How to Collaborate:
Procedural Knowledge in the Cooperative Development of Science

Paul Thagard
University of Waterloo
pthagard@uwaterloo.ca
Draft 4, August, 2005; comments welcome.

Introduction

A philosopher once asked me: "Paul, how do you collaborate?" He was puzzled about how I came to have more than two dozen co-authors over the past 20 years. His puzzlement was natural for a philosopher, because co-authored articles and books are still rare in philosophy and the humanities, in contrast to science where most current research is collaborative. Unlike most philosophers, scientists know how to collaborate; this paper is about the nature of such procedural knowledge.

I begin by discussing three related distinctions found in philosophy and cognitive science: knowledge how vs. knowledge that, procedural vs. declarative knowledge, and explicit vs. implicit knowledge. I then document the prevalence of collaboration in the sciences and its scarcity in philosophy. In order to characterize the sorts of procedural knowledge that make collaborative research possible and fruitful, I discuss how scientists collaborate, how they learn to collaborate, and why they collaborate. Contrary to some recent suggestions by philosophers, I will argue that knowledge how often does not always reduce to knowledge that, and that collaboration has many purposes besides the pursuit of power and resources. The relative scarcity of philosophical collaborations might be explained by the nature of philosophy, if the field is viewed as inherently personal or a priori. But I argue against this view in favor of a more naturalistic one, with the implication that the main reason why philosophers do not collaborate more is that they do not know how.

August 10, 2005

My account of collaboration is based on my own experience, published advice by practicing scientists, and interviews with a group of highly successful scientific collaborators who are members of the Social Psychology area of the University of Waterloo Psychology Department. For the past two decades, Waterloo's social psychology program has flourished, both in collaborative publication and in graduate training: their former Ph.D. students are now employed in many universities worldwide, including McGill, Toronto, Stanford, and UCLA. I interviewed 6 faculty members and 1 graduate student, asking them why they collaborate, whom they collaborate with, what they have learned about how to collaborate, and how they learned it.

Procedural vs. Declarative Knowledge

Gilbert Ryle (1949) introduced into philosophy the distinction between knowing how and knowing that. Knowing how is the ability to do certain sorts of things, which Ryle argued was different from knowledge of truths. Similarly, Michael Polanyi distinguished between tacit and explicit knowledge and argued that science depends heavily on tacit knowledge that cannot be made explicit (1958; 1967). In the late 1960s, researchers in the field of artificial intelligence introduced a distinction between declarative and procedural representations of knowledge, where the latter consisted of programmed functions for answering particular questions (Barr & Feigenbaum, 1981, 173- 179). The distinction between declarative and procedural knowledge was carried over into psychology by researchers such as John Anderson (1983), although his procedures consist of specifiable rules. In contrast, Ryle, Polanyi, and the AI proceduralists would reject the claim that procedural knowledge can be captured by explicit rules.

A related distinction was proposed by psychologists in the 1980s: implicit vs. explicit memory (Schacter, 1996). Whereas explicit memory involves conscious recollection of specific facts and episodes, implicit memory involves non-conscious change in performance on tasks in which people learn motor skills. There is psychological and neurological evidence that implicit memory operates differently from explicit memory. For example, there are brain-damaged patients with organic amnesia who are incapable of learning new facts but are capable of skill learning. Neuroimaging studies have suggested that explicit and implicit memory systems operate in different areas of the brain.   Hence we have empirical reasons to believe that there is something to the distinction between implicit/tacit/procedural/knowledge how, and explicit/declarative/knowledge that.   I will return later to the question of whether these kinds of distinctions stand up in the particular case of knowing how to collaborate.

Collaboration

In an earlier article, I documented the prevalence of collaborative research in science in 1992 (Thagard, 1997, pp. 244-245):

 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. 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.

A cursory glance at recent volumes of these journals shows that scientific research continues to be highly collaborative. In contrast, only 4 of 27 articles in the 1992 issues of the *Journal of Philosophy* were co-authored, and the 2004 volume had only 3 collaborative contributions.

In my 1997 paper, I distinguished 4 different kinds of collaborative relationships:

1. Employer/employee, in which a researcher tells a subordinate what to do.

2. Teacher/apprentice, in which a researcher trains a student to do research.

3. Peer-similar, in which researchers of similar knowledge, interests, and status decide to work together.

4. Peer-different, in which researchers from different disciplines but similar status collaborate.

I have been fortunate to have had collaborations from all 4 of these categories. All of them involve different kinds of procedures which I will now try to categorize.

Procedural Knowledge for Collaboration

Employer/employee

The most straightforward kind of collaboration is employer/employee. In the sciences, such relationships arise when people are hired to perform tasks such as running experiments and programming computers. Here the question of working effectively with coworkers is no different from what arises in any managerial situation. Being a good manager requires many interpersonal skills needed to improve the performance of employees. According to Goleman (1995, p. 152): "Many managers are too willing to criticize, but frugal with praise, leaving their employees feeling that they only hear about what they are doing when they make a mistake." Good managers know

how to motivate, encourage, and correct their employers, and effective research leaders must acquire the same kinds of procedural knowledge. Hill (2003) describes common mistakes made by new managers, such as not realizing that rewards have advantages over punishments in controlling subordinates' behavior. Other challenges to new managers is learning how to delegate, and figuring how to function as a coach as well as a conductor. Below I will discuss the extent to which this kind of knowledge consists of or can be translated into verbal rules.

Teacher/apprentice

The most prevalent and important kind of collaboration in science is teacher/apprentice. Unlike philosophy, where graduate students mostly conduct their research independently of their supervisors, students in the sciences are usually trained by close collaboration with faculty. In psychology, for example, new Ph.D. students are normally matched with advisors from the start, with the expectation that a student will immediately become part of the advisor's research team. Students typically start off doing a small part of the planned research, running a particular experiment and analyzing the resulting data. Normally they progress to the point where they can conduct experimental research on their own. Successful psychology students acquire a huge amount of procedural knowledge, such as how to design a good experiment, how to analyze the data insightfully, and how to write an effective research article. But there is another sort of procedural knowledge to be acquired, namely how to collaborate. Given the labor-intensive nature of experimentation in psychology and other sciences, no researchers are likely to be very successful unless they are able to recruit, train, and motivate graduate students who will contribute much of their research.

All the social psychologists I interviewed talked about the importance of collaborating with students, which is crucial both for the pedagogic purpose of training them how to do research and for the professional purpose of getting the professor's research done. The faculty members take very seriously their role as mentors, convinced in the tradition of the field that the best way (probably the only way) for students to become social psychologists is for them to work closely with someone who already knows how to do successful research. Emphatically, graduate students are not treated as research assistants who are merely assigned tasks; this role is more commonly filled by undergraduate students, who are often supervised by the graduate students. Instead, graduate students are treated as apprentices with the goal that by the time the they get their doctorates they will have become fully accomplished researchers. This kind of graduate education is a sort of symbiosis, in which the students benefit from training and mentoring, while the faculty benefit from having intelligent and industrious collaborators. Students also provide some skills that the faculty may be lacking, such as familiarity with the latest statistical and computational tools.

Two leading social psychologists, Mark Zanna and John Darley have written about how to manage the relationship between faculty and graduate students. They say (2004, p. 118) that "the central issue to be deal with initially by faculty members is how to identify a research project on which their graduate students will work." They recommend against simply handling graduate students a project, since graduate students are supposed to be more than research assistants:

Instead, we give our students an idea with, perhaps, some thoughts about how that idea might be turned into a study. The student, then, does some

background reading, develops some ideas of his or own, and together we design the study. Although we may have thought (perhaps a lot) about a design and a procedure beforehand, our students are genuinely engaged in working with us on the design and procedure. (pp. 118-119)

Zanna and Darley recommend having graduate students working simultaneously on at least two lines of research, one stemming from the students' interests and the other from those of the faculty member. Students should also be encouraged to work with more than one faculty member. Ideally, a doctoral thesis comes from an independent line of research initiated by the student, signaling the transition from student to autonomous researcher.

Based on their own experience with many graduate students, Zanna and Darley recommend that a faculty member should aim to have three or four graduate students spaced out at different stages in their careers. They also provide advice about how to lead students to discover for themselves how to design experiments and how to deal with questions of order of authorship on published papers. They also sensibly advise against faculty members forming overly friendly personal relationships with their students, and against advising too many students outside one's own area of expertise. One of the social psychologists I interviewed spoke of the importance of making sure that students are getting research done without monitoring them too closely.

These valuable suggestions might be encoded as a set of verbal rules, the ten or more commandments of working with graduate students. But there are many more subtle interpersonal skills that are not so easily encoded. The distinguished social psychologist Richard Nisbett learned about doing experiments from the reactions he got

in discussions with his supervisor, Stanley Schacter. Nisbett says (personal communication, Feb. 23, 2001) "He let me know how good my idea was by grunts: non-committal (hmmm...), clearly disapproving (ahnn...) or (very rarely) approving (ah!)." Interestingly, many of Schacter's students went on to become very successful social psychologists, as have many of Nisbett's students. Perhaps one of the most important skills that an advisor can attempt to pass on to students is taste: how to identify a topic or project that is actually worth doing.   In psychology, much dreary work consists of experiments about experiments that merely ape prevailing experimental methods, without promising much psychological insight.   Similarly, in philosophy too much writing is minor variations on overworked themes, with no awareness of why a problem might have been philosophically interesting in the first place.   Nonverbal communication via both positive and negative emotions may help direct students in useful directions.

One of the most delicate social pedagogical skills, as useful for parenting as it is for teaching, is how to correct people's errors without hurting their feelings or crushing their spirit.   Motivation is a major part of success in science or anything else, and nothing will destroy motivation faster than communication from a respected superior that one is stupid or useless.   On the other hand, nothing is learned from content-free positive encouragement.   In order to gauge the level of positive and negative advice warranted for a particular student, a teacher (or parent or manager) has to know how to estimate the emotional state of the student and convey information in ways that promote improvement but not discouragement.   This is a kind of emotional intelligence that is very difficult to put into words, as I discuss further below.

Peer-similar

Once a student has a Ph.D., he or she needs to get a research program underway. In the sciences, this can involve reversing roles and applying procedural knowledge picked up as a supervisee to acting as a supervisor. New professors who had good graduate advisors should have implicitly or explicitly picked up knowledge of how to initiate and conduct research that they can use with their own students, carrying on teacher/apprentice collaborations from the other direction. But it is also a good strategy to find peer-similar collaborators for the following reasons listed by Zacks and Roediger (2004, pp. 150-151).

Often, two heads really are better than one. Collaboration allows you to multiply your effectiveness, and some problems simply are not solvable without combining multiple sources of expertise. Collaboration stimulates new thinking: Two people thinking together often generate ideas that neither would come up with alone. Collaboration is fun: One of the pleasures of academia is its potential for collegial interactions, and rewarding research collaborations can be the most rewarding intellectual relationships one has. Finally, collaborating is strategic: Collaborating with other junior colleagues helps build a strong cohort. Collaborating with senior colleagues strengthens your ties to the department and makes it easier for the senior faculty to know how your research is going when it comes times for evaluations and an eventual tenure decision. Collaborating with colleagues at other universities integrates you into the

broader research community and … may provide access to resources or expertise not available at your institution.

Zacks and Roediger also point out possible downsides to joint work if one ends up working with unreliable or incompetent collaborators.

The social psychologists I interviewed have had many peer-similar collaborations, and provided much advice about how to make them work.   First, it is crucial that the collaborators respect each other and are capable of dealing with each others as equals. Disagreements inevitably arise in the course of conducting experiments and writing articles, and the collaboration will dissolve if one of the co-authors tries to take too much control of the research, or too much credit for it.   Second, the collaborators must be able to trust each other.   Even in a peer-similar collaboration, the researchers will bring somewhat different knowledge and skills to the joint enterprise, and there will be some division of labor.   Hence researchers must choose collaborators whom they can trust to carry out their share of the work reliably.   Third, the collaborators must be able to keep their egos out of the joint work, being confident about putting forward ideas without getting defensive when they are challenged.

The simplest kind of peer-similar collaboration is when two researchers do joint work, but many other arrangements are possible.   A favorite among the social psychologists I interviewed is for two faculty members to jointly supervise a graduate student.   The student gets advice from more than one expert and has supervision even when one of the faculty is on sabbatical.   The faculty members benefit from pooling their skills and expertise.   One strength of the Waterloo social psychology program is that students are strongly encouraged to work with more than one faculty member,

guaranteeing their exposure to different ideas and styles of research.   It also provides them with insurance that if one line of research falters empirically they have another line to pursue.   One of the social psychologist engages in very large scale international collaborations with many researchers, which I will discuss as a peer-different collaboration.

From my own experience and others I have observed, here are some tentative rules for working with peers.   Work with people who share your general intellectual stance, but who have areas or methods that complement yours.   Have fun: the interchange should be intellectually stimulating for both of you.   Work out a good division of labor and method of working that suits you both.   Some collaborators write articles together with both talking and one typing, but my own preference has always been for alternating drafts.   Levine and Moreland (2004) describe a long collaboration that shifted from writing together to doing successive drafts.   Unlike employer/employee and teacher/apprentice collaborations, there is no authority ranking in peer-similar collaborations, so the peers need to work out ways of dealing with intellectual and literary disagreements.   Here the normal sorts of interpersonal skills and emotional intelligence that are valuable in everyday life will serve the collaborators well.

One of the most fruitful collaborations between psychologists is the joint work of Daniel Kahneman and Amos Tversky, which has been very influential in psychology and economics.   Here is how Kahneman (2003, p. 723) describes the beginning of his collaboration with Tversky:

The experience was magical. I had enjoyed collaborative work before, but

this was different. Amos was often described by people who knew him as

the smartest person they knew.   He was also very funny, with an endless

supply of jokes appropriate to every nuance of a situation. In his presence,

I became funny as well, and the result was that we could spend hours of

solid work in continuous mirth.

They spent most of a year working together on their enormously successful 1974 *Science*

article on heuristics and biases, and three years producing the 1979 *Econometrica* paper

on prospect theory that founded the field of behavioral economics.   Writing together,

they would sometimes spend delightful hours advancing a paper by only a single

sentence.

Peer-similar collaborations are normally enjoyable and productive, but they can

also have negative effects on a research career.   Roediger and Balota (2004, p. 399)

warn:

Collaborate but watch out. You can get sucked into collaborations that

take an immense amount of time but do not lead to publishable products,

or, if the product is published, you do not get much credit. Watch out for

colleagues to whom "collaboration" means "be my research assistant and

carry out or oversee this project on which I'll be first author."

According to Taylor and Martin (2004, p. 370): "Too much collaboration with senior

people, too early in your career, places you at some risk."   The problem is that a junior

researcher may not get due credit because of the existing reputation of a senior

collaborator.   Best (2004, p. 350) advises: "One should pick collaborators carefully! …

Successful collaborations grow out of contacts with colleagues who share similar

interests, who have work styles that fit together well but are not necessarily the same,

and who respect and value the relationship. Usually, these sorts of relationships do not just happen – they must be carefully cultivated and nurtured."

Peer-different

Now consider what is potentially the most exciting kind of collaboration, peer-different, in which one allies with another researcher with different knowledge and skills.   My first experience with collaboration was of this nature, when I began working with the social psychologist Richard Nisbett.   We published three articles together, and joined with the computer scientist John Holland and cognitive psychologist Keith Holyoak to produce in 1986 a successful book, *Induction: Processes of Inference, Learning, and Discovery*. Nisbett had extensive experience of peer-similar and teacher/apprentice collaborations, and Holyoak had some experience of both. Nevertheless, we were all in uncharted territory in trying to integrate ideas from philosophy, computer science, and two branches of psychology.   Nevertheless, we persevered and along the way learned some useful things about peer-different collaboration.   Subsequently, Holyoak went on to have numerous peer-different collaborations with neuroscientists, and I have worked with students from many different backgrounds, including philosophy, psychology, computer science, and engineering.

What have I learned about how to do this kind of collaboration?   Probably the most important piece of relevant declarative knowledge is that people from different disciplines may vary a great deal in their methods and conceptual schemes, so that the same terms may have different meanings.   For example, my co-authors and I, in our initially meetings, were misled by the diversity of our uses of the word "schema".   The procedural lesson is that if you want to do interdisciplinary, peer-different research, then

you need to be willing to spend a great deal of time listening to and learning from your collaborators in order to find out how they think about things. Stating this as a rule does not fully capture the acquisition of the kind of patient, attentive, conversational style that is necessary to elicit mutual understanding.

One of the social psychologists I interviewed, Geoff Fong, is involved in a massive international collaboration with 19 investigators in many countries, all looking at ways to control tobacco use. The investigators combine 8 disciplines, and the success of the project depends on their being able to appreciate and respect the widely varying knowledge and skills of their collaborators. The primary investigators need to be able to lead, but also to give way when appropriate. In such a widely multidisciplinary investigation, respect, trust, and leaving your ego at home are even more important than in smaller, peer-similar collaborations.

I recently collaborated on a multiply authored, multidisciplinary paper. In order to write a paper called "Is the brain a quantum computer?", I assembled a group of 5 experts in neuroscience, physics, and computer science, with all of us feeding ideas to a talented undergraduate who wrote most of the paper (Litt et al, forthcoming). In one respect, this organization was modeled after what Richard Nisbett did in putting together the group that eventually produced the book *Induction,* but it is novel in having an undergraduate serve as the focal point of the group. Geoff Fong calls this the *Ocean's Eleven* model of interdisciplinary collaboration, after the movie in which a gang of people with different special talents is assembled in order to carry out a creative robbery.

According to Epstein (2005), the key elements to a successful interdisciplinary collaboration are attitude, communication, time, proximity, institutional climate, funding,

roles, appropriate topics, and publication. Interdisciplinary collaborators should know how to be open, communicative, efficient in managing time, and committed to working together. They also need to be concerned with arranging adequate research funding and operating in a supportive environment. Thagard (2005) describes the important contribution of institutions and organizations to fostering interdisciplinary research in cognitive science. Thus peer-different collaboration, like the other three kinds of cooperative research, requires a large amount of procedural knowledge.

Learning to Collaborate

How do scientific researchers acquire knowledge about how to collaborate? There are many possible sources of this information, ranging from explicit instruction to implicit role modeling. The most explicit kind of instruction is the kind of advice I have quoted from various chapters in Darley, Zanna, and Roediger (2004). But most researchers do not acquire their procedural knowledge from such general sources, but rather by interaction with supervisors during their graduate training. Sometimes instruction can be highly explicit, if supervisors vocalize their collaborative practices. For example, here is a rule I learned from Richard Nisbett: When research has progressed enough to be written up, the first draft is done by whoever is most capable of producing it, and that person becomes the first author. This rule is obviously much fairer than the practice of some researchers that the senior faculty member who initiated the project should be the first author. The structure of this kind of learning is: Experienced collaborator tells neophyte about a useful way of collaborating; so neophyte adopts this way of collaborating. We could characterize this as rule-based procedural learning.

Normally, however, learning to collaborate is much more subtle, and people new to the experience may not realize what they are picking up from working with more experienced collaborators. Without thinking much about their experiences, a graduate student may acquire models of the following sorts of practices crucial to future collaborations: how to get a student involved in a project, how to run a laboratory group meeting, how to divide labor in the conduct of a research project, how to write a joint paper, how to jointly revise a paper in response to reviewer's criticisms, and so on. The structure of this kind of learning is: Experienced collaborator displays a practice or behavior; neophyte stores memory of this practice or behavior; so when the occasion arises, the neophyte duplicates this practice or behavior. We could characterize this as case-based (analogical) procedural learning, because what the neophyte has acquired is not generalizations but cases, possibly repeated, of practices and behaviors. If pressed, the neophyte might be able to generate a rule describing what to do, but this is not required for the neophyte to apply what has been learned from the experienced collaborator. All that needs to be remembered are such situations as: THIS is what we did to design that experiment, or THAT is what we did to write that paper.

A third kind of procedural learning is even more subtle, because it depends on a kind of implicit social cognition. A neophyte may not have any conscious awareness or memory of the physical, intellectual, or social behavior of the experienced collaborator but nevertheless encode and eventually duplicate them. There are simple physical examples. Reportedly, many of Wittgenstein's philosophical students adopted some of his odd mannerisms, and I know a number of psychology graduate students who adopted some of the behaviors (smile, laugh) of their supervisors. More complex, and harder to

identify, would be cases where the neophyte acquires patterns of emotional interaction from a collaborator in the same way that children learn patterns of emotional interaction from their parents.   The structure of this kind of learning is: Experienced collaborator displays physical, emotional, and social behaviors; neophyte unconsciously acquires these behaviors. Thus for a neophyte researcher, an experienced collaborator may function as a role model operating at three different levels: explicit instruction, patterns of behavior and research practice, and patterns of social and emotional interaction.   Of course, there is also room for procedural creativity, as in the team I put together to write the paper on quantum computing I mentioned earlier.

Having senior role models is probably the main way in which scientists learn to collaborate, but they can also learn from their own experience and from observing that of others.   One of the social psychologists I interviewed told me that his style of working with graduate students changed in mid-career because he found that the fairly directive style he initially used was not successful, particularly with the most talented students. He eventually switched to a more indirect style that enabled the students to feel more responsible for their joint projects.   Another of the social psychologists described learning from the collapse of a previously successful collaboration, which died because the partners each thought they were doing too much work for not enough credit.   Thus scientists can improve their ability to collaborate based on mature experience as well as their early role models.

If graduate students work with more than one faculty member, as in the Waterloo social psychology program, they have the opportunity to learn different styles of doing research and interacting with collaborators.   It may be disconcerting for a student to

work with one advisor with a particular style and then have to adjust to another advisor with a different style, for example one that is more or less direct than that of the first advisor. Similarly, faculty members have to learn to adjust to the different personalities of their students and other collaborators, because working with someone who is brash and arrogant may require different interactions than working with someone who is reserved and diffident. Sadler and Woodie (2003) describe how people's interpersonal behaviors can alter the behaviors of their interaction partners. Anyone who collaborates effectively with more than one person must be prepared to learn to adjust to different personalities and research styles.

Procedural and Declarative Knowledge

The previous sections have made the case that collaboration requires procedural knowledge, knowing how as well as knowing that. But is there really a difference? According to Stanley and Williamson (2001), there are no linguistic differences between knowing-how and knowing-that sentences, so that procedural knowledge is a species of propositional knowledge. After briefly showing the flaw in their argument, I will discuss the extent to which procedural knowledge of collaboration is reducible to propositional, declarative knowledge. I will argue that knowledge how to collaborate falls into three categories, being either (1) already explicit and propositional, (2) capable of being made explicit, or (3) inherently implicit and incapable of being made explicit.

Bechtel and Abrahamsen (1991, p. 151) have a long and insightful discussion of how connectionist models clarify the nature of knowledge how as distinct from knowledge that. As a preliminary, they remark on the syntactic difference between knowing that Sacramento is the capital of California and knowing how to ride a bicycle:

"knowing that" requires completion by a proposition whereas "knowing how" is completed by an infinitive specifying an activity. They correctly state that "this linguistic distinction does not, however, settle the matter as to whether there are different psychological representations involved." Stanley and Williamson argue that from the perspective of deeper syntactic and semantic theories there are no significant linguistic differences between knowing-that and knowing-how sentences, and they conclude that knowing-how is not inherently different from knowing-that. But as Bechtel and Abrahamsen recognized, the issue here is psychological, not linguistic, and the structure of English is irrelevant to the issue of whether knowing how employs the same mental structures as knowing that. I have already mentioned the reasons why psychologists think that implicit memory involves different underlying mechanisms from explicit memory, and Reber (1993) provides an extended argument for why implicit learning is different from explicit learning of rules. Other arguments that implicit knowledge does not reduce to explicit propositions can be found in Dreyfus and Dreyfus (1986), Klein (1998), and Sahdra and Thagard (2003).

Sahdra and Thagard (2003) discuss procedural knowledge in molecular biology, focusing on the practices of biologists in using instruments to conduct experiments. Their discussion also applies to social knowledge:

> Decisions based on procedural knowledge are not analyzable into isolatable elements of a situation. Such decisions are intuitive in that they involve quick and effortless recognition of the key patterns of the situation. Procedural knowledge is intuitive recognition of the relevant patterns without recourse to analytical reasoning (p. 488).

Social procedural knowledge can be implicit in just this way, when persons are able to detect patterns in social situation unconsciously and react to them accordingly. For example, the leader of a research laboratory that constitutes a group collaboration may notice an emerging conflict between lab members and quickly move to defuse the situation. Such knowledge could, of course, be explicit, if the leader has formed rules about what to do in such situations, but often the leader simply acts without conscious deliberation thanks to perceptional/emotional recognition of an undesirable social pattern.

Hence it is reasonable to suppose that there is implicit knowledge distinct from explicit, but that does not settle the question of whether procedural knowledge about collaboration is implicit or explicit. Some of it does seem to be explicit and propositional, as summed up in rules such as the following:

Be fair about co-authorship.

Bring student collaborators along slowly.

Find collaborators who complement your own strengths.

A second category of procedural knowledge may not be explicit because it is stored as cases rather than rules, but it is easily made explicit, producing rules such as:

Have regular laboratory group meetings.

Make sure all group members report on what they have been doing.

Thus I am rejecting the claim made by the developmental psychologist Jean Mandler (2004) that procedural knowledge is always implicit, inaccessible to consciousness. In the skill of hitting a baseball, for example, some procedural knowledge can be made explicit (e.g., keep your eye on the ball), whereas other aspects of it remain implicit because they are encoded in perceptual and/or motor neuronal systems.

What are the social skills needed to carry out a successful collaboration? They are probably not much different from those required to be a successful parent or spouse. Here are some likely candidates. You need to be able to recognize and understand the emotional states of your collaborator. In an academic collaboration, there are some emotionally tricky issues, such as allocation of credit (usually reflected in co-authorship) and critical revision of each other's work. It is crucial to be able to recognize if one's collaborator is unhappy about some aspect of the project, in order to rectify whatever intellectual or interpersonal problem is causing this unhappiness. Otherwise, the collaboration may dissolve because of resentment or disinterest. As one of the social psychologists I interviewed said: "Collaboration *is* communication." Understanding and predicting the behaviors of others is not just a matter of being able to theorize about them, but sometimes requires the tacit ability to simulate their thinking (Goldman and Mason, forthcoming).

Another useful emotional skill is to be able to inspire one's collaborator by enthusiasm about particular ideas or methods, so that the collaborators come to share emotional values about them by means of emotional contagion (Hatfield, Cacioppo, and Ratson, 1994). Such contagion occurs when people automatically mimic and synchronize their movements with the facial expressions, voices, postures, movements, and instrumental behavior of others. This mimicry affects people's subjective emotional experiences, leading them to catch the others' emotions. Nothing is more enjoyable about a collaboration than mutual inspiration and contagious excitement. When I asked one of the social psychologists I interviewed how he motivated graduate students to be enthusiastic about their projects, he expounded: "I *am* enthusiastic." Of course,

emotional contagion can also transmit negative emotions, as when collaborators share critical reactions about their competitors.   Collaborators need to be able to monitor the personal relationships that make joint work possible, by paying attention to nonverbal cues, keeping collaborators happy, acting approving, backing off when necessary, and not talking too much. According to Hill (2003, p. 299), effective teams seek to engage in task conflict while avoiding emotional conflict, which requires close but usually tacit attention to the emotional states of team members.

Thus the most plausible candidates for aspects of procedural knowledge that would qualify as irreducibly implicit are those associated with interpersonal interactions that require perception and emotional interpretation.   I provided examples of cases where, whether working with students or peers, it is crucial to foster cooperation by tracking and responding to the emotional states of one's collaborators.   Emotions are an essential part of the scientific thinking of individuals (Thagard, 2002, forthcoming), and are even more important when scientific research takes place in the context of social interaction.   Sharing the emotions that motivate scientific research, such as curiosity, excitement, and the potential joy of discovery, is a crucial part of collaborative work. Equally important is avoiding or regulating the negative emotions that sometimes arise, such as disappointment and even despair.   My argument can be summed up as follows:

(1) Emotions are an essential part of scientific thinking.

(2) Interpersonal communication includes non-verbal interpretation and transmission of emotions.

(3) Such interpretation and transmission are rarely accessible to consciousness.

Therefore:

(4)   Some procedural knowledge essential to successful collaboration is inherently

implicit and not translatable into verbal rules or other propositional information.

Hence not all knowledge how is reducible to knowledge that.

Why Collaborate?

I have describe many aspects of the procedural knowledge about collaboration that scientists acquire, but have not addressed the question of why they bother to acquire this knowledge.   What is the point of collaboration in science, and why do the vast majority of philosophers not collaborate?

I elsewhere provided an epistemic answer to why scientists collaborate by showing how collaboration increases the overall reliability, power, fecundity, speed, and efficiency of research; these criteria were taken from Alvin Goldman's (1992) account of how epistemic practices can be evaluated.   A different answer to the question was suggested by Wray (2002) who provided a more pragmatic explanation:

> Collaborative research is becoming more popular in the natural sciences,
> and to a lesser degree in the social sciences, because contemporary
> research in these fields frequently requires access to abundant resources,
> for which there is great competition. Scientists involved in collaborative
> research have been very successful in accessing these resources, which
> has in turn enabled them to realize the epistemic goals of science more
> effectively than other scientists, thus creating a research environment in
> which collaboration is now the norm. (p. 151).

Although modern science certainly does involve a competition for resources, and scientists may sometimes join together in order to get large grants, Wray's explanation is

not very plausible. For myself and the researchers with whom I have collaborated, there are much more direct motivations for collaboration, especially being able to do more work and to do better work, including research that each of us would not have been able to do individually. There are also pedagogic motivations, in that many of my collaborations have come with students whom I have been happy to help train to do interdisciplinary research. Another personal motivation is simply that it is fun to work closely with clever people with similar goals. Of the 7 social psychologists I interviewed, *none* mentioned access to resources as one of the reasons they collaborated, although several of them have had collaborative projects funded. Rather, they collaborate for the reasons I mentioned: training students, accomplishing more and better research, and fun. As Epstein (2005) reports, interdisciplinary collaboration tends to be expensive, and, funding for interdisciplinary research is often harder to secure than more traditional disciplinary work. Hence Wray's broad functional explanation is not plausible as a psychological explanation of why scientists collaborate.

So, if collaboration is productive, creative, and fun, why do philosophers do it so rarely? The answer to this question is partly social and partly intellectual. The social explanation is that philosophical training rarely involves collaboration, so that philosophers beginning their careers have not learned about how to collaborate. But that raises the question of why collaborative training has not become part of the philosophical tradition, which is largely due to the nature of dominant forms of philosophical method.

Looking at philosophy from a very broad perspective, there appear to be four major approaches on the current scene: a priori reflection, logico-linguistic analysis, historical scholarship, and naturalistic theorizing. On the first of these, which assumes

that truths in epistemology, metaphysics, and ethics can be gained a priori, there is little reason to collaborate, since a priori knowledge (if there were any) is an entirely individual matter, stemming from the internal reflections of a single mind.   Hence philosophers who think that their intellectual enterprise is a priori have no reason to collaborate at all.

In contrast, advocates of logical or linguistic analysis and historical rumination might have some reason to collaborate, if they can get help with their projects in philosophical logic, philosophy of language, or the history of philosophy.   But the gains might be small, because of the inherent costs of collaboration, namely the extra time it takes to communicate with other people and develop common views.   In the terminology of the theory of parallel computation, the gains of collaboration within these traditions are likely to be sublinear: having n collaborators is likely to lead to less than n times extra work being accomplished.   In the sciences, collaborations of the employer/employee and teacher/apprentice kinds normally have sublinear benefits.

In contrast, the naturalistic approach to philosophy can potentially achieve a superlinear gain, because collaboration can produce work that would never have been done otherwise. Philosophical naturalism sees philosophy as intertwined with science, so that deliberations on the nature of reality, knowledge, and morals are intimately connected with developments in physics, psychology, neuroscience, and other branches of science.   Because of the difficulty of acquiring expertise in more than one field, it is invaluable to be able to develop peer-different collaborations with scientists who have expertise in fields relevant to some area of philosophical concern, for example with physicists for metaphysics, and with psychologists for philosophy of mind.   The ideal

result, as in peer-different collaborations in the sciences, is that research gets accomplished by a team of individuals that would never have come about otherwise. Under such circumstances, two heads really are more than twice better than one.

There is both anecdotal and statistical evidence that collaboration is more common among naturalistic philosophers than among adherents to a priori, logico-linguistic, or historical approaches. Philosophers I know who are atypically prone to collaboration include Alvin Goldman, Stephen Stich, Rob Cummins, and William Bechtel, all of whom take a collaborative approach. But perhaps there are collaborations in other approaches to philosophy that I simply am not aware of. It is therefore statistically useful to compare the 2004 volumes of other philosophy journal besides the *Journal of Philosophy* mentioned earlier. I picked the 2004 volume of the *Review of Metaphysics,* and contrasted it with the 2004 volume of *Philosophy of Science*. The latter has ADD% collaborative articles, whereas the former has only ADD%, even less than the *Journal of Philosophy*. Philosophers of science today tend to be much more naturalistic than other philosophers, so this finding is consistent with my hypothesis that philosophical naturalism is more compatible with collaboration than other approaches. An additional factor here is that, by virtue of working with scientists, philosophers can observe how scientists collaborate and learn from them how to do so. It is no accident that Goldman, Stich, and I all had an early association with the psychologist Richard Nisbett, an inveterate collaborator.

Conclusion

I therefore predict that if philosophy continues to become more naturalistic, then it will also tend to become more collaborative. Whether or not this happens, I am sure

that science will continue to be highly collaborative, and scientists will continue to learn how to collaborate from their supervisors and their peers. Transmission of procedural knowledge about how to collaborate is not always easy, since it requires time and mutual concern. But given the qualitative and quantitative gains in research productivity that collaboration can produce, the topic of collaborative knowledge will continue to be important for social epistemology.

References

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

Barr, A., & E., F. (1981). *Handbook of artificial intelligence*. Los Altos: Kaufmann.

Bechtel, W., & Abrahamsen, A. (1991). *Connectionism and the mind*. Oxford: Blackwell.

Best, D. L. (2004). Varieties of college and university experiences. In J. M. Darley, M. P. Zanna & H. L. Roediger (Eds.), *The compleat academic: A career guide* (pp. 345-359). Washington: American Psychological Association.

Darley, J. M., Zanna, M. P., & Roediger, H. L. (2004). *The compleat academic: A career guide*. Washington: American Psychological Association.

Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind Over Machine*. New York: Free Press.

Epstein, S. (2005). Making interdisciplinary collaboration work. In S. J. Derry, C. D. Schunn & M. A. Gernsbacher (Eds.), *Interdisciplinary collaboration: An emerging cognitive science* (pp. 245-263). Mahwah, NJ.

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

Goldman, A. I., & Mason, K. (forthcoming). Simulation. In P. Thagard (Ed.), *Handbook of philosophy of psychology and cognitive science*. Amsterdam: Elsevier.

Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (1994). *Emotional contagion*. Cambridge: Cambridge University Press.

Hill, L. A. (2003). *Becoming a manager: How new managers master the challenges of leadership* (2nd ed.). Boston: Harvard Business School Press.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: MIT Press/Bradford Books.

Kahneman, D. (2003). Experiences of collaborative research. *American Psychologist*(723-730).

Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, Massachusetts: MIT.

Levine, J. M., & Moreland, R. L. (2004). Collaboration: The social context of theory development. *Personality and Social Psychology Review, 8*, 164-172.

Litt, A., Eliasmith, C., Kroon, F. W., Weinstein, S., & Thagard, P. (forthcoming). Is the brain a quantum computer? *under review*.

Mandler, J. M. (2004). *The foundations of mind: Origins of conceptual thought*. Oxford: Oxford University Press.

Polanyi, M. (1958). *Personal knowledge*. Chicago: University of Chicago Press.

Polanyi, M. (1967). *The tacit dimension*. London: Routledge & Kegan Paul.

Reber, A. S. (1993). *Implicit learning and tacit knowledge*. New York: Oxford University Press.

Roediger, H. L., & Balota, D. A. (2004). Managing your career: The long view. In J. M. Darley, M. P. Zanna & H. L. Roediger (Eds.), *The compleat academic: A career guide* (pp. 393-407). Washington: American Psychological Association.

Ryle, G. (1949). *The concept of mind*. London: Hutchinson.

Sadler, P., & Woody, E. (2003). Is who you are who you're talking to? Interpersonal style and complementarity in mixed-sex interactions. *Journal of Personality and Social Psychology, 84*, 80-96.

Sahdra, B., & Thagard, P. (2003). Procedural knowledge in molecular biology. *Philosophical Psychology, 16*, 477-498.

Schacter, D. L. (1996). *Searching for memory: The brain, the mind, and the past*. New York: Basic Books.

Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy, 98*(8), 411-444.

Taylor, S. E., & Martin, J. (2004). The academic marathon: Controlling one's career. In J. M. Darley, M. P. Zanna & H. L. Roediger (Eds.), *The compleat academic: A career guide* (pp. 363-392). Washington: American Psychological Association.

Thagard, P. (1997). Collaborative knowledge. *Noûs, 31*, 242-261.

Thagard, P. (2002). The passionate scientist: Emotion in scientific cognition. In P. Carruthers, S. Stich & M. Siegal (Eds.), *The cognitive basis of science* (pp. 235-250). Cambridge: Cambridge University Press.

Thagard, P. (2005). Being interdisciplinary: Trading zones in cognitive science. In S. J. Derry, C. D. Schunn & M. A. Gernsbacher (Eds.), *Interdisciplinary collaboration: An emerging cognitive science* (pp. 317-339). Mahwah, NJ: Erlbaum.

Thagard, P. (forthcoming). *Hot thought: Mechanisms and applications of emotional cognition*. Cambridge, MA: MIT Press.

Wray, K. B. (2002). The epistemic significance of collaborative research. *Philosophy of Science, 69*, 150-168.

Zacks, J. M., & Roediger, H. L. (2004). Setting up your lab and beginning a program of research. In J. M. Darley, M. P. Zanna & H. L. Roediger (Eds.), *The compleat academic: A career guide* (pp. 135-152). Washington: American Psychological Association.

Zanna, M. P., & Darley, J. M. (2004). Mentoring: Managing the faculty-graduate student relationship. In J. M. Darley, M. P. Zanna & H. L. Roediger (Eds.), *The compleat academic: A career guide* (pp. 117-131). Washington: American Psychological Association.