

Regarding scientific significance

P.D. Magnus

June 8, 2007 – June 15, 2008

I put this paper on my website in 2007 and revised it somewhat over the following year. It has been over two years since then, and I do not foresee following it through to proper publication. When someone asked for permission to cite the unpublished paper, I decided to post this stable version. It is suitable for citation.

Magnus, P.D. 'Regarding scientific significance.' Unpublished. 2008.

Abstract

In *Science, Truth, and Democracy*, Philip Kitcher introduces significance graphs (structures that illustrate how and which questions are significant) and well ordered science (a norm defined by an imagined process of ideal deliberation). Jeremy Simon has argued that these two parts of Kitcher's account are intimately connected. In this paper, I argue that the connection between significance graphs and well-ordered science is rather more complicated. I survey three objections to Kitcher's account, two from Simon and a third by analogy with similar positions in ethics. This paper aims to show that Kitcher's account relies on some questions being ones about which we are curious. This has the consequence that ideal deliberation does not yield a precise agenda for basic research. Nevertheless, the ideal might draw our attention to features that make actual arrangements more or less well-ordered.

Philip Kitcher [Kit01] argues that there is no context-independent, objective way of characterizing scientific significance. Instead, which questions are scientifically significant for us depends inescapably on what we care about and what we are interested in. In order to explicate this point, Kitcher introduces *significance graphs*— structures that illustrate which questions are significant and how. He further specifies an ideal procedure for assigning resources to scientific projects, carrying them out, and exploiting the results. An arrangement is said to be *well-ordered* if it pursues projects in the way that the ideal procedure would select. This notion of well-ordered science is meant to provide a normative account of how scientific labor ought to be distributed.

I describe Kitcher on scientific significance in §1 and well-ordered science in §2. In §3, I consider Jeremy Simon's [Sim06] interpretation of Kitcher. Simon argues that the two parts of Kitcher's account are intimately connected. As he

understands it, the function of the ideal procedure is to determine the significance graph for the entire community. I argue that the connection between significance graphs and well-ordered science is neither as direct nor as simple as Simon supposes. I consider and answer two objections Simon makes against Kitcher's account. In §4, I consider a further objection motivated by similarities between Kitcher's notion of well-ordered science and full-information accounts of well being.

Throughout what follows—in understanding significance graphs and in answering some objections—I show how Kitcher relies upon a largely unanalyzed notion of natural curiosity. I ultimately argue that natural curiosity breaks down when considering well-ordered science, driving a wedge between Kitcher's account of significance and the normative ideal of well-ordered science.

1 Scientific significance

Science aims at discovering significant truths. But which truths are significant? In earlier work [Kit93], Philip Kitcher supposed that this was unproblematic. He no longer thinks so. Scientific significance is not an objective property, he now argues [Kit01]. Rather, it depends on the interests of the community doing the science.¹

It may help to consider why the alternatives face substantial difficulties. Suppose that one wanted to provide an interest-independent account of scientific significance. What, then, would make a truth significant? One might say that a truth is significant insofar as it facilitates the discovery of further truths. Yet a truth that facilitates cataloging trivia is not thereby deeply important. Facilitating further discovery is only a source of significance if the further discoveries are themselves significant. A regress threatens. One might nevertheless say that significance is a matter of discovering truths, and discovering more truths makes for greater significance. It is hard to count truths, however, and most enquiry could lead to the discovery of a countable infinity of truths. There is no way to weigh such infinities against one another unless the truths themselves are more or less significant. One might still insist that if a set of truths A is a proper subset of a set of truths B , then A must be less significant than B . This is a special case of more-is-better that does not require enumerating truths, but it is really no help in developing an interest-independent account of significance. First, few if any distinct scientific theories make claims that are a proper subset of the claims made by another theory. Second, an interest-dependent account of significance can accept this special case: More knowledge is better than strictly less. That dictum alone does not tell us *how much* more significant the extra knowledge is. It looks as if trying to account for significance in an interest-independent way fails to answer the general question of how to compare the significance of different truths.

Following up this line of argument is beyond the scope of this paper. I want instead to accept that scientific significance is in some sense interest dependent

¹For more on the difference between Kitcher's earlier and more recent work, see [Mag07].

and to consider Kitcher's claim that it depends on the community's interests. This is indisputable for research directed toward some application. For example, there is nothing intrinsically interesting about techniques for maximizing the yield of nuclear warheads. Weapons research is motivated by military objectives— or by commercial objectives, when there are military paymasters.

Also indisputably, some research is significant because it furthers other independently significant research. For example, before the discovery of the syphilis bacillus there were studies that asked how occurrences of syphilis vary with climate and season. This research was deemed significant at the time because it could perhaps provide a clue to the etiology of the disease and as to possible treatment [Fle79, p. 19]. Once the bacillus was discovered, questions of climatic and seasonal variation ceased to be significant. In order to understand the kind of significance those questions enjoyed, it is not enough to ask if an enquiry is significant *simpliciter*. Instead, it is essential to ask how such enquiry relates to other enquiries. Those other enquiries might themselves be significant because of their relations to still other enquiries. This web of inherited significance can be represented in what Kitcher calls a significance graph. He defines it in this way: "A significance graph is constructed by drawing a directed graph with arrows linking expressions, some of which formulate questions that workers in the field address, others that encapsulate claims they make, yet others that refer to pieces of equipment, techniques, or parts of the natural world" [Kit01, p. 78]. An arrow pointing into a node means that the question or enquiry inherits significance. For example, figure 1 is a very simple significance graph for the syphilis research described earlier. When the bacillus was discovered, the arrows pointing to 'How do occurrences of syphilis vary with climate and season?' went away, and the question ceased to be significant. Questions about the bacillus became significant instead.

Kitcher provides a partial significance graph for developmental biology [Kit01, fig. 1, p. 79]. The question of how genes are switched off is scientifically significant because the mechanisms of cell division are significant— those in turn are significant because organism development is significant. Kitcher cautions that this example should not be weighed too heavily. He calls it "extremely partial" [Kit93, p. 79] and constructs another graph that incorporates some of the same nodes; cf. [Kit01, fig. 2, p. 80]. We get different graphs if we adopt the perspective of a particular field of science or focus on the significance of a specific development. These different perspectives are "compatible and valuable for different purposes" [Kit01]. As such, we should be careful not to reify significance graphs. They are simply a way of illustrating how some questions inherit their scientific significance from others.

It is important to note that the inheritance illustrated in significance graphs cannot account for all scientific significance. The parent nodes (the ones with no arrows pointing into them) must derive their significance from some other source. Practical application is one obvious source— think of the nuclear warheads. If there were no other source, then we would arrive at a crude kind of pragmatism: Significant truths would only be significant insofar as they furthered our practical ends or contributed to a system that furthered those ends.

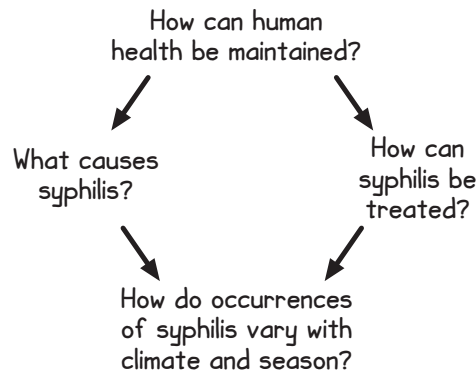


Figure 1: A partial significance graph for medical enquiry, prior to the discovery of the syphilis bacillus.

Although science would aim at significant truth, such truth would only be the power to exercise efficient control over nature. Truth would be eclipsed by significance.

This could be avoided if there were some objective source of significance, if the natural light revealed some questions to be of intrinsic epistemic importance. Bishop and Trout [BT05] suggest something like this solution. They advocate what they call *strategic reliabilism*, according to which belief forming strategies are praised for reliably producing true beliefs about significant matters. Significant truths, they say, are ones the knowledge of which either promotes human well being or has intrinsic epistemic value. They explain:

We take it that discovering the truth about the physical or social structure of the world is intrinsically valuable. So even if we can't be sure that it will lead to any practical results, the physicists at CERN and Fermilab have epistemic reasons (beyond their prudential reasons) for spending cognitive resources on trying to discover the Higgs boson. [BT05, p. 97]

They give no argument for the objectivity of epistemic reasons and no account to explain why epistemic reasons justify esoteric activities like trying to discover the Higgs boson. Kitcher, who used to believe in such epistemic reasons, now argues that there is no intrinsic epistemic significance. He advocates relativism about significance, in the sense that significance claims must always be indexed to particular enquirers.

Kitcher is, nevertheless, a realist. Relativism about significance is not supposed to ripen into relativism about truth. “What... is the ultimate source of epistemic significance?” he asks. “The answer is... commonplace and disappointing to those who expect a grand theory that will invest the sciences with overriding importance” [Kit01, p. 80]. There are some questions and discoveries which are just interesting: How do living things function and grow? What was the past like? What’s *out there*? (The latter works best when pointing into the night sky.) These are significant and interesting questions, just in themselves and not because the answers promise to materially remunerate us. As Kitcher explains,

We expect other people to see the point of such questions, and we describe those who don’t as lacking in ‘natural curiosity.’ Partly as a result of our having the capacities we do, partly because of the cultures in which we develop, some aspects of nature strike us as particularly salient or surprising. In consequence, we pose broad questions, and epistemic significance flows into the sciences from these. [Kit01, p. 81]

Some of the parent nodes in significance graphs are there just because of natural curiosity; there are parts of nature that we are naturally curious about, and it is good to learn.

Consider, for example, dinosaurs. It may be that the study of fossil dinosaurs serves some practical or instrumental purpose. It provides rhetorical material for debates about the history of the Earth which connect to broader religious and cultural agendas. It contributes to stratigraphy, perhaps, which in turn serves mineral and resource exploration. Yet none of these practical ends are in view when a child walks into a museum and looks up with awe at the maw of a fossil *Tyrannosaurus*. We are naturally curious about dinosaurs, and just knowing about them is a good thing. ‘What were dinosaurs like?’ is a significant question and can serve as the parent node of some significance graph. This is sufficient to show that some of the parent nodes in our significance graphs are there for other than practical reasons. Some questions answer to our natural human curiosity, and they are significant to us. This is sufficient to keep Kitcher’s account from collapsing into crude pragmatism.

One may still object that natural curiosity is carrying too much theoretical weight here; I will return to this worry below.

2 Well-ordered science

A significance graph alone is not sufficient to determine how scientific resources should be distributed among possible projects. First, because different individuals might be concerned with different basic questions. A single significance graph does not tell us how to aggregate the interests of community members. Second, because a significance graph does not contain any information about

relative levels of significance. It shows *that* some projects are significant but not *how* significant each is in comparison to the others.

In contrast, the notion of well-ordered science is meant to answer the question of how scientific resources should be spent. Kitcher asks us to imagine a process of ideal deliberation [Kit01, p. 118–122]. Community members each start with their own preferences. They meet together and learn about the preferences of their compatriots. They recognize that other members of community have as much right to pursue their aspirations as they do themselves. Once all interests have been illustrated and discussed, the group votes to draw up a list of priorities. This first phase of ideal deliberation addresses the first weakness with significance graphs: Individual preferences are refined and aggregated.

Once the deliberators have decided on a set of concerns, experts determine what projects could be pursued to address these concerns. These “possible agendas for inquiry” [Kit01, p. 120] are presented in a report to the community. Where there is uncertainty or disagreement, the experts acknowledge it and provide their best estimates.

The members of the community then discuss the potential costs and benefits of each agenda. They discuss any ethical issues that arise from the various proposed research programs. After everyone has been heard out, they vote on which agenda should be pursued. Research effort is then distributed in accord with the outcome of this vote. This second stage of ideal deliberation addresses the second weakness of significance graphs: The community decides which significant projects are worth pursuing and in what measure.

There is a third stage of deliberation, after research has been conducted, in which the community decides how results should be applied and disseminated.

Kitcher’s suggestion is not that we should implement this ideal process. Our actual social arrangements could never give us enough time, patience, or resources to do so. More importantly, Kitcher is not even suggesting that we should necessarily try to procedurally *approximate* the ideal process. Rather, Kitcher’s discussion of the ideal procedure is a roundabout way of specifying an arrangement of scientific resources. Any social arrangement that results in the distribution of scientific effort that the ideal process would count as ‘well-ordered.’ Kitcher explains that “there’s no thought that well-ordered science must *actually institute* the complicated discussions I’ve envisaged. The thought is that, however inquiry proceeds, we want it to match the outcomes those complex procedures would achieve. . .” [Kit01, p. 123].²

3 Simon’s interpretation

Jeremy Simon [Sim06] argues that Kitcher’s ideal procedure functions to determine the community’s significance graph and further that the community’s

²Kitcher does not say what actual system of organization would do best at producing well-ordered science. He speculates, and his sympathies are clearly with some democratic arrangement— but he does not say anything decisive. Perhaps well-ordered science would be most readily obtained by having a benevolent central office assign scientific resources.

significance graph exclusively defines what is scientifically significant. Simon writes:

For Kitcher, to determine whether a question is scientifically significant, one must refer to the graph, or, what amounts to the same thing, to the ideal deliberation that produced the graph, to see whether the question appears there. Nothing short of repeating ideal deliberation, or accessing its epitome, the graph, can yield a verdict on significance. [Sim06, p. 209]

As I understand Kitcher, there can be no question of “repeating ideal deliberation”; it is an ideal process that could never be done even a first time. More importantly, Simon’s reading of Kitcher (according to which a significance graph is the ultimate product of ideal deliberation) misunderstands the relation between significance graphs and well-ordered science.

The first stage of ideal deliberation identifies a set of shared concerns. As Kitcher notes, “Ideal deliberation must involve presenting the structure of significance graph, where the multiform sources of significance are revealed” [Kit01, p. 118]. This does not mean, however, that deliberators are free to draw any significance graph that they like. At this stage, they merely select questions and problems which count as intrinsically significant. These can be the parent nodes of a significance graph, the nodes that inherit none of their significance from other questions or problems.

Between the first and second stages of ideal deliberation, experts suggest projects and possible agendas. They determine which projects and questions would inherit significance from those that the ideal deliberators deemed ultimately significant. In effect, they draw the significance graph that results from the community’s choices. The graph does not result from the community’s choices. Experts draw the graph based on the best scientific knowledge available.

The second stage of ideal deliberation selects an agenda for allocating scientific resources; science is well-ordered if it allocates resources in the same way that ideal deliberation would. Although the selected projects must count as significant, the ultimate sources of significance and the connections between them are irrelevant. As such, any preferred allocation will be compatible with an indefinitely large set of different significance graphs. If the allocation includes a project to clone sheep, for example, this might be deemed significant because of its systematic connections with developmental biology, because of potential applications in ranching, because people find sheep interesting, or because of some complicated combination of these and other factors.

The third stage determines how results will be disseminated and applied. It does not have any obvious connection with significance graphs.

If I am right about the loose connection between significance graphs and well-ordered science, then many of Simon’s criticisms misfire. In the remainder of this suggestion, I consider two interesting objections raised by Simon that are independent of this point.

3.1 The demarcation reductio

Simon argues that Kitcher’s account of science yields the absurd consequence that whether a specific project counts as *scientific* or not may depend on what is happening in distant corners of the community. The argument can be formulated as a reductio: Kitcher claims that (a) science aims at significant truth and (b) significance is relative to the community. It follows from (a) and (b) that changes in the community could change whether or not something counts as science. Yet surely the status of some enquiry as *scientific* does not depend on whether the community cares about the outcome. To take the example of syphilis research, discussed above, the epidemiological data on the disease across climates and throughout the year might have ceased to be *interesting* once the pathogen was identified— but it did not cease to be science.

Simon suggests that the problem may be resolved by indexing significance to individuals rather than to communities. Yet Simon’s suggestion reduces to a similar absurdity, as he admits: “Two people could be simultaneously engaged in identical projects, but, if their motives are different, it may be that only one of them is doing science by pursuing a significant project” [Sim06, p. 205]. To continue with our example from §1, imagine an epidemiologist collecting data about the prevalence of syphilis in the tropics. It would be absurd to say that the data collection ceases to be science because, half a world away, other scientists discover the microbe responsible for the disease— that much would be avoided if significance were relative to the individual still toiling in the tropics. Yet it would be just as perverse to say that the data collection ceases to be science because, in a fit of ennui, the scientist no longer considers his own work to be of any significance— indexing significance to individuals makes the demarcation between science and non-science change with the motivation and mood of single investigators.

Simon puts the upshot this way: “[I]t would be odd for a laboratory to halt an experiment not because the problem had been solved or rendered irrelevant by a scientific revolution, but because, as an indirect result of a social revolution, the question was rendered unscientific” [Sim06, p. 203]. This oddity would still arise for Simon’s account if the scientists themselves were part of the social revolution. The problem inevitably arises for any relativist account of significance, so changing the locus of relativity is no help. It arises because Simon uses Kitcher’s characterization of science as a demarcation criterion. The problem can be resolved by recognizing that separating *genuine science* from *whatall else* may proceed along several dimensions.

First, one may distinguish between science and other non-scientific activities. This distinction is crucial to the Popperian effort to discredit some enquiries as pseudo-scientific (as in [Pop65].) Marxist history and Freudian psychoanalysis, Popper says, do not involve falsifiable claims and so are simply not science. Such demarcation does not admit of degrees. Once an interlocutor can be pilloried as not even doing science, there is no further debate to be had with them. Simon supposes that philosophers of science must provide a demarcation criterion of this kind. “Despite the failure of positivistic criteria for the demarcation of sci-

ence,” he writes, “the need to distinguish science from other endeavors remains” [Sim06, p. 213]. Simon thus takes Kitcher’s characterization of the aim of science to be a demarcation criterion like Popper’s. If ‘searching for significant truth’ defines science in this way, then changes in what counts as significant will move the line of demarcation.

Second, one may distinguish well-executed, successful science from shoddy, dreadful science. Since ‘scientific’ can function as an honorific term, there are contexts in which this distinction can be marked with the labels ‘genuinely scientific’ and ‘unscientific.’ Whereas pseudo-science is unscientific because it only purports to be science but really is not, shoddy science is unscientific because it tries to be science but does a miserable job at it. It makes a hash of scientific methods, offers bad arguments, and reaches unsound conclusions. Kitcher [Kit85][Kit93, p. 195] suggests that this distinction suffices to discredit so-called creation science.

Third, one may distinguish between scientific enquiry into matters of importance and enquiry into mere trivia. Imagine I make a bold conjecture as to the number of instances of the letter ‘e’ present in the university library. I begin a research program to investigate this matter using responsible and sophisticated methods. Yet, even though it is scientific in the first and second senses, one may still object that my wild endeavor is not real science. Again, this is because ‘science’ can be honorific—denying the label ‘science’ to my e-research is a way of condemning it as mere data fetishism. A similar denigration is suggested by Ernest Rutherford’s quip that all science is either physics or stamp collecting. The important questions, he suggests, are the ones asked and answered by physicists; the rest are only science in an anemic sense.

How should we understand Kitcher’s claim that science aims at significant truth? Kitcher [Kit85] sees no hope of finding a demarcation criterion of the Popperian kind; see also [Kit93, p. 196]. So he is not specifying *what science is* in the first sense.

If we understand Kitcher’s specification in the second sense—distinguishing real science from irresponsible storytelling—then it would entail that any enquiry aimed at discovering trivia would *ipso facto* employ unreliable or unreasonable methods. Surely this is absurd. The reliability of methods does not depend on whether they are used for important purposes.

If we understand it in the third sense—distinguishing real science from data fetishism—then it would entail this scenario: Significant social changes that lead us as a community to value different things might make old projects no longer seem important. Individuals who still pursued them would be doing science in the first and second sense even though their work had become trivial. Their work would not become pseudoscience or methodologically shoddy, but it would cease to be worthy of the appellation ‘science.’ There is nothing odd or worrisome about this scenario. So reading Kitcher in this way avoids the alleged absurdity.

3.2 The stultification objection

Simon worries that well-ordered science might fail to follow lines of enquiry which would turn out to be rewarding. He writes:

[H]istory warns us against requiring wide approval of scientific interests. There are many examples of scientists, such as Galileo, Boyle, and Semmelweis, who studied questions that, although of interest to few others at the time, ultimately came to be considered of broad and vital importance. [Sim06, p. 207]

Although the figures that Simon mentions were arguably concerned with issues central to scientific concerns of their time, examples are not hard to find. Gregor Mendel's work, for example, did not receive recognition until years after his death. Ideal deliberators in the nineteenth century would not have selected it to be part of the agenda for science. Yet, in retrospect, we recognize that Mendel's work was important for the advancement of science. Science was better off for having included it.

Specific examples like this illustrate the commonplace that basic science leads to significant consequences that could not possibly have been anticipated. If well-ordered science means only pursuing lines of research that have foreseeable payoffs, then it would stifle scientific development.

Simon suggests that this his approach, making scientific significance accrue to the interest of even a single individual, is the only way to avoid this stultification. He writes, "Only by expanding the province of legitimate scientific inquiry to questions of interest to anyone can we do justice to the intuition that in these cases... science was done before the recognition arrived" [Sim06, p. 207]. Simon's conclusion here turns partly on his conflating different notions of demarcation; the account of well-ordered science does not say anything about whether Galileo or Mendel were 'doing science' in the first or second sense. Moreover, we should consider ways in which such cases might be recognized as significant within Kitcher's framework.

First, one might say: The arrows in a significance graph reflect objective connections between the various nodes. The correct graph in the nineteenth century would have included an arrow from the nature of heredity to Mendel's peas, even if few in the nineteenth century would have drawn the graph that way. Admittedly, this would require some revision to Kitcher's account. He indexes significance graphs to particular times and means for them to reflect "what workers in the field know at the time to which the [graph] is indexed" [Kit01, p. 78]. Yet well-ordered already relies on various idealizations, and so revising it to include ideal significance graphs (the graphs that scientists ought to draw) might not seem like such a major change. However, this would unravel the notion of well-ordered science. The process of ideal deliberation is meant to yield an agenda for the scientific enquiry to be conducted. If deliberators were idealized so as to *already know* what the objective significance graphs are, then they would never decide at the end of the process to learn whatever facts are encoded in the graphs. Well-ordered science could never advise enquiry into

such matters, but surely those connections are often things that ought to be researched.

Second, one might reply: Ideal deliberators will be educated about the significance of each of the options, as judged by their community's best scientific account of the world. The scientists will also tell them that scientific research can lead to surprising consequences. As such (the reply concludes) ideal deliberators will select an agenda for science that includes some larks. Although Kitcher does not explicitly discuss this, there is nothing in the structure of ideal deliberation to preclude such an outcome. Yet, this is no help to the community trying to decide *which* seemingly-impractical projects should be funded.

Third: Abstruse basic research, although it does not inherit significance from any practical applications, typically does address big questions about which we are naturally curious. As we saw (at the end of §1) such natural curiosity is a source of scientific significance for Kitcher. This provides the community with some way of sorting out which basic research would be worth pursuing: viz., the projects that address questions about which we are naturally curious.

Again, one might worry about the work that natural curiosity is doing for Kitcher's account.

4 Problems with idealization

In this section, I develop an objection to Kitcher's notion of well-ordered science by comparing it to full-information accounts of well being. The latter are well-explored ethical positions.

In full-information accounts of well being, the good for an individual is defined to be the life that she would prefer from an idealized standpoint of knowing about all the possible lives that she might live. (For summaries of such views, see [Sob94], [Ros95], and references therein.) There are two significant differences between such accounts and Kitcher's account of well-ordered science:

First, Kitcher is concerned with the community rather than with an individual. Nevertheless, Kitcher's ideal procedure requires that we imagine each individual learning the preferences of all her fellows. Voting at the end ideal deliberation only makes sense because each (idealized) individual has participated. Kitcher briefly considers allowing proportional representation rather than direct democracy [Kit01, p. 123], but it must still be the case that each representative can be imagined to participate in ideal deliberation. If the process of ideal deliberation is incoherent, then well-ordered science is not well-defined.

Second, Kitcher is concerned with scientific significance, rather than with general well being. Ideal deliberators are asked to decide between scientific priorities and distributions of scientific resources. These choices are similar to decisions that we make in our ordinary, unidealized lives. Contrawise, full-information accounts imagine an idealized person choosing between different possible lives. Critics have objected that it is incoherent to imagine an individual seeing the intrinsic value in every possible life while at the same time being the person who has lived whatever life she has lived. David Sobel writes: "The

narrative unity of a life can provide the context to make sense of choosing one option over another, but this context is significantly dropped when we are choosing between lives rather than from within them” [Sob94, p. 808]. This problem does not arise directly for Kitcher, since he imagines ideal deliberators making decisions from the context of their own (idealized) lives.

Still, a version of the worry may arise for Kitcher.³ Consider a religious zealot who believes that all statements in the Bible are literally true, that the world was created some small number of years ago, and that science is the enemy of faith. In considering what arrangement of science would be well-ordered, we imagine this fanatic participating in ideal deliberation. In the first phase, he listens to each of his fellows explain what kinds of inquiry they consider to be important. We must imagine that he takes their testimony seriously and believes that community decisions should reflect their concerns as much as his own. This may already be hard to imagine, but it is incorporated into the very conception of well-ordered science; in the framework Kitcher proposes, “democratic ideals are taken for granted” [Kit01, p. 116].

At the second phase of ideal deliberation, we imagine the zealot reading reports from scientists about different possible research programs. If the zealot were to *actually* read such reports, then he would believe none of it. He would scoff at the expressions of scientific consensus and expert opinion. It is not clear why the zealot would react any differently in the ideal case, so what could he decide at the end of the second phase? Perhaps the zealot would recommend no scientific research whatsoever. Just as well-ordered science presumes a community with democratic values, it presumes a community that recognizes some value to enquiry. It does *not* address the broader question of how much effort should be directed toward science in the well-ordered society.

It is even more difficult to imagine the zealot in the third phase of ideal deliberation. In the third phase, deliberators decide how results of enquiry will be applied and disseminated. Yet, as Kitcher discusses [Kit01, ch. 12], scientific results sometimes lead religious folk into a crisis of faith. Are we to imagine the ideal zealot, experiencing such a crisis, deciding whether or not the results of enquiry should be made available to his actual crisis-free self? So perhaps the zealot is not admitted into ideal deliberations at all.

Consider a different problematic character: A devotee of astrology who thinks that our lives are dominated by the influence of the planets. Grant that she thinks that it is important for people with different opinions to be allowed to speak. Unlike the zealot, she can easily be imagined to participate in the first phase of ideal deliberation. After the first phase of deliberation, there is a selected set of scientific questions. Experts enumerate the different possible projects addressing these questions and report back to deliberators. Plausibly, some of the questions deemed important will be about stars and planets. Which experts should be consulted? Astronomers, certainly, but the devotee of astrology will want some astrologers included as well. If popular charlatans must

³The examples here are similar to Sobel’s discussion of an Amish person as idealized agent [Sob94, p. 801].

be consulted between the first and second phases of ideal deliberation, then the outcome might not end up being anything we would want to call *science*—well-ordered or otherwise.

This problem could be resolved directly if there were some clear demarcation criterion. Yet, as we have seen, Kitcher has no way of expelling astrology as simply not a science. (Using the list from §3.1, that would be demarcation in the first sense.) Nevertheless, he does think we can distinguish science done well from science done dreadfully (demarcation in the second sense.) Astrology gets dreadful marks when graded as a science. Its claims and arguments amount to bunkum. If the deliberating devotee asks for astrology to be addressed in the report, then experts will report that all the odds are against astrology. Yet, one might object, there is no guarantee of this outcome. If the experts in the community do believe in astrology, then their reports will list drawing up astrological charts among the possible projects for addressing questions about the stars and planets. Astrology might appear among the projects of ‘well-ordered science’ for that community. This objection is correct, I think, but it does not have any bite. Ideal deliberation and well-ordered science are just meant to separate serious enquiry from trivial enquiry (demarcation in the third sense.) If the community is in the thrall of enquiry that hardly rises above the level of dreadful science, then that community will deem such enquiry significant. Such scientific luminaries as Francis Bacon, Robert Boyle, and Isaac Newton thought alchemy was a significant and important enquiry, so the docket of well-ordered science for their communities would have included some alchemy. Today we recognize astrology and alchemy to be shoddy science at best, and so well-ordered science for us would not include either of them.

In the second phase of deliberation, then, the devotee of astrology reads reports offering a dismal assessment of astrology. In ordinary life, she dismisses scientific scepticism about astrology. Ideal deliberation requires her to take these reports seriously. We must ask what policies she would advocate, were she to take these lessons to heart. However, her dedication to astrology might be so central to her self-conception and life-defining projects that she would no longer count as *herself* if she abandoned astrology and took astronomical reports to heart. The devotee as ideal deliberator who accepts that Mercury is basically just a ball of hot rock would be a different person than she actually is.

As we are now imagining her, the devotee is not merely someone who reads the horoscopes or consults an astrologer. She is someone whose identity is so bound up with hokum that she cannot reckon with scientific results even in an imagined, ideal setting. This person, like the zealot we considered above, must be excluded from ideal deliberation. These cases show that some people simply cannot conceivably participate in ideal deliberation, because their identities are inextricably bound up with details that the conception would idealize away. For the zealot and the astrological devotee, this is because well-ordered science presumes a community of individuals at least somewhat concerned with democracy and at least willing to consider the advice of scientific experts. I do not think either of these constraints are terribly prohibitive, but we should ask whether less extreme personalities could survive idealization.

Both Sobel [Sob94, p. 803] and Rosati [Ros95, p. 309] point out that, as a psychological matter, the order and context in which information is presented effects how subjects respond to it and what they deem to be salient. For a full-information account, this means that the preferences of the ideal agent might depend on the order in which various possible lives are presented. Since there is no natural order in which to present the options, well being fails to be well-defined.

An analogous problem arises for Kitcher's process of ideal deliberation. In the first phase of deliberation, deliberators tell one another about their interests. In the course of this process their prior preferences change become "tutored" [Kit01, p. 118]. For some practical matters, this might be unproblematic. A deliberator begins with no real knowledge of the challenges facing agriculture, for example, but learns from fellow deliberators who are farmers. Although she had not even been aware of banana blight before deliberation began, she eats bananas and can recognize the blight as a problem. So her tutored preferences would include a concern for banana blight. At the end of the first phase of deliberation, she votes to have 'How is it possible to fight banana blight?' included among the list of significant questions.

This is less plausible when we consider questions which appeal to our natural curiosity. Labeling it 'natural' curiosity should not lead us to think that it is a reliable and invariant baseline. There is an indefinite list of things we could be curious about. Well-ordered science is only well-defined if there is a more-or-less determinate fact as to which concerns would pique deliberators' curiosity, but how curious people are and what they are curious about depends in subtle and complicated ways on their personal history. It is a familiar difficulty in education that exciting student curiosity depends on a complicated interplay between the presentation of information in the classroom and the temperament of students. Student curiosity may depend on factors even less relevant than the order of presentation: whether the class is before or after lunch, whether the classroom is warm or cold, whether it has windows or not, and so on. Ideal deliberation does not involve any specified meeting place, order of presentation, or seating arrangement—yet the vagaries of curiosity can depend on these things, and so ideal deliberation yields no determinate outcome. Ordinary folks could conceivably participate in ideal deliberation, unlike the zealot or the astrologer, but there would be no well-defined outcome that could define what arrangement of science would be well-ordered.

At the end of §1, I argued that natural curiosity is a legitimate and important source of scientific significance. My argument here does not undermine that. Rather, it shows that natural curiosity—although a source of *actual* significance—is too defuse to have any determinate influence in the *ideal* deliberations that Kitcher asks us to imagine. Since his account of scientific significance and his notion of well-ordered science are distinct, the argument I have given need not be fatal to either; but it does point to a tension between them.

Kitcher could insist that natural human curiosity—although important in actual enquiry—plays no rôle in constituting the ideal of well-ordered science. Without natural curiosity as a distinct source of significance, as we saw in §1,

scientific significance would only be instrumental and practical. Since the advancement of science requires basic research, such crudely utilitarian ordering would stultify science; recall §3.2. So well-ordered science could only be a normative constraint on technology, on science directed toward practical applications. Basic research would need to be left free and might be directed at whatever actual people happen to be curious about.

Alternately, Kitcher could heed advice from Rosati [Ros95, pp. 324–325] and abandon well-ordered science as a *constitutive ideal*, an ideal that tells us what constitutes the right distribution of scientific resources for a community. Instead, ideal deliberation can serve as a *regulative ideal*, an ideal that guides theoretical enquiry. Adapting Rosati’s suggestion: The ideal of well-ordered science bids us to construct theories of scientific significance that afford us a critical perspective on our lives and choices, one sensitive to relevant information and rationality, while reminding us to ground scientific priorities in a community’s concerns.

The ideal procedure need not define a specific outcome as *well-ordered* in order for us to recognize that some arrangements are poorly ordered. The notion of well-ordered science draws our attention to the factors which make configurations of the scientific community better or worse. This suffices for Kitcher’s intended applications of the ideal of well-ordered science [Kit01, esp. pp. 126–135, 192–197]. For example, we can criticize biomedical research on problems that have some mild impact on a few people if it crowds out research on problems that severely impact many people. We can criticize research which is likely to be applied in harmful ways. And we can criticize research that ignores applications or effects that are of interest to the broader community. At a more abstract level, we can ask whether centralized or decentralized funding would better promote projects that the community would deem significant.

This concession does mean that well-ordered science will not recommend a specific configuration of scientific effort. For example, it will not specify exactly how much ought to be spent on space exploration rather than particle physics. Even if the ideal procedure could be repaired so that it specified such details, our inability to implement the ideal procedure would mean that we could not know its precise recommendations. It would be metaphysically well-defined, but still epistemically indeterminate. We could still only appeal broadly to the kinds of considerations incorporated in the ideal procedure.

I have argued that Kitcher’s account relies on some questions being ones about which we are curious. This has the consequence that ideal deliberation does not yield a precise, univocal agenda for basic research. *Maximally* well-ordered science has no determinate content. Nevertheless, the ideal might draw our attention to features that make actual arrangements *more* or *less* well-ordered.

References

- [BT05] Michael Bishop and J.D. Trout. *Epistemology and the Psychology of Human Judgment*. Oxford University Press, 2005.

- [Fle79] Ludwik Fleck. *The Genesis and Development of a Scientific Fact*. University of Chicago Press, [1935] 1979. Translated by Fred Bradley and Thaddeus J. Trenn.
- [Kit85] Philip Kitcher. Good science, bad science, dreadful science, and pseudoscience. *Journal of College Science Teaching*, 14(3):168–173, December/January 1984/1985.
- [Kit93] Philip Kitcher. *The Advancement of Science*. Oxford University Press, 1993.
- [Kit01] Philip Kitcher. *Science, Truth, and Democracy*. Oxford University Press, 2001.
- [Mag07] P.D. Magnus. Distributed cognition and the task of science. *Social Studies of Science*, 37(2):297–310, April 2007.
- [Pop65] Karl R. Popper. *Conjectures and Refutations*. Harper & Row, 1963,1965.
- [Ros95] Connie S. Rosati. Persons, perspectives, and full information accounts of the good. *Ethics*, 105(2):296–325, January 1995.
- [Sim06] Jeremy Simon. The proper ends of science: Philip Kitcher, science, and the good. *Philosophy of Science*, 73(2):194–214, April 2006.
- [Sob94] David Sobel. Full information accounts of well-being. *Ethics*, 104(4):784–810, July 1994.