

KUHN'S RISK-SPREADING ARGUMENT AND THE ORGANIZATION OF SCIENTIFIC COMMUNITIES

One of Thomas Kuhn's profoundest arguments (alas, sadly neglected) is introduced in the 1970 "Postscript" to *The Structure of Scientific Revolutions* (Kuhn 1970). Kuhn is discussing the idea of a "disciplinary matrix" as a more adequate articulation of the "paradigm" notion he'd introduced in the first, 1962, edition of his famous work (Kuhn 1962). He notes that one "element" of disciplinary matrices is likely to be common to most or even all such matrices, unlike the other elements which serve to distinguish specific disciplines and sub-disciplines from one another. This is the element which he calls "values", which, as he notes (1970, 184), being common to a number of otherwise distinct disciplinary matrices, "do much to provide a sense of community to natural scientists as a whole". On the other hand, they also do much, and crucially in Kuhn's view, to promote and sustain a healthy diversity among the practitioners who share any specific disciplinary matrix. In particular, Kuhn claims (1970, 186) that "individual variability in the application of shared values may serve functions essential to science." This crucial point is worth unpacking. And it is, likewise, worth quoting Kuhn at length in order to unpack it. Here is the first of two crucial passages (Kuhn 1970, 186):

To a greater extent than other sorts of components of the disciplinary matrix, values may be shared by men who differ in their application. Judgments of accuracy are relatively, though not entirely, stable from one time to another and from one member to another in a particular group. But judgments of simplicity, consistency, plausibility, and so on often vary greatly from individual to individual. ... Even more important, in those situations where values must be applied, different values, taken alone, would often dictate different choices. One theory may be more accurate but less consistent or plausible than

another ... In short, though values are widely shared by scientists and though commitment to them is both deep and constitutive of science, the application of values is sometimes considerably affected by the features of individual personality and biography that differentiate the members of the group.

Of course, Kuhn is aware that these remarks, especially the last sentence, will be read in an unsympathetic way. He says (1970, 186): "I am occasionally accused of glorifying subjectivity and even irrationality", and, indeed, he was frequently accused of this in the circles I inhabited around the time of his writing – i.e. L.S.E. in the Lakatos era (Cp. Lakatos & Musgrave 1970 for some accusations of irrationalism). So he supplements what he has said already with the following remarks (1970, 186), which constitute the heart of what later came to be called "the risk-spreading argument", but which were not enough to deflect the charges of irrationalism.

But that reaction [that he, Kuhn, "glorifies" irrationality] ignores two characteristics displayed by value judgments in any field. First, shared values can be important determinants of group behaviour even though the members of the group do not all apply them in the same way. ... Second, individual variability in the application of shared values may serve functions essential to science. The points at which values must be applied are invariably also those at which risks must be taken. Most anomalies are resolved by normal means; most proposals for new theories do prove to be wrong. If all members of a community responded to each anomaly as a source of crisis or embraced each new theory advanced by a colleague, science would cease. If, on the other hand, no one reacted to anomalies or brand-new theories in high-risk

ways, there would be few or no revolutions. In matters like these the resort to shared values rather than to shared rules governing individual choice [which, Kuhn implies, dictate the same response for each individual subject to them] may be the community's way of distributing risk and assuring the long-term success of its enterprise.

As I've already said, this argument has been sadly neglected. (See however Kitcher 1990, Hoyningen-Huene 1993, D'Agostino 1993, and Rueger 1996). It is also incompletely articulated. Before moving on to some analogues to Kuhn's argument, I'd therefore first like to say more about what he's got in mind and perhaps even to amplify his argument.

It will be easiest, I think, to work with Kuhn's own contrast, specifically between a rules-driven and a values-driven assessment of and commitment to paradigm work in the sciences specifically, but, really and as Kuhn's own argument makes clear, in many different kinds of communities of enquiry (Cp. Wenger 1998).

Suppose that we have a collection of individuals, each engaged in work within a particular "disciplinary matrix". Suppose that they are producing and assessing variants of some paradigm achievement in that discipline. While there may be some differences among the variants that these individuals produce, there cannot, on the rules model, be any variation in the ways in which they assess these variants. (This may not be entirely correct; see below for a clarification.) So if A produces the variant α and B the variant β , then, if there are rules of assessment R to whose use both A and B are committed, then, short of there being a "tie" between α and β , either both must accept that α is better than β , or both must accept that β is better than α . And, in this case, if both are rational, they both must devote their future energies to the articulation and improvement of whichever of the variants is, according to R, the better of these two.

What's the matter with this? you may ask. It is, as Kuhn saw, risky. Why? Well, as I've tried to indicate elsewhere (D'Agostino 2000), there are two courses of action, leading to an improvement of variants, that are not available to A and B if, because they have arrived at the same conclusion

about the relative merits of α and β , they both prefer the same one of these variants.

First of all, they can't improve the better of the two variants through a *competition* between it and the worse of the two since both A and B will be working on the better variant (if they are rational).

Secondly, they can't *dual-track* the two variants with a possible longer-term reversal of their original judgment – i.e. the originally inferior variant being improved to such a degree that it is now superior (according to R). They can't do this, in particular, since both A and B must, if they are rational, work on the better of the two variants and, hence, neither can work to improve the worse of these two.

In each case, then, A and B "invest" all their community's resources of time, energy, and attention in the superior variant and hence lose opportunities to improve both it and its rival that they would have had if they'd been able (rationally) to work on both.

And how would a values, rather than a rule, orientation assist? Kuhn (1977, 324) puts the matter very clearly.

When scientists must choose between competing theories, two men fully committed to the same list of criteria for theory choice [i.e. to the same values] may nevertheless reach different conclusions. Perhaps they interpret simplicity differently or have different convictions about the range of fields within which the consistency criterion must be met. Or perhaps they agree about these matters but differ about the relative weights to be accorded these or other criteria when several are deployed together.

Actually, there are, according to Kuhn, two bases on which a values approach (as opposed to a rules approach) promotes diversity in judgments ... that, in turn, spreads risk and permits progressive courses of development that might not be facilitated by a rules approach.

First of all, values are not "self-interpreting" or "self-applying". What A and B share, if they share a commitment to the importance of simplicity in assessing rivals in some disciplinary matrix, is not so specific and so determinative that each

must, on pain of irrationality, interpret this value or criterion in exactly the same way in every situation.¹ To be sure, there may be (must be?) situations where both interpret the value and apply it in the same way (Cp. Dworkin 1986, p. 62 for an argument to this effect in the case of legal interpretation). But there may be situations where they don't and needn't. So A and B can look at the variants α and β in terms of their relative simplicity and, at least sometimes, A and B can reach different conclusions, consistently with both valuing "simplicity", about which is the simpler and, hence, the preferable variant (at least with respect to this criterion). And if they do reach different conclusions, then A might work on one variant and B on the other and, hence, they might, collectively, get the benefits of "dual-tracking" these variants – e.g. the benefits of competition that I outlined earlier.

Secondly, however, and as Kuhn clearly implies, the values that participants use to assess variants are plural and not always perfectly "aligned" in terms of their implications for those judgments of overall superiority on which a commitment is based. Let me explain.

As Kuhn points out, scientists have a number of values or criteria in mind when it comes to theory-choice. He lists simplicity, consistency, accuracy, and plausibility. As he also clearly recognizes, how a variant ranks with respect to one of these values (and relative to another variant) need not track its relative ranking with respect to others of these values. For instance, α might be more accurate but less simple than β . But, in this case, even if they agreed about all this (and, of course, they needn't, as I've already indicated), A and B might reach different conclusions about the overall merit of the two variants. Perhaps A thinks that accuracy is more important than simplicity and B thinks the opposite. In this case, even if A and B agree that α is more accurate but less simple than β , A might prefer α and B might prefer β . As Kuhn says, "they agree about these matters but differ about the relative weights to be accorded these ... criteria when [they] are deployed together". And, again, we get risk-spreading diversity of judgment and, hence, the potential for the two improvement cycles which I've sketched above.

Let me summarise. Thomas Kuhn developed a "risk-spreading argument" which showed, in

effect, how scientific "communities of practice" could achieve both solidarity among their members and diversity in the activities of these individuals. In particular, he demonstrated the value of this diversity to the enterprise in which this community was engaged. Because individuals can be united in their commitment to a value and yet divided in the ways in which they interpret and balance it against other values they are also commonly committed to, the community is able to develop variants of its paradigm achievement across a wider front than it would have access to if each individual were bound, by the canons of rationality, to reach the same conclusion about these variants as every other "paid up" member of the community was bound to reach (Cp. Barnes 2001, 20). Diversity ensures risk-spreading and risk-spreading permits the community to explore its domain of enquiry in an efficient and effective manner.

All this makes the scientific community sound a bit like "the market", at least as it is portrayed by some of its theorists and advocates. Allan Walstad, for instance, put it this way (2002, 5):

Unlike methodology, which seeks to prescribe the correct judgment, the market takes advantage of differing judgments. People act on the basis of their individual judgments. Different judgments lead to different choices. Diverse options are explored, and the results can be compared.

There is, nevertheless, a notable tension between this model of diversity-driven competition and the model of scientific consensus that is much-beloved of scientists, philosophers of science, and, perhaps especially, by "end-users" of their investigations. Surely, science is differentiated, say, from literary criticism (or even the market) precisely because or to the extent that it arrives, however circuitous the route, at consensus on "what the facts are". How or to what extent is this compatible with the idea that "individual variability in the application of shared values may serve functions essential to science"? There is, in fact, no particular difficulty, though there is an aspect of the solution to this conundrum that needs to be highlighted. Let me explain.

How could individuals, interpreting and

balancing values in different ways, come to the same conclusion about which of the available variants was the better or best? They could do so in the case of *dominance*, certainly—i.e. where there is a variant that is so much better than the others in certain respects that, no matter how (within limits) you interpret the values and no matter how you weight them relative to one another, this variant is the better or best overall. (This case is emphasized in Laudan & Laudan 1989). When this happens, scientific consensus is achieved despite the fact that individual scientists have indeed interpreted and balanced the values differently.

The crucial thing to notice about this situation is that, when we have dominance, we have agreement on which variant is better while, at the same time, we have the potential for risk-spreading diversity of assessments *in the future*. After all, the consensus does *not* occur because the individual scientists agree about how to interpret and balance the values they use to assess variants. Even as they agree that *this* variant is best (relative to whatever that evaluation is relative to), they disagree about how simplicity should be interpreted or about how consistency and accuracy should be traded-off or balanced against one another. It's just that all that disagreement is irrelevant in this particular case. (In the next section I will return to this rather cavalier claim of irrelevance).

So, while we have a consensus on the assessment of variants, giving us whatever "pay-offs" are supposed to be associated with this and, certainly, giving us a sense of "solidarity" within the scientific community, we also have a *residual divergence* in values that helps support that diversity in assessments that is, in turn, necessary to spread risk and, hence, to facilitate the efficient and effective exploration of a paradigm's potentialities.

Diversity and solidarity are reconciled in the case of dominance. We can agree while disagreeing. While it would be interesting to identify other, less restrictive cases where diversity and solidarity are reconciled (perhaps in some looser way than is possible with dominance), that is not my primary task here and it would be, in any event, beyond my powers. (Cp. however D'Agostino 1996, §34 where I try to show how

the dominance of a particular articulation of the paradigm can be recognized, progressively, across a broader and broader range of community members). What I *do* want to do is two-fold.

First of all, I'd like to sketch some other ways in which residual divergence can be promoted within communities of enquiry.

Secondly, I will address the issues raised by List and Pettit's recent work (2002) on the so-called "discursive paradox". The considerations I develop suggest a way of dealing with the conundrum they identify.

Extending the Risk-Spreading Argument

In extending the risk-spreading argument from values to "the concrete", as I propose to do, I am only, in effect, following Kuhn's own thinking. For, after all, he was at pains, though sometimes confusingly, to assert (1970, 10) that a paradigm, in his specific sense(s), is a "concrete scientific achievement".

Let me begin by saying that, as many others have argued and emphasised, the concrete – that which is embodied, "real" in that specific sense – is, in Stuart Hampshire's phrase (1983, 106), inexhaustibly describable; there is no a priori limit to what, even keeping strictly to the truth, we can say about it. Something concrete has aspects or characteristics and, because it also has parts and relations to other objects, and because these parts and related objects are themselves concrete, it has inexhaustibly many characteristics or aspects. There is more to be perceived, comprehended, and said about any concrete object than we can, within the limits of our own finitude, perceive, comprehend and say. (This is a source, I believe, the "frame problem", as it's called in Artificial Intelligence, i.e. the problem of specifying, in advance, what is and what isn't relevant collateral information that bears on a particular problem. See, for instance, Pylyshyn 1987. Because of inexhaustibility, there is simply too much that *might* be relevant.)

Why does inexhaustibility matter? Well, if there is more to say than we can say, then what we *do* say will necessarily be selective. And if our descriptions or even silent musings on a concrete achievement are selective, then, in principle, what one person says or thinks about it can differ

from what other people say or think about that very same achievement. There can, in other words and in a Kuhnian patois, be diversity in the understanding of concrete achievements, including those that constitute paradigms for scientific activity. Kuhn is quite explicit about this. Cp. Kuhn 1970, 44:

Scientists can agree that a Newton, Lavoisier, Maxwell, or Einstein has produced an apparently permanent solution to a group of outstanding problems and still disagree, sometimes without being aware of it, about the particular abstract characteristics that make those solutions permanent. They can, that is, agree in their identification of a paradigm without agreeing on, or even attempting to produce, a full interpretation or rationalization of it.

In any event, this diversity, like that in the application of shared values, may serve functions essential to science. How could this work?

Two enquirers, A and B, agree that the concrete achievement Π constitutes a model for their own activities. Because Π is concrete, it has inexhaustibly many features and anyone's understanding of it will therefore be selective. Hence, when A *interprets* Π , there is some prospect that he will be selective in identifying its exemplary characteristics in a different way than B will be when she interprets it. There will, potentially and quite frequently, be two different selective "images" of Π , Π_A and Π_B . Hence, while A and B remain united in their belief that Π provides a model for their own activities, they, as with the case of values, are in a position to develop this paradigmatic achievement in different ways and, thus, to spread the risk of developing the paradigm for the community to which they belong. Let me explain.

Suppose, for instance, that A's interpretation of Π is based on a selection of Π 's characteristics that includes α and β and that B's interpretation is based on the characteristics β and δ . (I've shown A and B as sharing an element of their interpretations in order to make explicit what should be obvious: their interpretations can't differ "too much" without difficulties for the dynamics of the community to which they both belong.) In this case, when it comes to exploring the relevance of

– to some problem on which both A and B are working, where A looks at the potential relevance of α to this problem, B, on the contrary, look at the potential relevance of δ . They thus explore a wider range of possibilities than they would, collectively, if both had made the same selection from the inexhaustibly many potentially relevant characteristics of the paradigm which both aim to emulate, or, to put it another way, if they hadn't had to *interpret* this concrete achievement but could just directly apply it.

So, again, as with values, we have diversity supporting a spreading of risk that facilitates a more efficient and effective exploration of the paradigm (and of the world). And, again, we also have solidarity—both A and B are committed to extending the reach of Π , and both of them are committed to doing so subject to the discipline of those shared but diversely interpreted and balanced values whose significance to their enterprise I have already explained.

Notice, furthermore, how values can mediate the progressive development of the shared paradigm. In particular, A and B will compare their two interpretations according to their (differing understandings of their) shared values and, if, say, Π_A proves to be dominant, with respect to these values, then both A and B will accept that interpretation of their shared paradigm as the favoured interpretation—indeed it will *become* the paradigm (for them)—and they will proceed thereafter on that basis—which, they agree, is a better basis for exploring the world. Of course, just as their agreement on this issue doesn't mean their agreement on the interpretation and balancing of shared values, neither does it mean their agreement on the *subsequent* interpretation of Π_A , i.e. the newly-contrived leading edge for their further explorations. For it too is a concrete achievement and hence it too is subject to interpretation—to a selection, and hence variable selections of its features for emulation.

We have, if you like, *a perpetual motion machine of the Kuhnian kind*. We resolve differences (through dominance) but never exhaust the differences that we require to support diversity in exploratory behaviour and, hence, to spread risk (by producing more differences which need to be resolved).

That values need to be interpreted and balanced (and can be differently by different

people) and that concrete paradigms need to be interpreted (and can be differently by different people) are two reasons why risk can be spread in the development of scientific practice. I said, earlier, that Kuhn's original distinction between rules-based and values-based practice was overdrawn, however valuable in leading to his genuine insights. In particular, rules too are not really self-interpreting or self-applying. That this is so shouldn't blind us, however, to an important distinction *within* the realm of rules. This is the distinction between prescriptive and proscriptive rules. (Cp. Hayek 1973-6). Let me explain the significance of this distinction for our concerns.

Suppose that A and B are subject to the prescriptive rule R. It is of the nature of such rules that they prescribe (relatively) specific behaviour – e.g. those subject to them are supposed, in a situation S, to perform the specific action ϕ . Bearing in mind the points I've already made about rules not being self-interpreting, A and B, subject to this rule, must both ϕ in S. This is quite different from the situation where they are subject, rather, to a proscriptive rule, such as R' which forbids the performance of some action, say ψ , in the situation S'. In this latter case, both A and B can be compliant with the rule they are subject to while performing quite different actions in the relevant situation. For it is consistent with neither of them ψ -ing that, say, A ϕ s and B χ s, and since performing these different actions might lead A and B in quite different directions, they will be free to explore more of the relevant territory they occupy in a proscription-based regime than they would be in a prescription-based regime. And this means diversity in the application of shared rules (both obey R') and, hence, risk-spreading with all that it entails for the efficient and effective advancement of their joint cause.²

The Essential Tension

We have rather a complicated picture of the scientific community à la Kuhn. We've in fact identified three distinct supports for the diversity of behaviour that, on Kuhn's account, promotes risk-spreading and, hence, the more efficient and effective exploration of the domain of enquiry. These are: (1) the reliance on values, rather than rules; (2) the importance of concrete and hence

(multiply) interpretable exemplars of achievement; and (3) the use of proscription rather than prescription where rules are indeed a mechanism of solidarity and coordination.³

One of the key ideas, indeed the most important idea that I've introduced is that of "residual divergence". Even when individual members of a community agree about something, they will, because of the mechanisms I've sketched, still have a basis in the future for those disagreements that are so crucial for spreading risk. Their agreement is shallow and their disagreement is deep, we might say. In cases of dominance, A and B agree about the "conclusion" (e.g. that α should be preferred to β) without agreeing, even, indeed, while disagreeing about the "premises" which support this conclusion. Let me explain.

Remember, there are two points which Kuhn made about values.

First of all, they are not "self-interpreting", as I've put it. Hence, it is possible that A might judge that α is simpler than β whereas B judges that β is simpler than α . How, then, could they nevertheless agree that α is, overall, superior to β ? There are, in fact, a number of different possibilities.

Perhaps they agree about the overall superiority of α because, despite their disagreement about simplicity, they agree that α is more predictive than β and because each values predictivity more highly than simplicity. So we might have, schematically, that:

CASE 1

	<i>Simplicity</i>	<i>Predictivity</i>	<i>Weighting</i>	<i>Overall worthiness</i>
A	$\alpha > \beta$	$\alpha > \beta$	P > S	$\alpha > \beta$
B	$\beta > \alpha$	$\alpha > \beta$	P > S	$\alpha > \beta$

Alternatively, perhaps they agree about the overall superiority of α because, despite their disagreement about simplicity, there is a compensating disagreement about (a) the relative predictivity of the two variants and (b) the relative weightiness of the two values in determining overall superiority. So we might have, schematically, that:

CASE 2

	<i>Simplicity</i>	<i>Predictivity</i>	<i>Weighting</i>	<i>Overall worthiness</i>
A	$\alpha > \beta$	$\beta > \alpha$	$S > p$	$\alpha > \beta$
B	$\beta > \alpha$	$\alpha > \beta$	$P > S$	$\alpha > \beta$

In each case, in other words, we have collective agreement on the "conclusions" of individual courses of reasoning which in fact diverge on relevant "premises".

This is an interesting point to arrive at, particularly in view of Alvin Goldman's discussion (2004), in the first issue of this Journal, of the so-called "discursive paradox" articulated by Christian List and Philip Pettit (2002). For it would seem that I am identifying as a virtue what List and Pettit say is a vice. In particular, I say that we reconcile solidarity and diversity, both of which facilitate the proper functioning of a community of enquiry, by fostering shallow agreements (in cases of "dominance") with deeper disagreements (about the bases for these shallow agreements). And List and Pettit say, in effect, that this is vicious. What they seem to require, when individuals reason about some issue to which they will have to commit themselves jointly is that there should be at least as much agreement, if I may put it that way, about the premises as there is about the conclusion. In particular, they identify (2002, 95-6), as a one of "two plausible demands that we might want to make on the aggregation of judgment. ... that in aggregating judgments a group should reach a collective set of judgments that is itself rational." And, although the particularities of Case 2 above differ from those of their own exemplars of collective judgment, Case 2 does not, it seems to me, exhibit a "collective set of judgments that is itself rational" by their standards. Or, in any event, it could easily be converted to a case which clearly does not, simply by adding a third member of the community resulting in the following profile of judgments:

CASE 3

	<i>Simplicity</i>	<i>Predictivity</i>	<i>Weighting</i>	<i>Overall worthiness</i>
A	$\alpha > \beta$	$\beta > \alpha$	$S > p$	$\alpha > \beta$
B	$\beta > \alpha$	$\alpha > \beta$	$P > S$	$\alpha > \beta$
B	$\beta > \alpha$	$\beta > \alpha$	$S \text{ ? } P$	$\beta > \alpha$

For we now seem to have precisely the sort of array of "premises" and "conclusions" which List and Pettit think is "irrational". For each of two bases for comparing alternatives, we have a majority judging β to be superior to α and yet we also have a majority judging α to be superior to β overall. How are we to reconcile this divergence between what the "collective" seems to think about the "premises" feeding into their judgments of overall superiority with what the "collective" seems to think about precisely this question of overall superiority?

As Goldman indicated (2004, 13), we can pose the problem in terms of a choice. Should we look at the premises or at the conclusion? If we look at the premises, we have, in Case 3, a majority in favour of the judgment that β is superior to α in both relevant respects and, hence we might infer that the collective should conclude that β is superior to α overall, since that conclusion follows from the premises that it, the collective, collectively "accepts". If we look at the conclusions, though, we have a majority in favour of the judgment that α is superior to β and, hence might infer that the collective should conclude, despite the confusing situation with the premises, that α is indeed superior to β .

For List and Pettit the choice between premises-driven and conclusion-driven approaches is a tortured one. For me, it is easy. A conclusion-driven approach *preserves residual divergence* and hence supports risk-spreading and, accordingly, facilitates the effective exploration of the domain of enquiry. Let me explain.⁴

The possibility that individuals might converge on the same conclusion from different sets of premises is one that is recognised if we take a conclusion-based approach to the so-called discursive paradox. While such an approach recognises, what the parties themselves recognise, namely that they agree about which concrete option to prefer, it does so without demanding that they should standardise on their reasons for preferring this option. It thus preserves the diversity of grounds for preference which is so important in facilitating that risk-spreading which, in turn, supports the wider *collective* exploration of the domain of enquiry. (And it does so without, as List and Pettit might fear, undermining the solidarity of the community of enquiry. That solidarity is

founded on a shallow agreement about values, rules, and concrete exemplars, it is not *propositionally* grounded, and hence isn't threatened by any supposed incoherence among the specifically propositional commitments of the community's members.)

On the other hand, to demand that, because a majority of participants agrees on argumentative premises, the collective as a whole should agree on those premises (and hence on the conclusion which *they* support) seems unnecessarily Procrustean. Certainly, it will have a tendency to reduce diversity in the interpretation and application of values and will thus work against the project of risk-spreading.

Alvin Goldman (2004, 11) aptly suggests that a proper "social epistemology" should, *inter alia*, "examine social practices in terms of their impact on knowledge acquisition". I claim to have made

a start at that project here. In particular, I have, relentlessly following Kuhn's teachings, identified three specific "social practices" which, I claim, facilitate the (efficient) acquisition of knowledge. These are: the predominance of values over rules as grounds for solidarity in the scientific community; the use of concrete and hence multiply interpretable paradigms as the exemplars which scientists are meant to extend and apply; and, drawing on Hayek, the predominance, where rules do play a role in enquiry, of proscriptive over prescriptive rules. All three "social practices" contribute to "residual divergence", as I have called it – i.e. to the persistence in the face of that shallow agreement that is necessary for group solidarity of that deep disagreement that is necessary if risk is to be spread and the domain of enquiry efficiently explored.

References

- Barnes, B (2001). "Practice as collective action". In T Schatzki, K Knorr Cetina & E von Savigny (eds), *The Practice Turn in Contemporary Theory*. London: Routledge.
- D'Agostino, F (1993). "Demographic" Factors in Revolutionary Science: The Wave Model'. *Methodology and Science*, vol. 26, pp. 41-52.
- D'Agostino, F (1996). *Free Public Reason*. New York: Oxford University Press
- D'Agostino, F (2000). "Incommensurability and Commensuration: Lessons from (and to) Ethico-Political Theory". *Studies in the History and Philosophy of Science*, vol. 31, no. 3, pp. 429-47.
- D'Agostino, F (2004). 'Liberalism and Pluralism', in G Gaus & C Kukathas (eds), *Handbook of Political Theory*. London: Sage.
- Dworkin, R (1986). *Law's Empire*. London: Fontana Press.
- Goldman, A (2004). "Group Knowledge versus Group Rationality: Two Approaches to Social Epistemology". *EPISTEME*, vol. 1, no. 1, pp. 11-22.
- Hampshire, S (1983). *Morality and Conflict*. Oxford: Basil Blackwell.
- Hayek, FA (1973-6). *Law, Legislation and Liberty*, Chicago: University of Chicago Press
- Hoyningen-Huene, P (1993). *Reconstructing Scientific Revolutions*. Chicago: University of Chicago Press
- Kitcher, P (1990). "The Division of Cognitive Labor". *Journal of Philosophy*, vol. 87, pp. 5-22.
- Kuhn, T (1962). 'The Structure of Scientific Revolutions', in *International Encyclopedia of Unified Science: Foundations of the Unity of Science*. Chicago: University of Chicago Press, vol. 2, no. 2.
- Kuhn, T (1970). *The Structure of Scientific Revolutions*, 2nd edn. Chicago: University of Chicago Press.
- Kuhn, T (1977). *The Essential Tension*. Chicago: University of Chicago Press.
- Lakatos, I (1970). "Falsification and the Methodology of Scientific Research Programmes", in I Lakatos & A Musgrave (eds), *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.
- Lakatos, I & Musgrave, A (eds.) (1970). *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.

- Laudan, R & Laudan, L (1989). "Dominance and the Disunity of Method." *Philosophy of Science*, vol. 56, pp. 221-33.
- List, C & Pettit, P (2002). "Aggregating Sets of Judgments: An Impossibility Result". *Economics and Philosophy*, vol. 18, no. 1, pp. 89-110.
- Pylyshyn, Z (ed.) (1987). *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Norwood, NJ: Ablex Publishing Corporation.
- Rueger, A (1996). "Risk and Diversification in Theory Choice". *Synthese*, vol. 109, pp. 263-80.
- Walstad, A (2002). "Science as a Market Process". *Independent Review*, vol. 7, no. 1, p. 5ff.
- Wenger, E (1998). *Communities of Practice: Learning, Meaning and Identity*, Cambridge: Cambridge University Press.

Notes

- 1 Since this point is true of rules as well as values, the *sharp* contrast between a values-based and a rules-based approach is actually overdrawn. We might view it as some Kuhnian scaffolding, that we can now dispense with, that enabled him to see the point about diversity of interpretations. (This is the qualification whose necessity I flagged above.)
- 2 Perhaps the "hard core" of a Lakatosian "research program" is an example of proscription in science. As long as each scientist preserves the hard-core assumptions (she is proscribed from doing otherwise), then she is free to explore the domain of enquiry in any way she pleases. (Or not. Certainly, Lakatos's ideas about the "heuristic" of the research program seem more prescriptive than proscriptive.) See for instance Lakatos 1970.
- 3 David Resnik (1996) certainly provides a persuasive argument that an ethics of research code might assist in solidifying the scientific community (and, indeed, in securing its functionality in relation to its epistemic aims). Although some of the twelve rules which he nominates are not grammatically proscriptive, all are sufficiently vague, abstract, or ambiguous to allow for multiple interpretation and, hence, for some risk-spreading. Resnik acknowledges both their vagueness (p. 580) and their potential to conflict with one another (p. 581) and, hence, the need to balance their sometimes competing demands. Nevertheless, Resnik seems to miss the opportunity to use these observations as a basis for advocating the kind of pluralism which I recommend. He says, for instance, that (p. 581) "resolving conflicts of rules is a messy business and ... I can offer no overall procedure for adjudicating conflict", and thus forgoes saying, as I would, that conflicts of rules is a basis for the diversification of judgments and hence for the spreading of risk.
- 4 If there is some "collective irrationality" about this approach, so be it. (The appearance of "collective irrationality" is, I believe, entirely an artefact of the way List and Pettit set up the problem and entirely disappears once we recognize the demands of a "deep pluralism" about values. See my 2004.)

Fred D'Agostino is Associate Professor of Philosophy and Program Director for Contemporary Studies at the University of Queensland. His research interests include philosophy of linguistics, philosophy of the social sciences, political philosophy, and scientific method. Fred held the co-editorship of the *Australasian Journal of Philosophy* for several years. His books include *Incommensurability and Commensuration* (2003); co-edited with Jerry Gaus *Free Public Reason* (1996); and *Chomsky's System of Ideas* (1986).