



Conference

Collective Knowledge and Epistemic Trust *Interdisciplinary Perspectives on Social Epistemology* Greifswald 6 – 8 May 2010

Abstracts

(in alphabetical order)

Max Albert

Methodology and Scientific Competition

Why is the average quality of research in open science so high? The answer seems obvious. Science is highly competitive, and publishing high quality research is the way to rise to the top. Thus, researchers face strong incentives to produce high quality. However, this is only part of the answer. High quality in science, after all, is what researchers in the respective field consider to be high quality. Why and how do competing researchers coordinate on common quality standards? I argue that, on the methodological level, science is a dynamic coordination game, with critical rationalism as a possible equilibrium.

Elizabeth Anderson

Democracy, Public Policy, and Lay Assessments of Scientific Testimony

Responsible public policy making in a technological society must rely on a great deal of technical scientific research. Legitimate public policy making in a democratic society requires that the reasons for such policies be open to discussion and evaluation by citizens at large. Yet lay citizens—those without technical training—generally lack the knowledge needed to directly assess the reasons scientists give for the conclusions they draw in the research that underlies public policies. Hence, there is a tension between these two demands—that public policies be both empirically responsible and democratically legitimate. The tension can be resolved by developing reliable second-order criteria lay persons can use to assess the trustworthiness of scientists who make claims about their research. Such criteria of trustworthiness are accessible to laypersons even if they lack expertise in the technical fields on which public policies are based. I offer some criteria and apply them to cases in which lay critics have raised doubts and controversy about scientific claims in areas such as global warming, genetically modified foods, the safety of nuclear waste disposal plans, and purported links between vaccines and autism.

Dieter Birnbacher

Distrust of Science and Technology in the Acceptance of Technological Risks

One thing that the psychology and sociology of risk have shown is that there is a sharp divergence between the laws governing the perception, assessment and acceptance of technological risks by laymen and by scientific and technological experts. Whereas the standards of the latter roughly correspond to a socially extended Bayesian decision theory that takes into account essentially (negative) utility and probability and assesses risks from a risk-neutral perspective, the assessment of technological risks by non-experts is oriented by aversion to risk as well as by a network of contextual factors including catastrophe potential, voluntariness, naturalness and perceived ability to control. A widespread view among scientists and engineers is that these factors, though explanative of the *acceptance* of technological risks, are of little relevance to their *acceptability* and mostly express purely

"emotional" reactions unrelated to the facts. Some even go further and argue that because of their irrationality, these emotional reactions should be ignored in technological planning. In the paper, it is argued that this criticism is only partly justified. It overlooks that emotions are not necessarily irrational and that the very concept of risk is inseparable from values inaccessible to a purely cognitive approach. Furthermore, even a Bayesian approach is bound to integrate a number (though not all) of the "qualitative risk factors" described by psychologists of risk given that they can be construed as negative psychological outcome values and should be included into a comprehensive calculation of risks.

Geoffrey Brennan

**The Division of Intellectual Labour and the Institutional Structure of Enquiry
(with some attention to the role of esteem)**

The point of departure in this paper is with the account of the role of the division of labour in human progress offered by Adam Smith in the first three chapters of *The Wealth of Nations*. Smith refers explicitly to the division of intellectual labour in this context—but he largely ignores the epistemic complications that trade itself brings; and he does not register the question as to how in general we can be confident that what others tell us is 'true' is indeed the case. Economists are familiar with the idea that ordinary exchange routinely creates "informational asymmetries" and generalized principal-agent problems. They are also familiar with the idea that institutional structures shape incentives and type-selection; though work on academic institutions specifically is relatively sparse—and tends (as does much economics) to deal with a "stripped-down" account of agent motivations that rather restricts the ways in which institutions can "get purchase" on behaviour. My ambition in this paper is to put these various strands of work together, including specifically allowance for the role of esteem in the operation of the institutions of enquiry.

Martin Carrier

**Values and Objectivity in Science:
Value-ladenness, Pluralism and the Epistemic Attitude**

I contrast a Baconian notion of objectivity, expressing a detached stance of scientists and a correspondence of their views to the facts, with a pluralist notion, involving a critical debate between conflicting approaches. These conflicts include substantive hypotheses or theories but extend to values as well. Many scientific judgments are value-laden in that value-commitments affect their assessment. Relevant values are epistemic and codify the features scientific knowledge is intended to exhibit, but also non-epistemic, in particular, pragmatic or utilitarian. Although the objectivity of science can be accomplished in large measure within such a pluralist framework, a shared epistemic attitude is required in addition. This epistemic attitude is epitomized in rules adopted by the scientific community concerning general principles of testing and accepting assumptions in science.

Paul Faulkner

Lies and the Problem of Trust

We frequently trust others to do things. And one of the things we trust others to do is tell us the truth. That others might have an interest in lying makes this trust risky; and, from one highly plausible perspective, this risk makes trust problematic. This paper argues that what renders trust non-problematic is the existence of certain norms of trust. These norms can underpin the rationality of accepting what others tell us, and cooperation more generally. This paper considers how this is so and how these norms of trust might be philosophically justified.

Gerd Gigerenzer

Trust Your Doctor: Paternalism and Collective Innumeracy in Health Care

The heuristic "If you see a white coat, trust it" guides much of patient-physician interaction, including cancer screening, which is the focus of my talk. Trust is an "ecologically rational"

strategy in a world where doctors (i) are free to practice evidence-based medicine without being penalized for doing so by malpractice litigation, (ii) are not surrounded by incentives that conflict with best treatment, and (iii) understand health statistics in the first place. However, trust is not ecologically rational in today's health care environment where these three properties are not satisfied. I provide evidence that (i) tort law and malpractice trials, believed to hold physicians accountable and thus improve the quality of care, can actually lead to harmful overtreatment and overmedication; (ii) conflicts of interest impede best care; and (iii) evidence-based medicine is virtually impossible given widespread statistically illiteracy in doctors, and that paternalism and geography-is-destiny operate in its place. When it comes to health, large sections of the general public, including devoted proponents of rational choice such as neo-classical economists, nevertheless prefer to trust doctors or avoid scientific evidence.

Sanford Goldberg

Trust in Oneself and Others

Agents are often in a position in which they rely on information from sources whose operations they cannot independently confirm. The paradigmatic example of this involves testimony from another source (speaker, newspaper or other media outlet, blog, web site, etc.) However, one's reliance on one's own cognitive processes can be seen as involving a kind of self-trust. In this paper I compare our trust of others to our self-trust: in what ways (if any) do these sorts of trust differ, and how do these differences affect the epistemological account we offer of each?

Kai Spiekermann

Courts of Many Minds

In *A Constitution of Many Minds* (Princeton UP 2009), Cass Sunstein argues that the three major approaches to constitutional interpretation all rely on some variation of a 'many minds' argument:

- "Traditionalists" would have judges defer to the "many minds" on previous courts;
 - "Populists" would have judges defer to the "many minds" in the electorate at large;
 - "Cosmopolitans" would have judges defer to the "many minds" on foreign benches.
- Here we assess each of these claims through the lens of the Condorcet Jury Theorem. In the cases of the Traditionalist and Cosmopolitan approaches we explore the implications of partial interdependence among past and foreign courts, respectively. In the case of the Populist approach, we consider ways of ensuring that the average voter is indeed better than random and that the views of a majority of voters should indeed be followed.

Russell Hardin

Collective Knowledge Acquisition

There are at least two ways to conceive collective knowledge. First, it is what everyone knows. Second, it is the sum of what each of us knows: it is what you know, plus what I know, plus what A knows, etc. I will deal exclusively with the latter and will refer to it as collective knowledge or even merely knowledge. On this conception, we might view our knowledge as the result of a division of labor. This makes the *acquisition* of collective knowledge a problem of collective action. I will focus on this problem of action and will assume that there is no difficulty in defining collective knowledge in the second sense above. If there is no hierarchical authority or control figure to parcel out different tasks, we might waste a lot of time and effort duplicating each other's contributions. If we manage this problem, there may still be overlapping efforts. If our two sets of knowledge overlap too much, however, there is little savings on efforts to establish knowledge through division of labor. Collective knowledge becomes especially valuable when it covers a vast territory well beyond anything an individual could master. Civilization and science depend on the massive acquisition of collective knowledge.

Rainer Hegselmann

Epistemic Networking: An Integrated Model of Opinion and Network Dynamics

In my talk I'll present an extension of the so-called bounded confidence model. The extended model will allow to analyse processes of the following type: (1) Agents modify their opinion on a certain topic in a social exchange process. (2) A true opinion exists, at least some of the agents are looking for the truth and to a certain degree they are successful in their quest for the truth. (3) The agents are networking. The social exchange process is confined to those others which are socially not too far away in terms of network distance. Taking (1) – (3) together we get a complicated and interwoven social-epistemic network dynamics. The talk will make use of ENSIM, a simulator that allows to experiment with these processes. The simulator especially allows to analyse cost and benefits of networking, measured in terms of distance to the truth.

Melissa A. Koenig

Learning from What Others Tell Us: Possible Mechanisms

When learning about the world, infants and young children depend deeply on others for information. Because systems of social transmission convey both reliable and unreliable information, it becomes crucial to evaluate the reliability of other sources. In this line of research, I investigate one important marker of a speaker's trustworthiness: whether she typically spoke truthfully in the past. Three studies examined 24-month-olds' sensitivity to the prior accuracy of the source and the way in which young children modify their word learning from inaccurate sources. In Experiments 1A, 2 and 3, toddlers interacted with an accurate or inaccurate speaker who linked a novel word with one of two objects. When tested in Experiment 1A, children performed less systematically in response to an inaccurate than to the accurate source. In Experiment 1B, children appropriately suspended their doubt when novel labels were taught by a different source, with no labeling history. Toddler's retrieval of the new word-object links was tested in Experiments 2 and 3. In Experiment 2, children responded randomly in response to a second speaker's comprehension tests when novel words were previously presented by an inaccurate source. In Experiment 3, toddlers responded randomly when their memory for words that had been trained by an inaccurate source was taxed by a brief delay period. Taken together, these findings suggest that toddlers attend to speaker reliability, treat inaccuracy as a feature of a particular individual and that the word-object representations formed as a result may be fragile and short-lived. Findings are discussed in terms of possible mechanisms by which children adjust their word-learning from problematic speakers.

Jennifer Lackey

Disagreement and Belief Dependence

At the center of work in the epistemology of disagreement is a debate regarding what is rationally required when one is faced with an epistemic peer with whom one disagrees about a given question. A and B are epistemic peers relative to the question whether p when A and B are evidential and cognitive equals with respect to this question—that is, A and B are equally familiar with the evidence and arguments that bear on the question whether p , and they are equally competent, intelligent, and fair-minded in their assessment of the evidence and arguments that bear on this question. While there is considerable dissent regarding the appropriate response to peer disagreement, there is nonetheless widespread consensus regarding a certain class of disagreements that rationally require no doxastic revision whatsoever. I call this thesis Independence and formulate it as follows: Independence: A's disagreement with epistemic peer B requires doxastic revision for A only if it is independent of other instances of disagreement that require doxastic revision for A. Despite both the widespread acceptance and intuitive plausibility of Independence, I argue in this paper that there is no interpretation of this thesis that turns out to be true. I proceed by distinguishing

several notions of belief dependence that may be at issue here—likemindedness, source dependence, testimonial dependence, non-independence, and collaborative dependence—and show that each fails to support the view that the amount of doxastic revision needed in the face of peer disagreement necessarily depends on belief independence. I then suggest that the considerations put forth in this paper provide further reason to accept the *justificationalist* view of the epistemology of disagreement that I have developed elsewhere, and reveal why such an account is preferable to rival views.

Gerhard Schurz

Meta-Induction and Epistemic Trust in the Evolution of Collective Knowledge

The justification of induction is of central significance for cross-cultural social epistemology. Different 'epistemological cultures' do not only differ in their beliefs, but also in their belief-forming methods and evaluation standards. For an objective comparison of different methods and standards one needs (meta-)induction over past successes. A notorious obstacle to the problem of justifying induction lies in the fact that the success of object-inductive prediction methods (i.e. methods applied at the level of events) can neither be shown to be universally reliable (Hume's insight), nor to be universally optimal. My proposal towards a solution of the problem of induction is *meta-induction*. The meta-inductivist applies the principle of induction to all competing prediction methods that are *accessible* to her. By means of mathematical analysis and computer simulations of *prediction games* I show that there exist meta-inductive prediction strategies whose success is universally optimal among all accessible prediction strategies, modulo a small short-run loss. The proposed justification of meta-induction is mathematically analytical. It implies, however, an a posteriori justification of object-induction based on the experiences in our world. In the final section I draw conclusions about the significance of meta-induction for the social spread of knowledge and the cultural evolution of cognition, and I relate my results to other results in evolutionary game theory which utilize meta-inductive learning mechanisms.

Michael Strevens

Secrecy and Sharing in Science: Resolving the Tension

Research programs regularly compete to achieve the same goal, such as the discovery of the structure of DNA or the construction of a TEA laser. The more the competing programs share information, the faster the goal is likely to be reached, to the benefit of society. But the "priority rule" —the scientific norm that enjoins that the first program to reach the goal in question receive all the credit for the achievement—provides a powerful disincentive for programs to share information. This paper will investigate ways in which the clash between social and self interest is resolved in scientific practice, considering extra-scientific solutions (e.g., the NSF's requirement that the results of research be made available in an appropriate way), self-policing within science (e.g., Merton's scientific norm of "communism"), and other, more private, arrangements.

Torsten Wilholt

Trust in Science: The Social Epistemology of Conventional Standards

In order to assess the trustworthiness of a research result, one would ideally have to consider all inductive risks taken in the course of the inquiry that ultimately led to the acceptance of the result. Many inductive risks are difficult to access, however, as they are implicit in practices of experimental design, data analysis, and other methodological choices. Therefore, a thorough appreciation of all inductive risks would turn the assessment of a scientific result into an enormously complicated task. It would help if the degrees of inductive risk were always to fall within certain limitations. However, as I will argue in this talk, *individual* rationality does not restrict the magnitudes of permissible inductive risks in scientific research. Instead, research communities *stipulate* standards that place explicit and implicit constraints on inductive risk. I suggest that these stipulations can best be understood as a

collective effort to enable and preserve epistemic trust in science, because they make a reliable assessment of the trustworthiness of research results practically feasible, both for members of the research communities themselves and for external “users” of scientific knowledge. This analysis contributes both to explaining the rationale behind apparently arbitrary methodological standards in the sciences and to consolidating the indispensability of social epistemology for the understanding of scientific methodology.

Marcus Willaschek

The Rationality of Trust: Epistemological Reflections on Jealousy and Evidentialism

The paper will argue for the importance of trust by investigating the epistemic standing of pathologically jealous persons. By trust, we mean an attitude towards a proposition p characterized by the following two features: (a) full acceptance: p is held to be true without any reservation; (b) conscious fallibility: the subject is aware of the fact that the available evidence for p does not guarantee that p is truth. Traditional epistemology has maintained that trust, thus defined, is irrational, since it violates the “evidentialist” principle according to which one ought to adjust the degree of belief according to the strength of the available evidence. Against this view, it will be argued that applying this principle to interpersonal relations would have the absurd consequence of turning pathological cases of jealousy into hallmarks of rationality and epistemological diligence. By contrast, if it is accepted that certain extreme kinds of jealousy are irrational, the evidentialist principle must be rejected.