PATRICK RYSIEW

INTRODUCTION

The papers in this issue address a wide range of topics and issues in social epistemology. The articles by Kornblith and Rupert treat foundational questions about the relation between individual cognitive agents and the larger groups and settings in which they operate. Those by Rescher and D'Agostino deal with issues concerning science considered as a joint enterprise, while Maki's focuses on the use of economic models in understanding science, and Schmitt's on the proper aims of education. The issue closes with an exchange between Sismondo and Brown on the nature of both science and the 'Wars' being fought over it.

In the opening article, Hilary Kornblith attempts to cast doubt on a tempting and traditional view of human reason. On this view, reason constitutes a 'neutral court of appeal' to which we might look in resolving intellectual matters. Against this, Kornblith cites evidence against the neutrality of individual reason, and argues that for its proper exercise individual reason requires the right kinds of social arrangements.

The traditional view of reason has going for it the phenomenology of individual reason: when I reflect on an intellectual matter, evaluating the evidence for or against a particular theory, it certainly seems as though I am simply examining the (de)merits of the theory itself, with nothing else intruding. But, Kornblith argues, the phenomenology here may be deeply misleading. (As Kornblith notes, this idea will be familiar to many feminist and Marxist theorists, for example.) Kornblith's point of departure is Frank Sulloway's work on the effect of something as apparently innocent as birth-order on one's intellectual temperament. According to Sulloway, firstborns tend to be significantly more intellectually conservative than later-borns in matters involving significant conceptual innovation. (Firstborns' conceptual conservativeness also tends to make them more professionally successful than later-borns.) If this is right, then if a panel of experts in which firstborns predominate is deciding the merits of some novel theory, the field is significantly slanted in favor of the going theory – not because of evidential matters, but simply because these people happened to have been born earlier than their siblings.

Biases such as these are not self-presenting, even if one is well-intentioned and generally epistemically responsible. Nor, in public discussion, can we count on biases to cancel each other out in. That all depends on the social distribution of the bias in question: if it is widespread and one-sided, public debate may serve only to leave certain views even more deeply entrenched. By the same token, however, a given bias need not be undetectable – it is unlikely, for example, to produce smooth distortions over the whole range in which it operates. The question is how to go about calibrating our methods of belief-formation so as to minimize bias' having any pernicious effects, epistemic or otherwise. There is no simple or merely formal solution to this problem, according to Kornblith. Most important, perhaps, will be for a community of inquirers to proceed in the knowledge that its work may be distorted by some common bias, and with a consequent willingness to take seriously specific charges to this effect.

While Kornblith aims to discredit naïve views of individual reason, Robert Rupert seeks to raise problems for those who think that talk of group minds or group mental states should be taken at face value. (See, for example, Margaret Gilbert's position on group belief, most recently articulated in her contribution to Issue 2 of this *Journal*.) While admitting that mentality should not be defined in such a way that it cannot genuinely

https://doi.org/10.3366/epi.2004.1.3.163 Published online by Cambridge University Press

occur at the group level, Rupert argues that current ideas about the mind do pose a serious challenge to proponents of group minds or mental states. (An endorsement of the latter entails that the former exist.) While we are far from having a complete science of the mind, certain ideas have emerged as our best current guides to mentality; and when we examine composite systems — collections of two or more minds – we find that they fail to exhibit the central features of minds as we know them best.

In general, talk of minds, and of specific mental states, can and does do serious causal-explanatory work. Indeed, for the majority of philosophers and cognitive scientists, it's not clear how else to conceptualize and explain large swathes of human behavior, if not in terms of certain types of internal, representational states. This is why cognitivism flourishes, and behaviorism has been abandoned. (Cognitivism also flourishes because we each have first-personal experience of our conscious states; there is no group analogue of this, Rupert thinks, though he sets this aside so as to focus attention on the causal-explanatory issue.) But it is not clear, Rupert argues, that there are any collective minds or group intentional states which might similarly do some real causal-explanatory work. Rupert agrees with Gilbert in rejecting what the latter calls a 'simple summative' account of group belief: a court's decision needn't be a majoritarian report of its members opinions, for example. Nonetheless, Rupert contends, that decision can be explained without essential recourse to group mentality, in terms of the minds and mental states of the relevant group's members. (By comparison, we understand only very dimly how those individuals' mental states and the laws governing them might be reduced to physical states and laws.)

Further, Rupert seeks to show that none of the going theories of mental representation (/representational content) apply in any straightforward way to group states. To this end, he presents a brief overview of a number of leading approaches to the understanding of representational content. In each case, he thinks, there are barriers to seeing the theory as applying smoothly to the group case. It may be, he suggests, that groups simply lack certain architectural features – for example, (quasi-)perceptual capacities – to which these theories assign a central place.

Rupert closes by pointing to a further *prima facie* disanalogy between the individual and group cases. Individual cognition is global: for a given topic or task, any of one's other mental states might in principle be relevant. (Hence the well-known 'frame problem' in artificial intelligence.) But in typical group systems the channels to access and affect representations are well-defined; if global considerations intrude, that is because individuals introduce them.

The next two articles deal with science – specifically, with issues of the distribution of goods which we think well-functioning science will share around. One of these goods is credit. In science, credit for discoveries leads to such things as enlarged opportunities and enhanced recognition. In the case of individual discoveries, it seems that how much epistemic credit should be given should be a function of two factors – the inherent importance or significance of the finding and the difficulty of arriving at it. The case is complicated where the discovery in question, like very many actual scientific discoveries, is the result of the conjoined efforts of more than one person – the product of a division of cognitive labor, in Philip Kitcher's phrase. Here, even if there is agreement as to how much credit ought to be given, there is still the question of how it should be distributed. This is the subject of Nicholas Rescher's article. According to Rescher, we must differentiate between two very different sorts of joint investigative efforts – distributive efforts and collective efforts. In the former, the parties involved engage in a coordinated effort towards some common epistemic goal; this is essentially a 'divide and conquer' strategy. Rescher's proposal is that, where a joint discovery results from merely distributive cooperation, the classic principle of 'fair-share proportionalism' applies: each contributor should get whatever credit properly belongs to their particular investigative contribution thereto.

Some cooperative investigative efforts don't lend themselves to a divide and conquer strategy, however. If two people are doing a crossword puzzle together, it won't be an effective strategy for one person to take responsibility for the

https://doi.org/10.3366/epi.2004.1.3.163 Published online by Cambridge University Press

vertical clues, the other for the horizontal ones, with the two combining their results at the end. The better strategy will incorporate interactive feedback between them as they proceed. Such feedback is the defining feature of what Rescher calls collective efforts. With regard to this type of cooperative investigation, Rescher's proposal is that credit belongs to the group members collectively and equally, since it is only as such a collective that they made the discovery at all. This does not violate fair-share proportionalism, however, since there are no competing individual claims: they did it together.

Whether inquiry is individual or cooperative, epistemic credit is due because of results obtained – our epistemic concerns are product-oriented, according to Rescher. This marks an important point of contrast with morality: our moral concerns are process-oriented; it is people's practices that matter, not any particular type of outcome. Rescher concludes by suggesting that this underlying difference in their teleologies or goal-structures underpins and sheds light on other important differences between the epistemic and moral domains – including why there is really no such thing as epistemic discredit, and why inadvertence cancels credit in the moral case, while epistemic luck does not.

Whereas in Rescher's article the focus is on the allocation of credit for a discovery, Fred D'Agostino's concerns the spreading-around of something else that is essential to well-functioning science: risk. Such is the thrust of Thomas Kuhn's 'risk-spreading argument', which D'Agostino sees as having been unjustly neglected. The core idea of this argument is simply that if all scientists reacted in the same way to a given new finding, science as we know it would cease. Most new ideas fail; on the other hand, some succeed; and even our best theories are surpassed in time. It thus seems that what we want is for there to be sufficient diversity among scientists that some will seek out and develop new lines of thought, while others will chisel away at safer, more well-entrenched positions. In short, we need diversity to ensure risk-spreading; and we need risk-spreading to ensure effective and efficient inquiry on the whole.

Thankfully, according to Kuhn, even if scientists agree in their values those values may themselves be a source of diversity. For one thing, those values are not self-applying – A and B might both prize simplicity as a virtue of a theory, while disagreeing about which of two theories is the simpler. Further, various values are multiply combinable – A and B might both prize simplicity and predictive power, while disagreeing about their relative weight, and so about the overall merit of a given theory.

In his article, D'Agostino suggests further ways in which a level of agreement or solidarity among scientists is compatible with a deeper diversity that is promotive of risk-spreading. One further source of diversity derives from the fact that anything is inexhaustibly describable — there are always more things to be said about a given thing than there is time to say them in. (Another characterization of the frame problem's entering wedge.) If so, then what one does say or think about a given thing will always be necessarily selective, in which case there will be room for diversity in the understanding of various scientific achievements, even those which constitute a successful research paradigm. Solidarity at the level of a shared paradigm remains, then, but not at the exclusion of diversity.

Kuhn himself regarded rules as hostile to diversity, and as therefore inhibitive of risk-spreading. That's not clearly so, D'Agostino points out – rules are no more self-applying than values, after all. But it is especially doubtful if we are mindful of the distinction between prescriptive and prohibitive rules. Individuals might behave in all sorts of divergent ways, while all complying with a given prohibition.

From the perspective of the risk-spreading argument, it is an eminently good thing for science overall that there can be situations in which individuals share certain core values but disagree on the overall merit of a given theory, or in which they are both operating within a given 'paradigm' or obeying a given rule, but elaborating or complying with it in different ways. It is, in short, a very good thing from this perspective that scientists' 'shallow' agreements, as D'Agostino puts it, can coexist with 'deeper' disagreements. For, once again, in such a case solidarity is preserved, but there remains residual diversity, the fuel for further risk-spreading down the road.

It is interesting to compare this picture of a well-functioning science with "the discursive paradox", as it has been termed by Christian List and Philip Pettit (and as discussed by Alvin Goldman in the inaugural issue of this *Journal*). For, as D'Agostino explains, situations such as those just described can give rise to cases in which there is an apparent divergence between individuals' collective conclusion on some question – e.g., the overall superiority of theory *a* to theory *b* — and the views of those involved on the 'premises' – e.g., for each criterion, the majority of individuals' judging theory *b* superior. Such collective judgments would not be rational, by List and Pettit's lights. Insofar then, as the type of diversity Kuhn and D'Agostino describe is essential to effective scientific inquiry, we must either revise List and Pettit's requirements on rational collective judgments, or see some measure of irrationality as among the costs of a well-functioning science.

Yet another source of risk-spreading within science, of course, is precisely the pursuit of credit: the promise of superior rewards can be an incentive to pursue ideas and lines of research that are off the beaten path. This is the sort of non-epistemic determinant of scientists' behavior that has led some to the conclusion – one extreme in 'the Science Wars' – that science really isn't about truth-seeking at all. Others have countered, however, that such supposedly 'grubby' motives can, in fact, have very good epistemic effects at the level of the larger scientific community (Kitcher, for example, as well as Alvin Goldman & Moshe Shaked). Just as (we're told) the behavior of self-interested individuals in a free market can maximize social welfare as an unintended consequence, it might be that credit-seeking on the part of scientists within a free market of ideas will have the effect of promoting superior truth-acquisition overall. Thus, economics serves as a resource for social epistemology – in this case, as a resource for defending science as worthy of the epistemic respect it has traditionally been afforded.

Uskali Mäki raises concerns about the use of economics within social epistemology, however. There may be no in-principle incompatibility between credit-motivation and truth-acquisition. But neither is economics bound to yield results which favor the 'pro-science' side in the Science Wars. For example, both informational cascades ('herd thinking') and externalities and path dependence – other ideas in the economist's toolbox – are mechanisms which can stifle diversity.

In addition to undercutting the idea that economic epistemology is inherently pro-science, Mäki discusses the important question of the nature and function of economic models themselves, whether or not they are applied to intellectual markets such as science. If, as Mäki suggests, such models provide only 'how-possibly' explanations, then no specific conclusions about how science actually operates can be drawn from them. There are, moreover, issues of reflexivity. First, as Mäki says, economics is itself a contested and heterogeneous discipline – much as science is. In order to assess the credentials of 'economic epistemology', then, we need an epistemology of economics, including a social epistemology thereof. Second, some researchers have suggested that repeated exposure to economic models which assume self-interested actors tends to encourage people to become such actors themselves. If this is right, then one wonders whether a specific model of scientific behavior might not also, in being current, tend to shape science itself in its own image. – All the more reason to think that, while economics is a valuable resource on which social epistemologists can draw, it must also be among the things they study.

As in Rescher's article, the notion of an activity's having certain proper aims figures prominently in the article by Frederick Schmitt, who adds to his contributions to social epistemology by addressing the question, what are the aims of education? ('Education' here refers to the social activity of teaching, involving at least one teacher and at least one student, inside or outside of an institutional setting.) As posed by Schmitt, this question does not ask after the motives of educators, the characteristic effects of education, or yet again which among those effects we take to be valuable. The question, rather, concerns the proper function of education – which effects education, by its nature, ought to produce.

Among Schmitt's central claims is that there is no nonvacuous intrinsic aim common to all types of education. Thus, 'practical education' – education in some trade, skill, practice, and so on – aims to produce students who are competent in that art, period. Practical education will therefore vary in its aim, depending on the subject matter (or art) in question. Nor, because it aims simply at producing one competent in the art, does practical education have as an intrinsic aim anything specifically epistemic – true belief, justified belief, rational belief, and so on.

In both of these respects, practical education differs from liberal arts education, according to Schmitt. For the latter does have a specifically cognitive-epistemic aim, and it has this aim regardless of its particular subject matter: besides acculturation beyond what's ordinarily supplied by one's upbringing, liberal arts education as such aims at the fostering of a certain sort of intellectual excellence, the possession of 'a sound mind', in the student.

But how should the latter be understood, exactly? Schmitt suggests that aiming to produce a sound mind means aiming to foster a student's abilities to perceive with accuracy and sensitivity within, to imagine possibilities about, to inquire into, and to make discoveries in, an arbitrarily selected specific subject matter. But is there some single epistemic aim – the production of true, rational, or justified belief – which might unify the items on this list? Goldman has suggested that the aim of education is the production of knowledge, in the sense of true belief. But Goldman's real concern, Schmitt suggests, is with the (epistemic) evaluation of various types of educational regimes and their effects, and so his 'veritistic' outlook is consistent with taking a different line as to education's intrinsic function (see above). And, in fact, Schmitt argues that the organizing epistemic aim of liberal arts education is the production of justified belief. Each of the abilities just mentioned are instrumental to, and in some cases (and in the right combination) constitutive of, justified belief. But not all are a matter of rational or critical thinking – certainly, accurate perception isn't. And while true belief might typically be a product of such abilities, whether it is depends on such things as the particular culture in question and features of the particular subject matter itself — a connection that's too tenuous, and too often broken, for the production of true belief to be the organizing aim of education. The value of these abilities, and the success of a given liberal arts educational regime (*qua* educational regime) which might promote them, is essentially tied to their producing justified belief, not true belief.

This issue closes with a discussion between Sergio Sismondo and James Robert Brown, centering on an issue at the heart of the Science Wars, mentioned above – the proper understanding of how science proceeds and progresses. In this dispute, 'rationalists' – Brown's term for the position he himself defends – hold that is it evidence and reason which drive science. On the other hand, social constructivists, adherents to David Bloor's 'strong programme', many of those engaged in science and technology studies, and so on, hold that it is social and other 'external', non-epistemic factors which are the more significant determinants of what goes on in science.

As Brown notes, disputes about the proper understanding of science go hand-in-hand with a similar meta-level dispute about how best to understand the efforts of those engaged with this issue. According to Sismondo, what makes the war metaphor apt is that much of what is going on in 'the Science Wars' is 'boundary work', in Thomas Gieryn's sense – i.e., an ideologically driven attempt to maintain a given discipline's epistemic authority about some subject matter. Such, Sismondo argues, is how we should see Brown's own recent book, *Who Rules in Science?*: it purports to be a defence of a certain view of science, and a critique of various non-rationalist interpretations of the scientific enterprise; but it is really an attempt to defend the terrain of traditional philosophy of science.

Brown disagrees that this is what he is up to, and he questions the extent to which the idea of its participants doing boundary work captures what is going on in the Science Wars. Brown contends that, whatever other motivations he may have had in writing his book, there are substantive issues on which he and his opponents genuinely disagree, and he intended his book to contribute to their resolution.

Whether it concerns the nature of science, or the nature of the Science Wars, the discussion is sure to continue.

https://doi.org/10.3366/epi.2004.1.3.163 Published online by Cambridge University Press

Taken together, the articles collected here give a real sense of just how numerous and diverse are the questions which fall under the heading of social epistemology, as well as how much very good work is being done in attempting to answer them.

Patrick Rysiew is Assistant Professor of Philosophy at the University of British Columbia. His principal research area is epistemology, including its points of intersection with philosophy of language and of mind. Among his current research interests is the relevance of both social and evolutionary factors to understanding such issues as our capacity for rational belief. His publications include "Goldman's Knowledge in a Social World: Correspondence Truth and the Place of Justification in a Veritistic Social Epistemology" (*Protosociology*, 2003); "Reid and Epistemic Naturalism" (*Philosophical Quarterly*, 2002); "The Context-Sensitivity of Knowledge Attributions" (*Noûs*, 2001); and "Testimony, Simulation, and the Limits of Inductivism" (*Australasian Journal of Philosophy*, 2000).

https://doi.org/10.3366/epi.2004.1.3.163 Published online by Cambridge University Press