SIX

REJECTING THE IDEAL OF VALUE-FREE SCIENCE

Heather Douglas

6.1 Introduction

The debate over whether science should be value free has shifted its ground in the past sixty years. As a way to hold science above the brutal cultural differences apparent in the 1930s and 1940s, philosophers posited the context of discovery–context of justification distinction, preserving the context of justification for reason and evidence alone. It was in the context of justification that science remained free from subjective and/or cultural idiosyncrasies; it was in the context of justification that science could base its pursuit of truth. Even within the context of justification, however, values could not be completely excluded. Several philosophers in the 1950s and 1960s noted that scientists needed additional guidance for theory choice beyond just logic and evidence alone. (See, for example, Churchman 1956, Levi 1962, or Kuhn 1977.) *Epistemic values* became a term to encompass the values acceptable in science as guidance for theory choice. Some argued that only these values could legitimately be part of scientific reasoning or that it was the long-term goal to eliminate nonepistemic values (McMullin 1983). By 1980, "value-free science" really meant science free of nonepistemic values.

But not all aspects of science were to hold to this norm. As the distinction between discovery and justification has been replaced by a more thorough account of the scientific process, the limits of the "value-free" turf in science have become clearer. It has been widely acknowledged that science requires the use of nonepistemic values in the "external" parts of science, that is, the choice of projects, limitations of methodology (particularly with respect to the use of human subjects), and the application of science-related technologies.[1] So the term *value-free science* really refers to the norm of epistemic values only in the internal stages of science. It is this qualified form of "value-free science" that is held up as an ideal for science.

Many assaults can and have been made on this ideal. It has been argued that it is simply not attainable. It has been argued that the distinction between epistemic and nonepistemic is not clear enough to support the normative weight placed on the distinction by the ideal. (I have argued this elsewhere [Machamer and Douglas 1999], as have Rooney 1992 and Longino 1996, more eloquently.) One can, however, take a stronger tack than the claim that value-free science is an unattainable or untenable ideal. One can argue that the ideal itself is simply a bad ideal. As I have argued in greater detail elsewhere, in many areas of science, particularly areas used to inform public policy decisions, science should not be value free, in the sense just described (Douglas 2000). In these areas of science, value-free science is neither an ideal nor an illusion. It is unacceptable science.

Rejecting the ideal of value-free science, however, disturbs many in the philosophy of science. The belief persists that if we accept the presence of values (particularly nonepistemic values) in the inner working of science, we will destroy science and set ourselves adrift on the restless seas of relativism. At the very least, it would be a fatal blow to objectivity. As Hugh Lacey has recently warned, without the value-free ideal for science's internal reasoning, we would lose "all prospects of gaining significant knowledge" (Lacey 1999, 216).

I disagree with this pessimistic prediction and instead think that rejecting the value-free ideal would be good for science by allowing for more open discussion of the factors that enter into scientific judgments and the experimental process. In this chapter, I will first explain why nonepistemic values are logically needed for reasoning in science, even in the internal stages of the process. I will then bolster the point with an examination of ways to block this necessity, all of which prove unsatisfactory. Finally, I will argue that rejection of the value-free ideal does not demolish science's objectivity and that we have plenty of remaining resources with which to understand and evaluate the objectivity of science. By understanding science as value laden, we can better understand the nature of scientific controversy in many cases and even help speed resolution of those controversies.

6.2 Choices and Values in Science

To make the normative argument that values are required for good reasoning in science, I will first describe the way in which values play a crucial decision-making role in science, which I will then briefly illustrate. The areas of science with which I am concerned are those that have clear uses for decision making. I am not focused here on science used to develop new technologies, which then are applied in various contexts. Instead, I am interested in science that is used to make decisions, science that is applied as useful knowledge to select courses of action, particularly in public policy.

One hundred years ago, science was little used in shaping public policy. Indeed, the bureaucracies that now routinely rely of scientific expertise in their decision making were either nonexistent in the United States (e.g., Environmental Protection Agency, Consumer Product Safety Commission, Occupational Safety and Health Administration, Department of Energy) or in their earliest stages of development (Food and Drug Administration, Centers for Disease Control). Now, entire journals (*Chemosphere*, *Journal of Applied Toxicology and Pharmacology*, *CDC Update*, etc.), institutions (e.g., National Institute for Environmental Health Sciences, Chemical Industry Institute of Toxicology, National Research Council), and careers are devoted to science that will be used to develop public policy. Although science is used to make decisions in other spheres as well (e.g., in the corporate world and nongovernmental organizations), I will draw my examples from the use of science in public policy. It is in this realm that the importance of scientific input is the clearest, with the starkest implications for our views on science.

In the doing of science, whether for use or for pure curiosity, scientists must make choices. They choose a particular methodological approach. They make decisions on how to characterize events for recording as data. They decide how to interpret their results.[2] Scientific papers are usually structured along these lines, with three internal sections packaged within an introduction and a concluding discussion. In the internal sections of the paper (methodology, data, results), scientists rarely explicitly discuss the choices that they make. Instead, they describe what they did, with no mention of alternative paths they might have taken.[3] To discuss the choices that they make would require some justification for those choices, and this is territory the scientist would prefer to avoid. It is precisely in these choices that values, both epistemic and, more controversially, nonepistemic, play a crucial role. Because scientists do not recognize a legitimate role for values in science (it would damage "objectivity"), scientists avoid discussion of the choices they make.

How do the choices require the consideration of epistemic and nonepistemic values? Any choice involves the possibility for error. One may select a methodological approach that is not as sensitive or appropriate for the area of concern as one thinks it is, leading to inaccurate results. One may incorrectly characterize one's data. One may rely upon inaccurate background assumptions in the interpretation of one's results.[4] When the science is used to make public policy decisions, such errors lead to clear nonepistemic consequences. If one is to weigh which errors are more serious, one will need to assign values to the various likely consequences. Only with such evaluations of likely error consequences can one decide whether, given the uncertainty and the importance of avoiding particular errors, a decision is truly appropriate. Thus values become an important, although not determining, factor in making internal scientific choices.

Clearly, there are cases where such value considerations will play a minor or even nonexistent role. For example, there may be cases where the uncertainty is so small that the scientists have to stretch their imaginations to create any uncertainty at all. Or there may be cases where the consequences of error are completely opaque and we could not expect anyone to clearly foresee them. However, I contend that in many cases, there are fairly clear consequences of error (as there are fairly well-recognized practices for how science is used to make policy) and that there is significant uncertainty, generating heated debate among scientists.

In general, if there is widely recognized uncertainty and thus a significant chance of error,[5] we hold people responsible for considering the consequences of error as part of their decision-making process. Although the error rates may be the same in two contexts, if the consequences of error are serious in one case and trivial in the other, we expect decisions to be different. Thus the emergency room avoids as much as possible any false negatives with respect to potential heart attack victims and accepts a very high rate of false positives in the process. (A false negative occurs when one rejects the hypothesis—in this case, that someone is having a heart attack—when the hypothesis is true. A false positive occurs when one accepts the hypothesis as true when it is false.) In contrast, the justice system attempts to avoid false positives, accepting some rate of false negatives in the process. Even in less institutional settings, we expect people to consider the consequences of error, hence the existence of reckless endangerment and reckless driving charges. We might decide to isolate scientists from having to think about the consequences of their errors. I will discuss this line of thought later. But for now, let us suppose that we want to hold scientists to the same standards as everyone else and thus that scientists *should* think about the potential consequences of error.

In science relevant to public policy, the consequences of error clearly include nonepistemic consequences. Even the most internal aspects of scientific practice—the characterization of events as data—can include significant uncertainty and clear nonepistemic consequences of error. An example I have discussed elsewhere that effectively demonstrates this point is the characterization of rat liver tissue from rats exposed to dioxin. (See Douglas 2000 for a more complete discussion.) In a key study used for setting regulatory policy completed in 1978, four groups of rats were exposed to three different dose levels of dioxin (2,3,7,8-tetrachloro-dibenzo-p-dioxin) plus a control group (Kociba et al. 1978). After two years of dosing, the rats were killed and autopsied. Particular focus was placed on the livers of the rats, and slides were made of the rat liver tissues, which were then characterized as containing tumors, benign or malignant, or being free from such changes. Over the next fourteen years, those slides were reevaluated by three different groups, producing different conclusions about the liver cancer rates in those rats. Clearly, there is uncertainty about what should count as liver cancer in rats and what should not.

What does this uncertainty mean for the decision of whether to characterize or not characterize a tissue slide as containing a cancerous lesion? In an area with this much uncertainty, the scientist risks false positives and false negatives with each characterization. Which errors should be more carefully avoided? Too many false negatives are likely to make dioxin appear to be a less potent carcinogen, leading to weaker regulations. This is precisely what resulted from the 1990s industry-sponsored reevaluation (see Brown 1991) that was used to weaken Maine water-quality standards. Too many false positives, on the other hand, are likely to make dioxin appear to be more potent and dangerous, leading to burdensome and unnecessary overregulation. Which consequence is worse? Which error should be more scrupulously avoided? Answering these questions requires reflection on ethical and societal values concerning human health and economic vitality. Such reflection is needed for those uncertain judgments *at the heart* of doing science.

One might counter this line of thought with the suggestion that scientists not actually make the uncertain judgments needed to proceed with science but, instead, that scientists estimate the uncertainty in any given judgment and then propagate that uncertainty through the experiment and analysis, incorporating it into the final result.[6] Two problems confront this line of thought. The first is purely practical. If the choices scientists must make occur early in the process, for example, a key methodological choice, it can be quite difficult to estimate precisely the effect of that choice on the experiment. Without a precise estimate, the impact on the experiment cannot be propagated through the experimental analysis. For example, in epidemiological studies, scientists often rely on death certificates to determine the cause of death of their subjects. Death certificates are known to be wrong on occasion, however, and to be particularly unreliable for some diseases, such as soft-tissue sarcoma (Suruda, Ward, and Fingerhut 1993) The error rates for rare diseases like soft-tissue sarcoma are not well known, however, and other sources of data for epidemiological studies are difficult or very expensive to come by. Expecting scientists to propagate a precise estimate of uncertainty about their source of data, in this case, through a study, would be unreasonable.

The second problem is more fundamental. To propagate the uncertainty, the scientist must first estimate the uncertainty, usually making a probabilistic estimate of the chance of error. But how reliable is that estimate? What is the chance of error in the estimate, and is the chance low enough to be acceptable? Making this kind of judgment again must involve values to determine what would be acceptable. Having scientists make estimates of uncertainty pushes the value judgments back one level but does not eliminate them. (This problem is first discussed in Rudner 1953 and, to my knowledge, is not addressed by his critics.) The attempt to escape the need for value judgments with error estimates merely creates a regress, pushing back the point of judgment further from view and making open discussion about the judgments all the more difficult. This serves to obscure the important choices and values involved, but it does not eliminate them.

Thus, *if* we want to hold scientists to the same responsibilities the rest of us have, the judgments needed to do science cannot escape the consideration of potential consequences, both intended and unintended, both epistemically relevant and socially relevant. This is not to say that evidence and values are the same thing. Clearly, logically, they are not. Values are statements of norms, goals, and desires; evidence consists of descriptive statements about the world. Hume's prohibition remains in effect; one cannot derive an *ought* from an *is*. This does not mean, however, that a descriptive statement is free from values in its origins. Value judgments are needed to determine whether a descriptive label is accurate enough and whether the errors that could arise from the description call for more careful accounts or a shift in descriptive language. Evidence and values are different things, but they become inextricably intermixed in our accounts of the world.

6.3 Scientists, Responsibility, and Autonomy

Although I hope to have convinced my reader by now that nonepistemic values do have a legitimate role to play in science and are needed for good reasoning, one still may wish to shield scientists from having to make value judgments as part of their work. There are two general and related objections to my position that can be made: (1) Scientists shouldn't make choices involving value judgments—they should do their science concerned with epistemic values only and leave determining the implications of that work to the policy makers, and (2) we should shield scientists from having to think about the consequences of error in their work in order to protect the "value neutrality" of the scientific process. I will address each of these in turn.

When the issue of values in science was raised in the 1950s by Churchman and Rudner, the response to their suggestion that values played an important role in science was that scientists do not need to consider values because they are not the ones performing the decisions for which consequences of error are relevant and/or they are simply reporting their data for the use of decision makers. The example of rat liver characterization choices from the previous section demonstrates the difficulty of holding to a "reporting data only" view of scientists' role in public policy. Even in the act of reporting "raw" data, some decisions are made as to how to characterize events in turning those events into raw data. (I also argued previously that reporting raw data with uncertainty estimates does not free the statements from relying in part on value judgments.) Those choices involve the potential for error and, in the example, clear and predictable consequences of error. Thus, even raw data can include judgments of characterization that require values in the process.

Scientists, however, rarely report solely raw data to public decision makers. They are usually also called on to interpret that data, and this is to the good. It would be a disaster for good decision making if those with far less expertise than climatologists, for example, were left with the task of interpreting world temperature data. Policy makers rarely have the requisite expertise to interpret data, and it is fitting that scientists are called on to make some sense of their data. Yet scientists' selection of interpretations involves selection of background assumptions, among other things, with which to interpret the data.

For example, in toxicology, there is a broad debate about whether it is reasonable to assume that thresholds exist for certain classes of carcinogens, or whether some other function (e.g., some extrapolation toward zero dose and zero response) better describes their dose-response relationship. There are complex sets of background assumptions supporting several different interpretations of dose-response data sets, including assumptions about the biochemical mechanisms at work in any particular case. Which background assumptions should be selected? Depending on which background assumptions one adopts, the threshold model looks more or less appropriate. In making the selection of background assumptions, not only epistemic considerations should be used but also nonepistemic considerations, such as which kinds of errors are more likely, given different sets, and how we weigh the seriousness of those errors. In short, we cannot effectively use scientific information without scientific interpretation, but interpretation involves value considerations. And because few outside the scientific community are equipped to make those interpretations, scientists usually must interpret their findings for policy makers and the public.

Still, to preserve the value-free ideal for useful science, one might be tempted to argue that we need to insulate scientists from considering the consequences of scientific error (the second objection). Perhaps we should set scientists apart from the general moral requirements to which most of us are held. Perhaps scientists should be required to search solely for truth, and any errors they make along the way (and the consequences of those errors) should be accepted as the cost of truth by the rest of society. Under this view, scientists may make dubious choices with severe consequences of error, but we would not ask them to think about those consequences and would not hold them responsible if and when they occur.

In considering this line of thought, it must be noted that, in other areas of modern life, we are required to consider unintended consequences of actions and to weigh benefits against risks; if we fail to do so properly, we are considering negligent or reckless. Scientists can be held exempt from such general requirements only if (1) we thought that epistemic values always trumped social values *or* (2) someone else could take up the burden of oversight. If we thought that epistemic values were a supreme good, they would outweigh social and moral values every time, and thus scientists would not need to consider nonepistemic values. If, on the other hand, someone else (with the authority to make decisions regarding research choices) were set up to consider nonepistemic values and social consequences, scientists could be free of the burden. If both of these options fail, the burden of responsibility to consider *all* the relevant potential consequences of one's choices falls back to the scientist. Let me consider each of these possibilities in turn.[7]

Do epistemic values trump other kinds of values? Is the search for truth (or knowledge) held in such high esteem that all other values are irrelevant before it? If we thought the search for truth (however defined, and even if never attained) was a value in a class by itself, worth all sacrifices, then epistemic values alone would be sufficient for considering the consequences of research. Epistemic values would trump all other values, and there would be no need to weigh them against other values. However, there is substantial evidence that we do not accord epistemic values such a high status. That we place limits on the use of human (and now animal) subjects for their use in research indicates we are not willing to sacrifice all for the search for truth. That our society has struggled to define an appropriate budget for federally funded research, and that some high-profile projects (such as the Mohole project in the 1960s[8] and the superconducting supercollider project in the 1990s) have been cut altogether suggests that in fact we do weigh epistemic values and goals against other considerations. That epistemic values are important to our society is laudable, but so, too, is that they are not held transcendently important when compared with social or ethical values. The first option to escape the burden of nonepistemic reflection is closed to scientists.

The second option remains but is fraught with difficulties. We could acknowledge the need to reflect on both social and epistemic considerations (i.e., the intended outcomes, the potential for errors and their consequences, and the values needed to weigh those outcomes) but suggest that someone besides scientists do the considering. We may find this alternative attractive because we have been disappointed by scientists' judgments in the past (and the values that shaped those judgments) or because we want to maintain the purity of science, free from social values.[9] The costs of nonepistemic research oversight by outsiders, however, outweigh the potential benefits.

For this option to be viable, consideration of nonepistemic consequences cannot be an afterthought to the research project; instead, it must be an integral part of it.[10] Those shouldering the full social and ethical responsibilities of scientists would have to have decision-making authority with the scientists, in the same way that research review boards now have the authority to shape methodological approaches of scientists when they are dealing with human subjects. However, unlike these review boards, whose review takes place at one stage in the research project, those considering nonepistemic consequences of scientific choices would have to be kept abreast with the research program at every stage (where choices are being made) and would have to have the authority to change those choices if necessary. Otherwise, the responsibility would be toothless and thus meaningless.

To set up such a system would be to dilute any decision-making autonomy the scientists have between the scientists and their ethical overseers. This division of authority would probably lead to resentment among the scientists and to reduced reflection by scientists on the potential consequences of research. After all, increased reflection would only complicate the scientist's research by requiring more intensive consultation with the ethical overseer. Without the scientists' cooperation in considering potential consequences, the overseers attempting to shoulder the responsibility for thinking about the consequences of science and error would be blind to some of the more important ones.

To see why, consider that scientists performing the research may in many cases be the only ones who are both aware of the uncertainties and potential for error and of the likely or foreseeable consequences of error. For example, before the Trinity test in 1945, several theoretical physicists realized there was a possibility a nuclear explosion might ignite the atmosphere. Hans Bethe explored this possibility and determined that the probability was infinitesimally small. Who else could have thought of this potential for error and followed it through sufficiently to determine that the chance of this error was sufficiently small to be disregarded? This is a dramatic example, but it serves to illustrate that we need scientists to consider where error might occur and what its effects might be. Few outside Los Alamos could have conceived of this possibility, much less determined it was so unlikely that it was not a worry. Only with the active reflection of scientists on the edge of the unknown can the responsibilities be properly met.

Thus, the responsibility to consider the social and ethical consequences of one's actions and potential error cannot be sloughed off by scientists to someone else without a severe loss of autonomy in research. We have no adequate justification for allowing scientists to maintain nonepistemic blinders on an ongoing basis. Because both epistemic and nonepistemic values are important, scientists must consider both when making choices with consequences relevant to both. To keep scientists from considering the consequences of their work would be a highly dangerous approach (for science and society), with risks far outweighing any benefits. However, some might still insist that the damage to the objectivity of science caused by accepting a legitimate role for nonepistemic values in scientific reasoning would be so severe that we should still attempt to shield scientists (somehow) from that responsibility. I will argue in the next section that objectivity is robust enough without needing to be defined in terms of the value-free ideal.

6.4 Implications for Objectivity and Science

Objectivity is one of the most frequently invoked yet vaguely defined concepts in the philosophy of science.[11] Happily, in recent years, some nuanced philosophical and historical work has been done to attempt to clarify this crucial and vague term. What has become apparent in most of this work is that objectivity is an umbrella concept encompassing a broad, interrelated, but irreducibly complex set of meanings. For example, in the philosophical literature of the past decade, several authors have pointed out that objectivity has, in fact, multiple meanings already in play (see, e.g., Lloyd 1995; Fine 1998). Historical work has suggested how this could come about, with detailed work tracking how the meaning of *objectivity* has shifted and accrued new nuances over the past three centuries (Daston and Gallison 1992; Daston 1992; Porter 1992, 1995). I will argue in this section that we can discard the value-free meaning of objectivity without significant damage to the concept overall. Despite the long association between "value free" and "objective," there is nothing necessary about the link between the two concepts.

Before embarking on a description of objectivity's complexity, I should make clear that not all of the other traditional meanings associated with objectivity are discussed here. Some of the meanings attached to objectivity are functionally unhelpful for evaluating whether a statement, claim, or outcome is, in fact, objective. For evaluating the objectivity of science, we need operationalizable definitions, definitions that can be applied to deciding whether something is actually objective. This restriction eliminates from consideration some of the more metaphysical notions of objectivity, such as an aperspectival perspective or being independent of human thought. Because we currently have no way of getting at these notions of objectivity, they are unhelpful for evaluating the objectivity of science or the objectivity of other human endeavors. I will not consider them here.

Even without functionally useless aspects of objectivity, there are seven distinct meanings for objectivity, aside from "value free"; that is, there are seven clear and accessible ways that we can mean "objective" without meaning "value-free." This result suggests that there are considerable resources inherent in the term *objectivity* for handling the rejection of the value-free ideal. Let me elaborate on these seven alternatives.[12]

Two of the senses of *objectivity* apply to situations where we are looking at human interactions with the world. The first is perhaps the most powerfully persuasive at convincing ourselves we have gotten ahold of some aspect of the world: manipulable objectivity. This sense of objectivity can be invoked when we have sufficiently gotten at the objects of interest such that we can use those objects to intervene reliably elsewhere. As with Ian Hacking's famous example from *Representing and Intervening*, scientists don't doubt the objective existence of electrons when they can use them to reliably produce images of entirely different things with an electron-scanning microscope (Hacking 1983, 263). Our confidence in the objective existence of the electron should not extend to all theoretical aspects connected to the entity—the theory about it may be wrong, or the entity may prove to be more than one thing—but it is difficult to doubt that some aspect of the world is really *there* when one can manipulate it as a tool consistently.

In cases where some scientific entity can be used to intervene in the world and that intervention can be clearly demonstrated to be successful, we have little doubt about the manipulable objectivity (sense 1) of the science. However, the controversial cases of science and policy today do not allow for a clear check on this sense of objectivity. The science in these cases concerns complex causal systems that are fully represented only in the real world, and to attempt to do the intervention tests in the real world would be unethical or on such long time scales as to be useless (or both). Imagine, for example, deliberately manipulating the global climate for experimental purposes. Not only would the tests take decades, not only would it expose world populations to risks from climate change, but also it still would not be conclusive; factors such as variability in sun intensity and the length of time needed to equilibrate global carbon cycles make intervention tests hugely impractical. It is very doubtful that we will have a sense of manipulable objectivity for cases such as these.

For some of these cases, there is another potentially applicable meaning for objectivity, one that trades on multiple avenues of approach.

If we can approach an object through different and hopefully independent methods and if the same object continues to appear, we have increasing confidence in the object's existence. The sense of objectivity invoked here, convergent objectivity (sense 2), is commonly relied on in scientific fields where intervention is not possible or ethical, such as astronomy, evolutionary biology, and global climate studies.[13] When evidence from disparate areas of research points toward the same result or when epistemically independent methodologies produce the same answer, our confidence in the objectivity (in this sense) of the result increases. (See Kosso 1989 for a discussion of the problem of epistemic independence.) We still might be fooled by an objectively convergent result; the methods may not really be independent, or some random convergence may be occurring. But objectivity is no *guarantee* of accuracy; instead, it is the best we can do.

In addition to these two senses of objectivity focused on human interactions with the world, there are senses of objectivity that focus on individual thought processes. It is in this category that one would place the "value-free" meaning of objectivity. As I argued previously, this sense of *objective* should be rejected as an ideal in science. It can be replaced with two other possibilities: detached objectivity or value-neutral objectivity. *Detached objectivity* refers to the prohibition against using values *in place of* evidence. Simply because one wants something to be true does not make it so, and one's values should not blind one to the existence of unpleasant evidence. Now it may seem that my defense of detached objectivity contradicts my rejection of value-free objectivity, but closer examination of the role of values in the reasoning process shows that this is not the case. In my preceding discussion and examples, values neither supplant nor become evidence by themselves; they do shape what one makes of the available evidence. One can (and should) use values to determine how heavy a burden of proof should be placed on a claim and which errors are more tolerable. Because of the need for judgments in science throughout the research process, values have legitimate roles to play throughout the process. But using values to blind one to evidence one would rather not see is not one of those legitimate roles. Values cannot act in place of evidence; they can only help determine how much evidence we require before acceptance of a claim. The difference between detached objectivity (sense 3) and value-free objectivity is thus a crucial one.

Value-neutral objectivity should also not be confused with value-free objectivity. In value-neutral objectivity (sense 4), a value position that is neutral on the spectrum of debate, a midrange position that takes no strong stance, is used to inform the necessary judgments. Value-neutral objectivity can be helpful when there is legitimate and ongoing debate over which value positions we ought to hold, but some judgments based on some value position are needed for research and decision making to go forward. Value-neutral objectivity has limited applicability, however; it is not desirable in all contexts. For example, if racist or sexist values are on one side of the relevant value spectrum, value neutrality would not be acceptable, because racist and sexist values have been rightly and soundly rejected. We have good moral reasons for not accepting racist or sexist values, and thus other values should not be balanced against them. Many conflicts involving science and society reflect unsettled debates, however, and in these cases, value neutrality, taking a reflectively balanced value position, can be usefully objective.

I have presented four alternative meanings for objectivity in addition to value free. There are three remaining, all concerned with social processes. The possibility of social processes undergirding objectivity has received increased attention recently, and in examining that body of work, I found three distinct senses of objectivity that relate to social processes: procedural objectivity, concordant objectivity, and interactive objectivity. Procedural objectivity (sense 5) occurs when a process is set up such that regardless of who is performing that process, the same outcome is always produced. (This sense is drawn from Megill 1994; Porter 1992, 1995.) One can think of the grading of multiple-choice exams as procedurally objective, or the rigid rules that govern bureaucratic processes. Such rules eliminate the need for personal judgment (or at least aim to), thus producing "objectivity."

Concordant objectivity (sense 6) occurs when a group of people all agree on an outcome, be it a description of an observation or a judgment of an event. The agreement in concordant objectivity, however, is not one achieved by group discussion or by following a rigid process; it simply occurs. When a group of independent observers all agree that something is the case, their agreement bolsters our confidence that their assessment is objective. This intersubjective agreement has been considered by some essential to scientific objectivity; as Quine wrote:

"The requirement of intersubjectivity is what makes science objective" (1992, 5).

Some philosophers of science have come to see this intersubjective component less as a naturally emergent agreement and more as the result of the intense debate that occurs within the scientific community (Longino 1990; Kitcher 1993; Hull 1988). Agreement achieved by intensive discussion I have termed *interactive objectivity* (sense 7). Interactive objectivity occurs when an appropriately constituted group of people meet and discuss what the outcome should be. The difficulty with interactive objectivity lies with the details of this process: What is an appropriately constituted group? How diverse and with what expertise? How are the discussions to be framed? And what counts as agreement reached among the members of the group? Much work needs to be done to fully address these questions. Yet it is precisely these questions that are being dealt with in practice by scientists working with policy-relevant research. Questions of whether peer review panels for science-based regulatory documents are appropriately constituted and what weight to put on minority opinions and questions of whether consensus should be an end goal of such panels and what defines consensus are continually faced by scientists.

I will not attempt to answer these difficult questions here. The point of describing these seven aspects of objectivity is to make clear that value free is not an essential aspect of objectivity. Rather, even when rejecting the ideal of value-free science, we are left with seven remaining aspects of objectivity with which to work. This embarrassment of riches suggests that rejecting the ideal of value-free science is no threat to the objectivity of science. Not all of the remaining aspects of objectivity will be applicable in any given context (they are not all appropriate), but there are enough to draw on that we can find some basis for the trust we place in scientific results.

6.5 Conclusion

Rejecting the ideal of value-free science is thus uncatastrophic for scientific objectivity. It is also required by basic norms of moral responsibility and the reasoning needed to do sound, acceptable science. It does imply increased reflection by scientists on the nonepistemic implications and potential consequences of their work. Being a scientist per se does not exclude one from that burden. Some scientists may object that their work has no implications for society and that there are no potential nonepistemic consequences of error. Does the argument presented here apply to all of science? My argument clearly applies to all areas of science that have an actual impact on human practices. It may not apply to some areas of research conducted for pure curiosity (at present). But it is doubtful that these two "types" of science can be cleanly (or permanently) demarcated from each other. The fact that one can think of examples at either extreme does not mean there is a bright line between these two types (the useful and the useless) or that such a line would be stable over time.[14] In any case, debates over whether there are clear and significant societal consequences of error in particular research areas would be a welcome change from the assertion that nonepistemic values should play no role in science. Understanding science in this way will require a rejoining of science with moral, political, and social values.

I would like to close this chapter by suggesting that opening the discourse of science to include discussion of nonepistemic values relevant to inductive risks will make answering questions about how to conduct good science easier, not harder. If the values that are required to make scientific judgments are made explicit, it will be easier to pinpoint where choices are being made and why scientists disagree with each other in key cases. It will also make it clearer to the science-observing public the importance of debates about what our values should be. Currently, too many hope that science will give us certain answers on what is the case so that it will be clear what we should do. This is a mistake, given the inherent uncertainty in empirical research. If, on the other hand, values can be agreed on, agreement will be easier to reach about how to best make scientific decisions (for example, as we now have clear guidelines and mechanisms for the use of human subjects in research) and about what we should do regarding the difficult public policy issues we face. If values cannot be agreed on, the source and nature of disagreement can be more easily located and more honestly discussed. Giving up on the ideal of value-free science allows a clearer discussion of scientific disagreements that already exist and may lead to a speedier and more transparent resolution of these ongoing disputes.

ACKNOWLEDGMENTS

My thanks to the University of Puget Sound (Martin Nelson Junior Sabbatical Fellowship) and the National Science Foundation (SDEST grant #0115258) for their support during the writing of this chapter. Also thanks to Harold Kincaid, Wayne Riggs, Alison Wylie, and Ted Richards for their detailed and insightful comments on earlier versions of this chapter. I gave one of those versions at the Eastern APA in December 2001, and the vibrant discussion that followed also helped clarify this work. All remaining muddles are mine alone.

NOTES

1. Rudner (1953, 1) noted the importance of values for the selection of problems. Nagel (1961, 485–87) and Hempel (1965, 90) also noted this necessary aspect of values in science. Rescher (1965) provides a more comprehensive account of multiple roles for values in science, as does Longino (1990, 83–102).

2. Richard Rudner (1953) made a similar point about the practice of science, although Rudner focused solely on the scientist's choice of a theory as acceptable or unacceptable, a choice placed at the end of the "internal" scientific process.

3. It can be difficult to pinpoint where scientists make choices when reading their published work. One can determine that choices are being made by reading many different studies within a narrow area and seeing that different studies are performed and interpreted differently. With many cross-study comparisons within a field, the fact that alternatives are available, and thus that choices are being made, becomes apparent.

4. Note that in reading a scientific paper with any one of these kinds of errors, it would not be necessarily obvious that a choice had been made, much less an error.

5. "Significant chance of error" is obviously a vague term, and whether it applies in different cases can be a serious source of debate. The fact that there is no bright line for whether a chance of error is significant does not mean that one need not think about that chance at all.

6. This is distinct from asking scientists to not consider consequences of error at all, to be addressed later.

7. I have argued these points in greater detail in Douglas 2003.

8. See Greenberg 1967, chapter 9, for a detailed account of Mohole's rise and fall.

9. Someone else may need to do some reflective considering in addition to scientists, but that would still leave the presumption of responsibility with the scientists.

10. There may be special cases where we decide to let scientists proceed without considering what might go wrong and whom it might harm, but these cases would have to be specifically decided, given the research context, and then still carefully monitored. What makes science exciting is its discovery of the new and the unknown. It is difficult to be certain at the beginning of a research project that no serious consequences (either of error or of correct results) lurk in the hidden future.

11. In comparison, consider the concept of truth. Although it, too, is often invoked, much effort has been spent trying to precisely define what is meant. With objectivity, in contrast, it is often assumed that we just "know" what we mean.

12. See Douglas 2004 for a more detailed discussion of these aspects of objectivity.

13. One can create controlled laboratory conditions for small-scale climate studies or evolutionary studies, but there is always a debate over whether all of the relevant factors from the global context were adequately captured by the controlled study.

14. The example of nuclear physics is instructive. Once thought to be a completely esoteric and useless area of research, it quite rapidly (between December 1938 and February 1939) came to be recognized as an area of research with immense potential practical implications.

REFERENCES

Brown, W. R. 1991. "Implication of the Reexamination of the Liver Sections from the TCDD Chronic Rat Bioassay," in M. Gallo, R. J. Scheuplein, and K. A. Van der Heijden, eds., *Biological Basis for Risk Assessment of Dioxins and Related Compounds*, pp. 13–26. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press.

Churchman, C. W. 1948. "Statistics, Pragmatics, and Induction." *Philosophy of Science*, 15, pp. 249–68.

Churchman, C. W. 1956. "Science and Decision-Making." *Philosophy of Science*, 22, pp. 247–49.

Daston, L. 1992. "Objectivity and the Escape from Perspective." *Social Studies of Science*, 22, pp. 597–618.

Daston, L., and P. Gallison. 1992. "The Image of Objectivity." *Representations*, 40, pp. 81–128.

Douglas, H. 2000. "Inductive Risk and Values in Science." *Philosophy of Science*, 67, 559–79.

Douglas, H. 2003. "The Moral Responsibilities of Scientists: Tensions between Autonomy and Responsibility." *American Philosophical Quarterly*, 40, pp. 59–68.

Douglas, H. 2004. "The Irreducible Complexity of Objectivity." *Synthese*, 138, pp. 453–73.

Fine, A. 1998. "The Viewpoint of No-One in Particular." *Proceedings and Addresses of the APA*, 72, pp. 9–20.

Greenberg, D. S. 1967. *The Politics of Pure Science*. Chicago: University of Chicago Press.

Hacking, I. 1983. *Representing and Intervening*. New York: Cambridge University Press.

Hempel, C. G. 1965. "Science and Human Values," in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, pp. 81–96. New York: Free Press.

Hull, D. 1988. *Science as a Process*. Chicago: University of Chicago Press.

Jeffrey, R. C. 1956. "Valuation and Acceptance of Scientific Hypotheses." *Philosophy of Science*, 22, 237–46.

Kitcher, P. 1993. *The Advancement of Science*. New York: Oxford University Press.

Kociba, R. J., et al. 1978. "Results of a Two-Year Chronic Toxicity and Oncogenicity Study of 2.3.7.8-Tetrachlorodibenzo-p-Dioxin in Rats." *Toxicology and Applied Pharmacology*, 46, pp. 279–303.

Kosso, P. 1989. "Science and Objectivity." *Journal of Philosophy*, 86, pp. 245–57.

Kuhn, T. 1977. "Objectivity, Value, and Theory Choice," in *The Essential Tension*, pp. 320–39. Chicago: University of Chicago Press.

Lacey, H. 1999. *Is Science Value Free?* New York: Routledge.

Levi, I. 1962. "On the Seriousness of Mistakes." *Philosophy of Science*, 29, pp. 47–65.

Lloyd, E. 1995. "Objectivity and the Double Standard for Feminist Epistemologies." *Synthese*, 104, pp. 351–81.

Longino, H. E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton, NJ: Princeton University Press.

Longino, H. E. 1996. "Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy," in L. H. Nelson and J. Nelson, eds., *Feminism, Science, and the Philosophy of Science*, pp. 39–58. Dordrecht: Kluwer.

Machamer, P., and H. Douglas. 1999. "Cognitive and Social Values." *Science and Education*, 8, pp. 45–54.

McMullin, E. 1983. "Values in Science," in P. D. Asquith and T. Nickles, eds., *Proceedings of the 1982 Biennial Meeting of the Philosophy of Science Association*, vol. 1, pp. 3–28. East Lansing, MI: Philosophy of Science Association.

Megill, A. 1994. "Introduction: Four Senses of Objectivity," in A. Megill, ed., *Rethinking Objectivity*, pp. 1–20. Durham, NC: Duke University Press.

Nagel, E. 1961. *The Structure of Science: Problems in the Logic of Scientific Explanation*. New York: Harcourt, Brace, and World.

Porter, T. 1992. "Quantification and the Accounting Ideal in Science." *Social Studies of Science*, 22, pp. 633–52.

Porter, T. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton, NJ: Princeton University Press.

Quine, W. V. 1992. *The Pursuit of Truth*. Cambridge, MA: Harvard University Press.

Rescher, N. 1965. "The Ethical Dimension of Scientific Research," in Robert Colodny, ed., *Beyond the Edge of Certainty*, pp. 261–76. Englewood Cliffs, NJ: Prentice-Hall.

Rooney, P. 1992. "On Values in Science: Is the Epistemic/Non-Epistemic Distinction Useful?" in D. Hull, M. Forbes, and K. Okruhlik, eds., *Proceedings of the 1992 Biennial Meeting of the Philosophy of Science Association*, vol. 2, pp. 13–22. East Lansing, MI: Philosophy of Science Association.

Rudner, R. 1953. "The Scientist *Qua* Scientist Makes Value Judgments." *Philosophy of Science*, 20, pp. 1–6.

Suruda, A. J., E. M. Ward, and M. A. Fingerhut 1993. "Identification of Soft Tissue Sarcoma Deaths in Cohorts Exposed to Dioxin and Chlorinated Naphthalen." *Epidemiology*, 4, pp. 14–19.