Conference "Epistemic groups and collaborative research in science"
Nancy, December 17-19, 2012

INVITED SPEAKERS

Denis Bonnay (University Paris Ouest, IRePh & IHPST)
"What people think: judgement aggregation and opinion research"

On one hand, it is common doxastic practice to attribute beliefs to groups of agents, on the other hand, we have learnt from judgment aggregation theory that aggregating opinions is anything but simple. Why is that group beliefs so easily spring into existence, when they should be so hard to achieve? In the recent philosophical literature on collective beliefs, most, if not all, answers to this problem implicitly or explicitly assume that group beliefs do not supervene on individual beliefs simpliciter. Group beliefs require something more, such as deliberate consistency maintenance or joint intentions. This something more is meant to explain why non-individual beliefs are possible and real in spite of impossibility results. In this talk, I want to argue in favor of a different account of collective beliefs, according to which a) collective beliefs supervene on individual beliefs simpliciter, b) coherence, stricto and lato sensu, is a guide rather than an enemy to collective belief attribution, c) doxastic groups are characterized as doxastic units displaying a high-level of doxastic coherence. This account will be based on a critical comparison between the take on collective belief which is congenial to judgment aggregation theory and standard methodology used in sociology and opinion research.

Bryce Huebner (Georgetown University, USA)
"Accountability and values in radically collaborative research"

In this paper I discuss a crisis of accountability that arises when scientific collaborations are massively epistemically distributed. I argue that social models of epistemic collaboration, which are social analogs to what Patrick Suppes called a "model of the experiment," must play a role in creating accountability in these contexts. I also argue that these social models must accommodate the fact that the various agents in a collaborative project often have ineliminable, messy, and conflicting interests and values; any story about accountability in a massively distributed collaboration must therefore involve models of such interests and values and their methodological and epistemic effects.

Christian List (London School of Economics, UK)
"Three kinds of collective attitude"

Erik Olsson (Lund University, Sweden)
"Probabilistic Updating in Epistemic Groups: The Laputa Model"

The talk describes a Bayesian model for updating degrees of belief and trust in epistemic groups. Various features of the model are presented, including qualitative updating rules that can be derived from the underlying probabilistic framework. The underlying model is complex and analytical results are correspondingly difficult to prove. For that reason, a simulation program, Laputa, has been developed which allows for effortless exploration of the model. In the talk it is demonstrated how various more complex holistic features of the model can be uncovered with the help of Laputa. For instance, it turns out inquirers in the Laputa model, just as real inquirers, are vulnerable to belief polarization. Finally, two applications of the model are discussed that involve assessing the epistemic goodness of social practices in the sense of Goldman (1999): determining the optimal threshold of assertion for group members and the optimal communication structure for the group.

Jan Sprenger (Tilburg University, The Netherlands)
“A socio-epistemic variant of the No Alternatives Argument”

In most scientific disciplines, several lines of research compete with each other. Only one of them (call it H) may achieve an important intermediate result, or comply with a set of constraints that are essential to the solution of a scientific research puzzle. Does this lack of alternatives provide a valid argument in favor of H? Inparticular, should we allocate our resources to this line of research, rather han to less explored competitors?
We analyze this question by means of the No Alternatives Argument (NAA) recently put forward by Dawid, Hartmann and Sprenger (2012). While our previous analysis focuses on the empirical adequacy of a scientific theory, I now transfer the structure of the argument to the social epistemology of science. It turns out that while NAA remains logicallyvalid in the socio-epistemic case, it is questionable whether the argument is strong enough to justify clear-cut preferences among different lines of research.

K. Brad Wray (State University of New York, USA)
"The Impact of Collaboration on the Epistemic Culture of Science"

I examine how collaborative research affects the epistemic culture of science. First, I argue that some groups of scientists hold views that are irreducibly the views of the group. I also address the following two questions: (i) what do appeals to collective knowledge explain?; and (ii) what is the epistemic significance of the phenomena that such appeals explain? Finally, I examine challenges that collaborative research raises for refereeing in science. I argue that journal editors and editorial boards are out of step with the changes occurring in science as a consequence of collaborative research.

Anouk Barberousse, Henri Galinon and Marion Vorms (University Lille 1, University of Clermont, University of Paris 1, France)
"Community Modeling Systems: New Wave Scientific Collaboration"

In an attempt at analyzing present-day collaboration practices within the climate-modeling community, we present three examples of "community modeling" which differ according to the intensity of the involved collaboration. We propose new epistemic tools in order to conceptualize the apparition of these large-scale collaborations. We try to evaluate the relative weights of the epistemic, technical, and operational constraints that are put on the elaboration of collaborative meteorological, climate and environmental models.

Thomas Boyer and Cyrille Imbert (Archives Henri Poincaré, France)
"Scientific groups: the reason for collaborating"

For a given scientific problem, scientists can either work on their own and compete against each other, or can unite their forces and collaborate as a team. What is the best strategy for them? Instead of assuming an arbitrary reward or efficiency profile, we propose a basic sequential model of the research process, which accounts from the micro agent scale for these macro distributions. Investigating the model enables one to understand group configurations and group dynamics. For instance, we study why teams may want to grow bigger and bigger, at the expense of the community's efficiency. And we thereby shed light on the role of the priority rule in the development of groups through various fields of modern science.

Rogier de Langhe (Tilburg University, The Netherlands)
"To Specialize or to Innovate? An Internalist Account of Pluralistic Ignorance"

Academic and corporate research departments alike face a crucial dilemma: to exploit known frameworks or to explore new ones; to specialize or to innovate? Here I show that these two conflicting epistemic desiderata are sufficient to explain pluralistic ignorance and its boom-and-bust-like dynamics. The robustness of this result suggests that pluralistic ignorance is an inherent feature rather than a threat to the rationality of epistemic communities.

A key insight in social psychology (Allport 1924) is that individuals' public behavior is a function not only of their preferences, but also of the actions of others. Others' actions might lead individuals to act differently from their intrinsic preferences, in turn increasing the social pressure on others to comply as well. This self-reinforcing process is responsible for the typical boom-and-bust-like pattern typical of situations characterized by pluralistic ignorance (see Bicchieri and Fukui 1999, Kuran 1991). Among other things, pluralistic ignorance poses a serious challenge to the rationality of epistemic communities. Many insights from epistemology and philosophy of science, areas with a traditional focus on the individual, might not be robust against a generalization to many agents. It is to this challenge this paper turns. Its resolution requires an internalist account of pluralistic ignorance, viz. an account that can explain pluralistic ignorance in scientific communities from epistemic factors only. Here I show that this can be done by explaining pluralistic ignorance as a result of the conflicting epistemic desiderata of innovation and specialization. Academic and corporate research departments alike face this essential tension: to exploit known frameworks or to explore new ones; to specialize or to innovate? I show that pluralistic ignorance is a robust consequence of this essential tension. The robustness of this result suggests that pluralistic ignorance is an inherent feature rather than a threat to the rationality of epistemic communities.

Pluralistic ignorance in science

Pluralistic ignorance is the collective acceptance of a *norm* that agents privately reject but publicly accept because they believe others accept it. My focus here lies with norms in science, more specifically with how norms enter science in the classification of the world. This is a conventionalist perspective to classification, following Carnap (1950) and Kuhn (1962). Conventionalism is the view that there is no unique and optimal classification of the world. It became a popular doctrine after the development of non-Euclidian geometries. On a conventionalist perspective, how we "slice up" the world is essentially a social convention. How the world is sliced up matters because this affects what questions are raised and what counts as a solution to those questions. A classificatory framework gives meaning to scientific results and provides the basis for their rational evaluation. As a consequence there is no theoretical basis for choice between alternative classificatory frameworks. They are not "true" or "false". Radical conventionalism as held by the Strong Programme take this to entail that *all* classification is underdetermined and that individual scientists can change conventions whenever they please. But other forms of conventionalism are more subtle and claim that choice for a framework can indeed be rationally grounded albeit not in the world but pragmatically, for example based on considerations of fruitfulness. Just because there is no unique, optimal classification of the world does not necessarily mean that it is arbitrary or that individuals can change it unilaterally. For example anything can be agreed upon as money (e.g. tobacco, cows, shells...) but once a framework is in place agents cannot unilaterally change or cancel the norm. So although it doesn't make sense to compare the intrinsic qualities of the standards, there are important differences in pragmatic value that are contingent on the (social) context, such as the fruitfulness of a classificatory framework. Philosophers have often resisted such classificatory relativism because it makes our classification of the world relative to the circumstances. But classificatory relativism needn't lead to skeptical conclusions as long as one has a model of the circumstances, and as long as those circumstances are epistemic. This is what I set out to do in this paper. I explain pluralistic ignorance in science as a result of the social dynamics of adoption driven by two competing epistemic desiderata, innovation and specialization. In section two I introduce this essential tension. In section 3 I use them to explain the characteristic boom-and-bust-like pattern of episodes characterized by pluralistic ignorance and in section 4 I spell out how they constitute an internalist account of pluralistic ignorance by drawing only on epistemic factors.

REFERENCES

Allport, Floyd (1924). *Social Psychology*. Boston: Houghton Mifflin

Bicchieri, Cristina, and Yoshitaka Fukui (1999) "The Great Illusion: Ignorance, Informational Cascades, and the Persistence of Unpopular Norms", *Business Ethics Quarterly*, 9, 127–55

Carnap, Rudolf (1950). *Logical Foundations of Probability.* Chicago: Chicago University Press

Kuhn, Thomas (1962). *The Structure of Scientific Revolutions.* Chicago: Chicago University Press

Kuran, Timur (1991), "Now Out of Never: The Element of Surprise in the East European Revolution of 1989." *World Politics* 44(1), 7-48

Jeroen de Ridder (VU University Amsterdam, The Netherlands)
"What Could Group Knowledge Be?"

We often talk as if groups and other collectives are capable of having knowledge: 'The government knows that structural changes to the labor market are required.' 'The committee knew which candidate they wanted.' For science, such talk seems particularly natural: 'We now know that the earth revolves around the sun.' 'Mathematicians have proved Fermat's theorem.' 'The so-and-so lab has discovered a cure for X.'

Nonetheless, the idea that groups can have knowledge in any literal sense remains controversial in mainstream epistemology. Part of the explanation for this may be that proponents (e.g., Margaret Gilbert, Frederick Schmitt, Raimo Tuomela, Kristina Rolin, Brad Wray, Bird) and critics (e.g., Philip Kitcher, Ronald Giere, Kay Mathiesen, Melinda Fagan, Jennifer Lackey) of group knowledge use different conceptions of it and aren't always clear on what it is supposed to be and when something qualifies as group knowledge. Hence, real progress can be made by analyzing in more detail the ways in which groups can be involved in instances of knowledge and whether one or more of these ways gives rise to something that is group knowledge in a literal sense. I will argue in the affirmative, and, in doing so, will furthermore offer an argument to the effect that that group knowledge in fact has multiple senses.

I start with a brief discussion of a quick dismissal of group knowledge that hinges on the issue that groups do not literally have minds and mental states of their own, so that any notion of group belief or knowledge is a non-starter. This argument is uncompelling, since it may well be that there are group states that are so much *like* a state of individual knowledge that they nonetheless qualify as group knowledge.

Next, I investigate so-called summative construals of group knowledge, according to which group knowledge is a sum of individual pieces of knowledge of the group members. I rehearse and develop some well-known arguments against these construals. Although I will show that summative construals have more potential for dealing with alleged counterexamples than is sometimes admitted, I do wind up agreeing with the criticism that any summative construal of group knowledge is unsatisfactory, since it does not essentially involve the group in its depiction of group knowledge. That, it seems to, is a minimal constraint on any plausible conception of group knowledge.

I then look at non-summative construals of group knowledge, the most prominent of which are joint acceptance accounts. Common to these accounts is that they try to develop a construal of group knowledge by first giving an account of group belief and then adding further conditions to arrive at group knowledge. While such accounts do a better job in making the group essential to group knowledge, they also face considerable difficulties. First, they need to specify a plausible mechanism or procedure that can generate group beliefs. Jennifer Lackey has recently argued that this problem is insurmountable, since joint acceptance accounts cannot distinguish between group beliefs on the one hand and group lies and bullshit (in Frankfurt's sense) on the other. I consider her argument and respond to it. Second, the conditions under which a group belief constitutes group knowledge must be specified. I will argue (i) that existing proposals for doing this are largely unsatisfactory, (ii) that it can nonetheless be done, but (iii) that it is unclear whether anything in the real world meets the conditions thus specified so that it remains an open question whether this construal of group knowledge has any instances.

I also look at Alexander Bird's recent proposal for understanding 'social knowing'. Bird argues that since scientific knowledge plays the same functional role in societies as individual knowledge does in individuals, we should say that scientific knowledge is group knowledge. I clarify the details of his proposal and then argue against it on the grounds (i) that it is both too strict and too liberal an account of group knowledge: it excludes what I think are clear instances of group knowledge and it includes things which aren't group knowledge and (ii) that the analogy Bird uses to support his proposal is flawed.

Finally, I develop a proposal of my own for understanding group knowledge. The core of this proposal is that there are instances of knowledge which are such that satisfaction of the justification (or warrant) condition on knowledge essentially involves a group rather than an individual. I will argue that this is a sufficient condition for such instances of knowledge being group knowledge. I develop this proposal in more detail and defend it against various possible objections.

The conclusion will be that 'group knowledge' has different legitimate senses and that there are different legitimate concepts of group knowledge available.

Meghan Dupree (University of Pittsburgh, USA)

"Valuable but not Viable: Collaborations as Knowledge Producing Communities"

According to feminist contextual empiricism, as developed by Helen Longino, knowledge is fundamentally social. Individual epistemic agents cannot autonomously obtain knowledge. Instead, knowledge is primarily attributed to a community and is obtained by individuals via their membership in this community. Scientific knowledge is produced by the activity of a scientific community, and individual scientists can only have knowledge of content that has been appropriately processed and accepted by the scientific community.

Longino does not claim that any community of individuals counts as the appropriate kind of community to produce scientific knowledge. She suggests four norms that separate scientific communities from non-scientific communities. First, the community must have an appropriate venue or set of venues for discussing and criticizing content produced by community members. Next, there must be an uptake of this criticism, not merely a toleration of it. Third, the community must have public standards that are used for evaluating and accepting or rejecting the evidence. Finally, there must be tempered equality inside of the community. Tempered equality implies that a diverse range of perspectives are represented and included in the venues and criticism, and these differing perspectives are received with equal interest from the community.

In order to make her empiricism similar in form to previous versions of empiricism, Longino reworks the notion of justification to a notion of epistemic acceptability: the acceptability of epistemic content by a community that practices the aforementioned norms. Moreover, instead of claiming that justified scientific knowledge is true knowledge, she claims the appropriate measure of success for scientific knowledge is *conformation*. Knowledge *conforms* if it fits together with the previously accepted body of scientific knowledge to predict, explain, or manipulate phenomena in the material world.

On Longino's view, the findings of an individual researcher do not count as knowledge until they have been processed and absorbed by the scientific community. But what about the findings of a collaborative research project? If the collaboration adheres to the norms prescribed by Longino, then can the members of a collaborative research project know their findings?

In this paper I suggest that certain collaborative research projects qualify as scientific communities according to Longino's norms. If a collaboration counts as a community according to the norms, and the collaborators agree on certain findings, it follows, according to Longino, that the researchers know their findings. An individual researcher cannot claim to know the findings of her research project prior to acceptance of these findings by the scientific community. Therefore, collaborators know the findings of their study and individual researches do not. At first glance, this looks like a good reason to fund, promote and/or participate in only collaborative research projects! Furthermore, it appears counter-intuitive: the scientific community often treats the findings of collaborative research projects and individual research projects with equal epistemic credibility.

However, this argument treads on a fallacy of equivocation. Though there may be a technical sense in which the collaborators have knowledge of their findings, it is not the kind of knowledge that the individual researcher desires but lacks. Although some collaborations count as a communities according to Longino's norms, most scientific collaborations are not isolated communities but *sub-communities*. Sub-communities aim to provide findings that are known by a larger community. Due to their constitutive aims, most scientific collaborations are *subordinate* to the greater scientific community.

I then sketch a picture of epistemic relations between dominant and subordinate communities within contextual empiricism. Even if we accept that Longino's norms pick out a scientific community, I argue that community is not necessarily *viable* -- capable of surviving independently. A community C is *viable* if the community aims to provide knowledge that conforms to the purposes and projects of C. If a community C aims to have content that conforms to another community,C_M, then C is *subordinate* to C_M.

I further argue that Longino's norms must be revised for subordinate communities. Collaborators who hope to have their research adopted by researchers outside of the collaboration must not only have venues for sharing their work and uptake of criticism within the community, but must engage in these projects with the community they intend their findings to conform to. Likewise, they cannot develop their own public standards and practices for the evaluation of evidence, but must adopt and share the practices of the group that they are subordinate to. Finally, I discuss what it means for a collaboration to practice tempered equality. This equality must not be measured only with respect to the collaboration itself but with respect to the community the collaborators wish their findings to conform to. I then discuss how Longino's definition of knowledge is affected by the modifications of such norms, and argue that collaborative researchers don't know the epistemic content of their research findings prior to the acceptance of that content by the scientific community they are subordinate to.

I conclude by discussing the broader implications of subordinate communities. I address two worries one might have regarding the viability of communities. First, are there any viable communities? I argue that there are, in fact, viable communities and the subordination relation does not continue ad infinitum. Finally, I defend my view against the objection that the scientific community (at large) is subordinate to the general public. If scientists aim to have their research findings adopted by the general public, then it appears they are constrained by the general public in the same way that collaborations are constrained by the larger scientific community. This is problematic for my view as it would weaken the rigor of science and make it subject to public opinion. I argue that it is not a consequence of my view that the scientific community must conform to the evidence standards of the general public because, even if they wish for their findings to be adopted by the general public, they do not aim for their findings to *conform* to the knowledge and practices of the general public.

Paul Égré (ÉNS, France) and Olivier Roy (Münich Center for Mathematical Philosophy, Germany)

"Is common knowledge needed for coordination?"

To what extent is the notion of common knowledge (of preferences, rationality, or actions) needed between agents to ensure coordination? The notion of common knowledge was originally proposed by D. Lewis as a key component of the official definition of convention retained in his seminal book on the topic (Lewis 1969). Because of that, it is commonly thought that some amount of common knowledge between agents is necessary to ensure coordination, and in particular to achieve equilibria in coordination games. Since Lewis's original definition, skepticism has been expressed about the need for common knowledge on at least two fronts: from the standpoint of evolutionary game theory, where coordination equilibria are seen as resulting from very low principles of rationality (Skyrms 2004); and from the standpoint of epistemic game theory proper, where common knowledge has been argued not to be necessary to achieve Nash equilibria in games in general (Aumann and Brandenburger 1995). The aim of our paper is to clarify the issue, focusing on both criticisms. We argue that despite appearances to the contrary, Lewis himself did not quite view common knowledge as a necessary condition for the emergence or even the maintaining of coordination equilibria. However, he characterized it as a stability or reliability condition. We propose to relate this characterization to Aumann and Brandenburger's epistemic characterization of Nash equilibria, where common knowledge is presented as only a sufficient condition for equilibria, yet as a `tight' condition, such that its absence makes room for coordination failure.

Mads Goddiksen (Centre for Science Studies, Department of Physics and Astronomy, Aarhus University, Denmark)
"Collaboration and Authorship"

Historians and sociologists of science have described in detail how science has grown increasingly collaborative up through the late 19th and the 20th century (Beaver & Rosen, 1978; 1979; 1979; Beaver, 2001; Price & Beaver, 1966). Today, most new knowledge claims in science are produced by groups in which several scientists collaborate and combine knowledge, manpower, materials and other resources (Wuchty, Jones, & Uzzi, 2007).

This collaborative form of knowledge production raises new questions about how well, or even whether, the kinds of links which can exist between knowledge producer(s) and the knowledge produced are conveyed in scientific publications. Historically, it has been seen as transparent and unproblematic to link, for example, both credit and responsibility for the work of an individual scientist to the single author of publication which communicated that work to the scientific community. Historically, they have been the same person. As summarized by the historian of science Biagioli, “until the emergenc of large‑scale multi‑authorship … it seemed plausible to think of the scientist as the person who had the idea, did the work, wrote the paper, and took credit and responsibility for it” (Biagioli, 2000, p. 92). For collaborative knowledge production, however, with results being communicated through multi‑authored publications, the connections between authorship and credit and the ascription of responsibility are much more intricate.

First, collaborations often have fuzzy boundaries (Katz & Martin, 1997). Assessment of who should and who should not be counted as making an important contribution varies among groups and among scientists. This impacts the reliability of scientist composed authorship statements as indicators of collaborative contribution. It is not always the case that all co‑authors listed for a paper have taken part in the collaboration which led to the knowledge reported in the paper and, *vice versa*, sometimes some of those who have taken part in the collaboration are left off the list of co‑authors of the reporting paper. Based on a study of a large number of research groups from two collaborative research centres, Laudel (2002) has found that co‑authorship is usually granted for collaboration involving a division of creative labor and for time‑consuming service collaboration while it was seldom granted solely for the provision of access to research equipment or transmission of know‑how.

Second, between the knowledge producer(s) and the knowledge produced there is a link concerning both attribution of credit for having contributed to the production of that particular piece of knowledge *and* the ascription of responsibility for its reliability. These have often been assumed to be two sides of the same coin: you automatically take responsibility for contributions for which you receive credit. However, several misconduct cases in recent years have shown that collaborators themselves make a distinction between the two, willingly accepting a share of the credit for a well‑received paper, but distancing themselves from responsibility if it later turns out there are severe problems with the paper.

Despite these complexities, many qualitative and quantitative studies of the practice of scientific collaboration seem to assume a simple link between the knowledge producer(s) and the produced knowledge, and that authorship can therefore provide reliable, empirical access for the study of the practice of scientific collaboration, though this assumption has also been criticized (Laudel, 2002; Melin & Persson, 1996). At the same time, largely because of some well-publicized cases of scientific misconduct, there have been several discussions in the scientific community of how to attribute credit and ascribe responsibility to the individual authors of multi-authored papers. One discussion in particular, led primarily by editors of major journals in science and biomedicine, has arisen around this issue of responsibility and misconduct and has given rise to an attempt to explicitly identify the contributions made by authors to collaborative research projects. This has been implemented by several journals through the practice of requiring contribution statements to accompany journal articles.

In this paper we shall focus specifically on the question of responsibility in collaborative research. First, we view the editorial and journal discussions on the challenges of multi-authored papers in relation to responsibility and on how to meet these challenges. Next, drawing on the notion of epistemic dependence as introduced in the work by Hardwig (1991; 1985) we clarify the issue of responsibility in collaborative research. We describe various forms that collaborations can take, applying critically the taxonomies of Laudell (2002) and of Rossini and Porter (1979), and analyze their implications for the distribution of responsibility among collaborators. Finally, we shall compare our analysis to the suggestions made by some journal editors to split authorship into contributors and guarantors and describe the difficulties inherent in the suggested definition of guarantors.

Reference List

Beaver, D. & Rosen, R. (1978). Studies in scientific collaboration. Part I. The professional origins of scientific co-authorship. Scientometrics, 1, 65–84.
Beaver, D. & Rosen, R. (1979). Studies in scientific collaboration Part III. Professionalization and the natural history of modern scientific co-authorship. Scientometrics, 1, 231–245.
Beaver, D. & Rosen, R. (1979). Studies in scientific collaboration. Part II. Scientific co-authorship, research productivity and visibility in the French scientific elite, 1799–1830. Scientometrics, 1, 133–149.
Beaver, D. (2001). Reflections on Scientific Collaboration (and its study): Past, Present, and Future. Scientometrics, 52, 365–377.
Biagioli, M. (2000). Rights or Rewards? Changing Contexts and Definitions of Scientific Authorship. Journal of College and University Law, 27, 83–107.
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.
Katz, J. S. & Martin, B. R. (1997). What is research collaboration. Research Policy, 26, 1–18.
Laudel, G. (2002). What do we measure by co-authorships? Research Evaluation, 11, 3–15.
Melin, G. & Persson, O. (1996). Studying research collaboration using co-authorships. Scientometrics, 36, 363–377.
Price, D. J. d. S. & Beaver, D. d. B. (1966). Collaboration in an Invisible College. American psychologist, 21, 1011–1018.
Rossini, F.A. & Porter, A.L. (1979). Frameworks for integrating interdisciplinary research, Research Policy 8, 70–79
Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The Increasing Dominance of Teams in Production of Knowledge. Science, 316, 1036–1039.

Genco Guralp (Johns Hopkins University, USA)

"Collaborative Research in Cosmology: Discovering the Acceleration of the Universe"

This past year's Nobel Prize in physics was awarded to two teams which, working independently, confirmed the striking fact that the expansion of the universe is accelerating. For many cosmologists, this prize marks another major point in the chain of successful results cosmology obtained in its relatively short history of being an "experimental science." In fact, modern cosmology prides itself for becoming a *precision science*, breaking sharply with its "speculative" past. In this study, I examine the dynamics of collaborative research in this context of modern experimental cosmology, focusing on the discovery of the acceleration of the universe as a case study.

Even though the usual textbook account cites the serendipitous discovery of the Cosmic Microwave Background Radiation in (1964) as the beginning of the precision era in cosmology, many authors still refer to Edwin Hubble's (1929) discovery of the expansion of the universe, on the basis of his observations of "extra-galactic nebulae," as the turning period in the history of observational cosmology. It is remarkable that the problem that Hubble had to deal with, namely, establishing a reliable method to measure the distances to astrophysical objects, still continues to be one of the central challenges for modern observational cosmologists. In fact, the two experimental research collaborations that I examine in this paper devised new experimental techniques and methods of analysis to overcome the obstacles that cosmological distance measurements presented since the time of Hubble.[1] The main goal of both teams was to use supernovae explosion events as standard candles for cosmological measurements.

My paper is structured as follows: I begin by tracing the history of the two experimental collaborations, namely, the *High-z Supernova Search Team* (hereafter, HSS) and the *Supernova Cosmology Project* (hereafter, SCP), with the aim of focusing on the dynamics that eventually led to both teams' receiving credit for the discovery of the acceleration of the universe. This discovery, which was later interpreted in terms of a still not completely understood notion of "dark energy," forms an integral part of the modern *concordance* model in cosmology. Combining oral history interviews with the analysis of several key publications[2], which I refer to as the *evidence papers* (following [Staley, 2011]), I urge that the history of both these collaborations shows that the collaborative work they engage in is best characterized as what I call a *generative collaboration*. After establishing the historical background, I distinguish three features of a generative collaboration which marks it as a specific form of an epistemic group:

1. A generative collaboration sets out to obtain measurement results that have potentially wide-ranging significance for the science in question. Thus, both these programs aim at constraining the values of cosmological parameters, a project which is consequential for all cosmological models that purport to describe our universe. This contrasts with a narrow-range experimental program which aims at measuring, say, the X-ray emission from one particular binary-star system. This does not mean that the latter program has less scientific merit, but simply that it would be binding for a much narrower range of theories or models. As I explain below, this gives the generative collaboration significant *epistemic potency*.

---

1 Throughout the paper, I use the terms "experimental" and "observational" interchangeably. Even though the relationship between these terms may be problematic from a general philosophy of science perspective, I believe the context of this paper permits me to pass over this issue. I should also note that in the cosmology literature, these terms are very often used interchangeably.

2 The two main ones being [Perlmutter et al., 1999] for the SCP collaboration and [Riess et al., 1998] for the HSS.

2. A generative collaboration engages in a long-term research program which produces a significant number of papers as opposed to a program which culminates in a single (or a few) publication(s). For example, the collaborative work by the members of the HSS team spans a time of approximately two decades, and the SCP publications goes back at least fifteen years. Both collaborations have published more than 10 papers. I will refer to this as the *long-termism* of generative collaborations.
3. Being a long-term experimental research program, a generative collaboration establishes many auxiliary results which are then used in subsequent publications. As I document in my paper, the SCP and the HSS collaborations rely heavily on their own previous publications in their respective *evidence papers*, both to incorporate data from earlier measurement results and also to refer the reader to techniques and analytical tools developed earlier that are being used subsequently. In other words, a generative collaboration sets up a system of *internal referentiality*.

Thus, in the evidence paper of the HSS collaboration, we read: "Our own High-z Supernova Search Team has been assiduously discovering high-redshift supernovae, obtaining their spectra, and measuring their light curves since 1995. The goal is to provide an independent set of measurements that uses *our own techniques and compares our data* at high and low redshifts to constrain the cosmological parameters" (p. 1011, [Riess et al., 1998], emphasis added). In a similar fashion, the SCP evidence paper, after explaining the primary aim of the project as determining the values of the main cosmological parameters, adds that

"Goobar & Perlmutter (1995) showed the possibility of separating the relative contributions of the mass density, and the cosmological constant, to changes in the expansion rate by studying supernovae at a range of redshifts. The Project *developed techniques, including instrumentation, analysis and observing strategies*, that make it possible to systematically study high redshift supernovae (Perlmutter et. al. 1995a). As of 1998 March, more than 75 type Ia supernovae at redshifts z = 0.18—0.86 have been discovered and studied by the Supernova Cosmology Project (Perlmutter et. al. 1995b, 1996, 1997a, 1997b, 1997c, 1997d, 1998a)"[3] (p. 566, [Perlmutter et al., 1999], emphasis added).

Hence, both collaborations worked within the boundaries of their own research programs to a very large extent. I show that this *internal referentiality* is in line with the general prescription that was put forward in the inception of their respective projects.

Moreover, the collaboration's wide-range aims provides a crucial versatility to the project. This point becomes salient once we examine the various publications that the groups produced throughout the years, in comparison with the general goal that brought them together in the first place. Above I mentioned that the main aim of both groups was utilizing supernovae explosions for cosmological purposes. Yet, as the following list of titles taken from the publications of the HSS group attests, there are many different possible ways of studying supernovae towards that aim:

• "Optical Spectra of Type Ia Supernovae at z=0.46 and z=1.2"
• "Tests of the Accelerating Universe with Near-Infrared Observations of a High-Redshift Type Ia Supernova"
• "Supernova Limits on the Cosmic Equation of State"
• "The High-Z Supernova Search: Measuring Cosmic Deceleration and Global Curvature of the Universe Using Type Ia Supernovae"

As one can discern from above—and the case with SCP is strikingly similar—the versatility of the group means that it can deploy its expertise across several problems in observational cosmology such as the curvature of the space, nature of supernovae explosions, the equation of the state of the

3 All references in this quote are internal references, that is to say, they all refer to SCP collaboration papers.

universe etc. Previously, I referred to this characteristic feature of the generative collaboration as its *epistemic potency*. Epistemic potency gives the group the ability to command a wide epistemic field on the basis of its expertise.
After establishing these claims, and substantiating the three distinguishing features of generative collaborations that I introduced above, I conclude the paper with a discussion of the epistemic advantages that this type of collaboration possesses. Here, I make two observations, on the basis of my historical data. First, by spreading the establishment of the final result over a large span of papers which all pass through the peer review process independently, the collaboration significantly reduces the epistemic risk that is undertaken by the final discovery claim, which forms the high point of the endeavor. Even though this form of risk-reduction can be realized, at least in principle, by individual scientists or other forms of epistemic groups, the versatile character of a generative collaboration means that the final discovery claim will have a high degree of significance and will be robustly supported by evidence.[4] Secondly, as a wide-range research program has more potential for new results than a narrow one, a generative collaboration will, in general, have higher chances of securing funding and attracting talented students, two key factors which are crucial to sustain the long-term goals of the collaboration.

References

Perlmutter et al. Measurements of ω and λ from 42 High-Redshift Supernovae. The Astrophysical Journal, (517):565–586, 1999.
Riess et al. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. The Astrophysical Journal, (116):1009–1038, 1998.
Kent Staley. The Evidence for the Top Quark. Cambridge University Press, 2011.
William Wimsatt. Re-Engineering Philosophy for Limited Beings. Harvard University Press, 2007.

Koray Karaca (University of Wuppertal, Germany)

"The data-selection process of the ATLAS experiment as a distributed cognitive system"

According to Hutchins' account of "distributed cognition" (1995), cognitive processes can be distributed in a *distributed cognitive system* in the following three different senses:

> cognitive processes may be distributed across the members of a social group, cognitive processes may be distributed in the sense that the operation of the cognitive system involves coordination between internal and external (material or environmental) structure, and processes may be distributed through time in such a way that the products of earlier events can transform the nature of later events. (2001, p. 2068)

Recently, Giere has imported Hutchins's notion of distributed cognition into the context of philosophy of science to characterize cognitive processes in large-scale experimental systems. Giere has argued that in such systems distributed cognition goes beyond collective cognition in that it contains not only humans, but also instruments and other artifacts as parts of the cognitive system. In this sense, Giere (2007) regards distributed cognitive systems as hybrid systems that include both physical artifacts and ordinary humans.

In the present work, drawing on Hutchins' and Giere's accounts, I shall seek to characterize the data-acquisition system of the ATLAS experiment as a "distributed cognitive system". Several research units undertake the management of the data-selection system of the ATLAS experiment,

4 Here, I use the notion of robustness in the sense that is discussed in Wimsatt [2007].

which is currently underway at the Large Hadron Collider (LHC) at CERN.[5] The *Level-1 Trigger* unit consists of three sub-units; namely, *Calorimeter* and *Muons* which respectively execute the level-1data-selection triggers for calorimeter and muon detectors, and *Central Trigger Processor* (CTP) which generates the level-1 decisions. The *High Level Trigger* Unit manages the software infrastructure of the high level trigger system. This unit also consists of several sub-units; namely, *Physics* and *Event Selection Architecture* (PESA), *LVL-2 Infrastructure* and *EF Infrastructure*; these units are respectively responsible for the design and execution of selection algorithms, level-2 and level-3 triggers. *Data Acquisition Unit* consists of three sub-units, called *Dataflow*, *Online Software* and *Detector Control System* (DCS), each of which consists of several research units which manage various processes ranging from the monitoring of various types of data and the configuration of various parameters to the collection of experimental data. The *Data Acquisition Unit* manages the data-flow, the online software and the detector control system. In addition to the above-mentioned units, the unit called *Cross-System Activities* coordinates various activities across the aforementioned units. It is also to be noted that the operation of the data-selection system of the ATLAS experiment requires significant interactions between the above mentioned research units and other sub-groups of the ATLAS experiment, such as detector operation and data-analysis groups.

The above discussion suggests that the various cognitive processes involved in the data-selection system of the ATLAS experiment are managed by different research units, thereby indicating a *division of cognitive labor* among these units. This in turn exemplifies the first sense in which Hutchins suggests cognition can be distributed in a cognitive system. That is, cognitive processes associated with the data-selection system of the ATLAS experiment are distributed over the various research units. As has been stated above, the research units operating the data-selection system of the ATLAS experiment are in close interaction with other parts of the ATLAS experiment. For example, *Detector Interface Group* (DIG) provides coordination between the data-selection system and the ATLAS detectors. Given that those interactions with the external systems also involve certain cognitive processes, it would be correct to say that some of the cognitive processes associated with the data-selection system of the ATLAS experiment extend beyond the research units comprising the system itself, in that they are coordinated and executed jointly with some other research units which are parts of the ATLAS experiment. This suggests that one cannot talk about a definite cognitive boundary of the data-selection system. Therefore, the above considerations illustrate Hutchins' remark that cognitive processes in a system can be distributed in such a way that the operation of some of the cognitive systems requires coordination with external structures/units. Finally, I would like to point out that the data-selection system of the ATLAS experiment exemplifies Hutchins' third sense of distribution of cognition, i.e., cognitive processes can be distributed in time. To see this, let us note that the data-selection at the ATLAS experiment proceeds in three separate levels of increasing complexity and delicacy. At each level, "interesting events" are selected out from the collision-events occurring inside the LHC. The data-selection proceeds in such a way that the output of the previous level is the input for the next level. Note also that as higher levels of data-selection are attained cognitive processes necessary for data-selection become more complex, resulting in much finer selections of data. These considerations indicate that cognitive processes associated with the data selection system are also distributed in time in the way Hutchins suggests; leading me to regard the data-selection system of the ATLAS experiment as a distributed cognitive system in Hutchins' sense.

The case of the ATLAS experiment also vindicates Giere's claim that distributed cognitive systems are hybrid systems in that they are partly (dynamic) physical, computational and human cultural systems. In the ATLAS experiment, the first level trigger system, as well as its interface with the higher levels, consists of fast electronic devices. Whereas, the high-level trigger system is

---

5

See the ATLAS technical design report, sec. 17.

computerized; it is a large computer farm that has been design to execute specialized data-selection algorithms. Recalling that computer systems are also physical systems according to Giere's account, the above discussion suggests that the data-selection system of the ATLAS experiment is partly a physical and partly a computational system. Given that the overall data-selection system is managed by various research groups, it is also a partly human cultural system. A number of partly *autonomous* research groups in this system collaborate to achieve a common objective, i.e., data-acquisition. It is important to note that this objective is desired by all research groups and achievable by none of them individually.

References

ATLAS Collaboration, 2003: "ATLAS Technical Design Report: High Level Trigger, Data Acquisition and Controls", ATLAS-TDR-016; CERN-LHCC-2003-022.
Giere, R., (2004): "The Problem of Agency in Scientific Distributed Cognitive Systems", *Journal of Cognition and Culture* 4 (3-4): 759-74.

Hutchins, E., (1995): *Cognition in the Wild*, Cambridge, Mass.: MIT Press.

Hutchins, E., (2001): "Cognition, Distributed", *International Encyclopedia of the Social & Behavioral Sciences*, pp. 2068-2072.

Nicolas Lechopier (University of Lyon 1, France)

"Epistemic Communities and Reciprocity. Case-study of an Environmental-Health Research Partnership"

Research ethics concerns three kinds of issues. First, it concerns research practices, like observation, intervention, data collection, presentation of results etc. For all that, it is worth asking how diverse stakeholders taking part in the same research come to understand one another, maintain appropriate relationships, and account for differential power and varying relationships to values. Apart from actual interactions in the research arena, ethics imposes a system of regulation, which is traditionally based on standards or principles dictated by committees that evaluate research projects according to certain gold standards. (This paper does not discuss this second kind of issues). Research ethics is ultimately a field of conceptual analysis that help to clarify the complexities of moral order arising from research activities and enriching conversations among researchers, participants and society in a broader sense. In this last sense, research ethics belongs to the broader field of science studies, in which it addresses the issues generated by the multidimensional nature of scientific activities, their intersection in various axiological or value-oriented registers (epistemic, ethical and social values), and the various positions of research communities with respect to competing social interests.

One key transversal issue concerns the participation process itself. In global health research, genuine collaborations and partnerships became an ethical imperative (Emanuel 2004). Building relationships based on trust, mechanisms for sharing information, areas for formal discussions about the research, the ability of the research protocol to be impacted by the participants' involvement ..., all these issues became a matter of strategies or social techniques, as can be seen for example in the concepts of "community engagement" (Tindana 2007). Going away from this instrumental approach, I question in this paper the epistemological dimension of partnership building, which actually has direct ethical implications. My aim is to show that research cooperations between scientists and participants-non-scientists suppose a certain kind of knowledge concerning the epistemic attitudes of each partners, a reciprocal knowledge that contributes to buid something like an “epistemic community”.

I examined the case of an applied research program conducted by an international team in the brazilian amazon region, about the emergent problems of mercury contamination and Chagas disease related to poor land-use. Three small-scale communities participated in this 4 years project within which various interdisciplinary subprojets were conducted simultaneously about a range of environmental, health and social issues. In 2010, I conducted a qualitative inquiry on the research process itself, observing a serie of field work activities and interviewing researchers and local participants (n=45). I looked for situations of ambiguities and epistemical tensions, and tracked their modalities of resolution. This resulting paper paper focuses on two steps of the research process : when data are collected ; and when the researchers and participants evaluate the “results”.

(1) Research practices are, at the stage of collection of data, a mainly practical matter : it's all about preparing instruments, organizing transportations, interacting with people, collecting and storing samples, etc. These practices are contiguous and sometimes overlap with mundane local practices like planting, tending, harvesting, daily communications, etc. It appeared that research as a practice in its own right has sometimes a complicated relationship to the custom or local practices. Resolutions of these issue have been found during field experiences (for example, a researcher and a local repearing together a measurement device, giving an occasion of co-learning).

(2) Concerning the step of sharing the results, a difficulty arose after a proposed intervention : farmers were encouraged to experiment with agroforestry techniques that differed from their usual slash-and-burn methods, on one of their own plots. This experiment, supported by the research team, was designed to be as realistic as possible, meaning as close as possible to their customary farming practices. Various factors (weather, land features, organizational models, etc.) altered the anticipated conditions and practices planned, resulting in the apparent failure of some of these trials (some plantings, not all, ‘yielded nothing’). This lead to the revelation of a discrepancy between the criteria of success for an experiment and the criteria for a successful practice, i.e. since it is an experiment, rewards can be reaped even when the planting itself has not borne any fruit. There was a tension between the conducting of an ‘experiment to learn’ and the practical reality of obtaining truly operational results (good agricultural output) identified as being positive by the community itself. Dealing with this tension needed a range of mutual learning.

These kinds of epistemic discrepancies reveal potential misconceptions – or at least a certain pluralism – about the interpretation of what it is to do “research”, and possibly threaten the research partnership. In this case study, these issues have been managed by a range of practical, verbal and contextual activities that contributed to the building of an epistemic community, that is to say a more lucid – a more epistemically virtuous – research partnership.

Indicative bibliography

Amin, A. & Roberts, J. Knowing in action: Beyond communities of practice. *Research Policy* 37, 353–369 (2008).

Emanuel, E. J., Wendler, D., Killen, J. & Grady, C. What Makes Clinical Research in Developing Countries Ethical? The Benchmarks of Ethical Research. *The Journal of Infectious Diseases* 189, 930–937 (2004).

Grasswick, H. E. Scientific and lay communities: earning epistemic trust through knowledge sharing. *Synthese* 177, 387–409 (2010).

Phillips, L. J. Analysing the dialogic turn in the communication of research-based knowledge: An exploration of the tensions in collaborative research. *Public Understanding of Science* 20, 80–100 (2011).

Tindana, P. O. *et al.* Grand Challenges in Global Health: Community Engagement in Research in Developing Countries. *PLoS Med* 4, e273 (2007).

Wynne, B. Public uptake of science: a case for institutional reflexivity. *Public Understanding of Science* 2, 321 –337 (1993).

Conor Mayo-Wilson (Carnegie Mellon University, USA)
"The Dynamics of Scientific Collaboration"

Some scientists (e.g., Kahneman and Tversky) form collaborative relationships that last the entirety of their careers. Others, however, collaborate with different researchers at different times; old collaborations break apart, and new ones are formed. The dynamics of collaboration are often driven by social factors, such as personality conflicts, geographical proximity, institutional affiliation, and so on. This raises the question, "is there any epistemic benefit to encouraging scientists to change collaborators over time?" I argue there is: discovery is hastened by dynamic collaborative relationships.

This paper develops a simple game theoretic model to investigate the relationship between collaboration and speed of discovery in scientific communities. The game has two players: Nature and the scientific planner. It is a strategic, zero-sum game in which Nature's goal is to slow scientists' acquisition of knowledge, and the planner's goal is to design scientific communities that maximize the speed of learning. In the Nash equilibria of the game, both players employ *mixed strategies*, and the planner's mixed strategy ensures that (i) a diversity of research methodologies are used in the scientific community, and (ii) collaborative relationships change over time.

The game operates as follows. Nature chooses a state of the world that is unknown to the planner and to a community of scientists. For simplicity, I assume the state of the world is just some unknown real number.

At the same time, the planner makes two choices. First, to each scientist in a fixed community, she assigns a *method* for investigating the world. Methods might represent different experimental techniques for measuring an unknown physical constant, or they might represent different approaches to modeling some phenomenon. Together, nature's choice of a world and a scientist's method determine a sampling distribution, which describes the probability that the scientist will observe different types of data about the world at each stage of inquiry. In any given world, some methods are unreliable and produce data with high variance; others reliably indicate the true state of the world.

Second, the planner also chooses a *scientific network*, which represents which scientists form collaborations and share their findings. Importantly, the planner can choose to design a dynamic scientific network that evolves over time. The speed with which scientists learn is determined by the data they observe and by the information they receive from their neighbors in the network on each stage of inquiry.

The model has two crucial assumptions. First, no method is optimal in every world. In other words, for any method, there is some state of the world in which the method produces data with high variance, thereby failing to reliably indicate the true state of the world. Second, scientists cannot form collaborations with all others: at any given time, they can share their findings with only some small subset of researchers in the network. When these assumptions are made sufficiently precise, it can be shown that the game has a mixed strategy Nash equilibrium. One way the scientific planner can **ensure** the worst-case payoff associated with the equilibrium is (i) to diversify the set of research methods employed by scientists, and (ii) to design a dynamic network, in which scientists change the researchers with whom they share their findings at different stages of inquiry.

The paper ends with a discussion of the limitations of the model. Three deserve special mention. First, the notion of collaboration in the model is fairly shallow: two researchers are said to collaborate when they share their knowledge and change their beliefs accordingly. Real world collaborations often involved more than two individuals and are quite a bit more complex. For instance, different scientists might complete different parts of a large project (e.g., one performs the experiment, another analyzes the data, and a third develops a theoretical model). Second, while the model allows collaborative relationships to be dynamic, it assumes that each researcher employs the same method from one stage to the next. However, early in their careers at least, scientists often change the experimental and theoretical techniques they use in their daily lives. Finally, the model provides little guidance to scientific institutions (e.g., the National Science Foundation in the United States) interested in funding research or providing incentives to scientists to pursue collaborative projects.

Ryan Muldoon (University of Pennsylvania)
"Investigating Competition and Cooperation in Science"

Previous models of the division of cognitive labor implicitly assume that scientists are always better off working on their own. This talk seeks to explore how scientists decide whether they are better off collaborating with others, or competing with them. I argue that these decisions are driven by the demands of the problems under investigation, and the skills constraints that labs face. Thus, the formation of epistemic groups is driven by the epistemic landscape.

Eddie Soulier and Elie Abi-Lahoud (University of Technology of Troyes, France and University College Cork, Ireland)
"Social Epistemology for Knowledge Fostering in Online Communities of Intelligence"

I. Introduction

"Philosophers and researchers have been long studying, under Epistemology, Knowledge and Justified Belief . Classical theories define Knowledge as Justified True Belief. They consider that the subject (Epistemic agent) relies on individual mechanisms such as perception, memory and reasoning in order to assess her beliefs and to generate or acquire Knowledge. Many reprove those theories for isolating a subject while reasoning about the truth of a given Belief instead of considering her living environment and social interactions. In particular, the Social Epistemology theory, promoted in the 80's mainly by Alvin Goldman, explores "the ways that human Knowledge can be increased by social transactions". Social Epistemology studies several cases of human-to-human Knowledge exchange models as opposed to the traditional human-to-world reasoning models.

Those so-called *social transactions* are in the heart of the ISICIL project serving as a ground for our paper. Information Semantic Integration through Communities of Intelligence onLine proposes to study and to experiment with the usage of new tools for assisting corporate intelligence tasks. These tools rely on web 2.0 advanced interfaces (social network, blog, wiki, social bookmarking) for interactions and on semantic web technologies for interoperability and information processing . As a test bed, ISICIL provides a platform aligning social and semantic web technologies for scientific information monitoring. Currently, four major software functionalities are implemented in the platform: expert lookup, collaborative document edition, collaborative terminology edition and activity tracing. Several user interfaces are available at this web address http://isicil.inria.fr/v2/index.php.

Scientific information monitoring is indeed a belief-centric task whereas a watchman is confronted on a regular basis to the classical belief dilemma: what to keep and what to reject. The following work consists of tackling this dilemma from a Social Epistemology perspective.

II. SOCIAL EPISTEMOLOGY

Social Epistemology is often presented as a social theory of Knowledge . It aims at studying the impact of social factors on Knowledge adoption. Its scope extends to the study and identification of valuable Knowledge-related practices as Goldman states: "*Social Epistemology would identify and evaluate social processes by which epistemic subjects interact with other agents who exert causal influence on their beliefs*" .

A. Genesis

Early philosophical efforts fighting obscurantism opposed Knowledge to Belief (episteme/doxa), the former being what the latter is not. Notions of Truth and Reality came afterwards to state that Knowledge is rue while Belief is false, Reality being the Truth regulator: True is in conformity with Real. When Truth deviates from what seems real, a "rational" process of Justification is used to evaluate a Belief, defining Knowledge as Justified True Belief. Three major evolutions explain the need to evolve from this classical view of Epistemology to Social Epistemology, among others a more Situated view of cognition and dissemination of Information Technology in society.

B. Foundations and Major Approaches

Social Epistemology is founded upon three major hypotheses on (i) Knowledge boundaries, (ii) Knowledge nature and (iii) Epistemic involvement.

Several approaches (or theoretical trends) of Social Epistemology are studied in literature. They differ by their degree of compliance with the aforementioned hypothesis. Veridistic approaches are Truth centric. They are the closest to traditional Epistemology. Aretist approaches add the notion of Trust between agents. Argumentative approaches substitute the notion of Truth with Agreement (or Disagreement). Pragmatic approaches focus on the social use of epistemic practices (to satisfy, to make happier). Holistic approaches consider Knowledge as a shared result of a social cooperation.

C. Goldman's V-value Framework

The rationale behind Alvin Goldman's veritistic social epistemology is to "*evaluate social practices in terms of their veritistic outputs, where veritistic outputs include states like knowledge, error and ignorance*" .
The central element to Goldman's framework is the concept/ notion of veritistic value designated by v-value. Goldman uses the v-value to evaluate impacts of social practices in terms of Knowledge acquisition.

D. An Implementation of the V-value Framework

In the scope of ISICIL, the following implementation of the v-value framework adds a functionality allowing a watchman to address questions to a set of experts, to evaluate their answers and to decide on whether to adopt or not their beliefs.

We identified the following Socio-Semantic variables inspired by Goldman's framework to which we add the aretist dimension (Trust): ***Asker's Interest, Asker's Expertise, Informant's Interest, Informant's Expertise, Guaranties on Information, Informant's Reliability.***

An electronic forum allows us to capture the structure of an interaction between agents. However, traditional forums lack semantic descriptions, which complicates Knowledge capture.

We rely on the shared social norms approach described by to capture Knowledge within an interaction between an agent and a group of experts over a forum. Shared social norms are sets of instructions and rules guiding the use of a tool. They allow implicit capture of Knowledge from an interaction. A common example of shared social norm is the wikipedia chart. Another common example is *facebook's like button.* Its use implies a trust contract between users. It allows a user to express a preference that the system catches implicitly and uses for several deductions.

E. Socio-Semantic variables quantification

In an on-going work, available at the ISICIL website, we describe the suggested techniques to quantify, in the context of ISICIL, the variables listed in the previous section.

III. CONCLUSION AND FUTURE WORK

In this short paper we introduced the Belief dilemma and stressed on its complexity. We reviewed afterwards Social Epistemology, its foundations and major approaches. Finally, we presented the basis of a veridistic model implementing Goldman's V-value Framework.
Future work includes combining the previously presented variables in an aggregation algorithm.

IV. BIBLIOGRAPHY

Steup, M. (2011). Epistemology. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy.
Goldman, A. (1999). Knowledge in a Social World. Oxford Universtiy Press.
Goldman, A. (2006). Social Epistemology. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy.
Goldman, A. (2009). Social Epistemology: Theory and Applications. Royal Institute of Philosophy.
Gandon, F., & al., e. (2009). ISICIL: Information Semantic Integration through Communities of Intelligence onLine. PRO-VE'09, WIVE09. Thessaloniki: Springer.
Bouvier, A., & Conein, B. (2007). L'épistémologie sociale : présentation. In L'épistémologie sociale : une théorie sociale de la connaissance (pp. 9-54). EHESS.
Longino, H. (2002). The fate of knowledge. Princeton University Press.

Alain Trognon and Martine Batt (University of Lorraine, France)
"From results on Social Psychology of Collective Induction in Experimental Groups to Hypothesis on Epistemology of Collective Induction in Scientist's Groups"

---

1. Introduction (and abstract)

In almost a century of research's works, experimental social psychology of problem's resolution in groups have gain an extensive set of data firmly covered, thin, robust, convergent, orderly and, cumulative, which constitute indisputable resources for practitioners of groups and soon for designers of these *multi-agent systems* which are rapidly developing in *Artificial Intelligence*, in *robotics* and in *social-technologies*. Now, we know why and how it may be relevant to use a dyad or a group for respectively reinforcing an acquisition and finding innovative solutions. We know the conditions that must be met so that a group and a *team*, which is a "higher" form of group (Forsyth 2010), are productive (Johnson & Johnson 1992; Quin, Johnson & Johnson 1997).

Are epistemologists interested in these findings? That is the problem we will explore after having summarized the findings of experimental social psychology of problems resolution in groups and showing their relevance to scientific research groups.

2. What experimental social psychology has teached us about thinking and reasoning in groups?

2.1. Groups are more powerful than individuals in all cognitive activities

In experimental social psychology (Blanchet & Trognon 2008; Trognon & Batt 2012; Trognon, Batt & Marchetti 2011), one knew very early (Hill 1982) that in all the *tasks of memory*, groups are more powerful than individuals, although they generally do not reach their level of potential productivity (Clark & Stephenson 1989 ; Hinsz 1990). Groups process more data, of better quality, more deeply. They adopt strategies inaccessible to individuals, for example transactive memory, built by distributing information according to competences of their members (Wegner 1986).

Concerning *problem solving activities*, one know since 1982 that groups of size *m* performed better than an *average individual* independent and since 2002 that groups of 3, 4 and 5 members perform on *collective problems* better than the best of an equivalent number of individuals whereas groups of two perform at the level of the best of two individuals (Laughlin & al. 2002, 2003, 2006 ; Trognon, Batt & Laux 2011).

2.2. Factors of groups' competences

What explains the superiority of groups on individuals in problem reasoning and *concept attainment* ? It is the fact that groups, not only *gather* (cf. *infra*) but also *organize the resources of their members*. So they can invent solutions and strategies which did not exist in these groups before interaction. Ten years ago, Moshman & Geil have supplied to this thesis a sticking experimental demonstration by establishing that groups of students composed of 5 and 6 members were between 8 and 9 times more efficient in the leading task of the experimental psychology of reasoning (the *Wason selection task* in its standard form 1966) than students working on their own; and in showing that this result is only obtained thanks to the implementation of a *collaborative reasoning*. Indeed, by comparing the response selected by individuals of the group before and after the interaction phase, the authors find that most people change positively and that "having a member who initially selected the correct cards was neither a necessary nor sufficient condition for group success" (Moshman & Geil 1998).

Because they are *social shareness* systems (Tindale & Kameda 2000 ; Kameda, Tindale & Davis 2003), groups are better *processors of information* than individuals (Hinsz & al. 1997 ; Laughlin 1999 ; Laughlin & al. 1991, 2003, 2006 ; Propp 1999). This shows that *biases* we sometimes observe in the data processing of the groups (the restriction of the processed information (cf. *supra*) to the information shared by members (Stasser 1992; Stasser and Stewart 1992; Stasser and Titus 1985, 1987; Stasser & al., 1989) and the *Group-Think* (Janis 1972, 1981) are often exeeded by groups. So groups learn, and more or less independently of their members (Argote 1993). These latter progress also, but not all (Laughlin & Sweeney 1977 ; Tindale 1989). Among those which progress, certain members profit more than of others (Laughlin & al 2003, 2006). Thus *the groups are sources of vicarious learning for their members, and also of inventive learning* (Moschman & Geil 1998). "Does effective group problem solving transfert to subsequent individual problem solving?"(Laughlin & al. 2006). Three types of experimental devices have been dreamt up to answer this question: the use of successive tasks (as in Laughlin & al. 2006); individually retesting the knowledge of the group's consensual response after its members have left it (Moschman & Geil 1998); individually retesting the performance in the field of the competence to be acquired as in *social psychology of genetic development* (Perret-Clermont & Nicolet 1988). In this last approach, the experimentation proceeds in three phases. A pre-test assesses the skills of the subjects in the field of the cognitive ability to be acquired. They are then distributed in an experimental group of *interactive training* or in a group control of *personal training*. They perform one or more post-tests individually by using the same test as in pre-test or a test that keeps going an explicit relationship with it. The results observed about many tasks of acquiring concepts, for instance the concept of *conservation* (of length, volumes, numbers, etc.) show that "in certain conditions a situation of social interaction, which requires that subjects coordinate their actions between them or they confront their points of view can lead to a subsequent modification of the individual cognitive structure "(Perret-Clermont 2001). Thus, a child who does not know the principles of the conservation can discover them by working in cooperation with some peers to a suitable task. It is often a hard, long-term, deep and effective learning since the child who has thus made a lot of progress often justifies his claims by using new arguments that neither he nor his mate have put forward during the interaction phase.

3. Experimental social psychology of working groups and Epistemology of knowledge emergence in scientific groups

3.1. Experimental social psychology and Epistemology have get both a common interest in *collective induction*

The epistemologists should be interested in these last findings. Indeed, as Laughlin wrote in 1999: "consider a scientific research group. The group members observe patterns, regularities, and relationships in some domain; propose hypotheses to account for them; and evaluate the hypotheses by observation or experiment. They map a distribution of group member hypotheses into a single group response by some social combination process. They test positive hypotheses that are expected to have the property of interest if the hypothesis is correct, or negative hypotheses that are not expected to have the property of interest if the hypothesis is correct. If the results of observation or experiment support their predictions, the hypotheses become more plausible; if the results fail to support the predictions, the hypotheses are rejected or revised. They may exchange hypotheses and evidence with other groups or individuals who are working on the same problem. These and many other cooperative groups, such as auditors, securities analysts, intelligence analysts, and art experts, engage in collective induction, the cooperative search for descriptive, predictive, and explanatory generalizations, rules, and principles » (1999 : 50-51).

The *collective induction* is the task which is studied in experimental social psychology of problem's resolution in groups[6]. These results should have an application in epistemology as well.

3.2. The main conditions of collective induction

According to Trognon, Batt & Marchetti (2011), collective induction have *four conditions*. (1) Cooperative groups produce collective inductions. A cooperative group is a group composed by interdependant individuals. (2) The tasks made by these groups favor the cooperation. (a) According to Steiner (1966, 1972) cited by Laughlin & al. (2003), those tasks are disjunctive-unitary and complementary one. In disjunctive-unitary tasks each member performs the same task

6

We will present a task built to study collective induction during the conference and the group selects the best solution, thereby performing at the level of the best member. Complementary tasks are tasks « in which the single individual performs only a part of a total task, while other persons, possessing different kinds of resources, perform the remaining parts."(Steiner 1966 : 280). Collective inductions have these two fundamental features. (b) Groups' members must agree to a solution (consensus tasks). (3) Groups perform cooperative decision making, which is « essentially a process of resolution of disagreement in formulating a collective group response » (Laughlin 1999 : 51).

Five processes manage to a collective group response: random selection among proposed alternatives, voting among proposed alternatives, turn-taking among proposed alternatives, demonstration of preferability of a proposed alternative, and generation of a new emergent alternative. « Rather than voting, turn taking, or accepting a demonstration, groups may generate a new response on which members with different beliefs or preferences may agree. If the emergent group response is better on some criterion of truth or value than any of the responses favored by the group members, this process may extend beyond a way of resolving disagreement to fulfillment of the Gesttalt maxim that « the whole is greater than the sum of its parts » (Laughlin, Vanderstroep et Hollinghead, 1991: 52).

Demonstration occurs with intellective tasks, which are the problems for which there are demonstrable correct responses. Then groups adopt a solution when at least one member proposes it with logical or eureka problems or when at least two members proposes it with knowledge problems. More usually, « the number of group members that is necessary and sufficient for a group response is inversely proportional to the demonstrability of the response » (Laughlin, Vanderstroep et Hollingshead 1991 : 52). (4) In these groups, members interact freely (Trognon, Batt & Marchetti 2011). They welcome the critical discussion (Walton 2007), controversies (Bromberg & Trognon 2005 ; Johnson & Johnson 1997 ; Trognon & Galimberti 1996) and different kinds of dialogic persuasion (Walton & Krabbe 1995 ; Van Eemeren & Grootendorst 2011 ; Trognon, Batt, Sorsana & Saint-Dizier 2011) and they fight against fallacies. All these properties contribute to the superiority of collective induction compared with individual induction.

4. Conclusion: The risks of generalizing the results of casual groups to scientific groups

Whereas Laughlin and his colleagues think that « the results should provide understanding of many groups who seek, propose, and act on collective induction » (1991 : 62), we must care not to assign the good results of casual groups doing collective induction to scientific research groups doing the same cognitive activity (although at an upper level). On the one hand, scientific research groups are not immunized against the different kinds of dysfunction which often occur in groups, for instance the « group think » (Janis, 1972). Perhaps they are even less preserved than casual groups. We all know a lot of scientists who were unfairly treated by the scientific establishment (cf. F. Dagognet : *Tableaux et langages de la chimie*). On the other hand what is good for casual groups is not necessarily good for scientific groups. For instance, Laughlin & al. (2003: 684) wrote : «crucial factor in this superiority of groups over independant individuals in problem solving is the demonstrability of the operations, strategies, and procedures leading to problem solutions ». Demonstrability have four conditions : « (a) a group consensus on a conceptual system ; (b) a sufficient information ; (c) incorrect members (who) are able to recognize the correct response if it is proposed ; and (d) correct members (who) have sufficient ability, motivation, and time to demonstrate the correct response to the incorrect members. As demonstrability increases, problem-solving groups perform increasingly better than individuals » (ibidem : 684). The group consensus on a conceptual system promotes objectivity (Trognon, Batt & Marchetti 2011). But Tindale (1995) demonstrated that groups rather use common data than uncommon one. This 'natural' weakness leads to biased judgments. Since common data are institutional data, scientific groups would rather promote established knowledge than innovating one. If our suppositions are right, then the scientific revolutions don't have their origins in the scientific research groups.

References[7]

Blanchet, A., Trognon, A. (2008, 2° ed.). *La psychologie des groupes*. Paris : Armand Colin.

Trognon, A., Batt, M. (2012). Group Dynamics and Learning. Norbert Sell (Ed.), Chapter 1863. *Encyclopedia of the Sciences of Learning*. New York : Springer.

Trognon, A., Batt, M., Marchetti, E. (2011). Le dialogisme de la rationalité dans l'ordre de l'interaction. *Bulletin de Psychologie*, Tome 64 (5), 515, sept-oct 2011, 439-455.

Trognon, A., Batt, M., Sorsana, C., Saint-Dizier, V. (2011). Argumentation and Dialogue (pp.147-185). In A. Trognon, M. Batt, J. Caelen & D. Vernant (Eds.), *Logical Properties of Dialogue*. Nancy : Presses universitaires de Nancy.

Trognon, A., Dessagne, L. (2012). Equipes. Les équipes de travail. In José Allouche (coordinateur), Encyclopédie des ressources Humaines (3° édition). Paris : Hatier.

Trognon, A., Dessagne, L., Hoch, R. Dammerey, C., Meyer, C. (2012). Communications : les communications dans les situations de travail. In José Allouche (coordinateur), Encyclopédie des ressources Humaines (3° édition). Paris : Hatier.

Anna Zielinska (University Paris Descartes, France)
"Ethics committees without bioethical theory. Decision procedure in the evaluation of research protocols in biomedicine."

Biomedical research on humans needs to be evaluated by local committees before it is properly launched. It is generally presumed that this will be based on a theoretical framework, and will use theoretical tools. These tools in the context of biomedical research (and in clinical practice, which is not our subject here) are usually supposed to come from bioethics. The latter, even when it is described as an enquiry or a study, is supposed to be more than that.
In the classical *Encyclopaedia of Bioethics* (1995 edition), Warren Thomas Reich defined "bioethics" as "the systematic study of human conduct in the area of life sciences and health care, insofar as this conduct is examined in the light of moral values and principles". Yet, it is clear that the founders of the discipline, Beauchamp and Childress, did not only want to "study" human conduct, but also to guide it; their fundamental project was a constructive one. Their first and major contribution to international bioethics were the so-called Four Principles (cf. Beauchamp and

7 Nous donnons uniquement la bibliographie de nos travaux correspondant à cette communication ; les références données dans le texte sont toutes dans notre bibliographie

Childress, 1979). “Our goal was to develop a set of principles suitable for biomedical ethics”, noted Beauchamp in 2007, where each principle would be “an essential norm in a system of moral thought, forming the basis of moral reasoning”. He cautiously added that “more specific rules for health care ethics can be formulated by reference to these four principles, but neither rules nor practical judgements can be straightforwardly deduced from the principles.” These principles are (1) respect for autonomy, (2) nonmaleficence, (3) beneficence and (4) justice. These principles are both well intentioned and controversial in biomedical practice, whether in clinical or in research domains.

Bioethics as a discipline provides both a development and a constructive criticism of this kind of attempts. Its aims remain clear: to find a moral theory to guide the actions of practitioners in a biomedical context. And yet, why would this aim be desirable? The very discipline is founded on the premise that important decisions concerning human life and well-being can be guided by a set of principles, either *a priori* or *a posteriori*, which would be crucial to any decision made in concrete circumstances. This appeal to principles is supposed to protect us from arbitrary decisions made emotionally, in a situation of conflicting interests, or in biased circumstances. Independently of biomedical context, this idea received harsh criticism from many contemporary philosophers (cf. Dancy 2004 for details). The moral sphere is essentially complex, whereas the very idea of moral principles pre-supposes simple ideal situations (at least at the hypothetical level), where these principles might apply. Moreover, our effective judgment should take into account the whole context of the situation if they want to be truly relevant and helpful. Philosophers who criticise principles make often appeal to moral intuitions. This quite simple device is certainly unsatisfying in a biomedical context. Yet it seems that there is no way back to moral principles.

Given this double failure, I will defend here the following thesis: *collective and interdisciplinary expertise about each case is the only way to deal with controversial and problematic cases in biomedicine, in both clinical research and practice*. Garrard and Wilkinson have recently suggested that "bioethics is better off with moral theory than without it" (2003). My suggestion is that certainly, we do need moral theories in any thinking about moral decisions. Yet, these theories do not need to form moral systems, and in consequence, there is no need for a systematic enterprise such as bioethics.

This programmatic question constitutes the first, theoretical, part of this project. The second part is a case study of one of the French ethical boards, *Comité de protection des personnes* (this institution does not make medical decisions, it reviews research protocols in biomedicine before they are allowed to launch). I would like to present the kind of arguments used during the evaluation of research protocols, and to analyse whether there are bioethical principles advanced in the deliberation. Indeed, a huge variety of situations involved in each protocol may raise a huge variety of questions and objections. A property that looks crucial in a given case may not have any properly moral meaning, and finally lead to abandon the study. On the other hand, some bioethical principles (such as autonomy) may simply be dismissed, in the case of research into new procedures that will concern patients in emergency treatment, who are unable to give an express, informed consent. The way in which the arguments are raised by people from different backgrounds, with different worries, and very different experiences shows that this kind of epistemic community is able to create a unique and autonomous procedure of “moral evaluation” without appealing to bioethics.

It might be objected that the fact that there is an ethical committee (*Comité de protection des personnes*, e. g.) by no means excludes the idea of bioethics: indeed, members of the group can make appeal to bioethical principles, and apply them in relevant cases. Nevertheless, this is precisely the kind of normativity that should be avoided. The application of bioethical principles can be either inspired by an *a priori* moral framework, or an empirical one. The first is easily undermined by meta-ethical objections against principles (Dancy, quoted above). The second, however, is maybe more challenging, but, at the end of the day, even more intellectually frustrating. At the same time, people try here to take field research seriously and see how (bio)medical decisions are really being made *and* to somehow apply former findings to new empirical cases, but

here, the newest case is not being taken seriously in its particularity.
The only good reason to justify bioethics as a discipline (which means accepting its oversimplified meta-ethical framework), is to take seriously worries about the possibility of the proper evaluation of difficult moral cases in a biomedical context. This paper aims to show that this worry is unjustified, given the remarkable capacity of a cross-disciplinary ethical committee to provide an exhaustive account of analysed cases.

References

Ashcroft Richard E., "The Ethics and Governance of Medical Research", in Ashcroft & McMillan (eds), 2007, pp. 681–687.
Ashcroft, Richard E., John R. McMillan (eds) Principles of Health Care Ethics, John Wiley & Sons, Ltd, 2007.
Beauchamp, T. (2007), "The 'Four Principles' Approach to Health Care Ethics", in Richard E. Ashcroft, John R. McMillan (eds) 2007, pp. 3-10.
Brody Baruch A. (1998) The Ethics of Biomedical Research: An International Perspective, Oxford University Press.
Dancy, J. (2004), Ethics without Principles, Oxford, OUP.
Devettere Raymond J. (2009) Practical Decision Making in Health Care Ethics: Cases and Concepts, Georgetown University Press.
Gallin John I. et Ognibene Frederick P. (2012) Principles and Practice of Clinical Research, Academic Press.
Garrard, Eve, Wilkinson, Stephen (2003) "Does Bioethics Need Moral Theory ?", in Häyry Matti et Takala Tuija, Scratching the Surface of Bioethics, Rodopi.
Murphy Timothy F. (2004) Case Studies in Biomedical Research Ethics, MIT Press.