Introduction: Creativity, Conservatism & the Social Epistemology of Science

Adrian Currie

Penultimate version, forthcoming in Studies in History & Philosophy of Science A.

Abstract

The special issue Creativity, Conservatism & the Social Epistemology of Science collects six papers which, in different ways, tackle 'promotion questions' concerning scientific communities: which features shape those communities, and which might be changed to promote the kinds of epistemic features we desire. In this introduction, I connect these discussions with more traditional debate in the philosophy of science and reflect upon the notions of creativity which underwrite the papers.

The best science is both well-grounded and ground-breaking. We care about epistemic values like trustworthiness, confirmation and systematicity—factors which ground science—but also care for the generation of novelty, creativity and radical inventiveness; ground-breakers. The relationship between these two sets of virtues, the conservative and the creative, is rich grounds for philosophical investigation: hence this special issue. Such a focus allows us to bring new philosophical methods and perspectives into dialogue with novel questions and older concerns. In this introduction, I'll start with a quick discussion of more traditional debates in the philosophy of science, before explaining the relevance of both social epistemology and the formal approaches to it which the papers collected here discuss. After a tour of the papers, I'll close by reflecting on the notions of creativity implied by and underwriting these discussions.

I'll play fast and loose with notions like 'novelty', 'conservatism' and 'creativity': these are polysemous, complex notions acting in differing roles across differing contexts. Although there are potentially relevant discussions of creativity, conservatism and novelty in virtue epistemology (Kieran forthcoming), aesthetics (Gout 2010), and at the intersection of philosophy and creativity research (Klausen 2010), I'll focus on the philosophy of science.

Philosophical discussion of scientific novelty is often concerned with confirmation. Is a theory's capacity to predict new phenomena more important than its capacity to accommodate existing data (e.g., Musgrave 1974, Mayo 1991)? Here, 'novelty' relates to evidence, conceived of as a relationship between some scientific representation and the world (or our descriptions of it). But novelty need not be connected with theory-confirmation: new ideas, new phenomena, new technologies, approaches and techniques have all both been aims in science, and provide fodder for philosophical, historical and sociological exploration of it[1]. Moreover, science is about so much more than testing theories: it is a complex, all-too-human, pluralistic set of practices. Instead of considering novelty in

---

[1] See, for instance, Franklin (2005) on exploratory experiments, and Currie (2017) on surprise.

evidential terms, then, we seek a conception that is broader yet amenable to informing and understanding scientific practices. As we'll see, social epistemology provides such a perspective.

Tensions between the creative and the conservative clash between those two poles of 20[th] Century philosophy of science: Thomas Kuhn and Karl Popper[2]. For the former, the majority of scientists are typically in the business of picking at the coalface of normal science; applying a set of pre-existing tools, perspectives and research questions to the world. The latter takes a more heroic view: good scientists generate bold hypotheses and steadfastly test them. Popper indeed complained that Kuhn's picture of science encouraged conservative drudgery. As Popper says, "In my view the "normal" scientist, as Kuhn describes him, is a person one ought to be sorry for... He has been taught in a dogmatic spirit: he is a victim of indoctrination" (Popper 1970, quoted in Rowbottom 2011, 119). For Popper, the normal scientist has failed to develop the critical attitude which he takes to be central to good scientific practice. Kuhn responds: "[Popper and company] argue that the scientist should try at all times to be a critic and a proliferator of alternate theories. I urge the desirability of an alternate strategy which reserves such behaviour for special occasion" (Kuhn 1970 quoted in Rowbottom 2011, 119).

Kuhn's pluralistic response and Popper's occasional approving appeals to 'dogmatism' suggest this dispute is founded on a false dichotomy: and again a social perspective makes this explicit. If science is a social process, we might not need heroic individual scientists for creative results; and if science is pluralistic we need not choose between Kuhn and Popper, but find a beneficial balance between them. As Darrell Rowbottom points out, there is no reason to think that all scientists must adopt Popperian critical attitudes, or all scientists should be puzzle-solvers. In different contexts, different scientists might adopt different strategies; at the social level, a mixture of approaches could be desirable.

Social epistemology, then, reminds us that knowledge-production is a group-level activity. Scientific knowledge is not simply a matter of putting hypotheses to the test, but also of communicating results in pedagogical, professional and public settings; science involves coordinating differently skilled groups towards common epistemic aims. As such, the organization of scientific communities as well as the infrastructure supporting them, and the connections between science and the wider world, are critical components for understanding how science generates knowledge, and its nature. An emerging approach in social epistemology makes use of formal modelling to represent such communities and explore how they operate (see, for instance, the papers collected in Boyer-Kassem, Mayo-Wilson & Weisberg 2018). Much of this work is founded upon just the tension between the conservative and the creative we have been discussing, and the papers collected here each draw or build upon that tradition.

A striking aspect of formal social epistemology is pluralism regarding the use of abstract tools, particularly agent-based models. These include the use of landscape and evolutionary models adopted from biology (e.g., Weisberg & Muldoon 2009, Thoma 2015, Currie & Avin forthcoming, O'Connor forthcoming); two-arm bandit models from probability theory and reinforcement learning (Zolman 2010); rational decision (or marginal utility) modelling from economics (Kitcher 1990, Strevens 1993) and abstract argument models from artificial intelligence (Borg et al 2017). These models are extracted from their original context and re-construed in terms of scientific communities. For instance, in evolutionary landscape models locations in the landscape are traditionally read as

[2] Similar tensions are also explored in the sociology of science. For instance, Bourdieu (1975) conceives of science as a struggle between homogenising *conservation strategies* and heterogeising *subversion* strategies.

genotypes, and landscape height as fitness. Social epistemologists of science have reconceived them: locations are research questions, and height is the significance of results.

Much recent work in formal social epistemology has tackled what I'll call the *value* question: which properties of epistemic communities are desirable? Most saliently, is it beneficial to have a mixture of scientists adopting conservative and risky strategies and, if so, under what conditions? Most of the papers herein target a slightly different question, albeit one presaged by the value question. Remco Heesen points out in his contribution that most social epistemologists of science "…say very little about how a socially beneficial distribution of researcher types, once determined, should be realized". We more-or-less assume that creativity in some forms plays an important role in science—as Shahar Avin puts it, that "scientific novelty, even in its more extreme "maverick" variety, is an important collective epistemic good" (Avin, forthcoming). But this still leaves open what I'll call the *promotion* question: what aspects of how scientific communities are structured and incentivized encourage various features? Does, for instance, peer review make it more difficult to do revolutionary science? And, if it does, what levers do we have—what possible interventions might we make—to ensure that science's ecosystem becomes friendlier to such work? The notion of a *well-adapted* science clarifies the relationship between the two questions: "A research program is well-adapted when the standards, incentives and expectations governing investigations are geared towards overcoming the challenges of the relevant epistemic situation" (Currie, this issue). To understand whether a research program is well-adapted we need to answer both the value question and the promotion question. The former tells us what kinds of features we want our communities to have, the latter tells us how to get them.

Answering the promotion and value questions, and bringing them together in well-adapted sciences, is a big, complex ask and, as such, it is best tackled from a variety of approaches. This methodological pluralism is reflected in the six papers here.

Finnur Dellsén engages with both the value and promotion questions. He tackles Kyle Stanford's argument that realism and anti-realism make a difference due to influencing attitudes towards revolutionary science (2006). Stanford claims that if realists think our best science is likely true, then they should be happy with relatively conservative science; if anti-realists think our best science is likely false, they should encourage a more diverse, creative science. Dellsén's argument, in effect, claims that Stanford has overly simplified the value realists will place on creative, diverse science

> … while realists will place a lower probability on a search for radically distinct alternatives being successful, realists will also place a greater value [than anti-realists] on both possible outcomes of such a search (success and failure). (Dellsén forthcoming)

This is because while realists think it unlikely that radical alternatives will be true, they will think it more important to demonstrate the falseness of those radical alternatives (as they more value confirming existing theories) and, if the radical theory proves successful, they will more likely think the new theory is successful due to its truth—and so this result will be valued more as well. Dellsén's argument turns on Stanford being right about the differing probabilities realists and anti-realists will assign to the success or failure of radical theories, but wrong about the kinds of value the two camps will assign to those outcomes. Splitting the probability of outcomes from their value is a mainstay of decision theory, which Dellsén puts to good use.

Also drawing on decision theory, Remco Heesen critiques how credit incentives are attributed the capacity to generate optimal distributions of research strategies within scientific populations. Heesen points out that incentives can only change a population's makeup if there are sufficient

members of that population with the required dispositions to take up the roles incentivized. If, as a matter of fact, only a small set of the population has the capacity to adopt a risky strategy, then incentives alone cannot produce riskier behaviours. Heesen makes explicit and explores a hitherto unacknowledged (or under-acknowledged) assumption or idealization in models of scientific communities: that the personal attributes and dispositions of agents do not ultimately affect the makeup of scientific communities.

To make his argument, Heesen considers a simplified model aimed at capturing a tradeoff between a discovery's impact (hence its credit), the probability of success, and the speed of research. Heesen shows how in his model it is the predisposition of the scientist—whether they prefer safety or impact—which decides the optimal strategy, not the amount of credit provided for that impact. The upshot is that setting credit incentives is not necessarily sufficient to shape a community: individuals with the right dispositions are also required. This focus on fairly stable behavioural dispositions undermines 'passive' views of the role of credit. As Heesen says,

> Where previous work has argued that the credit economy operates like an invisible hand, using individual self-interest to motivate socially beneficial decisions, I have argued here that this mechanism should not be expected to have particularly beneficial effects in the context of the distribution of researcher types. This is because there is no reason to expect the distribution of predispositions in a scientific community to match the optimal (or a particularly good) distribution of impact-seeking versus safety-seeking.

This doesn't, however, rule out more active approaches to the credit economy. In particular, it doesn't rule out intervening on credit incentives to stop credit from aligning with scientific importance. That is, Heesen suggests we might explore "deviating from credit proportional to impact to get the right level of impact-seeking." Perhaps, then, a less *laissez faire* approach to scientific incentives is called for.

Cailin O'Connor uses an evolutionary model to explore the relationship between the inheritance of research strategies between labs (being more or less radical in one's approach, for instance), and the level of competition in science. O'Connor argues that the assumption that high competition encourages conservatism is mistaken (e.g., Stanford 2015, Currie forthcoming). However, the differential heritability of creative or conservative approaches to science might undermine the capacity of creative scientists to form flourishing research lineages.

> Successful risky science is often hard to repeat. While more conservative scientists will be able to train scientists capable of continuing their successful projects, and so create thriving lineages, successful risky science may not be the kind of thing one can easily pass on. (O'Connor forthcoming)

If O'Connor is right, it might be heritability rather than competition which sometimes pushes scientific conservatism. If so, I suspect competitive mechanisms might still be diversity-reducing in other ways. O'Connor's evolutionary mechanism works by, in effect, selecting those labs which have the highest combination of success and anti-conservatism—creating a kind of elite. Typically actual elites are gathered from rather small, highly biased subsections of populations, and there's reason to think this comes with epistemic costs (O'Conner & Bruner 2017), not to mention moral or political costs: presumably we want a just science too.

A perennial question regarding the models favoured by formal social epistemologists concerns their idealization: given the enormous distance between the complexity of real-world scientific systems and the simplicity of our models, how are we supposed to take guidance from them? The promotion

question makes this especially pressing. Three papers in the special issue take separate approaches on this issue.

Audrey Harnagel argues that peer-review funding practices "can hinder the ability of the scientific community to successfully explore the landscape and generate significant scientific knowledge" (Harnagel forthcoming). She extends previous landscape models (particularly Avin 2017) by incorporating empirical information about actual citations to construct the shape of the landscape. Harnagel showcases what she calls 'mid-level models' which shift from an explanatory or exploratory mode to a predictive one by combining models with empirical data. Incorporating data and thus "providing predictive information about real-world communities, thus acting as a tool for informing policy decisions… formulate[s] a methodology to move from how-possibly towards more predictive models". Her results suggest that reduced reliance on peer review—and turning instead to lotteries—is likely to promote more creative scientific communities.

Sticking with funding via peer-review, Shahar Avin synthesizes a complex set of arguments chasing similar conclusions as Harnagel, and compares these to the few recent real-world examples available. He thus provides both a history and an integrative argument drawn from the increasing number of diverse voices who are both worried about the cost and effect of peer-reviewed funding, and optimistic about the possibility of the role of lotteries and other chance-introducing factors. Avin notes an important discrepancy between these calls and actual introductions of lotteries to funding:

> One clear difference between the implemented policies and the arguments in the literature is that the policies only apply for a small subset of R&D, namely transformative research; this contrasts with the arguments which all see themselves as applying to all R&D, or at least basic research (Avin, forthcoming).

Although it is too early to tell whether these implementations have the desired effect, the relationship between peer review and scientific creativity is an important locus for consideration of the promotion question. In contrast to the arguments which Avin summarizes, I've attempted to tackle the promotion question in a more context-sensitive way (Currie, forthcoming). I table a variety of what I think are conservatism-boosting features of scientific practice and argue that they make some research programs—those targeting existential risks, for instance—difficult. On this view, answering the promotion question should involve "a thick description of what I call a scientific endeavour's *epistemic situation*: roughly the conditions of knowledge generation and the challenges facing it."

It may be that a good answer to the promotion question looks like a toolbox, or a series of levers, which can influence scientific communities towards various kinds of strategies, organizational features, and constitutions. And it might turn out that what the best tools or arrangement of levers are turns crucially on both the kind of epistemic task we're faced with, and with what we care about producing.

With the exception of my own contribution, the papers do not directly worry about what 'novelty' or 'creativity' in science might be exactly, but reflecting on the papers a cluster of related ideas emerge. There are three components worth mentioning.

First, we see a decoupling between the creativity of an individual and that of a group. Consider O'Connor:

> Even if competitive environments promote conservative decision-making by individual scientists, they seem to have the opposite effect on the selection of conservatism. (O'Connor forthcoming).

Splitting individual behaviours and motivations from those of the groups they form is a mainstay of social epistemology, and creativity is no different. This, I think, puts pressure on the insistence (best articulated by Berys Gaut) that creativity be considered an agential property. That argument relies on intuitions regarding what is required to distinguish true creativity from mere random searching and non-purposive action:

> The kinds of actions which are creative are ones that exhibit at least a relevant purpose (in not being purely accidental), some degree of understanding (not using merely mechanical search procedures), a degree of judgment (in how to apply a rule, if a rule is involved) and an evaluative ability directed to the task at hand. (Gaut 2010, 1040).

It is an open question whether Gaut's intuitions and the account implied here can be reconciled, either by going pluralist or finding non-agential versions of the conditions he highlights.

Second, the relationship between creativity and exploration is critical. Consider Harnagel on peer review:

"[the] bias of peer review… can hinder the ability of the scientific community to successfully explore the landscape and generate significant scientific knowledge" (Harnagel, forthcoming).

The notion that creativity ought to be understood in terms of a search-space or landscape has most affinity with Margaret Boden's 'computational' notion of creativity, in particular what she calls 'exploratory' creativity. Boden introduces the notion of 'conceptual spaces' which are "structural styles of thought" within which new thoughts might be had. Scientists (or philosophers for that matter) operate within not just styles of thought but by using particular tools, questions, technologies and so forth. Regardless, there is at the very least a useful metaphor here: understanding scientific creativity as exploring the potential of some space of possibilities.

> Exploring creativity is valuable because it can enable someone to see possibilities they hadn't glimpsed before. They may even start to ask *just what* limits and *just what* potential, this style of thinking has (Boden 2004, 5)

Creativity as implied in the papers collected here diverges somewhat from Boden's notion: as noted, they go beyond styles of thought and conceptual spaces, often considering problem spaces as well, but are also less explicitly psychological. Although human psychology plays an important role in answering the promotion question, the question is fundamentally about how a community ought to be organized to maximize its potential. On the conception implied here, then, creativity is considered a process or—better still—a strategy for exploration (Klausen 2010). Further, psychological factors might be necessary but insufficient to answer questions concerning creativity in scientific communities.

Third, there is a connection between creativity and risk throughout the papers. It is worth noting that the connection between risk and exploratory creativity is contingent on how risk is conceived. The risk is, typically, risk of failure—"the investigation may fail to be published" (O'Connor). Dellsén draws on a similar conception:

> …it is claimed that contemporary scientists are being steered towards 'safe' research that builds upon, and thus assumes, large parts of previous science; accordingly, new scientific

research is becoming less likely to disrupt existing scientific paradigms. (Dellsén forthcoming).

O'Connor's conception of risk is related to whether the science will pay dividends for the scientist, while Dellsén's concerns the risk of going outside well-confirmed background theory. Heesen takes creative, 'maverick', scientists to target high risk, high reward investigations. Here, reward is taken to involve credit, and risk to involve two components:

> First, the project may fail, as most attempts to revolutionize an area of science do. (In contrast, less world-shocking projects often have a greater chance of success.) Second, projects that aim for high impact tend to take more time to complete. (Heesen forthcoming).

Generally speaking, failure and success—and thus risk and safety—can be understood in many ways. We might fail to make an interesting discovery, or get published, or have a career, or whatever. Whichever notion is salient depends upon context and, crucially, perhaps we have some measure of control over that context. It might be that a successful answer to the promotion question would show us how to engender a successful creative science that involves little risk—but a lot of apparent failure.

Bibliography

Papers in the Special Issue:

Avin, Shahar (forthcoming). Mavericks and Lotteries. Studies in the History and Philosophy of Science.

Currie, Adrian (forthcoming). Existential Risk, Creativity & Well-Adapted Science. Studies in the History and Philosophy of Science.

Dellsén, Finnur (forthcoming). Should Scientific Realists Embrace Theoretical Conservatism? Studies in the History and Philosophy of Science.

Harnagel, Audrey (forthcoming). A Mid-Level Approach to Modeling Scientific Communities. Studies in the History and Philosophy of Science.

Heesen, Remco (forthcoming). The Credit Incentive to Be a Maverick. Studies in the History and Philosophy of Science.

O'Connor, Cailin (forthcoming). The Natural Selection of Conservative Science. Studies in the History and Philosophy of Science.

Other Citations:

Avin, S. (forthcoming). Centralized Funding and Epistemic Exploration. The British Journal for the Philosophy of Science.

Boden, M. A. (2004). The creative mind: Myths and mechanisms. Routledge.

Borg, A., Frey, D., Šešelja, D., & Straßer, C. (2017). Examining Network Effects in an Argumentative Agent-Based Model of Scientific Inquiry. In International Workshop on Logic, Rationality and Interaction (pp. 391-406). Springer, Berlin, Heidelberg.

Bourdieu, P. (1975). The specificity of the scientific field and the social conditions of the progress of reason. Information (International Social Science Council), 14(6), 19-47.

Boyer-Kassem, Mayo-Wilson & Weisberg (eds) (2018). Scientific Collaboration and Collective Knowledge. Oxford University Press.

Currie, A., & Avin, S. (forthcoming). Method Pluralism, Method Mismatch & Method Bias. Philosophers' Imprint.

Currie, A. (2017). The argument from surprise. Canadian Journal of Philosophy, 1-23.

Franklin, L.R., (2005), Exploratory Experiments, Philosophy of Science. 72: 888-899

Gaut, B. (2010). The philosophy of creativity. Philosophy Compass. 5(12), 1034-1046.

Kieran, M (forthcoming) Creativity as an Epistemic Virtue. In: Battaly, HD, (ed.) The Routledge Handbook of Virtue Epistemology. Routledge Handbooks in Philosophy . Routledge , London, UK . ISBN 9781138890206.

Kitcher, Philip. (1990).The Division of Cognitive Labor. Journal of Philosophy, 87: 5–22.

Klausen, S. H. (2010). The notion of creativity revisited: A philosophical perspective on creativity research. Creativity Research Journal, 22(4), 347-360.

Kuhn, T. S. (1970a). Logic of discovery or psychology of research? In Criticism and the growth of knowledge. Lakatos & Musgrave (eds) (pp. 1- 23). Cambridge: CUP

Mayo, D. G. (1991). Novel evidence and severe tests. Philosophy of Science, 58(4), 523-552.

Musgrave, A. (1974). Logical versus historical theories of confirmation. The British Journal for the Philosophy of Science, 25(1), 1-23.

O'Connor, C., & Bruner, J. (2017). Dynamics and diversity in epistemic communities. Erkenntnis, 1-19.

Popper, K. R. (1970). Normal science and its dangers. In Criticism and the growth of knowledge. Lakatos & Musgrave (eds) (pp. 51–58). Cambridge: CUP

Rowbottom, D. P. (2011). Kuhn vs. Popper on criticism and dogmatism in science: a resolution at the group level. Studies in History and Philosophy of Science Part A, 42(1), 117-124.

Stanford, P. K. (2006). Exceeding our grasp: Science, history, and the problem of unconceived alternatives. Oxford University Press.

Stanford, P. K. (2015). Unconceived alternatives and conservatism in science: The impact of professionalization, peer-review, and big science. Synthese, 1-18.

Strevens, M. (2003). The role of the priority rule in science. The Journal of philosophy, 100(2), 55-79.

Thoma, J. (2015). The epistemic division of labor revisited. Philosophy of Science, 82(3), 454-472.

Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. Philosophy of science, 76(2), 225-252.

Zollman, K. J. (2010). The epistemic benefit of transient diversity. Erkenntnis, 72(1), 17.