NOÛS 31:2 (1997) 242–261

*Collaborative Knowledge**

PAUL THAGARD
University of Waterloo

Abstract. Collaboration is ubiquitous in the natural and social sciences. How collaboration contributes to the development of scientific knowledge can be assessed by considering four different kinds of collaboration in the light of Alvin Goldman's five standards for appraising epistemic practices. A sixth standard is proposed to help understand the importance of theoretical collaborations in cognitive science and other fields. I illustrate the application of these six standards by describing two recent scientific developments in which collaboration has been important, the bacterial theory of ulcers and the multiconstraint theory of analogy, and by arguing that philosophy should become more collaborative.

In April 1994, a group of 450 physicists centered at Fermilab presented evidence for the existence of the top quark, an important theoretical construct of the Standard Model of particles and forces. Although the size of this group is unusual, the collaborative nature of their work is not. In the natural and social sciences, it has become much more the norm than the exception to have work produced by two or more cooperating scientists. This paper examines collaborative knowledge from both descriptive and prescriptive viewpoints. How prevalent is collaborative knowledge? Why do scientists collaborate? What kinds of collaboration are most productive? Why is collaboration in the humanities much rarer than in the sciences, and does it need to be?

After briefly reviewing the extent and nature of collaborative work in the sciences, I shall discuss collaboration from the perspective of the epistemic standards that Alvin Goldman has proposed for evaluating social practices. Although his standards need to be reframed somewhat in order to be helpful in understanding the advantages of collaborative scientific work, they capture many important aspects of collaborative knowledge. I shall, however, propose an extension of Goldman's standards to capture the contribution that collaboration can make to the explanatory goals of science.

1. The Prevalence of Collaboration

In recent years, researchers in many fields have paid increasing attention to social aspects of the development of knowledge. In philosophy, *social epistemology* has arisen to address numerous questions about how knowledge develops in social contexts (Fuller, 1988; Giere, 1988; Goldman, 1992; Hardwig, 1985; Hardwig, 1991; Kitcher, 1990; Solomon, 1994; Thagard, 1993, 1994). Psychologists are paying increasing attention to *distributed cognition,* examining knowledge not merely as the possession of individual minds, but as also dependent on social and physical environments. (Galegher, Kraut, & Egido, 1990; Resnick, Levine, & Behrend, 1991; Salomon, 1993) In computer science, there is growing interest in *distributed artificial intelligence*, which concerns how expertise can be investigated in networks of cooperating computers rather than in individual ones. (Bond & Gasser, 1988; Durfee, Lesser, & Corkill, 1989; Gasser, 1991; Hewitt, 1991). Finally, sociologists have made strong claims about the social production of scientific knowledge. (Barnes, 1985; Bloor, 1991; Latour, 1987; Latour & Woolgar, 1986). Surprisingly, however, there has been little discussion in any of these fields of the nature of the collaborations that have become prevalent in modern science.

This prevalence is easily documented. Figure 1 graphs the percentage of multiauthored papers in the physical and biological sciences, the social sciences, and

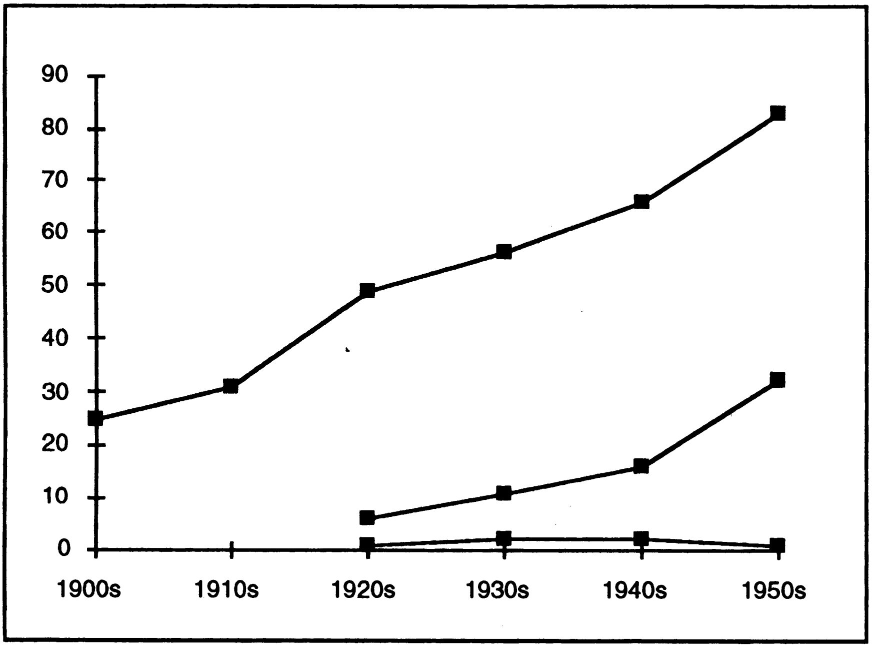

Figure 1. Percentage of multiauthored papers in the physical and biological sciences (top line), social sciences (middle line), and humanities (bottom line). Based on data from Merton (1973), p. 547.

the humanities up to the 1950s. By then, 83% of papers in selected journals in the physical and biological sciences were collaborative efforts. For social sciences, the number was 32%, while the humanities remained relatively constant at 1 or 2%. The trend has continued in recent decades: Figure 2 shows the percentage of multiauthored papers for selected journals in different fields for 1992. In *Physical Review Letters*, 88% of papers were multiauthored, and in *Cognitive Psychology* 75% of papers involved collaboration.

Research in the sciences is much more collaborative than work in the humanities. Although philosophers gain greatly from talking with each other, philosophical writings are rarely collaborative. In the 1992 volume of the *Journal of Philosophy*, for example, only 3 of 27 papers have 2 authors, and only 1 has 3 authors; the rest are single authored. The 1993 volume of the same journal has only 1 collaborative article. In contrast, of the 16 papers in the 1992 volume of *Cognitive Psychology*, only 4 are single authored, while 6 have 2 authors and the rest have from 3 to 6 authors. Similarly, of the 161 papers published in the Proceedings of the 1994 Conference of the Cognitive Science Society, only 52 are single authored, while 71 have 2 authors and 38 have between 3 and 8 authors.

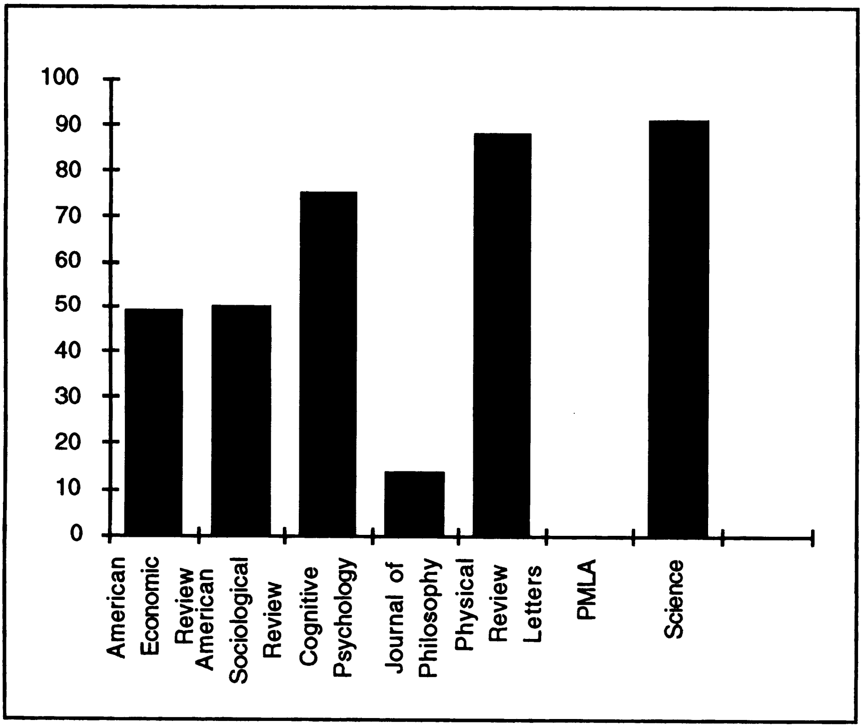

Figure 2. Percentage of multiauthored papers in selected journals in 1992. PMLA is the *Proceedings of the Modern Language Association*, whose 1992 volume had no collaborative papers.

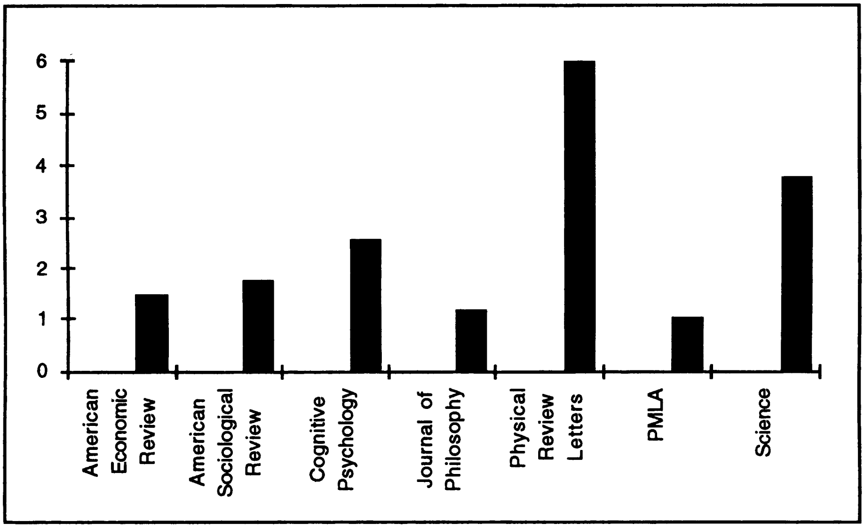

Figure 3. Mean number of coauthors in selected journals for 1992.

Even more extreme, in the January-April, 1992, volume of *Physical Review Letters*, only 67 of 558 contributions are single authored, around 12%. 168 papers have 2 authors, and 254 have from 3 to 5 authors. 6 papers have more than 100 authors, with the largest total being 291. Figure 3 shows the mean number of coauthors in journals in different fields. By looking more closely at different kinds of collaboration, we can start to explain why joint work is so common in the sciences but rare in the humanities.

2. The Nature of Collaboration

Not all collaborations are alike. There are at least four different kinds of collaboration, reflecting the different backgrounds and roles of the collaborators.

1. Employer/employee. This is the weakest form of collaboration, in which an employer simply tells an employee to perform a task that the employer knows how to do but does not want to spend the time on. Examples of such tasks include running experiments, writing computer programs, constructing apparatus, and so on. Technicians and research assistants do not normally make sufficient contribution to be considered as coauthors, but a talented assistant may become an apprentice as described in the next category.

2. Teacher/apprentice. This kind of collaboration is similar to the previous one in that there is asymmetry of knowledge and status, but it has a different goal. Apprentices do not merely perform work that the instructing researchers lack time for, but also aim to acquire the skills that will enable them to do the work themselves. For example, experimental psychologists typically work closely with graduate students who help to design and run experiments. Designing experiments and statistically interpreting the results are complex skills that are not learned by reading books or taking classes, but by working on projects with experienced researchers.

3. Peer-similar. Sometimes researchers of similar knowledge, interests, and status find it advantageous to work together. Perhaps the most famous collaboration of this century is the work by Francis Crick and James Watson on the structure of DNA. In psychology, there have been such productive duos as Allan Newell/Herbert Simon and Daniel Kahneman/Amos Tversky. Of course, "similar" does not mean "identical": any two researchers even in the same field will have somewhat different knowledge and skills to bring to a collaboration. But we can place in this category collaborations that involve people whose training has been substantially alike.

4. Peer-different. Cross-disciplinary research is more likely to bring together researchers with similar goals but with different knowledge and skills. In cognitive science, a typical collaboration involves a psychologist and a computer scientist (for other examples, see section 4.4 below). The former has expertise in theoretical and experimental psychology, including knowledge *how* to do experiments, while the latter has computational expertise including knowledge *how* to build programs that simulate aspects of thinking. Collaborations within physics may involve combinations of theoretical and experimental physicists who have very different kinds of skills.

Of course, the boundaries between these four kinds of collaboration can blur. A clever employee can turn into an apprentice, and a successful teacher/apprentice relationship should gradually become closer to a peer-similar collaboration. Researchers from disparate fields may start out as peer-different but become more similar as each learns more about the other's field. But these four different kinds of collaboration provide a start at addressing the question of what makes collaboration worthwhile.

3. Goldman's Standards for Epistemic Appraisal

The prevalence of collaboration strongly suggests that scientists must have good epistemic reasons for working together, but what are these? In the novel *Cantor's Dilemma,* written by a Stanford University chemist, a research scientist begins to suspect that his star post-doctoral fellow has been fabricating data on a very important experiment (Djerassi, 1989). This story illustrates one of the perils of collaboration, which can increase error as well as productivity. In considering the merits of different kinds of collaboration, we need to assess the occurrence of losses as well as gains.

Alvin Goldman has developed a set of standards for assessing epistemic practices. He advocates *veritism* as the principal approach to social epistemology, taking the goal of truth to be central to all intellectual pursuits. All his standards of appraisal for evaluating a social practice are concerned with truth (Goldman, 1992, p. 195):

1. The *reliability* of a practice is measured by the ratio of truths to total number of beliefs fostered by the practice;
2. The *power* of a practice is measured by its ability to help cognizers find true answers to the questions that interest them;
3. The *fecundity* of a practice is its ability to lead to large numbers of true beliefs for *many* practitioners;
4. The *speed* of a practice is how quickly it leads to true answers;
5. The *efficiency* of a practice is how well it limits the cost of getting true answers.

Before proceeding to apply these five standards to the four kinds of collaboration listed in section 2, it is useful to reframe the standards in less veritistic terms. Many scientists would blanche at describing their findings as "truths", since the truth of scientific claims only gets sorted out in the long run, as experiments and theories accumulate. Hence if our goal is to understand why scientists collaborate, we need to describe what they do according to more short-term goals than truth. As an alternative vocabulary, let us describe scientists as seeking *results*, which can include both empirical results consisting of experimental or observational findings, as well as theoretical results that consist of the development of theories that explain the empirical results. The criteria for counting something as a result are less stringent and metaphysical than those for counting something as a truth; as a first approximation, we can count an empirical or theoretical claim as a result if it is acceptable by a scientist's peers. Unanimous acceptance by one's peers is not required; perhaps a minimal requirement for a result is that it should be publishable in a good, peer-reviewed journal. Ultimately, we want to the results to be true, but in understanding everyday scientific practice we do not want to have to wait years or decades that might be required for full validation.

The opposite of a result is an *error*, an experimental or theoretical claim that would tend to be rejected by well-informed peers. We can now reframe Goldman's standards as follows:

1. The *reliability* of a practice is measured by the ratio of results to total number of results and errors fostered by the practice;
2. The *power* of a practice is measured by its ability to help cognizers find results that answer the questions that interest them;
3. The *fecundity* of a practice is its ability to lead to large numbers of results for *many* practitioners;
4. The *speed* of a practice is how quickly it leads to results;
5. The *efficiency* of a practice is how well it limits the cost of getting results.

The connection between these standards and Goldman's original veritistic ones is that what I call results are what scientists generally take to be true, and what I call errors are what scientists generally take to be false. From the perspective of scientific realism (e.g. Thagard, 1988, ch. 8), results often *are* true, and errors often *are* false. I agree with Goldman that science seeks and sometimes achieves truth, so the title of this paper is legitimately "Collaborative Knowledge" rather than just "Collaborative Belief." But for understanding the epistemic value of collaboration, we need shorter-term, more readily assessable standards than veritistic ones.

The question now becomes: How do the different kinds of collaboration affect the reliability, power, fecundity, speed, and efficiency of scientific research?

4. Why Collaborate? Gains and (Occasional) Losses

4.1. Employer/Employee

When a scientific researcher hires an employee such as a laboratory technician, research assistant, or computer programmer, it is probably unreasonable to expect increased reliability. Unless the employee has esoteric skills and tasks, most of what he or she does could probably be done at least as well by the researcher. So reliability is not the most relevant standard for appreciating this kind of collaboration. The researcher presumably cares enough about reliability not to want an employee whose work is dramatically increasing the error rate, but accepting a somewhat higher error rate will normally have to be an acceptable tradeoff for not doing everything by oneself. It is possible, however, that some tasks will actually be done more reliably by an employee than by the researcher, who may, for example, not be as good a computer programmer as a young assistant.

With good employees, potential losses in reliability are more than compensated for by gains in power, speed, and efficiency. Division of labor in which employees such as technicians do simpler or more time consuming tasks allows researchers more time to work on experimental or theoretical projects. The effect should then be that the researcher gets more desired results (power), and gets them faster (speed). Hiring an employee does increase the cost of research, thereby potentially reducing efficiency, but not nearly as much as hiring an additional researcher. Improvements in power and speed are not a sure result of hiring an employee, since the researcher has to spend time training and supervising the employee. Initially, the time spent may exceed the time that a researcher might have required to do a task alone, but in the long run the time and effort required to monitor the employee should drop well below the amount of researcher time and effort expended. Fecundity, the question of getting many results for many people, does not seem to be relevant to assessing employer/employee collaborations.

4.2. Teacher/Apprentice

In many of the natural and social sciences, graduate students are an essential part of the conduct of research. In the humanities, graduate students typically pursue projects unconnected to their advisor's research, who accordingly treat them at best with benign neglect. In contrast, in fields such as experimental physics and psychology, graduate students are often a crucial part of the research team with primary responsibility for collection of empirical results. Students may work with advisors on experimental design and put in long hours collecting data. As with employees, researchers find it worthwhile to collaborate with graduate students because the gains in power and speed potentially compensate for possible losses in reliability and efficiency. Reliability can suffer because newly trained students may not know as much as established researchers about how to avoid mistakes, and the cost of research is increased by the need to fund the students. But effective graduate students, who assume time-consuming tasks that would otherwise have to done by a researcher, can greatly contribute to more and faster results. This contribution involves power (how much gets done) as well as speed (how fast it gets done). Having an apprentice to perform such labor-intensive tasks as running experiments and writing computer programs can enable researchers to complete tasks that they would never have attempted otherwise.

Teacher/apprentice collaborations differ from employer/employee ones in a crucial respect. Researchers work with graduate students not only to increase their own productivity but also to *train* the students. Training in experimental work is much more complicated than imparting knowledge of the sort available in print. Effective experiments and their statistical analysis usually involve a wealth of techniques that can only be acquired by working with someone who has already had experience with them. Doing science requires much more than knowledge *that*; it requires knowledge *how* to design experiments, construct apparatus, and interpret complex data statistically. The goal of apprenticeship is not simply to enhance an advisor's career, but to bring students along to the point where they can do effective research on their own. In Goldman's terms, teacher/apprentice collaborations have the potential to increase fecundity, since they produce new researchers who can go on to get results of their own.

Teacher/apprentice collaborations are the most common type in the sciences, but they are rare in the humanities. One reason for this discrepancy is that researchers in the sciences often have grant money that they can use to hire graduate students as research assistants. Humanities graduate students are in contrast typically funded (if at all) by teaching assistantships that do not involve working closely with a supervisor. A second reason is that the humanities do not obviously lend themselves to the kind of division of labor that is natural in sciences where students can be assigned time-consuming tasks in data collection. Many projects in the natural and social sciences are decomposable in ways that make it possible to apportion different parts of them to different people. Some of the parts, such as running experiments or writing computer programs, can be handled by students. A third reason why collaboration is rare in the humanities is simply tradition: young assistant professors never worked collaboratively with their advisors, so they do not expect to work collaboratively with their students. Effective collaboration requires communication and organizational skills needed to establish a useful division of labor and maintain progress. In sciences where collaboration is entrenched, these skills can be learned implicitly as part of graduate training, when an effective advisor provides a role model for how students can conduct collaborative research when they have their own students.

The best apprenticeships turn into full-fledged collaborations, as students develop into equals. In physics, it is not unusual for students to continue to work with their advisors after graduation, as part of the large research teams increasingly found in that field. In contrast, psychologists need to cut loose from their students since otherwise the students will not establish a strong enough research record on their own to get tenure and grants. Young psychologists, unlike young physicists, are expected to establish their own track record, whereas in physics the costs of research are often so great that independence would be too much to expect of a recent Ph.D.

The issue of how independent students need eventually to be from their advisors exemplifies a difficult ethical issue in collaborations: the apportionment of credit and blame. Obviously, of the 450 names on a physics paper, not all researchers made equal contribution to the work described. Some may have made not any intellectual contribution at all, but are included simply because they are part of the team or managed the enterprise. It becomes difficult to know whom to reward for desired results, or whom to blame for the production of error. Notoriously, credit for work done jointly by students and advisors goes unduly to the established researchers. Merton terms this the "Matthew effect", from the Gospel according to St. Matthew: "For unto every one that hath will be given, and he shall have abundance: but from he that hath not shall be taken away even that which he hath" (Merton 1973, p. 445). Many psychology departments demand that candidates for tenure have developed a research program independent of that of their Ph.D. supervisor. Although this requirement can be helpful in making possible an assessment of a researcher's independent capabilities, it sometimes leads to premature termination of fruitful collaborations. The assumption is that young researchers cannot earn the sort of credit they need for promotion if they are working with their original supervisors. This assumption is clearly too strong, since there are collaborations between students and advisors where it is clear that the student made most of the contribution. Problems of apportionment of credit also arise with collaborations among equals.

My discussion has been largely from the perspective of established researchers. Graduate students and junior researchers may have different and less pleasant reasons for engaging in collaborations, for example to secure funding and research resources. Collaboration of the employer/employee and teacher/apprentice types sometimes involves power issues that differ from the epistemic ones discussed in this paper.

4.3. Peer-Similar

Whereas employer/employee and teacher/apprentice collaborations involve people with substantial differences in knowledge and status, a less common kind of collaboration brings together established researchers with similar knowledge and interests. What would two researchers gain by working together rather than independently? It might be expected that the time spent in coordination and communication would merely subtract from time that could be spent on individual projects. In the branch of computer science concerned with parallel computation, *n* processors working together in a network are expected to produce less than an *n*-times increase in efficiency. Improvement is expected to be sublinear because of the efforts required for communication and coordination. If researchers have similar backgrounds, what can they gain by working with each other?

Surprisingly, computer simulations have shown that the sublinear expectation does not hold for groups of complex agents working on tasks that require some degree of intelligence (Clearwater, Huberman, & Hogg, 1991). The task used in the simulations was cryptarithmetic, which requires decoding letters into numbers in a way that makes true mathematical equations such as WOW + HOT = TEA. Clearwater et al. developed a computer system that has 100 agents working on such tasks cooperatively. Cooperation takes place by having each agent that is randomly generating and testing solutions, announce any progress it is making to the other agents by communicating a "hint" to the other agents. The interesting result was that *n* agents communicating in this way could together solve cryptarithmetic problems more than *n* times faster than all the agents working alone. The cause of the superlinear improvement seems to be that hints effectively reduce the size of the search space: having agents start off at different locations increases the likelihood that some will find hints worth communicating to other agents to reduce their subsequent search. Two heads working together thus can be more than twice as good as two heads working alone. Goldman's standards provide a way of seeing how something similar can hold in scientific research.

First consider reliability. Because it is easier to identify blunders in others than in oneself, peer-similar collaborations can improve reliability by virtue of members of a team noticing mistakes that would get past them working alone. Reliability can occasionally suffer, however, if it leads to increased sloppiness based in overconfidence in one's collaborator's. Work on decision making has identified the phenomenon of *groupthink*, that sometimes members of a group can end up with more confidence in a decision than each member would have alone (Janis, 1982). Similarly, one researcher's confidence in a result may be buttressed by the confidence of a collaborator which in turn is based partly on the confidence of the first researcher, so that confidence is more a function of group hysteria than of the validity of the result. Reliability may also suffer if one member of the team is weak but no one notices. There have been cases of scientific fraud in which one researcher fabricated data but the collaborators were strongly motivated to deny it (LaFollette, 1992). As Hardwig (1985, 1991) has pointed out, we are very much epistemically dependent on one another: much of what each of us professes to know depends on information that we have acquired from others that we trust. The cost of epistemic dependence of the sort especially notable between collaborators is that mistakes can enter and propagate within the system because of collaborators who are inept or corrupt. Hence collaboration between equals may decrease as well as increase reliability.

In the Clearwater et al. experiment, the different processors were virtually the same, but they developed different approaches to a given cryptarithmetic problem because of random generation of hypotheses. Similarly, even if two researchers are very similar, they will not pursue exactly the same solution to a problem because of subtle differences in what they know and in what they are exposed to. Collaboration can thus lead to increased power and speed as researchers working together produce more results in less time. There is no guarantee, of course. Loss of some power and speed in teacher/apprenticeship collaborations can be justified because of the long term fecundity benefit that a new researcher is being trained. With a peer-similar collaboration, it is an open question whether progress justifies the time spent working together, but the examples mentioned earlier (Crick and Watson, Newell and Simon, Kahneman and Tversky) suggest that great progress can be made. Fecundity and efficiency (cost) do not seem to be issues for peer-similar collaborations.

It is important to recognize a relative difference between the kinds of results that may accrue from peer-similar collaborations compared to the first two kinds I discussed. Since employees and graduate students are more likely to be able to do the time-consuming routine work that experiments require than they are to make theoretical contributions, the gains associated with the first two kinds of collaboration are most likely to concern empirical results. In contrast, comparable collaborators have most to gain from each other conceptually, making progress toward theoretical results, although they can also benefit from working together to produce novel experimental designs. Molecular biologists frequently help each other out using analogies between experiments: when a researcher's experiment is having problems, another researcher can describe a similar experiment that suggests a way of overcoming the problems (Dunbar, 1995). Peer-different collaborations are even more strongly directed toward theoretical rather than empirical results.

In addition to those suggested by Goldman's epistemic standards, there may be other reasons for collaboration. In the first place, it can be fun for reasons that are independent of power and speed. Having a collaborator means having someone you know is interested in discussing your research, thereby alleviating the loneliness of the long-distance scholar. Some researchers find it easier to develop new ideas in conversation rather than by individual thinking or writing. In addition, sociologists such as Latour (1987), who view science as an aggressive process of building alliances, might see collaborations as ways of accumulating the political power to have one's ideas become dominant. Enlisting collaborators is one way of increasing the competitiveness of one's research program (Durfee 1992). It is unlikely, however, that merely having a collaboration, as opposed to having a collaboration that produces good results, increases one's scientific success.

4.4. Peer-Different

Interdisciplinary fields such as cognitive science are the obvious places to look for collaborations between researchers with very different backgrounds. Here is a brief selection of cross-disciplinary collaborations in cognitive science, in alphabetical order:[1]

- Robert Abelson (psychology) and Roger Schank (artificial intelligence);
- Patricia Churchland (philosophy) and Terry Sejnowski (computational neuroscience);
- Allan Collins (psychology) and M. R. Quinlan (artificial intelligence);
- Jerry Fodor (philosophy) and Zenon Pylyshyn (psychology);
- Ken Forbus (artificial intelligence) and Dedre Gentner (psychology).
- Kris Hammond (artificial intelligence) and Colleen Seifert (psychology);
- Geoffrey Hinton (artificial intelligence), Jay McClelland (psychology), David Rumelhart (psychology), and Paul Smolensky (physics, artificial intelligence).
- John Holland (artificial intelligence), Keith Holyoak (cognitive psychology), Richard Nisbett (social psychology), and Paul Thagard (philosophy, artificial intelligence);
- Mark Johnson (philosophy) and George Lakoff (linguistics);
- Daniel Osherson (psychology) and Scott Weinstein (philosophy)
- Michael Posner (psychology) and Marcus Raichle (neuroscience)

Schunn, Okada, and Crowley (1995) surveyed papers presented at the annual meetings of the Cognitive Science Society and found a high frequency of interdisciplinary collaborations.

Because cross-disciplinary researchers use different methodologies (e.g. psychologists' experiments with human subjects versus AI researchers' computer simulations), we should not expect empirical results to be the primary benefit of peer-different collaborations. But there are huge gains to be made in the number and rate of theoretical results. These gains in power and speed can come about because cross-disciplinary collaboration brings together previously isolated theoretical ideas that can produce fruitful combinations. Gains, however, are typically not the immediate result of cross-disciplinary work, since much time and effort is usually required for people from different fields to begin to understand each other. Peer-similar collaborators with the same kinds of intellectual background can expect to understand each other's work quickly, but extensive cross-disciplinary education will be required for people from different fields to be able to work together productively.[2] Once these barriers are overcome, however, there is a great potential gain in fecundity, since collaborative results can be developed and used by many people in many different fields.

Cross-disciplinary collaboration might contribute to reliability, through triangulation of methods that lead to more robust results. An ideal cognitive science collaboration might be one that combined experiments on human behavior with computer simulations of that behavior with brain scans of how that behavior is implemented in the human brain. The behavioral, computational, and neurological experiments would ideally provide a way of converging on valuable empirical and theoretical results. On the other hand, reliability can suffer from interdisciplinary collaboration if people from different fields have no way of critically evaluating the results of unfamiliar methodologies. Peer-different collaborators are exceptionally epistemically dependent on their coworkers, since they typically lack the skill to validate work done in a different field. For example, most psychologists know little about the pitfalls of computer modelling, just as most artificial intelligence researchers know little about the design of experiments involving human subjects.

5. Explanatory Efficacy

Although Goldman's standards of reliability, power, speed, efficiency, and fecundity help us to understand why cross-disciplinary collaboration can be a valuable epistemic strategy, they neglect an important aspect of scientific thought. The growth of scientific knowledge is not just a matter of the quantity and reliability of results: some theoretical and empirical results are much more qualitatively important than others. Importance depends on the goals of inquiry. Scientists and ordinary people do not strive simply to accumulate true beliefs. In everyday life, we want to acquire true beliefs relevant to our goals. In science, we want to acquire true beliefs that are relevant to the goals of science, which include explanation and technological application as well as truth. The greatest explanatory accomplishments of science are unifying theories such as quantum theory and relativity theory in physics and evolutionary theory and genetics in biology. Cognitive science has not yet had its Newton, Darwin, or Einstein to provide a unified theory that applies to all kinds of thinking, but many of the collaborations that have arisen in cognitive science have been important because of the steps toward theoretical unification they have provided.

A mature scientific field should not just be a list of unconnected results, but should rather be unified by a common explanatory framework. Coherence is greatest when a small number of theoretical principles serve to explain a large number of empirical observations. These principles, such as Darwin's central claim that species evolve as the result of natural selection, assume great importance because of their capacity to explain diverse observations, and the observations increase in importance when they can unified with others by means of the theory. Scientific revolutions occur when an old theory is replaced by a new one with greater explanatory coherence (Thagard, 1992).

Many collaborations in cognitive science, including both peer-similar and peer-different ones, derive their importance from the thrust toward unified theories. No single dominant theory has emerged, but collaborators have developed competing views of mind that have tied together diverse phenomena. For example, John Anderson (1983, 1993) and Allan Newell (1990) have each worked with many collaborators to show how diverse aspects of thinking can be viewed from the perspective of rule-based systems. Similarly, David Rumelhart and James McClelland (1986) have worked with a host of collaborators to show the applicability of a different explanatory framework, parallel distributed processing.

Collaboration can increase explanatory coherence in two ways. First, collaboration, especially across disciplines, can produce conceptual combinations that establish new theoretical frameworks. Ideas about rule-based and PDP systems have depended on integration of psychological and computational inspiration, as have ideas about analogical reasoning (Holyoak and Thagard, 1995). Second, assembling a broad range of experimental results to be unified by a theoretical framework requires the participation of large groups of experimenters. Collaboration therefore greatly aids the production of unifying theories with demonstrably broad scope. Explanatory relations among the theoretical principles and the experimental results transform a set of independent results into a coherent whole (Thagard, 1992). I want accordingly to add to my version of Goldman's five standards the following:

6. The *explanatory efficacy* of a practice is how well it contributes to the development of theoretical and experimental results that increase explanatory coherence. Appreciation of the practice of collaboration among researchers of equal stature depends on seeing how collaboration can be aimed at and can contribute to explanatory efficacy of the results obtained. With this addition, Goldman's standards of social epistemic appraisal can shed considerable light on the advantages and disadvantages of collaboration in science, as I will now show by considering their application to several case studies.

6. Applications

The discussion so far has been programmatic and abstract, but can easily be made more concrete by showing how the six standards for evaluating the benefits of collaboration apply in specific cases. The first two are taken from very recent science, concerning the collaborative development of (1) the bacterial theory of ulcers and (2) the multiconstraint theory of analogy. I then argue that naturalistic philosophy could benefit from increased collaboration.

6.1 Marshall and Warren: Ulcers

Two Australian physicians, Barry Marshall and J. Robin Warren, initiated a dramatic recent shift in medical beliefs concerning the etiology of peptic ulcers. In 1981, when they began their collaboration at Perth Royal Hospital, the dominant medical opinion was that ulcers are caused by excess acidity. By 1995, it had become widely accepted that the primary cause of most gastric and duodenal ulcers is infection with bacteria, *Helicobacter pylori*. I have elsewhere described the psychological, physical, and social processes underlying the development and acceptance of the bacterial theory of ulcers (Thagard, forthcoming). My concern here is much more narrow, to evaluate the collaboration among Marshall, Warren and numerous colleagues according to the epistemic standards of reliability, power, fecundity, speed, efficiency, and explanatory efficacy.

Collaboration was essential for the development of the bacterial theory of ulcers because of the involvement of several different medical specialties. Warren, a pathologist, noticed unusual spiral bacteria in gastric biopsies he examined in 1979, but he was unsure whether the bacteria had any medical significance. In 1981, he began working with Barry Marshall, a trainee in gastroenterology. Together they devised an experiment that found an association between the bacteria and peptic ulcers (Marshall and Warren, 1984). Subsequently, they were able to show that ulcers can often be cured using antibiotics (Marshall et al., 1988), a result since confirmed by a wealth of other studies. In addition to Warren and Marshall's expertise in pathology and gastroenterology, the experiments required the know-how of specialists in microbiology, electron microscopy, and pharmacy. Marshall et al. (1988) has a total of nine co-authors drawn from four medical specialties.

Barry Marshall won the prestigious 1995 Lasker award for the discovery that *H. pylori* causes peptic ulcer disease, but it is obvious that he could not have produced the important development on his own. As a gastroenterologist, he lacked the expertise to do the pathology and microscopy work required to identify the presence of the bacteria. Similarly, Warren lacked the expertise to do endoscopy to get the stomach biopsies. Having non-experts using instruments would have drastically reduced the reliability of the experimental results. The power and fecundity of collaboration is also evident in this case, since Marshall and Warren acting without each other and without microbiologists and other assistants would never have been able to answer the question of whether *H. pylori* is a causal factor in ulcer disease. Speed and efficiency are lesser benefits of this collaboration, although it is clear that without a team of people to perform such tasks as biopsies, bacterial cultures, and statistical analyses the research would have required much more time.

Collaboration can enhance explanatory efficacy in two ways, by contributing to theoretical and experimental results. We have already seen that the experimental results explained by the hypothesis that bacteria cause ulcers depended on collaboration. Although Marshall was mainly responsible for the formation and promulgation of this hypothesis, Warren contributed crucial links to his chain of reasoning, informing Marshall that the bacteria are associated with gastritis. When Marshall read that gastritis is associated with ulcers, he conjectured that the bacteria might be associated with ulcers (Marshall, 1989). Hence collaboration contributed theoretically as well as experimentally to explanatory coherence. In sum, the collaboration among Marshall, Warren, and their colleagues scores very high on the standards of reliability, power, and explanatory efficacy.

6.2 Holyoak and Thagard: Analogy

In our 1995 book *Mental Leaps:* Keith Holyoak and I defended a theory of analogical thinking derived from more than a decade of collaborative work that combined Holyoak's expertise as an experimental psychologist and my background in philosophy and computational modelling. The preface of the book thanks more than thirty additional collaborators, including, for example, psychology graduate students who worked with Holyoak at Michigan and UCLA, and programmers who worked with me at Princeton and Waterloo. Figure 4 summarizes some of the

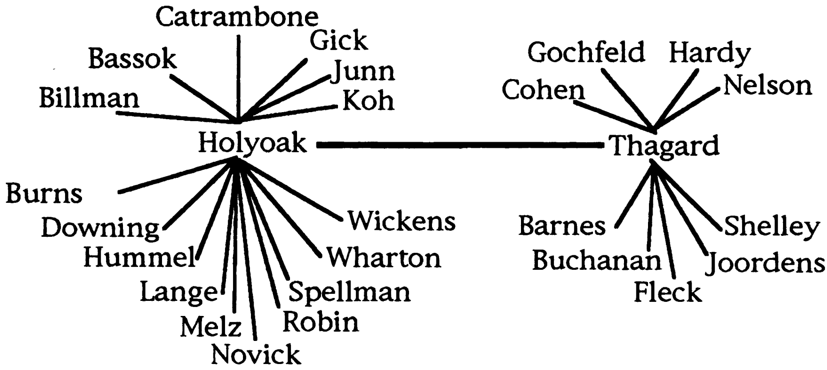

Figure 4. The collaborative nature of Holyoak and Thagard (1995). Names above Holyoak's are his Michigan collaborators; those below worked with him at UCLA. Names above Thagard's are his Princeton collaborators; those below worked with him at Waterloo. All collaborators named were co-authors of published papers.

collaborations that made possible the theoretical, experimental, and computational work that went into our book.

Like many projects in cognitive science, this work combined several different methodologies that surpassed any individual's expertise. While Holyoak and I worked jointly on theoretical issues, only he had the expertise to conduct rigorous psychological experiments, but he did not have sufficient computational background to write the computer programs that were essential to testing our theoretical claims against the psychological evidence. If Holyoak had tried to do the computer simulations, or if I had tried to do the psychological experiments, there would have been a dramatic loss in reliability. Instead, by combining experimental and computational skills, we were able to generate answers to many interesting questions concerning how people think analogically. It would be presumptuous to speak of the truth of our findings, but publication in respected journals such as *Cognitive Science* and *Artificial Intelligence* legitimates talk of results; hence, given my weakened version of Goldman's standards, we can say that our collaboration contributed to power and fecundity. The wealth of other collaborators also greatly enhanced the power of our research project, along with its speed and efficiency. In principle, Holyoak could have done all the experiments and I could have done all the simulations, but both enterprises are extremely time consuming and neither of us would have been able to do more than a fraction of what was eventually produced. Thus experimental results, and the computational simulations that tied them with theory, were greatly fostered by collaboration.

Theoretical developments were more narrowly collaborative, involving mostly Holyoak and me. Holyoak had the initial inspiration in 1987 that grew into our theory that analogical thinking fundamentally involves three kinds of constraints (structure, similarity, purpose), but subsequent theoretical contributions that emerged during joint writing of numerous articles and the book are so entwined that it would be virtually impossible to disentangle them. Hence on the presumption that it is legitimate to talk of theoretical and experimental results in this case, it is clear that explanatory efficacy was greatly aided by multiple collaborations. As in the ulcers case, collaboration contributed to the development of the multiconstraint theory of analogy according to all of the six epistemic standards, but most notably power and explanatory efficacy.

6.3 Why Philosophers Should Collaborate More

In section 4.2, I sketched some of the reasons why teacher/apprentice collaborations are so rare in the humanities: absence of grant money, lack of natural division of labor, and cultural tradition of solitary work. Despite these barriers, I am convinced that at least in philosophy there is much to be gained by increased collaboration, involving both peer-similar and teacher-student combinations. I will argue that the increasing influence of naturalistic approaches to philosophy points toward the need for increasing collaboration, and that training of philosophy graduate students would also be improved by increased collaboration.

Many areas of philosophy are now imbued with naturalism, seeing philosophical issues as continuous with scientific investigations. Epistemology is increasingly naturalistic, tying traditional concerns of justification and truth with empirical issues about human psychology and sociology. Metaphysics can be viewed as an extension of science, for example when issues in the philosophy of mind are integrated with developments in psychology and neuroscience. Psychological and computational results are even impinging on recent work in ethics and logic.

Because serious naturalistic philosophy requires knowledge of work in cognate fields, and since acquiring deep knowledge of fields such as psychology is a difficult and time-consuming task, philosophers can greatly benefit from collaboration with experts in those fields. The potential epistemic benefits of such peer-different collaborations include reliability, since working closely with an expert in another field increases the likelihood that the naturalistic philosopher will bring appropriate results to bear on philosophical problems. Potentially, there is also on increase in power, if empirical expertise can help to answer philosophical questions. (I avoid the question of whether philosophy achieves "results" or "truths".) Explanatory efficacy can also potentially be a benefit of increased collaboration between philosophers and other researchers, if theoretical progress is made by combining philosophical concerns with theoretical issues in psychology, linguistics, physics, or other fields.

Naturalistic philosophers can also benefit from collaboration with each other. It is becoming increasingly common for individual philosophers to know a lot about a particular field such as psychology, or linguistics, or artificial intelligence, or neuroscience, or physics. It is rare, however, for this extra-philosophical expertise to extend beyond a single field, so it should be natural for two philosophers, each with detailed knowledge of two related fields such as linguistics and psychology, to collaborate with each other, just as linguists and psychologists sometimes collaborate with each other.

Student/apprentice collaborations can also benefit from the fact that graduate students in philosophy often have backgrounds that complement the knowledge of their professors. Working collaboratively with graduate students in philosophy at the University of Waterloo, I have benefited from students' knowledge on topics with which I was comparatively unfamiliar, such as the history of mathematics, mythology, dynamic systems theory, and empathy. Hence a student/apprentice collaboration can take on the characteristics of a peer-different one.

Increased collaboration between philosophy professors and their graduate students can also improve training in philosophical research. In addition to gaining the intellectual benefits of work with established researchers, graduate students can acquire practical knowledge such as how to write, revise and submit papers and deal with the editorial process. As in the sciences, one benefit of collaboration at the beginning of a graduate student's career is early publications, no small matter in an increasingly competitive job market.

But my current argument is primarily epistemic: within the context of naturalism, collaboration between philosophy professors and graduate students has the potential to increase the reliability, power, and explanatory efficacy of philosophical research. To determine the extent to which this potential can be realized, it will be necessary to conduct a social experiment in which philosophy professors engage their students much more directly in their research projects.

The epistemic contribution of collaboration to philosophy is still largely a matter of potential, but the bacterial theory of ulcers, the multiconstraint theory of analogy, and innumerable theoretical and experimental results of current science show that scientific collaboration is a valuable epistemic practice. In this paper I have distinguished four different kinds of collaboration and shown how they affect six different standards of epistemic appraisal. Perhaps future work on the epistemology of collaboration will be more collaborative.

Notes

*This research is supported by the Social Sciences and Humanities Research Council of Canada. Thanks to Kathleen Gorman for research assistance. I have benefited from conversations with Dedre Gentner and Ziva Kunda, and from a suggestion by Sharon Derry. Gary Olson and Ed Hutchins and two anonymous referees provided valuable comments on an earlier draft. Correspondence and requests for reprints should be sent to the author at Philosophy Department, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1. Email: pthagard@watarts.uwaterloo.ca.

[1]For an overview of cognitive science, see Thagard (1996).

[2]For example, John Holland, Keith Holyoak, Richard Nisbett and I spent more than a year meeting and talking regularly before we considered doing a book together.

References

Anderson, J. R. (1983), *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. (1993), *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Barnes, B. (1985), *About science.* Oxford: Blackwell.

Bloor, D. (1991), *Knowledge and social imagery* (second ed.). Chicago: University of Chicago Press.

Bond, A., & Gasser, L. (ed.), (1988), *Readings in distributed artificial intelligence.* San Mateo: Morgan Kaufmann.

Clearwater, S., Huberman, B., & Hogg, T. (1991), "Cooperative solution of constraint satisfaction problems", *Science 254:* 1181-1183.

Djerassi, C. (1989), *Cantor's dilemma.* New York: Doubleday.

Dunbar, K. (1995), "How scientists really reason: Scientific reasoning in real-world laboratories". In R. J. Sternberg & J. Davidson (eds.), *Mechanisms of insight.* Cambridge, MA: MIT Press, pp. 365-395.

Durfee, E., Lesser, V., & Corkhill, D. (1989), "Cooperative distributed problem solving". In A. Barr, P. Cohen, & E. Feigenbaum (eds.), *The handbook of artificial intelligence, Volume IV.* Reading, MA: Addison-Wesley, pp. 83-147.

Durfee, E. H. (1992), "What your computer really needs to know, you learned in kindergarten". In *Proceedings of the Tenth International Conference on Artificial Intelligence* . Menlo Park: AAAI Press/The MIT Press, pp. 858-864.

Galegher, J., Kraut, R. E., & Egido, C. (ed.), (1990), *Intellectual teamwork: social and technological foundations of cooperative work.* Hillsdale, NJ: Erlbaum.

Gasser, L. (1991), "Social conceptions of knowledge and action: DAI and open systems semantics.", *Artificial Intelligence 47:* 107-138.

Giere, R. (ed.), (1992), *Cognitive models of science.* Minneapolis: University of Minnesota Press.

Goldman, A. I. (1992), *Liaisons: Philosophy meets the cognitive and social sciences.* Cambridge, MA: MIT Press.

Hardwig, J. (1985), "Epistemic dependence", *Journal of Philosophy 82:* 335-349.

Hardwig, J. (1991), "The role of trust in knowledge", *Journal of Philosophy 88:* 693-708.

Hewitt, C. (1991), "Open information systems semantics for distributed artificial intelligence", *Artificial Intelligence 47:* 79-106.

Holyoak, K. J., & Thagard, P. (1995), *Mental leaps: Analogy in creative thought.* Cambridge, MA: MIT Press/Bradford Books.

Janis, I. L. (1982). *Groupthink: Psychological studies of policy decisions and fiascoes* (2nd ed.). Boston: Houghton Mifflin.

Kitcher, P. (1990), "The division of cognitive labor", *Journal of Philosophy 87*: 5-22.

LaFollette, M. C. (1992), *Stealing into print: Fraud, plagiarism, and misconduct in scientific publishing.* Berkeley: University of California Press.

Latour, B. (1987), *Science in action: How to follow scientists and engineers through society.* Cambridge, MA: Harvard University Press.

Latour, B., & Woolgar, S. (1986), *Laboratory life: The construction of scientific facts.* Princeton, N.J.: Princeton University Press.

Marshall, B. J. (1989). "History of the discovery of C. pylori." In M. J. Blaser (Eds.), *Campylobacter pylori in gastritis and peptic ulcer disease* (pp. 7-22). New York: Igaku-Shoin.

Marshall, B. J., Goodwin, C. S., Warren, J. R., Murray, R., Blincow, E. D., Blackbourn, S. J., Phillips, M., Waters, T. E., & Sanderson, C. R. (1988). "Prospective double-blind trial of duodenal ulcer relapse after eradication of campylobacter pylori." *Lancet,* 2(8626/8627), 1437-1441.

Marshall, B. J., & Warren, J. R. (1984). "Unidentified curved bacilli in the stomach of patients with gastritis and peptic ulceration." *Lancet, 1*(8390), 1311-1315.

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

Newell, A. (1990), *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Olson, J. S., Card, S. K., Landauer, T. K., Olson, G. M., Malone, T., & Legett, J. (1993), "Computer supported cooperative work: Research issues for the 90s", *Behavior and Information Technology 12*: 115-129.

Resnick, L., Levine, J., & Behrend, S. (ed.), (1991), *Socially shared cognitions.* Hillsdale, NJ: Erlbaum.

Salomon, G. (ed.), (1993), *Distributed cognitions.* Cambridge: Cambridge University Press.

Schunn, C. D., Okada, T., & Crowley, K. (1995). Is cognitive science truly interdisciplinary? The case of interdisciplinary collaborations. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* . Mahwah, NJ: Erlbaum. 100-105

Solomon, M. (1994), "Social empiricism", *Nous 28*: 325-343.

Thagard, P. (1988), *Computational philosophy of science.* Cambridge, MA: MIT Press/Bradford Books.

Thagard, P. (1992), *Conceptual revolutions*. Princeton: Princeton University Press.

Thagard, P. (1993), "Societies of minds: Science as distributed computing", *Studies in History and Philosophy of Science 24*: 49-67.

Thagard, P. (1994), "Mind, society, and the growth of knowledge", *Philosophy of Science 61:* 629-645.

Thagard, P. (1996). *Mind: Introduction to cognitive science.* Cambridge, MA: MIT Press.

Thagard, P. (forthcoming). "Ulcers and bacteria: I Discovery and acceptance; II Physical and Social Interactions." *Studies in History and Philosophy of Science.*