



Organizing Inquiry

HOW SHOULD INQUIRY BE ORGANIZED so as to fulfill its proper function? At the dawn of modern science, both Bacon and Descartes attempted to address the question. They saw the history of attempts to understand the natural world as dominated for two millennia by the faulty doctrines and methods of Aristotelianism, and resolved that such stagnation should never occur again. Offering a common diagnosis, both saw the need for a method of discovery future inquirers could follow, whatever questions they were concerned to resolve, and a method of justification that would specify exactly when answers should be accepted; in focusing on the latter project, they hoped to forestall the premature adoption of views that would distort future research. Of course, their proposals for answering questions about discovery and justification differed. Descartes emphasized the virtues of rational analysis, both to direct the mind in discovering solutions to problems and to frame the options experiments might discriminate. By contrast, Bacon stressed the importance of unprejudiced observation and the patient accumulation of empirical data.

Neither of these proposals wins contemporary acceptance, although both have left their mark. Yet Descartes and Bacon continue to shape—and to limit—contemporary discussions of “proper” inquiry. From the seventeenth century to the present, their main questions—What is the right method of discovery? What is the logic of justification?—have dominated reflections about the sciences. In the twentieth century, the first question fell out of favor as an influential group of philosophers argued that there was no general method of scientific discovery. Recently, emphasis on the serendipity of discovery has given way to a more refined appreciation of the methodical ways in which investigations typ-

ically proceed, and scholars have turned to new formal techniques to resurrect the concept of a method of discovery.

For the past seventy years, however, the central normative question about science has concerned the logic of justification. How can we identify the conditions under which statements—particularly universal generalizations and statements about entities remote from observation—are properly accepted? The dominance of this question suggests a way of thinking about inquiry. To a first approximation, philosophers of science have written as if inquiry would fulfill its proper function provided all those who engage in it live up to the standards set by an ideal methodology, accepting only those statements that are genuinely justified in light of a correct logic of justification.

It is remarkable just how much this conception leaves out. Besides issues about methods of discovery (now regaining some respectability), there are other obvious omissions. We typically evaluate an enterprise by considering how well it achieves, or can be expected to achieve, its aims. As some philosophers have seen, an assessment of the sciences that ignores the fact that we risk error in hopes of gaining information cannot be right. Unless the logic of justification goes beyond the popular project of trying to specify the conditions under which particular types of conclusions are likely to be true, it will fail to see that what matters in well-organized inquiry is the standard of gaining significant truth. Even when discussions of the sciences appreciate the elementary point that the most reliable means of ensuring that the statements one accepts are not false is to accept nothing whatsoever, the resultant approach to issues of justification frequently supposes there is some context-independent good thing for which it's worth risking error. Philosophers conjure up general measures of content, or explanatory power, or information, and then puzzle about how, in general, we should weigh the chance of being mistaken against the abstract benefits they favor. But what inquiry seeks is significant truth, and, as we've seen, significance is thoroughly context-dependent. There is, then, no general problem of trading significance against chances of truth, and, accordingly, no general solution. The right strategy is to frame the problem, at the start, in terms of how well-functioning inquiry promotes the acquisition of significant truth, explicitly acknowledging the variability of standards of significance.

Even more obviously limiting is the traditional focus on individuals. Unlike Descartes, Bacon already recognized that inquiry is a collective affair, and his ideas about the social character of proper inquiry were influential in the formation of the Royal Society. (Some of those ideas will occupy us briefly in chapter 11.) The community of inquirers cannot simply be viewed as a magnified version of the individual. Indeed, when we appreciate how changes in the sciences stem from the combined efforts of individuals, we may be led to adjust our views concerning how single investigators should behave. Suppose methodology in the individualistic tradition has succeeded in picking out the right rules

for accepting hypotheses (or assigning probabilities to hypotheses) on the basis of evidence. A community of scientists, each of whom follows these rules and all of whom have access to exactly the same evidence will be homogeneous in its opinions. Yet, possibly for Millian reasons (canvassed in the last chapter), one might doubt if homogeneity is the best epistemic policy. Maybe the collective attainment of truth (or significant truth) would be advanced if some members of the community were to disagree. So an account of what individuals should do won't automatically tell us when collective inquiry is working properly.

The discussion of the last chapter identified a third major omission. Suppose we broadened the traditional perspective of concentrating on the attainment of truth by individual inquirers, considering the collective pursuit of significant truth (where significance is understood in the way suggested in chapter 6). Even that would not take into account the possibility that the search for significance can conflict with other important values. To assess the proper functioning of scientific inquiry we must consider if collective research is organized in a way to promote our collective values in the most encompassing sense. To enclose scientific research so that the standard to which it is held is the collective acquisition of epistemically significant truth is to acquiesce in the myth of purity, and, as the last chapter tried to show, that approach will allow projects that are morally suspect.

Traditional philosophy of science has offered welcome clarifications of important concepts and principles, and we should value its insights. But, as I've complained, the dominant perspective has provided a very narrow normative perspective. Perhaps that is why the classical view of science has attracted criticism, generating the polar pair of unacceptable images described in chapter 1. What follows is an attempt to put back some of the considerations that have been slighted.

When we lived in California, Bertie, the much-loved family dog, would come and go freely between the house and the fenced backyard. One afternoon, two or three hours before sunset, when coyotes become active, someone came to read the meter and left the gate to the yard open. When Bertie went out, he spotted the chance of adventure and set off to explore. Discovering his absence we had to formulate a plan for finding him (quickly, since he's too small to take on coyotes). We normally walked him along one of two routes, which we assumed were the places he'd be most likely to go. One route we estimated to be a more probable path than the other. There were four of us. How to proceed?

The problem, of course, is underdescribed, for I haven't specified the probabilities of finding him along either route, the chances of his doing something different, the number of people needed to scout each route thoroughly, and so forth. But one point ought to be clear. If the four of us had stayed together, even if we went over the most likely route, that would probably have been a very bad

strategy. What we'd like to have known, of course, provided it could be identified with negligible costs of time to calculate it, is the strategy that would maximize our chances of finding Bertie within the next two hours. We probably didn't manage that. We did avoid the really bad strategy of staying together, though, and it's quite possible that our approach yielded a probability of success close to the maximum value. (In any event, we found Bertie quite safe, and brought him home before dusk.)

This homely example not only shows clearly how a well-organized inquiry can require investigators to do different things but also suggests a criterion for appraising strategies for inquiry. Given the information available to us, we want our efforts to be organized so as to maximize our chances of attaining our goal. Similar problems arise in scientific contexts, and, at least *prima facie*, they are associated with a similar criterion of success and sometimes allow explicit analysis. Suppose, for example, that a community of chemists hopes to fathom the structure of a very important molecule (VIM). Two methods are available. One is assessed as highly likely to succeed if pursued with sufficient vigor, and as not necessarily slow. The other is risky: it might deliver the answer quickly or might be quite inadequate to the task. If the community contains enough members, then, just as in the case of the search for Bertie, the best strategy is to divide the labor.

Our dog-finding efforts began with a discussion whose upshot was to designate an agreed-upon role for each of us. We might have reached the same division of effort in a different way: imagine the family returned home sequentially, and each person left a note explaining the situation and the searches in progress so far. Scientists, of course, do not typically assemble to agree on roles for attacking a problem (although this type of explicit cooperation is not unprecedented, especially in conditions of war). But it is not hard to see how a community might mimic the second way of distributing the effort among its members: as newcomers enter the field they take stock of the ways in which central problems are currently being pursued and adjust their own research to maximize the chance that the community will succeed. Yet this requires a type of high-mindedness that may be rare. Can we expect inquirers to be devoted to the common project of advancing knowledge in the same way the family is dedicated to reclaiming its lost dog? That seems implausible. Scientists do not often declare in public that they intend to pursue some unpromising line of inquiry because their doing so will advance the community project of solving a problem—and I doubt if private commitments of this kind are more common. But this may not matter.

For there are other ways in which a community can reach a satisfactory distribution of its research efforts. Suppose everybody in the group of chemists recognizes that whoever fathoms the structure of VIM will receive great kudos, maybe even win a much-coveted prize. For each individual scientist the desire

to be the one to solve the problem (and perhaps to win the prize) is the dominant motivation. Now we can imagine some of the chemists surveying the current distribution of effort, and, believing that there is too much competition among the followers of one method, switching to the other. The latter may be more risky, but, because of the lower competition, the scientist's chance of winning the race for the solution would go up. So the community can avoid the disastrous state of homogeneity and even come close to the optimum.

The most important moral of my story so far is the clear delineation of a possibility: we can have a community well-designed for the attainment of epistemic goals in which social institutions we might have viewed as irrelevant, even contrary, to those goals (attributions of credit, prizes) are tailored to motivations typically viewed as antithetical to the goals (desire for personal glory). We might hope to go further. If we could identify a recurrent set of scientific predicaments, then we might be able to embed this particular analysis in a broader study, one that would show the effects of various types of social arrangements, given prevalent human motivations, and pick out the package that would best promote the community's attainment of truth. Finally, we could hope to expand the set of goals beyond the epistemic, repeating the same style of analysis with respect to this broader conception of value. So we'd arrive at an explicit account of the proper functioning of inquiry.

This is far too optimistic. The examples I've offered were tractable because it was possible to specify a *local* epistemic goal. Our family wanted to find the dog before dusk, the chemists wanted to reach a state in which one of them recognized the structure of VIM. Trouble sets in when we try to think more globally.

If there's no context-independent notion of significance, then any attempt to develop a measure of epistemic value—the quantity inquiry is to be designed to maximize—will embody current ideas about epistemic significance. As the notion of epistemic significance evolves between the time at which the possible achievements are assessed and the time at which new knowledge is gained—almost certainly with important surprises—the measure assigned retrospectively may differ quite radically from that proposed in advance. Perhaps however we tried to organize inquiry we'd be doomed to regret our choices. Or maybe almost any decisions would generate future states in which we were happy with what had been attained.

Even were we able to offer definite specifications of our epistemic goals, posing and solving the pertinent optimization problem would still be difficult. Focusing on a recurrent decision situation, such as the one in which the community has two available methods for tackling an important problem, might allow us to show that a particular social institution, like the public awarding of credit to the first solver, would lead to a good distribution of effort. But this success might be offset by comparable failures if the circumstances were slightly differ-

ent, and without detailed knowledge of the likely frequencies with which situations of various kinds arise, there would be no basis for viewing the institution in question as beneficial. Worse still, without analyzing the impact of the institution across the total range of contexts in which it might play a causal role, it's impossible to discern its overall effect. The dangers of local optimization are familiar from evolutionary biology (and other areas of inquiry). Moreover, as the last chapter already suggested, the pressure to earn public acclaim, which might prove valuable in promoting cognitive diversity, can also have much less salutary consequences—for example leading researchers to leap to conclusions that resonate with popular prejudices.

These points underscore the difficulty of elaborating a fully general social methodology. We would like the sciences to be well organized so we could learn as much as possible about the world as efficiently as possible, but it may not be possible to formulate any serious proposals about optimal organization until our knowledge of nature is vastly richer than it is; perhaps in order to identify the recurrent predicaments inquirers face and to estimate reliably the chances of success of various methods, we already need to know most of the things we want our inquiries to disclose. For all that, reflections on social methodology are far from useless, in that they can show us how to avoid really bad strategies and sometimes reveal to us problems that we wouldn't otherwise have seen or benign consequences where we might have harbored suspicions. The best is too much to hope for, but we may aspire to improve our situation.

Consider a sequence of questions. First: what are good policies for individuals to adopt if they want to learn epistemically significant truths? Second: what are good ways for communities to organize their efforts if they want to promote the collective acquisition of epistemically significant truth? Third: what are good ways for communities to organize inquiry if they want to promote their collective values (including, but normally not exhausted by, the acquisition of epistemically significant truth)? The first of these questions is closest to the philosophical tradition of trying to clarify the methods of inquiry, and, because of the context-dependence of the notion of epistemic significance it is already hard to treat both formally and generally. We can look for formal approaches to special instances of it or for a more general approach that aims at a qualitative characterization of promising strategies and pitfalls. Matters are similar, as we have just seen, with respect to the second question. On the face of it, the third question—the issue with which we really ought to be concerned—appears even harder.

For in the first two cases we think we know what we are trying to attain: epistemically significant truth, either for the individual or for the community. When we broaden the perspective to encompass our “collective values,” the goal becomes much more nebulous. How does this goal relate to the actual wishes

and preferences of the members of a society? How are we to integrate the preferences of different people? Can very different types of value be brought under a single measure? Is it even possible to undertake the local analyses that are available in the more limited epistemic projects? These are all serious concerns, and they will require our attention.

There are two very obvious ways of approaching the problems of the last paragraph. One is to suppose that whatever preferences people actually have, whatever they think about what it would be good for them to pursue, either individually or collectively, some ends are objectively worthy and there are objective relations among these ends. Call this general perspective *objectivism about values* (*objectivism*, for short). Objectivism can concede that there are many different kinds of values, some of them practical, some epistemic, some present, some future. It may even countenance human diversity, supposing some packages of good things are better for some people, different packages for other people. But objectivists think there's a right way of trading the epistemic against the practical—and, more generally, a right way of trading various different types of values for one another—a right way of balancing the present against the future, and a right way of integrating the objective interests of different individuals.

Here's an example (not, I hope, a particularly implausible one) of an objectivist position. It is objectively good for people to develop their talents and to enjoy as much liberty as possible in deciding on and pursuing their chosen goals. With respect to each assignment of various freedoms and resources (food, shelter, education, and so forth) at the different stages of the person's life there is an objective level of value for that person (not necessarily the same for all people). Further, there's an objective way of aggregating the levels of value for each person into a measure of the collective value achieved: let's suppose it consists in adding the individual levels of welfare subject to a function that discounts them if there are pronounced inequalities among individuals. So, with respect to different social arrangements and endeavors, there will be an overall measure of expected value, one that depends objectively on the expected contributions of those arrangements and endeavors to the levels of value achieved by individuals at the successive stages of their lives, that integrates those levels of value across the person's life-span and that aggregates the expectations for the population by adding the individual expectations, subject to the discounting that represents the costs of inequality. Relative to a specified collection of other social arrangements and endeavors, we could now say that a particular way of conducting inquiry within a social context fulfills the proper function of inquiry just in case it furnishes the maximal level of collective expected value attainable within that social context. (Plainly, we'd be happy to settle for ways of conducting inquiry that failed to maximize but that came relatively close.) In other words, inquiry functions well when it can be expected to lead to states in which people have, at the different stages of their lives, the resources and free-

doms they need (according to the basic account of what's valuable for the individual), and when the total level of lifelong value, calculated across the society, is high and not marked by large inequalities.

Perhaps others will see how to articulate a conception of this general sort. I don't. My doubts rest on the difficulty of divorcing what is good for a person from that person's own reflective preferences and the kindred problem of ignoring personal preferences in understanding the ways in which different distributions of goods across the stages of a person's life yield overall value. Further, I think that the general problem of understanding how to aggregate individual levels of well-being into a measure of collective welfare, in the ways objectivists propose, is extremely difficult. So I propose something more modest. Individual preferences should form the basis for our understanding of the personal good that inquiry (among other social institutions) is to promote. In moving from the individual to the measurement of value for the society, we should explicitly limit our discussions to societies that honor certain democratic ideals. Hence my approach to the fundamental question, "What is the collective good that inquiry should promote?" will start from a subjectivist view of individual value (using personal preferences as the basis for an account of a person's welfare) and will relate the individual good to the collective good within a framework in which democratic ideals are taken for granted.

The task for the next chapter is to deliver an answer to the fundamental question. I don't pretend that what I offer is the unique best answer—or even that the approach to which I've just committed myself is the preferred way of tackling the question. As I've indicated throughout this chapter, traditional discussions of scientific inquiry adopt a much narrower conception of the standard against which the proper functioning of inquiry should be assessed. The obvious difficulty in attempting to think more broadly about the role of the sciences within society is the lack of any clear conception of what the wider task of inquiry might be. In providing *an* answer I hope to respond to skeptical concerns that there is no coherent wider conception, and also to delineate the area in which a plausible answer may be taken to lie. It would be folly, however, to pretend that I have given convincing reasons for distinguishing my particular proposal from rival members of the family, or that my account of the details is likely to be correct. Others, perhaps, may be able to improve on it.



Well-Ordered Science

THERE'S A VERY SIMPLE WAY TO DEVELOP the idea that properly functioning inquiry—well-ordered science—should satisfy the preferences of the citizens in the society in which it is practiced. Projects should be pursued just in case they would be favored by a majority vote. Call this “vulgar democracy.”

Vulgar democracy doesn't require actual voting. Rather it offers a standard against which we can assess rival schemes for deciding which endeavors are to be undertaken. The idea of calling together the citizenry to cast ballots on each occasion of decision is evidently absurd, but vulgar democracy is only committed to seeking social arrangements (committees of representatives, for example) that we might expect to do well at mimicking the outcomes of the expression of individual preference. Nevertheless, as its name suggests, vulgar democracy is a very bad ideal.

The most obvious deficiency, of course, lies in the fact that people's preferences are often based on impulse or ignorance and thus diverge from favoring what would actually be good for them. Only a moment's reflection is needed to see that the most likely consequence of holding inquiry to the standard of vulgar democracy would be a tyranny of the ignorant, a state in which projects with epistemic significance would often be dismissed, perceptions of short-term benefits would dominate, and resources would be likely to be channeled toward a few “hot topics.” Because these consequences plainly diverge from the promotion of collective well-being, vulgar democracy is a bad answer to our question.

How can we do better? I offer a homely analogy. Imagine a family with a free evening and a strong shared wish to spend it together in some form of entertainment. They begin with a number of different proposals, explaining to one

another their preferences, the strength of the preferences, the considerations that move them. Each family member learns new things about the character of the various options, and each learns how the others view the possibilities. Nobody wants to do anything that any of the others regards as too unattractive, and they end up with a plan that reflects their collective wishes. Those collective wishes can't be conceived as the wishes that would emerge from a simple vote in the initial state of mutual ignorance; rather, they are the wishes that would be produced by a more intricate negotiation.

In order to use the analogy to articulate an ideal procedure for well-ordered science, we need a clear understanding of the kinds of decisions that will be needed. Let's conceive of ideal inquiry as divided into three phases. At the first phase, decisions are made to commit resources, such as investigators and equipment, in particular amounts to particular projects. The second phase pursues those projects in the most efficient way, subject to moral constraints that rule out certain physically possible options. At the third phase, the results of the various investigations are translated into practical consequences. So there are three different decisions to be made: How are resources initially to be assigned to projects? What are the constraints on morally permissible investigation? How are the results of the investigation to be applied? As we'll see, the first and the third decisions can be approached within a similar framework.

Begin with the first decision. I envisage individuals with different initial preferences coming together, like the family, to discuss the available courses for inquiry to pursue. The first thing to recognize is that, unlike the family, they are likely to begin from a very partial understanding of the possibilities. An obvious remedy for their ignorance is to insist on transmitting information so that each deliberator becomes aware of the significance, epistemic and practical, attaching to potential lines of inquiry. Ideal deliberation must involve presenting the structure of significance graphs, where the multiform sources of significance are revealed. Once this has been accomplished, the deliberators revise their own initial preferences to accommodate the new information. Specifically, I imagine that each considers how possible inquiries might bear on goals that were antecedently adopted. The product of the consideration is a collection of lists of outcomes the deliberators would like scientific inquiry to promote, coupled with some index measuring how intensely they desire those outcomes. Personal preferences have given way to *tutored* personal preferences.

The next step is for ideal deliberators to imitate the imaginary discussion of the family. They exchange their tutored personal preferences, explaining why they want particular outcomes to particular degrees and listening to the explanations given by others. In this process, I assume each is moved by respect for the preferences of others and aims to arrive at a consensus list in which none of the others is substantially underrepresented. The deliberators are committed to seeing the others as having, like themselves, a claim to realize their aspirations,

and thus to take seriously the others' descriptions of their preferences and predicaments and the rationales they provide for choosing as they do. Ideal deliberators thus recognize that they are engaged in a long-term sequence of interactions with people whose situations and fundamental wishes may be quite different from their own, and that such people cannot be expected to sacrifice their desires to the preferences of others.

At the end of this exchange, the preferences of each ideal deliberator are again modified, this time to absorb their recognition of the needs of others. The next step is for them to attempt to draw up a list that represents their priorities concerning the outcomes to which inquiry might prove relevant. One possibility is that there is consensus. After coming to understand both the current characteristics of significance graphs and the tutored preferences of other deliberators, each party formulates the same list, assigning exactly the same value to each outcome. If this is so, then the resulting list expresses the collective preferences, and no further accommodation is needed. A second possibility is that some deliberators favor different lists but each is prepared to accept a set of lists as fair and the intersection of the sets is nonempty. Under these circumstances, if the intersection contains a unique member, then that expresses the collective preferences. If not, then the ideal deliberators must decide by vote which of the lists in the intersection is to be preferred. Finally, if the intersection of the sets of lists deliberators accept as fair turns out to be empty, collective preferences are determined by vote on all candidates drawn from the union of these sets of lists.¹

At this point, our deliberators have formulated the issues they'd like inquiry to address and have indicated the relative weight to be given to these issues. Their formulation, tutored by a clear understanding of the sources of significance for scientific endeavors already completed as well as those that might now be undertaken, can be expected to recognize possibilities for satisfying curiosity as well as opportunities for practical intervention, long-term benefits as well as immediate payoffs. The next step is to assess the possibilities that particular scientific ventures might deliver what the ideal deliberators collectively want. Given a potential line of inquiry that might bear on some items on the collective list, we require an estimate of the chances the desired outcomes will be delivered, and it's appropriate to turn at this point to groups of experts. How are the experts to be identified? I suppose that the ideal deliberators can pick out a

1. It's clear from the extensive literature on social choice theory stemming from Kenneth Arrow's famous impossibility theorem that any procedure like the one described here may lead to counterintuitive conclusions under certain hypothetical circumstances. I don't assume that my suggestion is immune to such problems, but I hope they only arise in sufficiently *recherché* situations to make the standard I attempt to explicate appropriate for the purposes of characterizing well-ordered science. For lucid discussions of Arrow's theorem and its significance, see Amartya Sen, *Collective Choice and Social Welfare* (San Francisco: Holden Day, 1970).

group of people to whom they defer on scientific matters generally, that this group defers to a particular subgroup with respect to questions in a particular field, that that subgroup defers to a particular sub-subgroup with respect to questions in a particular subfield, and so forth. Further, it's assumed the experts identified are disinterested—or that any members of a group whose personal preferences would be affected by the project under scrutiny are disqualified from participating in the process. If matters are completely straightforward, there'll be consensus (or virtual consensus) at each stage concerning the appropriate people to consult, and these people will agree on exact probabilities with respect to the outcomes of research. In that case, the output of the search for probabilities is just the collection of chances assigned by the groups singled out at the ends of the various chains of deference.

Complications can arise in any of three ways.² First, there may be disagreement about to whom deference is warranted. Second, the experts may be unable to do more than assign a range of probabilities, possibly even a wide range. Third, the experts may be divided on which probabilities (or ranges of probabilities) should be assigned. I deal with all these complications in the same general way, namely through being inclusive. If any of them arises, the output of the search for probabilities is no longer a single set of values, but an explicit record of the verdicts offered by different groups, coupled with the extent to which those groups are supported by deliberators who have full information on the current state of inquiry and the past performances of the groups in question. Hence, instead of a simple judgment that the probability a scientific project will yield a particular desired outcome is a certain definite value, we may have a more complex report to the effect that there are several groups of people viewed as experts, that deliberators in full command of the track records of the various groups select particular groups with particular frequencies, that the groups divide in their judgments in specified proportions, and that these judgments assign specified ranges of probability values.³

At the next stage, we suppose a disinterested arbitrator uses the information about probabilities just derived, together with the collective wish list, to draw up possible agendas for inquiry. The arbitrator begins by identifying potential levels of investment in inquiry (possibly an infinite number of them). With respect to each level, the task is to pick out either a single assignment of resources to scientific projects best suited to the advancement of the deliberators' collective wishes, given the information about probabilities, or a set of such assignments

2. One might envisage a fourth. Mightn't it turn out that all the people most competent to assess a venture have an interest in the outcome? In the real world, this is an obvious possibility. But, for our purposes, it's enough that there are ideal experts, who share all the knowledge of actual judges but have no personal stakes in the line of inquiry.

3. I am indebted to Stephanie Rupy for a discussion that brought home to me the potential complexities of the appeal to experts in this context.

representing rival ways of proceeding that cannot be decisively ranked with respect to one another. In the simplest case, when the arbitrator is given point probability values, the decision procedure can be specified precisely: with respect to each budgetary level, one identifies the set of possible distributions of resources among scientific projects compatible with the moral constraints on which the ideal deliberators agree, and picks from this set the option (or set of options) yielding maximal expected utility, where the utilities are generated from the collective wish list and the probabilities obtained from the experts. (Although the process of deciding on moral constraints on inquiry hasn't yet been considered, we'll see below that it's quite independent of the decision to be made by the arbitrator.) If there are disagreements about who the experts are, disagreements among the experts, or latitude in the responsible assignment of probabilities, then the arbitrator must proceed by considering the distributions of resources that would meet both budgetary and moral constraints, subject to different choices for probability values, picking that set of distributions that best fits the views of the majority of those who are most often considered experts.⁴

The last stage of the process consists in a judgment by the ideal deliberators of the appropriate budgetary level and the research agenda to be followed at that budgetary level. Perhaps there is consensus among the ideal deliberators about which level of support for inquiry should be preferred, and perhaps the arbitrator assigns a single distribution of resources among lines of inquiry at that particular level. If that is not so, then the final resolution must be reached by majority vote. The result (whether it comes from consensus or voting) is the course of inquiry that best reflects the wishes of the community the ideal deliberators represent.

As already noted, the procedure just outlined presupposes some agreed-upon constraints on inquiry. I suppose these take a normal form, always stemming from the recognition that a particular way of pursuing inquiry would violate the rights of some individual or group. Any disagreements among the ideal deliberators are thus traceable to different conceptions of rights, perhaps to rival ideas about which individuals have rights, perhaps to alternative attributions of context-independent rights, perhaps to divergence about which rights that normally accrue can be suspended in the context of inquiry. (Here I have in mind the idea that, as with war or politics, a distinctive "public morality" might apply to scientific investigations.) Imagine, then, ideal deliberators exchange information about the putative bearers of rights and the strategies of inquiry those rights are supposed to debar, and they attempt to defend their conceptions by appeal to principles. In response to such interchanges, they

4. There are several ways of treating this problem formally. For my present purposes, I suppose that the arbitrator proceeds in the way we take disinterested, intelligent, and reasonable people to make their judgments when confronting divergent "expert" opinions.

modify their ideas about which moral constraints are appropriate. At the end of the process, they may find themselves in one of three situations: consensus, no consensus but agreement on one view as the fair representation of different points of view, no consensus and disagreement about how to represent the collective attitude. In the last case, once again, the issue is settled by majority vote.

No doubt there are possible societies—very likely actual societies—in which even ideal deliberation would end in irreconcilable disagreement. People might well fail to reach accord on permissible experimental procedures because they harbor different views about the moral status of animals or fetuses. Even when disagreements persist, the awareness of the sources of divergence in opinion, among parties committed to mutual respect, can affect the ways in which decisions about inquiry are made. So, if it's known that a minority favors more stringent constraints than the majority demands, that fact can be weighed in the choices of a distribution of resources to scientific endeavors, and even in the formulation of the collective wish list. If the deliberators see that a particular outcome could only be achieved by using methods the minority would view as impermissible, they may respond by decreasing the value assigned to the outcome or by committing extra resources so that a strategy acceptable to all may be pursued.

Let's now turn to the third phase of inquiry, the translation of results into applications. Note first that some of the achievements are likely to yield epistemic benefits, to contribute to answering questions that spring from human curiosity; the only issue that arises about these achievements is how to disseminate them both among the scientists to be educated in the next generation and among the wider public. With respect to the achievements that have practical significance, it's natural to think the decision has already been taken in the framing of the research agenda, and the appropriate procedure is simply to follow the policies instituted there. That, however, would be to overlook the possibility that changes in the significance graphs, the consequences of the research that has been undertaken, may modify judgments of relative significance. So I suggest that the ideal procedure at this stage is to mimic the decision-making process of the first phase, with the emphasis now on garnering specified practical benefits. In light of the new knowledge, our ideal deliberators revise their collective wish list, the experts update their views about the probabilities of satisfying various wishes, the arbitrator offers a set of options for gaining particular benefits at various levels of cost, and the ideal deliberators pick a policy for making use of the new information. We may think of that policy as reflecting their newly tutored collective wishes.

The question with which we began—Under what conditions is the science of a society well-ordered?—can now be answered. For *perfectly* well-ordered science we require that there be institutions governing the practice of inquiry within the society that *invariably* lead to investigations that *coincide* in three re-

spects with the judgments of ideal deliberators, representative of the distribution of viewpoints in the society. First, at the stage of agenda-setting, the assignment of resources to projects is exactly the one that would be chosen though the process of ideal deliberation I have described. Second, in the pursuit of the investigations, the strategies adopted are those which are maximally efficient among the set that accords with the moral constraints the ideal deliberators would collectively choose. Third, in the translation of results of inquiry into applications, the policy followed is just the one that would be recommended by ideal deliberators who underwent the process described.

Perfectly well-ordered science is surely too much to hope for. What we would like is, I suggest, a feasible approximation. In setting up structures for the funding and oversight of research, contemporary affluent democracies try, in rather haphazard ways, to come reasonably close to an important ideal. I propose the notion of perfectly well-ordered science as the ideal at which they are aiming.⁵

Before proceeding further, it will be useful to leaven the rather abstract presentation of my answer with some clarifications and illustrations. First, just as I absolved vulgar democracy of the charge that it required actual voting on scientific projects, so too 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 at the points I've indicated. Quite probably, setting up a vast populationwide discussion that mimicked the ideal procedure would be an extraordinarily bad idea, precisely because transactions among nonideal agents are both imperfect and costly. So the challenge is to find institutions that generate roughly the right results, even though we have no ideal deliberators to make the instantaneous decisions we hope to replicate.

Second, like vulgar democracy, the ideal procedure attempts to incorporate the views of every member of the pertinent society. It's an open question as to whether the collection of ideal deliberators contains distinct idealized representatives for each citizen or whether we can assume that people divide into groups whose members are sufficiently similar that they can be represented en bloc. In the latter case, we can suppose that the ideal deliberators proportionally represent the groups with shared perspectives (that is, if one group has twice as many actual members, then it has twice as many ideal representatives). The procedure I've outlined is indifferent as to whether we suppose one-to-one representation or proportional representation of groups with a common perspective.

A third obvious worry about my ideal is its dependence on the values, possibly quite erroneous, of particular societies. This is a direct consequence of my

5. As I suggest in the next chapter, much of the literature on science policy has been handicapped by any clear recognition of what the intended ideal might be.

decision, at the end of the last chapter, to retreat from giving an objectivist answer. Yet, it's natural to think that the only acceptable normative perspective is one that doesn't make science hostage to current beliefs about what things are worth pursuing. My conception of well-ordered science can easily be seen as implicitly recommending that inquiry ought not to lead us to improved views about what is valuable, a prospect many thinkers have taken to be important and liberating.

In response, various things need to be pointed out. First, it's a familiar fact that we can often appraise an activity from either of two perspectives, one that probes its actual success and the other that considers whether an agent did as well as he could, given the limitations of his view. My normative notion of well-ordered science belongs to the latter family. Moreover, as is apparent from my description, the construction of the collective wishes from the individual preferences involves, from the beginning, reflective transformation of those preferences, so there's no danger of holding inquiry hostage to capricious and irrational desires.

Further, if there is indeed a defensible version of objectivism, then it shouldn't be hard to see how to move from my conception of well-ordered science to something stronger. Let's say that the science of a society is well-ordered in the weak sense if it conforms to my criterion, well-ordered in the strong sense if, in addition, its collective values conform to the objective good. When science is well-ordered in the weak sense but not in the strong, then something is amiss. The error is, quite properly, traced to a failure to recognize what's objectively worth doing. But, to recapitulate the point of the last paragraph, the society is still organizing inquiry as well as could be expected, given its limited understanding of the objective good.

Nonetheless, the worry is quite right to spot a danger of conservatism if the only standard for appraising the practice of inquiry is by appeal to my notion of well-ordered science. Conformity to actual values might foreclose investigations that would reveal prejudices and transform aspirations. We'll confront this issue later, in chapters 12 and 13.

A different criticism of my relativization to societies charges that my treatment implicitly focuses on the wrong group. I have written throughout of the practice of science "within a society" and have conceived of the decisions of well-ordered science as representing the wishes of members of that society. The natural interpretation is to suppose that the societies I have in mind are the affluent democracies in which most scientific research is done, and that, for a particular democracy, well-ordered science requires conformity to the idealized wishes of the citizens of that democracy. An obvious defense of that way of proceeding is to invoke the idea that the resources to be committed are those of the society in question, and ultimately of its citizens, so that the citizens have a special right to say how such resources ought to be distributed. Both the interpre-

tation and defense are vulnerable to charges of myopia. Can we really overlook the fact that the kinds of inquiries undertaken have an effect on the well-being of billions of people outside the society? A decision to pursue a line of inquiry that could help in treating diabetes (say) might foreclose opportunities for malarial research. Contemplating examples like this, one can easily conclude that the appropriate group to be represented in the ideal deliberation isn't the citizenry of a particular society (for example some rich democracy) but the entire human species.

We should distinguish a number of positions. One extreme takes the form of the ideal deliberation to be a process that represents only the citizenry of a particular democratic society from which resources will be drawn to pursue inquiry and requires the ideal deliberators to focus only on the needs and aspirations of other members of the society—like the imaginary family, they restrict their attention to one another. At the opposite pole, we can envisage a similar process involving representatives of all members of our species. Two intermediate views are also worth considering. One would continue to restrict membership in the group of ideal deliberators to representatives of the citizens of the society which is to support the inquiry, but would require them to take the preliminary step of acquainting themselves with the needs of people who belong to different societies. In effect, there would be an extra step in the process of ideal deliberation, so that the exchange of views about priorities would include representatives from groups not represented at other stages of the decision. A second possibility is to broaden the class of deliberators to include representatives of other groups whose preferences and opinions count at all stages of the process.

The issues here are complex, and I shall be brief and blunt. Neither of the polar positions seems defensible. Although one might argue that the decisions about inquiry should be left entirely to those who will support it—so that the citizens of an affluent democracy have the right to declare how their funds and talents should be employed—that surely doesn't entail that it's permissible for them to ignore the plight of outsiders, especially when we reflect that their ability to dedicate some of their resources to inquiry may stem from accidents of their society's history or even from past injustices towards those whose priorities are now being excluded. On the other hand, the ideal of a deliberation among parties who cannot be expected to share common democratic ideals or to view one another as participating in joint enterprises looks hopeless: either the tutoring will be entirely inadequate to engender common understanding, or it will effectively transform the ideal deliberators so that their ability to represent the outsiders is highly suspect. Thus I favor one of the intermediate options, and suggest the better choice is that which restricts membership to representatives of the citizens but requires them to become acquainted with the preferences of others, on the grounds that this accom-

modates a wider spectrum of viewpoints but allows for a shared democratic framework among the deliberators. We might note that, in my original formulation, the ideal deliberators are supposed to take thought for future generations and to consider the implications for members of their society (yet unborn) who will reap the consequences of lines of inquiry now set in motion; by the same token, we might think that the deliberators can and should have an understanding of the consequences for people outside their own society, and it should figure in their deliberations even though the outsiders (like the unborn) do not vote. Clearly, however, much more needs to be said, and I don't pretend to have provided a compelling defense of my preferred way of developing the normative standard.

The last objection I'll consider here is the criticism that the standard for well-ordered science is toothless. Virtually anything, it may be suggested, can be approved. But although well-ordered science might plausibly countenance a number of courses of inquiry, it surely will not allow all. Investigations that bring large benefits to one segment of the society while harming others will not measure up to the standard. Neglect of practical concerns in favor of a sole focus on the epistemic will typically fall short—as will an exclusive concern with the practical (except, perhaps, when practical needs are extremely urgent). The more interesting challenge is to try to understand how far the current practice of the sciences lies from well-ordered science. Can we tell how well (or how badly) we are doing? Can we use the ideal of well-ordered science to improve our situation? To these questions I now turn.

Had we but world enough and time, we could follow a direct approach to designing an ideally well-ordered science. We would review all possible institutions, all possible contexts over which they might operate, formulate an optimization problem, and solve it. This is an impossible dream. We have no realistic prospects of canvassing social institutions and reviewing their entire range of effects across all the situations in which they might be employed—indeed, I think it likely that, in order to assess those effects, we'd already have to resolve many of the issues for which we hope to design inquiry. But we can still scrutinize our own practice from the perspective supplied by the standard.

Consider, for example, the actual ways in which research agendas are constructed. The channeling of research effort is subject to pressures from a largely uninformed public, from a competitive interaction among technological enterprises that may represent only a tiny fraction of the population, and from scientists who are concerned to study problems of very particular kinds or to use the instruments and forms of expertise that are at hand. Actual deliberations (as we'll see shortly) often involve agents who depart from the ideal in two different ways: potential consumers who have a highly incomplete understanding of the range of options and of their consequences, and inquirers who are strongly mo-

tivated to present research projects in ways they think will appeal to a much broader public.⁶ So, in this context at least, there are grounds for pessimism.

Contrast the appraisal of agenda-setting with that of the moral constraining of inquiry. Here, from the relatively mundane and unperturbing potential transgressions (scientific fraud, plagiarism, and so forth) to the truly disturbing cases (experiments that damage human subjects without their consent), we can point to a core set of moral constraints that are close to being universally acknowledged, and serious attempts, at least, are made to ensure that researchers abide by them. Partly because of terrible abuses of inquiry in the past—the examples of the Nazi doctors and the Tuskegee experiment—widely shared views about human rights have inspired systematic oversight of experimentation involving human subjects, which quite deliberately involves people with a variety of perspectives. Further, when members of contemporary societies hold fiercely opposing attitudes, as with research on human embryonic tissues or the use of nonhuman animals, the existence of a lively debate about the moral standing of the pertinent entities has created fora that are plainly intended to approximate something like the types of deliberation and negotiation I've described.

In what follows, I'll attempt to identify some likely problems for our current practices, thus indicating loci where we might hope to do better. I'll start with three different types of concern about the setting of research agendas and the use that is made of scientific results: one charges that the preferences of large segments of the public are consistently neglected, a second alleges that inquiry is distorted because the untutored preferences of outsiders lead to the neglect of problems of real epistemic significance, while a third suggests that the coherent systematization of widely shared preferences would recommend different priorities. The first often takes the form of complaints that the sciences don't take into account the needs of women, children, members of minorities, and people in developing countries; the second typically comes from scientists who have been disappointed by the lack of support for a project that fascinates them; the third usually comes from those who oppose a scientific project on the grounds that a systematic interest in the values professed by the champions of the project would lead to quite a different assignment of resources. There's no doubt the actual processes that shape our research agenda and convey the results of inquiry to the public give disproportionate emphasis to the predilections of people belonging to particular subgroups while members of other subgroups don't

6. Although there are obvious instances in which this discrepancy arises, it would be wrong to maintain that the interests of scientists are always antagonistic to those of the broader public. Part of Donald Stokes's thesis in *Pasteur's Quadrant* is that considerations of *perceived* epistemic significance and *perceived* practical import may sometimes coincide. Even here, however, there may be a deeper problem. For the scientists are responding to their perceptions of what society needs, and that may not coincide with the outcome chosen after the ideal deliberation I've outlined.

participate directly at all. Yet one might think that the agenda and the applications must be sensitive to a wide variety of preferences because at least some people who do play a direct role in the decisions—administrators of government agencies, manufacturers, and other entrepreneurs—are answerable to the public. Optimists hope there will be some type of invisible hand, so that apparently unrepresented minorities, however small, can offer a niche to which businesses and politicians will want to appeal, and thus affect the character of inquiry. If a practical problem is urgent for a subgroup then, within the public sphere of assigning resources to inquiry, elected representatives should find it advantageous to encourage research that addresses the problem, and, within the private domain, there will be room for commercial exploitation. Unfortunately, it's very easy to show that there are conditions under which a set of rational agents—whether bureaucrats or entrepreneurs—will do better to ignore the problems of small minorities: if the distribution of constituencies across the voting population assigns each electoral district a dominant group interest while the minority is thinly spread, and if the costs of starting to develop the pertinent technology are sufficiently high, the minority won't be worth bothering about. Further, given the account of the evolution of inquiry that I have offered, there are further reasons for thinking that an initial decision to favor the interests of one group may be self-perpetuating. A line of solution to a practical problem, and its associated research projects, may be suboptimal for a subgroup of the population *relative to a class of options that were never offered*, even though it's the best of those that continue to be available, precisely because of an original neglect of the preferences of the subgroup.

The root idea is that a decision to extend the significance graph in a particular direction may make it easier to continue in the same direction, perhaps by decreasing costs, perhaps by increasing chances of success. I'll illustrate it with a popular (but controversial) example, the dedication of resources to devising effective means of birth control. Let's assume that, prior to the research that led to the Pill, men would have preferred a pill that could be taken by women, women a pill that could be taken by men.⁷ Suppose further that the initial decision ignored the preferences of the large majority of women. For many people, men and women, avoiding conception in sexual intercourse is an important goal, and hence the provision by biotechnology (*avant la lettre*) of the female pill seemed an excellent solution—at least until there were concerns (possibly unjustified) about increasing rates of cancer and heart disease. Yet, given our assumptions, women were offered a choice that didn't completely reflect their

7. This strikes me as a lot more complicated than is often supposed, primarily because of issues about control. Of course, preferences have evolved in light of changing sexual mores and because of the frequency and severity of sexually transmitted disease.

preferences, the female pill vs. much cruder forms of contraception, when they might have been given the choice of the female pill vs. the male pill. Optimists think that, once the broader spectrum of possibilities is recognized, there will be pressure on inquiry to respond. But the invisible hand fails. Because the female pill has a head start, even if *both* projects are now pursued with roughly equal resources, it will be expected that the choice will be between a female pill with costs C (measured in terms of side effects as well as money) and a male pill with costs C^+ (where, as the notation suggests, $C^+ > C$). Even if women prefer a male pill to a female pill, with equal costs, they may continue to prefer a female pill at lower cost to a male pill at higher cost. There's a *Nonrepresentational Ratchet*: because of the initial neglect of female preferences, women never receive the choice they want.

The story I've sketched may be right—or it may not. Detailed sociological work is needed to decide that issue. My aim at this stage is to canvass possibilities, and the first of these is

The Problem of Inadequate Representation

A group is inadequately represented when the research agenda and/or the application of research results systematically neglects the interests of the members of that group in favor of other members of society. Because of the Nonrepresentational Ratchet an early problem of inadequate representation in a field may be self-perpetuating.

One can't show that the problem of inadequate representation exists simply by noting that members of a particular group aren't sufficiently represented in decisions about inquiry—but, by the same token, nor can one suppose that some invisible hand will operate to forestall the problem.

The second problem I'll consider stems from the fear that representation of perspectives outside science works too well. Because the preferences of the vast majority of citizens are untutored, areas of science that depend heavily on public funding can be shaped by governmental decisions that respond to widespread ignorance, with the result that practical projects whose significance can easily be appreciated are overemphasized, with concomitant neglect of questions of large epistemic significance. Although there are public discussions about budgets for scientific research, discussions that sometimes afford inquirers an opportunity to campaign for their favorite epistemic endeavors, these may be quite inadequate to solve the problem. Despite the compelling testimony scientists present to elected representatives, those whom they hope to convince are responsible to constituents whose preferences run strongly counter to the appeal for funding. The mirror image of the problem of inadequate representation is

The Problem of the Tyranny of the Ignorant

Epistemically significant questions in some sciences may systematically be undervalued because the majority of members of society have no appreciation for the factors that make those questions significant.

Once again, it's important not to conclude too quickly that we are confronting an instance of this problem. Just because scientists don't succeed in obtaining the resources required for their favorite projects, it doesn't follow that we aren't in a state of well-ordered science—after all, even if the public preferences were tutored, they might still oppose the line of inquiry envisaged.

Scientists' awareness of potential problems of the tyranny of the ignorance spawns a further chance of departing from the state of well-ordered science. Favored lines of inquiry can be promoted by advertising them as catering to the wishes of large segments of the community. Even if it should turn out that citizens' tutored preferences would fortuitously accord with the recommended agenda and the envisaged applications, misleading accounts of what can be expected introduce further political constraints on research (as scientists must make gestures at accommodating the expectations they have raised) and can also reinforce attitudes that oppose research responsive to tutored preferences. In a nutshell, public misperceptions of the rationale for research and applications don't foster the stable pursuit of inquiry that would correspond to tutored preferences.

The defenses of the genomes project exhibit very clearly why this is so. In some quarters, the important reason for mapping and sequencing the human genome is taken to be the opportunity to protect American leadership in biotechnology. Public testimony to Congress on behalf of the project, widely reported in the media, presented a very different line, emphasizing the biomedical breakthroughs that were likely to occur. Partly because the project's champions sometimes said as much, partly because it was what the public expected, the biomedical advances were understood in terms of readily achievable strategies for the prevention and treatment of disease. But, as I noted in chapter 1, a sober review of the relief afforded by enhanced understanding of the molecular bases of diseases offers a very mixed picture, with a few partial successes and some cautionary failures. It was not pointed out to Congress, or to the general public, that the immediate practical consequence of mapping and sequencing would be an enormously enhanced ability to offer genetic tests, typically without being able to give much advice for addressing health risks: that this would be likely to reveal painful information to patients, especially in a probable majority of situations in which genetic counselling would be ineffective; that it might well serve as the basis for new forms of discrimination; and that it would result in the proliferation of prenatal tests, which in turn could be expected to multiply, possibly by a significant factor, the number of abortions.

For most scientists there has always been a far deeper motivation. Developing sequence technology and applying that technology to selected nonhuman organisms is expected to make biologists of our century better able to explore the large questions of physiology, developmental biology, and even evolution (genomic analysis will shed light on evolutionary relationships and reveal the kinds of changes involved in speciation). The scientists involved believe that the work in which they're engaged will ultimately translate into an enormously richer and more complete view of physiology and development—although they would concede that this is likely to take a very long time—and that there will be consequent medical benefits.

This is eminently justifiable and probably correct. The envisaged strategy is akin to that adopted by early twentieth-century geneticists who self-consciously sought to resolve the most fundamental issues by working with tractable organisms rather than tackling the questions of human medical genetics head-on. If the significance graphs for the pertinent fields were clearly articulated and their historical development explained, it is quite possible that the wisdom of the strategy would be evident and that the inquiries envisaged would accord with the tutored preferences of the citizens whose taxes support the genomes project. Although the research agenda actually pursued adequately represents the preferences people would acquire as the outcome of ideal deliberation, the way of achieving this goal is unreliable, and that unreliability has serious consequences. Funding flows because biotechnology is viewed as a continuing source of jobs for Americans and because medical benefits are believed to be around the corner. The entrenchment of these beliefs causes trouble for inquiry because there's pressure to produce *some* kind of short-term "solutions," and it also distorts the applications of results by concealing the social problems that the real products of the project are likely to bring. After all, who needs to worry about inadequate counselling, lack of insurance coverage, and genetic discrimination when cures are just around the corner? Hence we currently confront

The Problem of False Consciousness

A research agenda may conform to the tutored preferences of the majority not because the public reasons for the agenda are those that would figure in an ideal deliberation, but because those reasons misrepresent the agenda in ways that cater to the actual (untutored) preferences of the majority. Because these preferences are not tutored, there may be harmful constraints on the pursuit of inquiry and serious threats to the proper application of its results.

Faced with the prospect of the tyranny of the ignorant, false consciousness may provide a way of reaching a better outcome. Yet it's not a feature of well-ordered science.

The last general problem I'll discuss has to do specifically with the applications of research. Sometimes a course of inquiry can be defended and publicly supported because its champions advertise its promotion of a goal that is highly valued in the community, even though it's not made clear (quite possibly because the advocates don't recognize the point) that a different course of inquiry or application would promote that value more completely or more justly. This problem, too, can be exemplified by the genomes project. One important reason for mapping and sequencing the human genome is the possibility of preventing the births of people who would suffer from devastating genetic diseases—the extension of the benign programs of prenatal testing begun in the attack on Tay-Sachs disease. Ideal deliberation of the scientific research agenda and of the translation of scientific knowledge into technology and public policy would thus invoke the goal of diminishing the number of children whose lives are doomed to have abysmally low quality.

But once this goal is recognized, as it should be, then our ideal deliberators ought to consider all the available ways of advancing toward it, reviewing possible projects that might improve the quality of children's lives, especially in those instances where we know that the expected quality, without intervention, is low. Consider now the fact that something of the order of a million American children live in apartments where they are exposed to toxic levels of lead. There's no current program for clearing up their environments and preventing them (and future children) from suffering severe consequences (sometimes damage to bodily organs, sometimes impairment of mental function). An obvious criticism of our actual practice of extending and applying science is that our research agenda would be ideally supported by a principle that would also favor the application of technologies we already have (techniques for removing lead) and inquiries designed to develop new technologies to fulfill the same function more efficiently. In other words, we face

The Problem of Parochial Application

An actual research agenda and a practice of application may be ideally supported by a principle that would licence forms of research not currently undertaken or applications of previous research that are not pursued.

Parochial application often occurs because it goes hand in hand with inadequate representation. In the genetic example, it seems plausible that the failure to apply the principle of improving the well-being of children by launching a program of lead removal results from neglecting the interests of people who are likely to live in the pertinent inner-city dwellings, to wit, members of ethnic minorities.

I've delineated some general problems in an effort to show that, even without detailed optimality analyses, we can sometimes identify ways in which the

practice of the sciences is likely to diverge from well-ordered science.⁸ I'll close with a more detailed discussion of whether we would be better off if there were more public input into decisions about which inquiries should be pursued.

An obvious thought is that involving representatives of diverse perspectives in decisions about prospective inquiries and about the uses of existing knowledge would be likely to modify the aims pursued to bring them closer to those that would emerge under conditions of ideal deliberation. Contrast a number of ways in which research agendas might be set. One, *internal elitism*, consists in decision-making by members of scientific subcommunities. A second, *external elitism*, involves both scientists and a privileged group of outsiders, those with funds to support the investigations and their ultimate applications (call these people "paymasters"). A third, *vulgar democracy*, imagines that the decisions are made by a group that represents (some of) the diverse interests in the society with advice from scientific experts. The fourth, *enlightened democracy*, supposes decisions are made by a group that receives tutoring from scientific experts and accepts input from all perspectives that are relatively widespread in the society: in effect, it fosters a condensed version of the process of ideal deliberation I've outlined.

I take it that the status quo in many affluent democracies is a situation of external elitism that groups of scientists constantly struggle to transform into a state of internal elitism. Vulgar democracy is, as I've insisted, likely to be a bad idea. The interesting question is whether enlightened democracy would be preferable to either form of elitism.

There are three influential arguments that incline people to dismiss the possibility. The first, already hinted at above, invokes the idea of an invisible hand. Consider the incentives for paymasters under external elitism. To achieve their ends they must respond to the preferences of the main constituencies within society, and their decisions must thus take into account the heterogeneous preferences of the citizenry. Explicit representation of those preferences isn't needed, and may well prove inefficient.

Earlier, I suggested we have no reason to believe in the invisible hand. Here, I want to add three points. First, even if the paymasters' decisions respond to the preferences of citizens, they won't reflect the transformation of those preferences that would occur under tutoring. It's thus unlikely that the agenda will be

8. In my judgment, I've only noted the most obvious problems. Once the ideal of well-ordered science is recognized, there's an important need for a political theory of science that will consider the various ways in which the interests of actors and social institutions might easily divert us from the outcomes that would be reached in a state of well-ordered science. In *Between Politics and Science* (Cambridge: Cambridge University Press, 2000), David Guston offers a framework that might be valuable in developing a theory of this kind, although, as with many other writings in this area, Guston's work seems to me to suffer from a lack of the ideal of well-ordered science.

set in a way that even approximates the best pursuit of inquiry. Second, even if there were some pressure to respond to the untutored wishes of major constituencies within the society, we can expect minority concerns will be slighted—even in those instances in which ideal deliberation would have produced a response to them. The failure to represent the interests of people beyond the society is likely to be even more severe. Finally, it's highly likely that paymasters will prefer to manipulate the preferences of those to whom the products of inquiry are to be offered in whatever directions will maximize profits, and there's every reason to believe that these will not coincide with the products of ideal deliberation. Reflections on the current biotechnology market in the United States are hardly encouraging.

The second argument in favor of elitism suggests that introducing citizens into the deliberations that shape research agendas would reduce the pool of available strategies by making it less attractive for paymasters to support research and for talented people to engage in scientific careers. Here we encounter familiar advertisements for the benefits of deregulation. Yet the fact that representatives of citizens with different perspectives would be involved in discussions about agenda-setting and about applications hardly affects the prospects of paymasters; indeed, under many circumstances, those who underwrite research and development are happy to pay for information about the constituencies to whom they intend to appeal. Perhaps, however, there are disincentives for the scientists whose views about what lines of inquiry are most interesting might be swamped by the demands of the outside majority.

Here too the existence and strength of the disincentive could be questioned, for it's worth recalling that brilliant people are sometimes prepared to spend their lives in carrying out the research projects directed by commercial (or governmental) concerns. More important, the argument assumes, without justification, that there's no way to organize broader deliberation so as to tutor the preferences of all discussants, generating a research agenda acceptable both to the discussants and to those who are to carry out the research. What lurks behind the suspicion is, I think, the thought that any attempt at democracy must be sufficiently close to vulgar democracy that the tyranny of the ignorant will be inevitable.

This is the last, and I think the most powerful, reason for defending elitism. Enlightened democracy would try to tutor the raw preferences of representatives of different perspectives within the society, would admit expression of the needs and perceived interests of all groups, and would thus conduct informed deliberations. Skeptics pose a dilemma: either the processes that precede agenda-setting are impossibly cumbersome and time-consuming or they fail to shift the views of the participants sufficiently to produce a genuine departure from vulgar democracy.

The best response to that dilemma would be to delineate, clearly and specif-

ically, a mechanism for enlightened democracy and to show, on the basis of sociological research and mathematical analysis, that the expected results are better than those that elitism would yield. It's not hard to do the mathematics (indeed, it can be developed in similar ways to those used in formal treatments of the problems of the last chapter), but the sociological information required to build realistic models is currently not available. Hence, I must settle for a weaker response. Democratic proposals within other areas of politics and political philosophy are always vulnerable to charges that the incorporation of preferences would be too cumbersome. We rightly reject *a priori* skepticism until we've explored whether the democratic suggestions can be made to work, either by acquiring empirical information or by trying various possible schemes (possibly on a limited scale). Unless, or until, sociological research shows that the project of approximating tutored collective preferences is hopeless, we have no basis for concluding that some form of elitism must be superior.

Doubts may linger. Perhaps they can be allayed by a brief look at three policy formulations that have had some influence on the institutional arrangements within which the sciences are practiced.