the community. However, the assumptions that scientists do not have perfect knowledge of what each other is doing and that they are not perfectly altruistic are exceedingly plausible.

of epistemically pure scientists might fail to behave rationally despite each member of the community behaving in a way which is individually rational.

One might worry that if tensions can arise between what is rational for the scientific community as a whole and what is rational for individual scientists when those individual scientists are motivated solely by the desire to increase the community's knowledge, they can surely arise in the situation we find ourselves in where scientists are motivated by a variety of factors only some of which are "pure". The desire for discovering truth is only one thing that motivates actual scientists because, like the rest of us, they are "epistemically sullied agents" (Kitcher 1993). They are also motivated by their desire to receive credit for making scientific discoveries and all of the goods which come along with that credit. As a result, one might worry that when we look at science as an epistemic system what we find is an irrationality—a system which falls far short of dividing its cognitive resources in a reasonable manner.

As we have seen time and time again, assuming that science fails to perform admirably well is a mistake. It simply does not follow that because a more "purely" motivated society of scientists fails to properly distribute cognitive labor that a more "sullied" community like we have will similarly fail. In fact, "from the community perspective, it is likely that sullied scientists will do better than the epistemically pure. This is because a pure community heads toward cognitive uniformity", but "the sullied community hedges its bets" by diversifying its cognitive labor (Kitcher 1993, pp. 310–311). Recognition of the fact that scientists are motivated by more than simply the desire to generate scientific knowledge has led to many economic models of the structure of science and a focus on the various incentive structures of science.⁶ One of the chief results of examining such models is the recognition that "the very factors that are frequently thought of as interfering with the (epistemically well-designed) pursuit of science—the thirst for fame and fortune, for example—might actually play a constructive role in our community epistemic projects, enabling us, as a group, to do far better than we would have done had we behaved as independent epistemically pure individuals" (Kitcher 1993, p. 351). Quite simply,

⁶See, for example, Brock and Durlaf (1999), Goldman (1999), Goldman and Shaked (1991), Hagstrom (1965), Hull (1988), Kitcher (1993), Latour and Woolgar (1979), Merton (1973), Rescher (1990), and Strevens (2003). For criticisms of economic models of science see Hands (1995, 1997), Muldoon (2013), Muldoon and Weisberg (2011), Weisberg and Muldoon (2009), and Wray (2000). There are other ways of modeling science such as Thagard's (1993) distributed A.I. approach, the ecological model, which Muldoon and Weisberg (2011) and Weisberg and Muldoon (2009) draw from "hill climbing" models in computer science, and various consensus models such as those put forward by Hegselman and Kraus (2006), Lehrer (1975), Lehrer and Wagner (1981), Wagner (1985), and Zollman (2010). Although these other models are interesting and worthy of careful consideration, we will limit our focus to the economic approach. The reason for this is twofold. First, such models are by far the most widely accepted approaches to modeling science and its distribution of cognitive labor. Second, the other primary ways of modeling science and the distribution of cognitive labor seem to agree with the economic models on the crucial points that hedging its bets is the best way for science to proceed. They also agree with economic models that there are social structures in place which help ensure that science has a diversity of cognitive labor.

“a profit motive can *discipline* one’s search for truth rather than distort it” (Goldman 1999, p. 260).

The central point illuminated by economic models of science (and other models) is that a beneficial division of cognitive labor is often achieved because “social institutions within science might take advantage of our personal foibles to channel our efforts toward community goals rather than the epistemic ends that we might set for ourselves as individuals” (Kitcher 1993, p. 351).⁷ In order to see this it is worth taking a look at a particular social institution present in science which helps to use scientists’ desire for credit as a way to ensure a useful division of cognitive labor.⁸ We will examine what Robert Merton (1973) dubbed the “priority rule”. The priority rule is the social institution that we find in science where all of the credit for a particular discovery goes to the first researchers to make that discovery. According to the priority rule, scientific discovery is a winner-takes-all matter where only first-place counts—everyone else loses when it comes to credit even when the independent discoveries of a particular fact made by different scientists are merely days apart. The priority rule is a very long-standing institution in science. Merton traces the priority rule back to such luminaries as Galileo and Newton—basically to the beginning of modern science. Interestingly, Merton maintains that the priority rule is one that most all scientists tend to view as a good thing. In fact, he claims that scientists believe that violations of this rule are *morally wrong*.

Michael Strevens (2003) argues that it is rational for credit to be awarded in accordance with the priority rule because society benefits from initial discoveries in science, but not from subsequent discoveries of the same fact. The recognition of this feature of how society as a whole benefits from scientific discoveries helps to explain why scientists view the priority rule as a just rule.⁹ Nonetheless, one might wonder how this social institution helps to ensure a favorable distribution of cognitive labor.

The priority rule promotes a division of cognitive labor because when scientists realize that only the first to make a discovery will receive credit it will prompt some to pursue projects other than those pursued by other scientists. After all, “two equally talented, equally industrious scientists may receive different rewards under the priority system just because one of them, but not the other, is lucky enough to select a research program that achieves its goal (and does so before any rival)” (Strevens 2003, p. 74). So, the reasonable thing in many cases may be to pursue a different project than one’s peers. Consider if Xenia knows that Xavier is

⁷Also see Hull (1988), Solomon (1992), and Thagard (1988).

⁸There are, of course, a number of such social institutions present in science. We will limit our focus here to just one of the most prominent ones because doing so is sufficient for our current illustrative purposes.

⁹Strevens (2003) suggests that the priority rule is perhaps simply a clear representation of a much more general reward scheme which we find in society as a whole. Namely, a reward scheme that rewards in proportion to the benefit provided to society as a whole. The priority rule exemplifies this rule because when it comes to scientific discovery only the first to make the discovery provides a benefit to society, so only the first to make a discovery gets rewarded.

investigating the cause of some phenomenon Y, then she knows that the only way she can receive credit for discovering the truth about Y is if she makes her discovery before Xavier. In some cases this may simply come down to luck—particularly if they are employing similar methods. Additionally, Xenia might be at a disadvantage if Xavier has been working on this project longer than she has. Xenia might decide that her best bet for getting credit for some important scientific discovery or other is to pursue a project related to phenomenon Z instead of investigating Y because there is less work being done related to Z. Thus, awareness of the priority rule can lead to a division of cognitive labor by capitalizing on scientists' desire for credit.

Of course, there are other ways of encouraging division of cognitive labor in science. Some of these ways will also appeal to our more “sullied” motivations; others will simply be a matter of resource allocation. Some of these social institutions are relatively indirect motivators like the priority rule; others are more direct methods for diversifying cognitive labor. For example, one way to directly encourage scientists to pursue a diversity of research projects is to simply offer increased funding for research projects which are underrepresented. “Very little of modern science can be conducted without funding, and as long as individual scientists at least propose to use different methods, funding agencies are in a position to encourage diversity by financially supporting it” (Goldman 1999, p. 257).

Admittedly, some of the motivators employed to encourage division of cognitive labor, and thus to increase the odds of science generating more and more knowledge, may be misguided.¹⁰ Also, it is likely that the current set of social institutions providing incentives in science is not optimal (Kitcher 1993; Muldoon and Weisberg 2011). Despite all of this, exploration of the social institutions present in science for encouraging division of cognitive labor helps to demonstrate that the sort of division of cognitive labor which is necessary for significant scientific progress can be achieved even by, and perhaps especially by, epistemically “sullied” scientific communities like ours. In fact, what we have seen here should remind us that we should not think that we can “identify very general features of scientific life—reliance on authority, competition, desire for credit—as epistemically good or bad” because “particular kinds of social arrangements make good epistemic use of the grubbiest motives” (Kitcher 1993, p. 305).

¹⁰For example see Merton's (1973, 1988) discussion of the “Matthew effect” where more well-known scientists receive more credit than less well-known scientists do for the same achievements. The Matthew effect is a social institution in science which many believe is a misallocation of credit, and so, a negative side effect of credit-based motivational structures in science. Though see Strevens (2006) for persuasive arguments for thinking that the Matthew effect does in fact distribute credit fairly.

16.3 Conclusion

In this chapter we have moved beyond our study of the individualistic characteristics of scientific knowledge by looking at science as an epistemic system. We have seen that the thoroughgoing social nature of science leads to some characteristics which make it particularly well suited for adding to the store of scientific knowledge. In particular, the social nature of science leads to a division of cognitive labor. We have seen that this division of cognitive labor both makes it so that trust plays an integral role in the generation of scientific knowledge and so that scientific progress is enhanced by the scientific community hedging its bets by scientists pursuing a wide variety of research projects utilizing a variety of methods. Although the individual scientists who make up the scientific community are not perfect, various social institutions in science help to make good use of their baser motivations. We have seen that when it comes to the epistemic system that is science: “Flawed people, working in complex social environments, moved by all kinds of interests, have collectively achieved a vision of parts of nature that is broadly progressive and that rests on arguments meeting standards that have been refined and improved over centuries” (Kitcher 1993, p. 390). Science may not be perfect, but it is an epistemic system that has been tremendously successful at generating knowledge of the world around us.

References

- Adler, J. (2014). Epistemological problems of testimony. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2014 Edition). <http://plato.stanford.edu/archives/spr2014/entries/testimony-episprob/>
- Brock, B., & Durlaf, S. (1999). A formal model of theory choice in science. *Economic Theory*, 14, 113–130.
- Cohen, P. A., Kulik, J. A., & Kulik, C. C. (1982). Educational outcomes of tutoring: A meta-analysis of findings. *American Educational Research Journal*, 19, 237–248.
- Durkheim, E. (1893/1997). *The division of labor in society*. New York: The Free Press.
- Fricker, E. (1994). Against gullibility. In B. K. Matilal & A. Chakrabarti (Eds.), *Knowing from words* (pp. 125–161). Dordrecht: Kluwer.
- Fuchs, D., Fuchs, L. S., Mathes, P. G., & Simmons, D. C. (1997). Peer-assisted learning strategies: Making classrooms more responsive to diversity. *American Educational Research Journal*, 34, 174–206.
- Gleick, J. (2004). *Isaac Newton*. New York: Vintage Books.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63, 85–110.
- Goldman, A. I. (2011). A guide to social epistemology. In A. Goldman & D. Whitcomb (Eds.), *Social epistemology: Essential readings* (pp. 11–37). New York: Oxford University Press.
- Goldman, A. I., & Shaked, M. (1991). An economic model of scientific activity and truth acquisition. *Philosophical Studies*, 63, 31–55.
- Hagstrom, W. (1965). *The scientific community*. New York: Basic Books.

- Hands, D. W. (1995). Social epistemology meets the invisible hand: Kitcher on the advancement of science. *Dialogue*, 34, 605–621.
- Hands, D. W. (1997). Caveat emptor: Economics and contemporary philosophy of science. *Philosophy of Science*, 64, 107–116.
- Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy*, 82, 335–349.
- Hegselmann, R. & Krause, U. (2006). Truth and cognitive division of labour: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation*, 9. <http://jasss.soc.surrey.ac.uk/9/3/10.html>
- Hull, D. (1988). *Science as a process*. Chicago: University of Chicago Press.
- Hume, D. (1739–1740/1978). *A treatise of human nature*. Oxford: Clarendon Press.
- Keil, F. C. (2006). Doubt, deference, and deliberation: Understanding and using the division of cognitive labor. In T. Z. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology: Volume 1* (pp. 143–166). Oxford: Oxford University Press.
- Keil, F. C., Stein, C., Webb, L., Billings, V. D., & Rozenblit, L. (2008). Discerning the division of cognitive labor: An emerging understanding of how knowledge is clustered in other minds. *Cognitive Science*, 32, 259–300.
- Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. New York: Oxford University Press.
- Kuhn, T. S. (1977). *The essential tension: Selected studies in scientific tradition and change*. Chicago: University of Chicago Press.
- Latour, B., & Woolgar, S. (1979). *Laboratory life: The social construction of scientific facts*. Beverly Hills: Sage.
- Lehrer, K. (1975). Social consensus and rational agnology. *Synthese*, 31, 141–160.
- Lehrer, K., & Wagner, C. (1981). *Rational consensus in science and society*. Dordrecht: Reidel.
- Longino, H. (2013). The social dimensions of scientific knowledge. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2013 Edition). <http://plato.stanford.edu/archives/spr2013/entries/scientific-knowledge-social/>
- Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development*, 73, 1073–1084.
- Merton, R. (1973). *The sociology of science*. Chicago: University of Chicago Press.
- Merton, R. (1988). The Matthew effect in science, II: Cumulative advantage and the symbolism of intellectual property. *Isis*, 79, 607–623.
- Muldoon, R. (2013). Diversity of the division of cognitive labor. *Philosophy Compass*, 8, 117–125.
- Muldoon, R., & Weisberg, M. (2011). Robustness and idealization in models of cognitive labor. *Synthese*, 183, 161–174.
- Rescher, N. (1990). *Cognitive economy: The economic dimension of the theory of knowledge*. Pittsburgh: University of Pittsburgh Press.
- Sarkar, H. (1983). *A theory of method*. Berkeley: University of California Press.
- Shapin, S. (1994). *A social history of truth: Civility and science in seventeenth-century England*. Chicago: University of Chicago Press.
- Shatz, D. (2004). *Peer review: A critical inquiry*. Lanham: Rowman & Littlefield.
- Smith, A. (1776/1904). *An inquiry into the nature and causes of the wealth of nations* (5th ed.). London: Methuen & Co.
- Solomon, M. (1992). Scientific rationality and human reasoning. *Philosophy of Science*, 59, 439–455.
- Strevens, M. (2003). The role of the priority rule in science. *Journal of Philosophy*, 100, 55–79.
- Strevens, M. (2006). The role of the Matthew effect in science. *Studies in History and Philosophy of Science*, 37, 159–170.
- Strevens, M. (2010). Reconsidering authority: Scientific expertise, bounded rationality, and epistemic backtracking. In T. Z. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology: Volume 3* (pp. 294–330). Oxford: Oxford University Press.
- Thagard, P. (1988). *Computational philosophy of science*. Cambridge, MA: MIT Press.
- Thagard, P. (1993). Societies of minds: Science as distributed computing. *Studies in History and Philosophy of Science*, 24, 49–67.

- Wagner, C. (1985). On the formal properties weighted averaging as a method of aggregation. *Synthese*, 25, 233–240.
- Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76, 225–252.
- Wray, K. B. (2000). Invisible hands and the success of science. *Philosophy of Science*, 67, 163–175.
- Zollman, K. J. S. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72, 17–35.