Towards a Political Philosophy of Science.

John Dupré
The Centre for the Study of Life Sciences (Egenis)
University of Exeter

Introduction

There was a time when philosophers of science treated science simply as a means of accumulating truth or at any rate justified belief. Its central question was how it achieved this excellent result, and what differentiated it from more benighted human practices with less respectable claims to knowledge—religion or philosophy, for instance.

Nowadays, though most philosophers still see science as the pre-eminent source of knowledge of the world, enthusiasm is often more nuanced. Scepticism is common concerning the limits of scientific truth—perhaps there are important areas of human concern that science cannot reach, and perhaps there are still important things to be said for other kinds of knowledge or wisdom. A related concern that will be the focus of this paper is whether even if science delivers truth, this is enough to count it as a good thing. This is one of the questions to which Philip Kitcher has turned his attention over the last fifteen years or so, and his path-breaking work on this and related questions have helped to bring political issues concerning science to the forefront of philosophical attention.

A good entry point into the topic is Kitcher's (2001) distinction between truth and significant truth. There is an infinite number of truths: I have x hairs on my head, I have fewer than x +1 hairs on my head, I have fewer than x + 2 hairs on my head… For some x all of these are true, but none is very significant. And of course there are many facts that have much greater significance than this for some—my childrens' birthdays, say, are important facts for me—that are of very little interest to science. There is surely not time enough for science to enumerate all the truths there are, and even if there were there are some that we would very much like to know sooner rather than later. Some of these are among the significant truths. How do we decide which these are, and thus direct our finite scientific effort to finding out the things that we have some reason to want to know?

This question perhaps seemed less pressing when many philosophers believed that science was a unified whole. This belief was often interpreted in terms of what has come to be known as the 'layer-cake' model, in which sciences were ordered in a hierarchy in which physics was the most fundamental, chemistry was to be derived from the laws of physics, biology from the laws of chemistry, and so on through psychology, sociology and so on. The most significant truths were fairly obviously the laws that articulated this hierarchical structure, and the more fundamental, i.e. the nearer to physics, the more significant. Also highly significant were the descriptions, for example of chemical entities in terms of physical constituents and their relations, which made possible the derivations from more to less fundamental levels.

I think this view of a unified science is no longer defensible (Dupré 1993, Cartwright 1999), a position that is also central to Kitcher's recent work. Kitcher belongs to a growing number of philosophers of science who see science as pluralistic, as consisting of particular theories or models designed to address particular classes of question. Scientific models abstract from the complexity of nature and aim to focus on a relatively small set of properties that are more or less decisive in generating a certain kind of phenomenon. Unlike the unified picture, pluralism offers no internal account of which truths are significant; this is something that must be decided by some parallel process. Moreover, for a pluralist there is little reason to suppose there is any limit to the possible topics that might be pursued scienctifically. New interests will make possible new sciences, and there is no reason to anticipate any limit to the interests we might acquire. Indeed, there is little reason to assume that the sciences we have allow only a finite number of truths to be discovered. So deciding what questions to address becomes an unavoidable part of the scientific process, and how these decisions are to be made is something on which philosophers of science have had little to say. To the extent that science is a public enterprise, funded by the public for the general good, this must be a political problem. Hence, Kitcher argues, we need a political philosophy of science.

Well-Ordered Science

Well-ordered science is the concept Kitcher (2001) introduced to refer to the way that institutions for the collection and storage of knowledge should fit into the organisation of a democratic society. In addition to questions already mentioned about what scientific research should be supported, this includes questions about certification—when a scientific claim should be taken to have been established—and access—who should have access to what parts of the accumulated, certified scientific knowledge within the state. And, especially important, how should scientific knowledge be applied to making decisions of policy?

It is clear that current democratic states have yet to achieve a well-ordered science. At least, it seems evident that neither individuals nor states reliably act in the ways that seem clearly mandated by sound scientific knowledge. Children die from measles because one thoroughly discredited scientific paper has persuaded millions that a safe and effective vaccine causes autism; in the most affluent nation on earth a majority of the population reject overwhelming evidence that we evolved over hundreds of millions of years from simpler organisms through natural processes, and believe that our existence is to some extent the reflection of the intentions of an all-powerful supernatural being.[1]

[1] According to a recent survey by the Pew Research Centre, 60% of Americans believed that humans have evolved over time, as opposed to having existed from the beginning of time in their present state. However, only 32% of Americans believed that this had happened through natural processes rather than as part of God's means to the creation of humans

And it is hard to argue that the allocation of scientific effort is optimal. Vast resources are devoted to amelioration or cure of the diseases that afflict the old in the richest countries, while little research is done on possibly quite simple measures that might massively reduce the devastating impact of infectious diseases, generally on the young, in the developing world. And finally, even where scientific knowledge has been acquired with obvious political implications, it is not always appropriately applied. An overwhelming scientific consensus predicts catastrophic climatic changes resulting from carbon emissions, yet little is done to reverse this disastrous process.

Kicther's *Science, Truth and Democracy* and *Science in a Democratic Society* (2001; 2011a) represent a systematic attempt to address these vital issues. Here I shall focus especially on two of the central issues: how is it to be decided what scientific research should be undertaken, and how should democratic decisions be made about the application of science to public policy?

Ill-Ordered Science: What is to be done?

What scientific research should be undertaken? A view that is probably held by many scientists is that this is something that they are best able to decide, and the ideal situation would be for all scientists to be free to address whatever questions they considered important and perhaps tractable. Whether or not this is right, and I shall say a bit more about it below, it is pretty clearly not very relevant. In a world of finite resources in which much of science is paid for by states, public decisions will need to be made about which projects are funded. How should this be done?

A modified version of the naïve first answer is that the relevant decisions be made by the community of scientists. In fact this seems close to what happens in practice in most democratic states. Scientists send proposals for research projects to funding agencies, and peers evaluate them and decide which should be supported. Unfortunately this just moves the problem up a level, to communities of scientists. If someone submits a proposal for funding on the mating behaviour of the Spangled Drongo, it will be assessed, if not by experts on the Spangled Drongo, at least by experts on bird behaviour or perhaps just animal behaviour. These experts may decide the project is less worthy than one on, say, foraging behaviour among wombats. But they are unlikely to decide that less or no funding should be allocated to any study of the behaviour of wild animals, a decision requiring an implausible kind of professional suicide. Whether limited resources are better devoted to animal behaviour or inorganic chemistry, say, does not seem to be something for which any particular kind of scientist has particular expertise.

Kitcher, in fact, argues that some scientific projects are better not done at all, and certainly should not be publicly funded. The example that he considers in detail in *Science, Truth and Democracy* is that of racial science, the exploration of

---

(http://www.pewforum.org/2013/12/30/publics-views-on-human-evolution/; accessed 17 May, 2014).

differences between people of different races. He discusses *The Bell Curve*, the notorious book by Richard Herrnstein and Charles Murray in which they argue, first, that economic class generally reflects talent: people are poor because they are relatively dumb; and, most notoriously, that the great over-representation among the poor of African-Americans reflects the fact that African-Americans are on average less intelligent than Americans of European descent. Suppose this is true? What would be the benefits of knowing it? As Kitcher argues, there are few obvious benefits and some obvious harms. Centrally, the position of a widely disadvantaged group is likely to be substantially worsened as those who have continued a centuries old tradition of discrimination against African-Americans, for example in employment, feel justified in their discrimination and are (even) less motivated to end the practice. Worsening the position of the already badly off is widely agreed to be a very bad thing.

There will no doubt be many who will respond to this argument by insisting that we should want the truth whether or not it may hurt. And, after all, the discovery claimed by Herrnstein and Murray really may show that alleged discrimination in employment isn't what it seems. Perhaps it is just a reflection of the systematically lower qualifications of African-Americans. Just as women are (appropriately it may well be claimed) under-represented in occupations requiring upper body strength, so African-Americans are appropriately under-represented in jobs that require exceptional intelligence.

At this point we need to be a little more critical of the assumption that the research in question is likely to generate truth. In the first place, research of this kind assumes that there is something being measured, intelligence, that is somehow a purely biological property, independent of upbringing and education. This is highly contentious. If, on the other hand, we recognise that measured intelligence reflects the outcome of a developmental process influenced at least as much by education as by any natural endowment, then we see that the outcomes of intelligence tests are likely to be a symptom of inequality rather than a measure of the cause of inequality. Moreover, this misinterpretation, if taken seriously, is likely to justify the continuation of the unequal treatment that in fact causes the perceived differences. Thus this research may not only be taken to justify unequal treatment, but it may help to perpetuate the phenomena that form the basis of its misguided interpretation.

It is, in addition, highly debatable whether a further premise of the most notorious aspect of the research in question is justified, namely that there are two kinds of people—Americans of African descent, and Americans of European descent—that may turn out to differ systematically in their properties. As is well-known, genetic diversity is much greater within standardly distinguished racial groups than between groups, and given the extent of human interbreeding it is difficult or impossible to provide sharp definitions of who belongs to which group. Hence we find such bizarre conventions as the 'one drop of blood' rule that defined anyone with detectable African ancestry as black. There is at least serious debate about whether racial categories have any ontological validity.

These concerns about the categories in which this research is framed raise a further concern. Quite apart from ethical questions about undertaking the research in question, values are in fact already embedded in these categories. Consider intelligence. Typical intelligence tests address various competencies, for example ability to handle and manipulate numbers, to visualise the relationships between shapes, or to deploy vocabulary effectively. Why these, rather than, say, the ability to recognise the makes of cars or to play fast reaction video games? If such tests are not, as is often said, merely a measure of peoples' abilities to do intelligence tests, it must be supposed that they are correlated with other skills, perhaps those supposed to be useful in succeeding in human life. Such correlations could no doubt be investigated. But what is success here? Clearly at some point this is going to require a normative decision. This is not a problem of objectivity: there could be a perfectly objective measure of the abilities to recognise cars or shoot down virtual vampires. We choose to measure intelligence because we value outcomes that this supposed capacity helps us to achieve.

Similar issues arise even more obviously for racial categories. There is nothing ontologically defective with the category of people with at least one drop of African blood (or, let us say, one ancestor native to Africa within the last four hundred years). But why anyone would be interested in this as a category for scientific research is another matter. Only a normative explanation, whether based on racial hostility or historical reparation, seems possible. At any rate, the fact that social, political or ethical values are embedded in much research from the outset makes even more pressing questions about the desirability of the research.

I don't want to make a detailed argument that research on psychological differences between races is epistemically misguided, though I think it is. The point is rather that the premises that underlie this research are at best controversial, and hence so are the results of the research. The combination of research the outcome of which is likely to be both dubious and harmful provides a paradigm for research it would be better not to do.

I don't assume that both these conditions must be present to make research undesirable. There is a lot of epistemically sound research that should not be carried out for obvious ethical reasons. It is argued that some Nazi research on issues such as hypothermia, while sound in principle, is so morally repugnant in its methods that it has been intensely debated whether it is even morally acceptable to make use of its outcomes. Perhaps exploding atomic bombs in earthquake faults would be a good way of learning about tectonics, but few people would advocate pursuing this line of enquiry.

The preceding remarks illustrate a central theme in Kitcher's work: the traditional idea of science as value free is indefensible. Values are unavoidably implicated not only in decisions about what topics we decide to investigate, but in the concepts in terms of which we formulate the questions we try to answer. These questions, in turn, cannot be answered apart from decisions about the goals we would like our science to serve. This is not, as is still sometimes

supposed, an argument that science is subjective, or the acceptance of its results a matter of taste. It is an argument that we cannot understand science properly without attending to these fundamental normative aspects.[2]

It is, at any rate, uncontroversial that there should be limits on what scientific research should be undertaken at all, let alone publicly funded, and Kitcher's contribution in his discussion of the example of racial difference is valuable especially for exploring the wide range of important respects in which this research may be highly undesirable even if from a scientific point of view, in terms of its likelihood to discover truths, it were perfectly sound.

The question then inevitably arises, how we should decide what research should be undertaken. Kitcher's answer, very roughly speaking, is that such decisions should be made democratically. The democracy he has in mind, however, is not the 'vulgar' democracy of popular referenda but a more Millian conception that recognises the importance and value of expertise. I shall consider some aspects of this solution in the next section. Here I note only that Kitcher does not advocate the simple and obvious solution of banning research that is deemed undesirable, though presumably the processes that he advocates for decision-making would at the least make public funding of, for example, the research just discussed on racial difference very unlikely.

I will confess, in passing, that I am somewhat tempted to a more coercive view. Democratic decision-making, vulgar or sophisticated, is likely mainly to affect questions of public funding. The Millian perspective that Kitcher largely endorses sets a high bar against limiting the intellectual activities of private individuals and this certainly applies to scientific inquiry. However, concerns about harmful science today apply not to private individuals in their garages or back rooms, but to vast corporations sometimes dwarfing nation states in their resources. This is not the place to consider in detail whether such corporations are in fact engaged in research that violates the constraints of well-ordered science. But to the extent that they are, it is hard to see how they could be restrained from doing so other than by banning relevant domains of enquiry. If a well-ordered democratic state would decide that research on racial differences should not be publicly funded, I'm not sure why it should not prevent such research being done outside the confines of state institutions.

Kitcher does have an argument for not banning research, namely that this may well give the impression that it is banned because if it were carried out it would produce the results which, it is feared, will do harm In the present case, banning research on racial difference may lead people to assume that it is banned because it is supposed that it will demonstrate that black people are less intelligent (2001, 105-7). . But though this argument is plausible enough, it is only one consideration among many that need to be balanced. As Kitcher regretfully notes, this argument may well apply as much to removal of public

[2] The role of values in science has been quite widely discussed in recent philosophy of science. See, e.g., Douglas (2009); Kincaid, Dupré, and Wylie (2007).

funding as to an outright ban. My suggestion here is just that a democratic process that effectively, and without overwhelming negative consequences, decided that research of a certain kind should not, for ethical reasons, be publicly supported, might very well have equally good grounds for saying that it should be altogether proscribed.

Democratic Science: What is to be done?

As I mentioned at the outset of this essay, the most important symptom of an ill-ordered science is the failure to employ science to improve individual lives and social policy. Often the relevant boundaries between the individual and the social here are open to debate. In the case of vaccination it is widely understood as a paradigmatic case for individual choice what medical technologies people choose to employ for themselves or their children. Yet vaccination is only the most obvious technology that problematizes such a view. Many people deciding not to treat their children with the MMR vaccine is already causing serious health risks not only to themselves but to others. One solution would be for states to enforce mandatory vaccination. But while this may be justifiable in the end, it would seem much preferable if people were able to understand and respond to sound scientific information so as to take sensible precautionary health measures without coercion. To do this they must have either the ability to assess for themselves the value of scientific research or have a high level of trust in scientific experts. But the first option seems unrealistic in the foreseeable future, and the second seems both frequently absent and anyhow problematic in various ways. This brings us to the heart of the problem Kitcher addresses in most detail in the second of the books under discussion here, *Science in a Democratic Society*. What is the proper relation between democracy and expertise?

Kitcher approaches the problem through what he calls 'the division of cognitive labour'. Various people, including Immanuel Kant, have been described as the last person to know everything worth knowing, but there is no doubt that that is a feat far beyond the reach of anyone currently living. PubMed, an index of biomedical publications contains almost 24 million citations at the time of writing, and a new one is added about every minute. Even if 95% of these have nothing very interesting to say, this still leaves a million or so worth reading, and this is just one major area of scientific knowledge. Hundreds of other databases can be found listing tens or hundreds of thousands of resources on topics from Japanese history to gardening, from astrophysics to philosophy (PhilPapers now lists over one million books and articles). Of course only a fraction of all this should probably count as 'worth knowing', but then the problem is to find out what fraction. The only solution to this problem is a division of cognitive labour: for many different areas of knowledge there are some people who know a good deal about what is known or credibly believed; these are the experts. If we need to know something about an area on which we are not an expert, rather than dive into this ocean of more or less reliable verbiage, we find an expert.

So far, so good. Now return to the central questions for political philosophy of science: how do we decide what science should be done, or at least funded, and

how do we apply science to real practical problems? Two problems arise. First, there is a wide consensus that decisions on public policy should be to some extent democratic. Placing decisions wholly in the hands of an elite caste, whether they be politicians, philosopher kings, priests or scientific experts, notoriously leads to abuse and oppression. But, on the other hand, democracy is likely to lead sometimes to policies that scientific evidence suggests will be disastrous. Second, there are areas in which, arguably, no expertise exists. I have already suggested that there are no scientific experts on which science should be supported, as all scientists, by virtue of the cognitive specialisation that is their job, are bound to be biased on this issue. And Kitcher adds an additional interesting if controversial claim, that there are no experts in ethics (2011b, 286). If ethics is a technology for promoting altruism and social consensus, there can perhaps be facilitators of this process, but it is up to the citizenry at large to reach their agreements.

How, then do we steer between the Scylla of epistemic equality and the Charybdis of Platonic authoritarianism in making these judgements about the content and application of science? Since these are, of course, normative questions, it will be relevant to mention Kitcher's (2011b) account of ethics. Ethics is, for him, a social technology. One may or may not be convinced by the evolutionary story that Kitcher tells about the emergence of technology as a solution to the social coordination problems faced by early humans, but the conclusion the basic idea to which this leads, that ethics is constituted by a set of social practices and institutions to regulate behaviour in the interest of social harmony seems compelling.

Implicit in Kitcher's treatment of the present issue is his treatment of science and democracy as social technologies in the same sense as ethics (see D'Agostino 2013). This perspective provides surely the right perspective on the questions under consideration, though it is important at least to gesture at the substantial and complex bodies of academic work that have attempted to disclose the workings of both of these institutions. Democracy, very crudely, is a technology for pubic decision-making, which, more or less aims to reflect the views of citizens, or the majority of citizens in decisions about public policy. As a technology, it is constituted by procedures such as voting, parliamentary government, referenda and, importantly, various more informal channels through which citizens come to express their views in ways that can become known to politicians and influence their behaviour. I emphasise this last feature, because it is arguably the quite recent proliferation of channels of communication, notably through the internet, that has brought this clash between democracy and science to the fore. Citizens today have access to quantities of information inconceivable a few decades ago, though perhaps not accompanied by comparably effective tools for assessing its quality. They also have multiple means of expressing their opinions, through blogs, social media, email petitions, and so on.

Science as a social technology has been studied intensively by sociologists and by scholars in Science, Technology and Society (STS). Important elements include university departments, laboratories, training programmes, academic journals,

and hierarchies of power and authority. It is a technology of knowledge production, but also for the production of practical technologies for the satisfaction of human needs and wants and for interacting with the world. It is also the home of expertise, of people taken to be qualified to provide authoritative opinions on a wide range of issues.

It is easy to see how these institutions can come into conflict in the context of public policy decisions. Consider, for instance, the debate in the UK and elsewhere in Europe on Genetically Modified (GM) foods. Emerging from a remarkable explosion of knowledge of, and technical capacity in interaction with, genomes, these were developed as products intended to meet human needs for food[3]. Within the scientific community they were widely agreed to be an excellent pathway to produce many desirable characteristics of crops—reduced competition from weeds, resistance to insect attack, extra nutrients, and so on—and hence to a more efficient and productive agriculture. Though a fair amount of due diligence was seen as necessary as in any major modification of the human food chain, scientists generally saw little reason to anticipate serious risks to health, and claimed possible benefits for the environment, including reduced needs for herbicides and pesticides.

Public reaction was another matter. Under the inspired label of 'Frankenfoods', GM crops were widely portrayed as a God-like interference with the natural order, likely to involve unacceptable threats to human health and to the environment. This perspective seemed quite unaffected by declarations by scientific experts that such fears were groundless, and political pressure driven by the strength of public hostility led to the large scale abandonment of the technology. In some sense this must presumably be seen as a success for democracy: the public on balance rejected this technology, and the democratic government responded to their preferences.

This appears also, however, to be a departure from well-ordered science. Assuming the experts are right—and there seems little compelling reason to doubt it in this case—an opportunity for a valuable technology, not to mention a leading position in the development of the technology as an internationally marketable commodity, appeared to have been passed up. Why did the UK public, and publics in much of the rest of Europe, take so strongly against this technology? The story is, unsurprisingly very complex, involving a variety of interests and arguments.[4] There is a history that goes some way to explain the background of public suspicion of scientists, for example in the UK the then recent fiasco over the management of the Bovine Spongiform Encephelopathy (BSE, or 'mad cow disease') outbreak, in which scientists had informed the government there was no risk to public health, something that turned out to be false. There was widespread and understandable suspicion of the motives of the

[3] They were also, of course technologies intended for the enrichment of corporations and their shareholders, a factor that no doubt played some part in generating the public suspicion directed towards them.

[4] Much more detailed discussion of some of these arguments can be found in Barnes and Dupré 2008, chs. 6 and 7.

corporations, notably Monsanto, who were major funders of GM research. It is also clear that this outcome was highly contingent, as no comparable public opposition emerged to GM technology in the US, where millions of acres of GM corn and other crops are cultivated without anyone being much concerned.

A disturbing issue raised by this case for well-ordered science is the question how informed public perceptions of science are even possible given the communications revolution alluded to above. Consider, for instance, the averagely well-educated reader of a news story about the bacterium, *Klebsiella Planticola*. According to a story that circulated widely a few years ago this was an organism that had been genetically modified in the 1990s with the perhaps excellent intention of helping to turn plant waste into alcohol—this at a time when growing motor fuel in fields still seemed a plausibly good idea. According to the story, however, a heroic amateur scientist, Elaine Ingham, had discovered that this organism had the potential to provide a film over the roots of plants that would generate a lethal dose of ethanol. Potentially, the story went, this could potentially wipe out plant life across the planet, and we would of course follow soon enough. If the reader were concerned whether there might be environmental risks to genetic modification of organisms this might seem to be compelling evidence for a pessimistic conclusion.

About ten years I encountered this story and looked up 'Klebsiella Planticola' on Google. Hundreds of websites reported this horrifying brush with disaster, and none that I could find raised any doubts about the credibility of the research. I Indeed it is still widely cited today in support of the alleged riskiness of genetic modification. Repeating this search today (May 14, 2014), the large majority of hits still report this result, often with headlines such as 'The Bacterium that Nearly Ate the World'. However there are also a few blogs that report the discrediting of the research. With some perseverance a source that is likely to be more impressive to the scientific reader than a blog of unknown provenance, *Nature Biotechnology* (Fletcher 2001), reports that Ingham has subsequently apologized to the New Zealand government for submitting false claims about the ecological impact of genetically modified organisms, claims that were backed up by a scientific reference that didn't in fact exist. Note, however, that this is a resource behind a paywall that is unlikely to be passable by the non-academic reader.

In fact, moreover, the modification that caused this furore is the copy multiplication of a gene present in wild strains of this organism, and general reflection on microbial evolution suggests that if it were useful for the organism to duplicate this gene it would probably have done so. Killing the organisms with which it associates—all or most plants—probably would not be a selectively advantageous strategy, and the extrapolation from the ability of this modified bacterium to kill a plant in a confined environment in a laboratory to its likely spread across the planet is, to put it politely, biologically unlikely. But, and this is my main point, to the reader with limited scientific knowledge and limited access to scientific resources (should he or she even think of looking for them) the conclusion that GM technology had almost wiped out life on Earth might seem a reasonable one.

Parallels with the case of the MMR vaccine are obvious. The notorious paper by Andrew Wakefield (Wakefiel et al. 1998) has been fully discredited in the scientific community, disowned by Wakefield's collaborators, and withdrawn by the Lancet. Dr Wakefield has been struck off the medical register by the British Medical Council. Yet again, consulting Google on this topic it is clear that in the wider public a debate still rages. Wakefield is said by many to have been smeared by the medical establishment, certainly one possible interpretation of what has happened to him. A particularly telling comment in the internet discussion is the following: "why don't you let parents just make their own decisions? Do your research and make whatever decisions you think are best for your OWN children. Isn't that the very essences of parenting?" This seems well to reflect the reality in which many parents in fact decide not to vaccinate their children, and the consequent rising incidence of disease.

Kitcher's Solutions

In the previous section I portrayed the current much discussed flood of information as presenting a problem for anyone, except perhaps the relevant expert, in deciding what to believe. One aspect of the well-ordered science that Kitcher advocates in *Science in a Democratic Society* is a public institution of 'certification' through which established results get into the 'repository' of socially accepted knowledge. Clearly Kitcher would like to see the consensus of climate scientists on global warming, of biologists on safe ways of improving crops, our best evolutionary theory and so on, certified for the repository, and the views of climate change deniers, GM scaremongers, or intelligent design theorists excluded. I agree with him. But how is this to be done in a way that is democratic and broadly acceptable to a democratic citizenry? Reflection on the broad (democratic?) discussion made possible by the internet are not altogether encouraging.

Kitcher does not endorse a naïve or 'vulgar' democracy. Certainly decisions on scientific matters—what science is to be funded, what results are to be certified, and how they should be applied—are not to be decided by referendum. Kitcher does base his answers to these questions on discussion and deliberation, but not of the anarchic kind that the internet represents. What is wrong with vulgar democracy? One problem is the venality of some participants. 'Resisters'—climate change sceptics, intelligent design theorists, and so on—who oppose rationally indisputable scientific results are in fact operating on the basis of different, and generally concealed, values. Climate science deniers are often associated with the fossil fuel industry; scientists arguing for the safety of cigarettes were frequently found to be paid by the tobacco industry. The proper kind of transparent discussion will reveal these values, and the democratic majority at least will find that they do not share them. It is not that the competent scientists do not have values of their own, but generally these will be values that, again with proper explanation and transparency, the citizenry will endorse.

Kitcher actually has two very different conceptions of the relevant conversation. The one that ultimately drives the argument is a Rawlsian one, a conversation of ideal deliberators who aim to reach consensus in a reflective equilibrium. This, it is supposed, will eliminate values that are not sustainable in such an ideal conversation. On the other hand, Kitcher describes actual conversations between experts and representative lay participants, the latter to be thoroughly briefed by the former, so as to be able to reach informed decisions subsuming sustainable values. It should also be noted that while the former conception appears to offer a criterion for what the normatively correct (ideal) outcome should be, this needs to be reconciled with a naturalistic account of ethics as an actual social technology that evolved from actual discussions between agents concerned to resolve social discord. It is not easy to see how these two conceptions are to be harmonised.

A somewhat cynical view is that "the notion of ''ideal deliberators'' often seems like little more than a philosophical delivery van for Kitcher's policy proposals" (Brown 2013, 395). Slightly less cynically, it is difficult to see what would make a deliberator ideal if, as Kitcher insists, there are no ethical experts. Not cynically at all, while I have a lot of sympathy for the model of an ideal conversation as an intellectual tool, it seems to me somewhat tangential to the most interesting feature of Kitcher's work in this area, which is the aim of understanding and addressing the discordance between two social technologies that Kitcher, and I expect most of his readers, admire, science and democracy. For this, surely, is a technological problem, largely independent of any theoretical, meta-ethical account of what would be the normatively desirable outcome of a well-ordered science.

This may seem wrong for the following reason. Surely we need some explanation of why we thought science was disordered in the first place, and doesn't this require an account of value against which the present situation can be judged wanting? I think that this worry does point to a real tension in Kitcher's writing. It seems clear that his concerns do begin with some firm convictions that things are amiss, for example the unwillingness of democratic states to adopt serious measures to combat climate change. And I don't doubt that there are good reasons for this. Personally I am more persuaded by formulations of these reasons in terms of evidence and consequences than in terms of the views that would be reached by ideal deliberators, though of course Kitcher aims to do both, and the formulations do answer different questions. And of course both evidence and consequences involve normative assumptions.

But surely if one is serious about democracy, even a sophisticated democracy free of all the familiar forms of vulgarity, one cannot prejudge the question what any suitable democratic process will decide about any positional policy issue. Kitcher is an optimist, and he tends to believe that a proper, non-vulgar democratic system will produce decisions more or less of the sort that sensible people like he or I would prefer. But it surely conceivable that a democratic society might reach a fully articulated view that without major changes in consumption patterns the world would be uninhabitable in two hundred years time, that Bangladesh would be under water in a hundred, and so on, and

democratically agree that they would rather keep their gas-guzzling cars and air-conditioned homes. One may conclude, so much the worse for our descendants and the Bangladeshis, or so much the worse for democracy, but I don't think the problem can be solved by yoking democracy to a theoretical account of what should be done. It might be, of course, that there are ethical experts after all, and Kitcher is one of them. But in that case shouldn't we try to devolve our decisions to Kitcher and his likes rather than to an unreliable democratic process?

So this brings me to the more practical kind of discussion that Kitcher considers. Part of the practical resolution between democratic and expert opinion is to be achieved by something like citizens' juries, including both a range of relevant experts and representatives of the public, the latter selected to cover as wide a range as possible of the diversity of perspectives within a population, and especially representing groups most likely to be affected by the decisions being taken. The problems of vulgar democracy will be addressed by making sure that all the facts as seen by experts are presented, and time has been allowed to raise all questions or doubts in the minds of the lay members of the panel. If all goes right, first the panel will reach some consensus on the matter at hand, and second the process will have sufficient general legitimacy that its decisions will command widespread respect.

It is important to say that this is the right *kind* of solution: it is an attempt to mesh the two conflicting social technologies in a way that will reconcile their conflicts. But I must confess to being a bit sceptical as to whether such a system is likely to achieve the benefits Kitcher hopes for. Will the fundamentalists, racists, climate change deniers, Christian scientists, alien abductees, and so on have a privileged place at the table as groups likely to be most affected by the reasonable consensus of these committees? Will they be convinced? Even if they are convinced, will they not be perceived as traitors by their wider communities? More generally, will these committees be perceived as genuinely democratic, or merely as the opportunity for the scientifically minded to generate propaganda for their views with the help of ambitious stooges from the wider public? Given the experience that many people are happy to draw their opinions from sources wholly opposed to almost universal scientific consensus, it is hard to know why they should change their views in the face of even the most well-meaning Quango.[5]

All this is not to deny that groups of this kind can be a good idea. Indeed they are demonstrably so. One of the most effective quangos, the UK Human Fertilisation and Embryology Authority (HFEA), which regulates research involving human gametes and embryos, is a very widely respected group, in many respects very much the kind of entity Kitcher advocates.[6] This body includes scientists, doctors, women who have experienced fertility treatments, a bishop, a lawyer, and even a philosopher. Unfortunately, its decisions do not appear to change the

[5] A quango, in the UK, is a quasi non-governmental organization.
[6] Sufficiently well-respected that it survived the so-called 'bonfire of the quangos' in which the Cameron administration attempted to dispose of as many as possible of such bodies.

minds of those who oppose them. For example when the HFEA first licensed therapeutic human cloning in 2004, there was outrage from a range of 'Pro-Life' groups. Perhaps most telling was the comment attributed to Josephine Quintavalle, of the pro-life group Comment On Reproductive Ethics: "It is very worrying indeed. We have decisions of this magnitude being taken by an unelected government quango".[7] The absence of vulgar democracy may be used as a weapon.

I have mentioned that there may be a problem in harmonising the ideal and practical conceptions of conversation. This is most obvious in relation to the question of the scope of the problems that are to be addressed. Recall that the kind of social technology that is the model for addressing ill-ordered science, the ethical project, is taken to have originated in discussions within small proto-human tribes. The problem of climate change, in contrast, is global. The possibility of small groups sitting round a table to hammer out a relatively local problem is appealing, perhaps even the adult male citizenry of a Greek polis might gather in the agora to similar effect. Representation of all the peoples of the Earth, and all the different perspectives within each nation, is a different matter, and perhaps only accessible to the ideal conversation. One trouble with the ideal conversation when it is more than some kind of rationalisation of a practical process is that it seems suspiciously like a job for an expert ethicist of whom, Kitcher has told us, there are none.

The difficulties with a conversation over climate change do not end here. Those most affected, on all but the most pessimistic scenarios, have yet to be born. Perhaps there are many possible people who will not be born unless we do something serious to address this problem. Who will speak for the unborn or the possibly never to be born? Kitcher thinks that among peoples' central life goals are the well-being of their children, grandchildren and, perhaps by extension, future generations generally. I'm not sure. I think most people care a lot about the well-being of their children and many, but by no means all, care about their grandchildren. But future generations long after their own deaths? I'm much less sure. Perhaps again they will need experts on ethics to speak on their behalf.

Conclusion

I have been somewhat sceptical about the proposals Kitcher sketches for reconciling science and democracy. He is, I have said, an optimist, and in this domain I am more pessimistic. Kitcher thinks that scientists discredit themselves by making excessive claims for their expertise and by acting on values other than those that should legitimately underlie their professional work. If institutions could be constructed that would expose these flaws and open their work to rational and civilised discussion, these flaws could be removed and well-ordered science would regain the deserved trust of the citizenry.

[7] See http://news.bbc.co.uk/1/hi/health/3554474.stm.

I fear that the sources of these defects go deeper, and that they will be even harder to remove than Kitcher supposes. I have described Kitcher's project as a political philosophy of science, and I wholeheartedly agree that that is something we need. The politics, however, is largely limited to the commitment to an admittedly sophisticated conception of democracy. This is, no doubt, a good thing to be committed to. But many of the problems that Kitcher is concerned with arise not merely from failures of democracy, but also from the intrinsic problems with the liberal, or neoliberal, framework within which most current democracies exist. In a social system that forefronts competition between individuals it is hard to imagine scientists who don't have their own agendas, even if these are no worse than the quest for personal success by doing good science. More problematically, a system that encourages the accumulation of wealth in large competitive corporations, and encourages these corporations to fund scientific research, will inevitably produce research infected with the values of the corporate funders. These are issues that, I fear, can only be addressed at a more systemic level than even the most well-constructed institutional add-ons for enlightened public debate.

Having expressed these doubts, I must nonetheless reiterate my endorsement of the importance of Kitcher's project, and the gratitude we should feel for the work he has done to open up the philosophy of science to these absolutely fundamental questions. The two books I have been considering on well-ordered science, together with his account of the ethical project, constitute a systematic attempt to address the political question of the role of science in society, and his account is full of valuable insights that should remain part of this debate.

My suspicion, however, is that the attempt to construct a democratic science may ultimately be impossible without more integration into the problem of constructing a democratic society. In 2012, the world's 100 richest people became $241 billion richer. They are now worth $1.9 trillion: just a little less than the entire output of the United Kingdom. [8] The problem that Kitcher is discussing is democracy of voice not equality of resources. But as resources become ever more unequal, democracy of voice becomes ever more unrealistic. It would be nice to see a democratic socialist account of well-ordered science, though perhaps given the distance we are from democratic socialism, Kitcher's account, despite inevitable weaknesses, will be more useful.

Bibliography

Barnes, Barry and John Dupré (2008). *Genomes and What to Make of Them*. Chicago: University of Chicago Press.

Brown, Mark B. (2013). Review of Philip Kitcher, *Science in a Democratic Society*. *Minerva* 51: 389-397.

[8] Monbiot, Guardian, Jan 14.

Cartwright, Nancy (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

D'Agostino, Fred (2013). Review of Philip Kitcher, *Science in a Democratic Society*. *Analysis* 73: 593-594.

Douglas, Heather (2009). *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.

Dupré, John (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.

Fletcher, Liz (2001). 'New Zealand GMO debacle undermines green lobby'. *Nature Biotechnology* 19: 292.

Herrnstein, Richard J. and Charles Murray (1994). *The Bell Curve: Intelligence and Class Structure in American Life*. New Yourk: Simon and Schuster.

Kincaid, Harold, John Dupré, and Alison Wylie (eds.) (2007). *Value-Free Science: Ideals and Illusion*. New York: Oxford University Press.

Kitcher, Philip (2001). *Science, Truth and Democracy*. Oxford: Oxford University Press.

Kitcher, Philip (2011a). *Science in a Democratic Society*. New York: Prometheus Books.

Kitcher, Philip (2011b). *The Ethical Project*. Cambridge, MA: Harvard University Press.

Wakefield, A. J., Murch, S. H., Anthony, A., Linnell, J., Casson, D.M., Malik, M., Berelowitz, M., Dhillon, A. P., Thomson, M. A., Harvey, P., Valentine, A., Davies, S. E., Walker-Smith, J. A. (1998). 'Ileal-lymphoid-nodular hyperplasia, non-specific colitis, and pervasive developmental disorder in children.' *Lancet* 351: 637-41. (Retracted in *Lancet*, 2010, 375:445.)