

fallisthagardcomments.pdf
10/31/05

The Epistemic Costs and Benefits of Collaboration

Don Fallis
School of Information Resources
University of Arizona
fallis@email.arizona.edu

Southern Journal of Philosophy, **44**, S, (2006): 197-208

The Epistemic Costs and Benefits of Collaboration

Abstract

In “How to Collaborate,” Paul Thagard tries to explain why there is so much collaboration in science, and so little collaboration in philosophy, by giving an epistemic cost-benefit analysis. In this paper, I argue that an adequate explanation requires a more fully developed *epistemic value theory* than Thagard utilizes. In addition, I offer an alternative to Thagard’s explanation of the lack of collaboration in philosophy. He appeals to its lack of a tradition of collaboration and to the *a priori* nature of much philosophical research. I claim that philosophers rarely collaborate simply because they can usually get the benefits without paying the costs of actually collaborating.

1. Introduction

There is a large and increasing amount of collaborative research in science (cf. Thagard 1997, 243-245, Wray 2002, 151-152). In “How to Collaborate: Procedural Knowledge in the Cooperative Development of Science,” Paul Thagard categorizes the various types of collaboration in science (viz., employer/employee, teacher/apprentice, peer-similar, peer-different). He then looks at how such collaboration can be carried out effectively. This is a critical task for the project in *social epistemology* that Alvin Goldman (1999) has emphasized.¹ According to this project, social epistemologists should try to identify practices that have good epistemic consequences. Thagard seems to have done just this with respect to collaboration.²

In addition to addressing the question of how to collaborate, Thagard addresses three additional questions of significant philosophical interest. Since I agree with the vast majority of what Thagard says about how to collaborate effectively, I will focus here on some important disagreements about these further questions:

- First, is knowledge of how to collaborate propositional knowledge? Thagard claims that the answer is *no*. However, I will argue that Thagard dismisses too quickly Jason Stanley and Timothy Williamson’s (2001) argument that procedural knowledge is a type of propositional knowledge.
- Second, why is there so much collaboration in science? Thagard tries to explain this phenomenon in terms of the epistemic benefits of collaboration. However, I will argue that a more fully developed *epistemic value theory* is needed in order to successfully give such a cost-benefit analysis. In addition, I will argue that Thagard dismisses too quickly Brad Wray’s (2002) explanation of why scientists collaborate.
- Third, why is there so little collaboration in philosophy? Thagard tries to explain this phenomenon in terms of the *a priori* nature of much philosophical research and the lack of a tradition of collaboration. However, I will argue that Thagard’s explanation is not successful. I will offer my own explanation for why philosophers rarely collaborate.

2. Is knowledge of how to collaborate propositional knowledge?

Thagard wants to argue that knowledge of how to collaborate is an example of something that is not propositional knowledge. However, Stanley and Williamson have argued to the contrary that procedural knowledge is a type of propositional knowledge. On their analysis, I know *how* to **R** (e.g., how to collaborate) if and only if, for some **X**, I know *that X* is a way to **R** (cf. Stanley and Williamson 2001, 430). Thagard criticizes Stanley and Williamson's claim that knowing-how is a type of knowing-that on the following grounds. He claims that there are important psychological differences between knowing-how and knowing-that. In particular, he claims that, in many instances, procedural knowledge is "irreducibly implicit."

These claims seem to me to be correct. However, they do not seem to me to be in conflict with what Stanley and Williamson say. On their analysis, knowing how to ride a bicycle is knowing that **X** is a way to ride a bicycle. But Stanley and Williamson (2001, 433) admit that, even if you have this propositional knowledge, you may not be able to explicitly state the relevant proposition. At least, you may not be able to do so in "nonindexical involving terms." In other words, you may only be able to point to your performance and say "*That* is a way to ride a bicycle."³ This suggests that, for Stanley and Williamson, even propositional knowledge can be "irreducibly implicit."

Of course, even if knowing-how is a type of knowing-that, the psychological differences that Thagard points to certainly indicate that knowing-how is a very special type of knowing-that. And these differences may be particularly important to the social epistemologist who wants to identify practices that have good epistemic consequences. However, the primary issue of interest for such a social epistemologist is not whether there is procedural knowledge that is "irreducibly implicit." The more critical issue is what the most effective means of transmitting various pieces of procedural knowledge are. Now, if a piece of procedural knowledge is "irreducibly implicit," then we are not going to be able to transmit it by making it explicit. We are going to have to use other techniques (e.g., allowing people to learn by example or from experience). But as Thagard suggests, making it explicit may not be the most effective means even for procedural knowledge that *could* be made explicit.

However, in many cases, it will be best to make procedural knowledge explicit. For example, if you can make a piece of knowledge explicit, you can write it down. And it is easier to disseminate recorded knowledge amongst people and across time (cf. Diamond 1997, 215-238, Goldman 1999, 162). Thus, an important virtue of Thagard's paper is that it provides many *explicit* recommendations for how to collaborate successfully. In addition, there is a lot of research on how to support collaboration. For instance, Xiaolong Zhang and George Furnas (2005) are trying to design virtual environments that will help people to collaborate effectively. These designers have to have fairly explicit knowledge of how collaboration works. This is so even if the goal is to design a tool that does not require its users to have such explicit knowledge (cf. Winograd and Flores 1987).

3. Why is there so much collaboration in science?

The most obvious way to try to explain why there is so much collaboration in science is with a cost-benefit analysis. In other words, we can try to show that the benefits of collaboration outweigh the costs. In particular, we can try to show that the *epistemic* benefits outweigh the costs.⁴ A number of philosophers (e.g., Maher 1993, Maddy 1997, Goldman 1999, Fallis 2002) have used this sort of cost-benefit analysis to explain methodological choices of scientists and mathematicians. Thagard tries to explain why so many scientists have chosen to collaborate in this same manner.⁵

Now, it is possible that scientists collaborate even though the costs outweigh the benefits. For example, it is possible that scientists are irrational or that they are simply ignorant of the costs of collaboration. But I take these to be explanations of last resort (cf. Fallis 2002, 376). I think that the prospects are good for showing, as Thagard would like to do, that the benefits of collaboration often outweigh the costs.

There are many costs associated with collaboration that the researcher working alone does not have to pay. One of the costs that is emphasized in Thagard's article is that researchers have to learn how to collaborate, but there are a number of others.⁶ For example, as Thagard points out, researchers who collaborate have to spend a lot of time simply communicating with each other. In particular, such researchers often have to spend time resolving disagreements about the intellectual content of their research.⁷ In addition, if such disagreements cannot be resolved, fewer and more "conservative" results may end up being published (McDonald 1995, 36). Also, researchers who collaborate may not get the credit that they deserve for their contribution to a research project (cf. Thagard 1997, 250). In physics especially, individual researchers may easily get lost in huge author lists that often "include technicians as well as scientists" (McDonald 1995, 36).⁸ But there may be epistemic benefits to collaboration that outweigh these additional costs.

As Thagard correctly points out in his article, "collaboration can produce work that would never have been done otherwise." In other words, some research projects can only be carried out as collaborations. In particular, there are research projects that require expertise from several different disciplines that no single individual is likely to have (cf. Wray 2002, 162). For example, a few years ago, I wrote a philosophy article that relied on the truth of a mathematical claim that I did not have the expertise to prove on my own. As a result, I had to collaborate with a mathematician, Gerrard Liddell, who did have the requisite expertise (see Fallis and Liddell 2002). However, the mere fact that a research project could only have been carried out as a collaboration does not show that the project was worth carrying out in the first place.

But there is empirical evidence that collaboration does tend to produce very valuable research. First, collaborative research tends to be more highly cited than other research (cf. Wray 2002, 156). And the citation counts of scientific and academic articles have consistently been found to be correlated with the quality of the articles (cf. Garfield 1972,

Lee et al. 2002). Second, in an empirical study on scientific research, Robert Merton and Harriet Zuckerman (1973, 547-548) found that Nobel laureates tend to collaborate more than other scientists. Presumably, this also indicates that there is some connection between collaborative research and quality.

Furthermore, there is a plausible explanation of how collaboration leads to valuable research. The basic idea is simply that two heads are typically better than one. First, in collaborative research, more smart people are contributing ideas to the project (cf. Thagard 1997, 252). Second, with more people taking a critical look at the research, it is more likely that any mistakes will be uncovered and corrected (cf. McDonald 1995, 36, Thagard 1997, 251). Third, when the labor is divided up in a collaborative research project, people can work on what they are best at.

It is also possible to identify some of the specific epistemic values that collaboration is likely to promote. For example, several philosophers (e.g., Kuhn 1977, 322, Kitcher 1981, 508) have discussed the epistemic value of having theories that unify a wide range of phenomena. In his earlier article, Thagard (1997, 254-255) plausibly suggests that interdisciplinary collaboration is especially likely to lead to such knowledge. Thagard (1997, 255-258) gives two actual case studies of scientific research that support this conclusion.

In addition to leading to more valuable research, there are at least two other important epistemic benefits of collaboration. First, collaboration can lead to the discovery of a *greater amount* of knowledge.⁹ Researchers who collaborate tend to be more productive (cf. Wray 2002, 157-158). Also, when teachers collaborate with students, they are creating more and better researchers for the future (cf. Wray 2002, 158). Second, collaboration can lead to *greater dissemination* of knowledge.¹⁰ With more people working on the project, there are more people to disseminate the research results. Also, the results are much less likely to be forgotten by the scientific community (cf. Wray 2002, 157).

It is important to note, however, that it is not enough that the epistemic value of a collaborative research project tends to be greater than the epistemic value of individual research projects. In order to account for collaboration with a cost-benefit analysis, we have to show that the epistemic value of a collaborative research project tends to be greater than the *combined* value of the individual research projects that the collaborators could have completed working separately. In other words, we need to show that collaboration is a more effective use of a fixed amount of resources.

Now, Thagard is probably correct that there are many circumstances where “two heads really are more than twice better than one.” In other words, it seems likely that collaboration is often a more effective use of a fixed amount of resources. But in order to show that this is the case, it is not enough to know what the various epistemic benefits of collaboration are. Since different courses of action can yield different benefits, we also need to be able to compare (i.e., determine tradeoffs between) the various epistemic benefits. (For example, how does the value of a higher quality collaborative article

compare with the value of several slightly lower quality articles written by individuals?) In other words, in order to definitively establish that the epistemic benefits of collaboration outweigh the costs, we need a full-fledged *epistemic value theory* (cf. Paterson 1979, Maher 1993, 208-244, Goldman 1999, 87-100, Fallis 2004). However, Thagard has not made use of such a theory.¹¹ Further work needs to be done to develop an epistemic value theory and to apply it to the issue of collaboration.

In addition to Thagard, Brad Wray (2002) has also offered an explanation for why there is so much collaboration in science. Wray points to competition over funding as an explanation of this phenomenon. Thagard criticizes Wray by pointing out that there are also many *epistemic* benefits of collaboration that motivate researchers to collaborate. But I do not think that Wray is denying this. In fact, it is an important part of Wray's argument that collaboration has numerous epistemic benefits.

Wray is essentially telling a natural selection story: Researchers who collaborate have greater epistemic success. Greater epistemic success leads to these researchers getting more funding. This funding allows them to do more research. (Also, other researchers are motivated to start collaborating in order to have greater epistemic success and in order to get more funding.) By contrast, researchers who do not collaborate do not have as much epistemic success. As a result, they are less able to get funding and less able to carry out research. In other words, competition over funding provides selection pressure in favor of researchers who collaborate. This sort of story is consistent with arguing that (and, in fact, requires that) the epistemic benefits of collaboration outweigh the costs.

4. Why is there so little collaboration in philosophy?

While collaboration is common in the sciences, it is rare in many other disciplines such as philosophy. As Thagard points out, there is some collaboration in philosophy. In particular, research projects in philosophy that require expertise from other disciplines are often collaborative. My aforementioned collaboration with a mathematician is an example of this. But there is not a lot of collaboration in philosophy. And, in particular, there is not a lot of collaboration just between philosophers whereas there is, as Thagard points out, a significant amount of peer-similar collaboration in the sciences. In addition to an explanation for the prevalence of scientific collaboration, we would like an explanation for the lack of philosophical collaboration.

Wray (2002, 159) thinks that his natural selection story easily provides an explanation of this phenomenon as well. In order to explain why philosophers rarely collaborate, Wray simply points to the fact that there is less competition over funding in philosophy. However, even if he has provided a successful explanation of why there is so much collaboration in science, I do not think that Wray provides a successful explanation of why there is so little collaboration in philosophy.

Admittedly, there is not as much competition over funding since less funding is needed to carry out philosophical research. However, there are other selection pressures that reward epistemic success in philosophy. For example, successful philosophers typically

get more time to do research (e.g., they get grants to buy out their teaching or they get positions with lighter teaching loads). Thus, if philosophers who collaborate had significantly greater epistemic success, we would expect philosophers who collaborate (just like physicists who collaborate) to be selected for. But we do not see this sort of selection, which suggests that philosophers who collaborate do not generally have significantly greater epistemic success (or that they pay much greater costs for collaborating). Wray provides no explanation for why this might be the case.

In order to explain why philosophers rarely collaborate, Thagard essentially claims that they pay greater costs for collaborating. According to Thagard, “the main reason why philosophers do not collaborate more is that they do not know how” (cf. Thagard 1997, 249). Since there is not a tradition of collaboration in philosophy, philosophers would have to pay the cost of learning how to collaborate from scratch. In other words, there are significant *switching costs* associated with moving to more collaborative research (cf. Shapiro and Varian 1999, 11).

However, while such switching costs are undoubtedly part of the explanation, they cannot be the *main* reason that philosophers do not collaborate. First, if collaboration were especially beneficial in philosophy, we would expect that the epistemic benefits would often outweigh these switching costs. After all, other sciences have moved to more collaborative research despite the switching costs. Second, Merton and Zuckerman (1973, 549) plausibly suggest that scientists do not collaborate as much in the later stages of their careers because they “often turn to broader “philosophical” or “sociological” subjects of a kind that have little place for collaboration.” But, unlike philosophers, these scientists do know how to collaborate. This suggests that the lack of collaboration in philosophy has more to do with what philosophical research is like than with the fact that philosophers do not know how to collaborate.

In fact, Thagard thinks that at least part of the explanation for why philosophers rarely collaborate has to do with the nature of philosophical research. In particular, he claims that many philosophers do not collaborate because they would not reap any epistemic benefits from collaborating. As Thagard points out, the philosophers involved in the aforementioned interdisciplinary collaborations tend to be those who take a naturalistic approach to philosophy. However, many philosophers take philosophy to be an *a priori* discipline. In other words, they think that philosophy is something that can be done sitting in an armchair working alone (cf. Descartes 1996). Thagard claims that, for these philosophers, “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.”

However, a prioricity cannot explain the lack of collaboration in philosophy. For one thing, even if a philosophical claim can be known a priori, it may be that no single individual is able to construct the requisite arguments to establish the claim. In other words, it may require collaboration. For example, I collaborated on an article on ethics with two other philosophers and that is exactly how it played out (see Frické et al. 2000). Each of us contributed a part of the argument that had not occurred to the rest of us.

But there is an even better reason to think that a prioricity cannot be the explanation for the lack of collaboration in philosophy. Mathematics is at least as a priori a discipline as is philosophy.¹² But, despite this fact, there is almost as much collaboration in mathematics as there is in the physical sciences (cf. Grossman and Ion 1995). For example, 63% of the articles published in the *Annals of Mathematics* in 2004 had more than one author (as compared with only 11% of the articles published in the *Journal of Philosophy*). In fact, collaboration is sufficiently prevalent in mathematics that it is common for mathematicians to calculate their *Erdős number*. You have Erdős number 1 if you collaborated with the prolific Hungarian mathematician, Paul Erdős, and you have Erdős number $n+1$ if you collaborated with someone who has Erdős number n . For example, as a result of my collaboration with a mathematician, my Erdős number is 6.¹³

Since switching costs and a prioricity cannot explain the lack of collaboration in philosophy, I would like to float my own explanation. This explanation is still in line with Thagard's desire for a cost-benefit analysis.¹⁴ The basic idea is that there are mechanisms that give philosophers the benefits of collaboration without their having to pay the costs of actually collaborating. From the time of Socrates, philosophers have typically done philosophy by spending a lot of time talking with (or corresponding with) other philosophers.¹⁵ For example, many philosophers have shared ideas in informal discussion groups, such as the Wise Club (founded by Thomas Reid) and the Vienna Circle.¹⁶ In addition, what I will refer to as the *philosopher's strategy* extends to formal as well as informal interaction. When philosophers give formal talks, engagement and criticism from the audience is encouraged. In fact, in a move that is fairly unique to philosophy, someone (viz., a commentator such as myself) is often officially assigned the job of coming up with objections and suggesting revisions. As a result of these mechanisms, a philosopher can easily benefit from the ideas and criticisms of other philosophers without a formal collaboration.

While this strategy is traditional in philosophy, this strategy is certainly available to other disciplines as well. After all, most disciplines have conferences and colloquia that provide opportunities for people to give talks about their research. However, I think that there are reasons why the philosopher's strategy will not obviate the need for collaboration in these other disciplines.¹⁷ First, in many scientific disciplines, the research described in such talks could not have been done in the first place without collaboration (cf. Wray 2002, 156). Second, in most scientific disciplines, such talks serve primarily to disseminate the research results. By the time that the talk is given, the data has already been collected and feedback from the audience can typically only influence the design of future experiments. By contrast, in philosophy, such feedback can be put to immediate use in shaping the research just like input from a collaborator can.

Mathematics is the other academic discipline that would probably reap the most benefits from the philosopher's strategy. But mathematicians can also reap more benefits from actual collaboration than philosophers can. For example, it is much easier to divide up the work among collaborators in mathematical research (cf. Thagard 1997, 249). The

“10,000 Page Proof” of the classification theorem for finite simple groups is probably the most famous example of such a division of labor (cf. Fallis 1997, 171-172). By contrast, since it is more difficult to divide up the work in philosophical research, philosophers are not forgoing much benefit in terms of division of labor by not collaborating. Also, mathematicians pay fewer costs from actually collaborating than philosophers do. For example, it is much easier to resolve intellectual disputes among collaborators because the correctness of mathematical arguments tends to be more cut-and-dried than the correctness of philosophical arguments.¹⁸

Much more certainly needs to be said about the properties of philosophical research that allows this strategy to significantly reduce the need for philosophical collaboration. But, hopefully, I have said enough here to indicate why philosophers so rarely collaborate. However, it should be noted that I do not mean to criticize Thagard’s call for more collaboration in philosophy (cf. Thagard 1997, 258-259). Philosophers may very well collaborate less than it would be rational for them to do. In fact, if they were to follow Thagard’s excellent recommendations for how to collaborate effectively, the epistemic benefits of philosophical collaboration would undoubtedly often outweigh the costs.

Acknowledgements

I would like to thank Kay Mathiesen, George Pappas, Paul Thagard, and the other participants in the *Spindel Conference on Social Epistemology* for helpful suggestions.

References

- Bhopal, Raj, Judith Rankin, Elaine McColl, Lois Thomas, Eileen Kaner, Rosie Stacy, Pauline Pearson, Brian Vernon, and Helen Rodgers. 1997. "The Vexed Question of Authorship: Views of Researchers in a British Medical Faculty." *British Medical Journal* 314:1009-12.
- Craig, Edward, ed. 1998. *Routledge Encyclopedia of Philosophy*. London: Routledge.
- De Millo, Richard, Richard Lipton, and Alan Perlis. 1979. "Social Processes and Proofs of Theorems and Programs." *Communications of the Association for Computing Machinery* 22:271-80.
- Descartes, René. 1996. *Meditations on First Philosophy*. ed. John Cottingham. Cambridge: Cambridge.
- Diamond, Jared. 1997. *Guns, Germs, and Steel*. New York: W. W. Norton.
- Fallis, Don. 1997. "The Epistemic Status of Probabilistic Proof." *Journal of Philosophy* 94:165-86.
- Fallis, Don. 2002. "What Do Mathematicians Want?: Probabilistic Proofs and the Epistemic Goals of Mathematicians." *Logique et Analyse* 45:373-88.

- Fallis, Don. 2004. "Epistemic Value Theory and Information Ethics." *Minds and Machines* 14:101-17.
- Fallis, Don and Martin Frické. 2002. "Indicators of Accuracy of Consumer Health Information on the Internet." *Journal of the American Medical Informatics Association* 9:73-79.
- Fallis, Don and Gerrard Liddell. 2002. "Further Results on Inquiry and Truth Possession." *Statistics and Probability Letters* 60:169-82.
- Frické, Martin, Kay Mathiesen, and Don Fallis. 2000. "The Ethical Presuppositions behind the Library Bill of Rights." *Library Quarterly* 70:468-91.
- Garfield, E. 1972. "Citation Analysis as a Tool in Journal Evaluation." *Science* 178:471-79.
- Goldman, Alvin I. 1987. "Foundations of Social Epistemics." *Synthese* 73:109-44.
- Goldman, Alvin I. 1999. *Knowledge in a Social World*. New York: Oxford.
- Grossman, Jerrold W. and Patrick D. F. Ion. 1995. "On a Portion of the Well-Known Collaboration Graph." *Congressus Numerantium* 108:129-31.
- Kitcher, Philip. 1981. "Explanatory Unification." *Philosophy of Science* 48:507-31.
- Kuhn, Thomas S. 1977. "Objectivity, Value Judgment, and Theory Choice." Pp. 320-339 in *The Essential Tension*, Chicago: University of Chicago.
- Lee, Kirby P., Marieka Schotland, Peter Bacchetti, and Lisa A. Bero. 2002. "Association of Journal Quality Indicators With Methodological Quality of Clinical Research Articles." *JAMA* 287:2805-8.
- Maddy, Penelope. 1997. *Naturalism in Mathematics*. New York: Oxford.
- Maher, Patrick. 1993. *Betting on Theories*. New York: Cambridge.
- McDonald, Kim A. 1995. "Too Many Co-Authors?" *Chronicle of Higher Education* 41(33):35-36.
- Merton, Robert K. and Harriet Zuckerman. 1973. "Age, Aging, and Age Structure in Science." Pp. 497-559 in *The Sociology of Science*, Chicago: University of Chicago.
- Paterson, R. W. K. 1979. "Towards an Axiology of Knowledge." *Journal of Philosophy of Education* 13:91-100.
- Shapiro, Carl and Hal R. Varian. 1999. *Information Rules*. Boston: Harvard Business School.

- Stanley, Jason and Timothy Williamson. 2001. "Knowing How." *Journal of Philosophy* 98:411-44.
- Thagard, Paul. 1997. "Collaborative Knowledge." *Nous* 31:242-61.
- Winograd, Terry and Fernando Flores. 1987. *Understanding Computers and Cognition: A New Foundation for Design*, Reading, Massachusetts: Addison-Wesley.
- Wray, K. B. 2002. "The Epistemic Significance of Collaborative Research." *Philosophy of Science* 69:150-168.
- Zhang, Xiaolong and George W. Furnas. 2005. "mCVEs: Using Cross-Scale Collaboration to Support User Interaction with Multiscale Structures." *Presence: Teleoperators & Virtual Environments* 14:31-46.

¹ See also Goldman's article in this issue.

² His recommendations fit well with my own personal experience with collaboration (on six research articles as of 2005).

³ Or, to use Thagard's own example, "THIS is what we did to design that experiment, or THAT is what we did to write that paper."

⁴ Scientists are often motivated by non-epistemic goals (such as the drive for scientific credit) as well as epistemic goals. However, as Goldman (1999, 260-263) points out, achieving epistemic goals is typically the most effective means for scientists to achieve such non-epistemic goals. Thus, it seems fairly safe to simply focus here on the epistemic benefits of collaboration.

⁵ Thagard (1997) gives an even more detailed cost-benefit analysis in an earlier article.

⁶ Thagard makes a number of recommendations for how to reduce many of these costs.

⁷ I have certainly experienced this cost in my own collaborations.

⁸ In order to deal with this situation, the editors of several medical journals have tried to impose a fairly restrictive notion of authorship (cf. Bhopal et al. 1997). As a result, many articles include a section that describes the contributions that "need acknowledging, but do not justify authorship" (see, e.g., Fallis and Frické 2002, 79).

⁹ Goldman (1987, 128) refers to this as the *power* of an epistemic practice.

¹⁰ Goldman (1987, 129) refers to this as the *fecundity* of an epistemic practice.

¹¹ In his earlier article, Thagard (1997) makes use of the five epistemic standards described by Goldman (1987, 128-129). But neither Thagard nor Goldman say how to resolve conflicts between these standards.

¹² As a matter of fact, I have essentially argued in previous work that mathematics is not all that a priori (cf. Fallis 1997, Fallis 2002). But here I am appealing to what is clearly the standard line in the philosophy of mathematics (viz., that mathematics is the epitome of an a priori discipline).

¹³ As a result of his many collaborations, Paul Thagard's Erdős number is 7.

¹⁴ Just as with scientists, I do not think that we need to appeal to the irrationality or ignorance of philosophers to explain their methodological choices.

¹⁵ There is a sense in which philosophers are collaborating when they do this. But Thagard clearly has in mind collaborating in the sense of co-authoring research articles.

¹⁶ See the entries for "Reid, Thomas" and the "Vienna Circle" in Craig (1998). Philosophical *blogs* are the latest method for such informal sharing of ideas.

¹⁷ The reasons may be different for different disciplines.

¹⁸ Admittedly, since proofs are getting longer and more complicated, it is not always a trivial matter to resolve such disputes even in mathematics (cf. De Millo et al. 1979, 272).