THE MORAL RESPONSIBILITIES OF SCIENTISTS

THE DEBATE AMONG PHILOSOPHERS OF SCIENCE in the 1950s concerning values in science hinged on the proper role of scientists in a modern democracy. Should scientists be giving advice to decision-makers? And should they, when giving this advice, consider the context of use and the potential consequences of error when deciding what to say? Or should scientists decide which empirical claims are adequately supported with no thought to the importance of these claims to society? These questions fundamentally concern the moral responsibilities of scientists as scientists. If, with Rudner and Churchman, one thinks that scientists should consider the potential consequences of error when deciding which claims to make, then values have an unavoidable place in scientific reasoning. If, with Levi and McMullin, one thinks that scientists should not be considering the potential consequences of error, then scientists can safely exclude social and ethical values from the heart of scientific reasoning.[1]

Which is the correct view? Should scientists consider the potential consequences of error in their advising? This question involves two general aspects. First, there is the question of whether all of us, as general moral agents, have a responsibility to consider the consequences of error when deliberating over choices, and in particular when deciding upon which empirical claims to make. I will argue here that we do have a general moral responsibility to consider the consequences of error, based on our concern over reckless or negligent behavior. Second, there is the question of whether scientists share this burden with the rest of us, or whether they have a special moral exemption from such considerations. In other words, we must ask whether scientists have a special professional status which means they should not consider the consequences of their work as scientists. This is a view that has been supported by some in the recent past and must be addressed seriously. In the end, I will argue that scientists do have a moral responsibility to consider the consequences of error in their work, but that this responsibility places no burden of special foresight on the scientists. We cannot expect scientists to be seers. Indeed, the very errors we are concerned with here mean that we cannot expect *perfect* foresight and prediction. But we should expect *reasonable* foresight and care from our scientists. Being a scientist provides no special exemption from this expectation.

Moral Responsibility and the Consequences of Error

The literature on moral responsibility has developed much in recent years, but it has mostly focused on three general issues: competence (when is a person morally capable of making decisions and thus is responsible for them), coercion (what kind of forces on a person make their choices not their own, but rather due to someone else, thus shifting moral responsibility), and causation (what conception of causality allows for both enough free will and enough foresight so that we can be considered responsible for the outcomes of our actions).[2] While the discussion around these issues is fascinating, none of it is very illuminating for our concerns here. With respect to competence, scientists are generally capable moral agents, as capable in their daily lives as the rest of us. No one has seriously argued that scientific training somehow impairs moral reasoning or moral sentiments. With respect to coercion, scientists are not usually under threat of force to *not* consider certain moral issues. Coercion would be considered as pathological in science as it would anywhere else. And the issue of causation applies just as much for scientists as for anyone else; either we have a causal structure that allows for moral responsibility, or we do not. More relevant to our concerns here are what is moral responsibility in general and what are our responsibilities with respect to consequences we do not intend to cause. Let me address each of these in turn.

What do we mean when we say a person is morally responsible for some action or for an outcome of an action? One basic distinction is between causal responsibility and moral responsibility. I may be causally necessary, and thus partially causally responsible, for the eventual actions of my great-grandchildren, but few would suggest I was morally responsible for them, that I should be praised or blamed for what they do. Or if this book, for which I am causally responsible, serves to block a bullet meant for a reader of the book, I am surely not morally responsible for having saved that person's life, and no one should praise me (on that basis) for having written this particular book. Being part of a causal chain is not sufficient for moral responsibility. Where does moral responsibility begin? One crucial marker for moral responsibility is the giving of praise or blame. When we hold someone morally responsible for something, we want to give praise to them or lay blame on them. We think that they could have done otherwise, that they chose a particular course, and we hold them responsible for that choice.[3] Thus, the attribution of praise or blame is central to moral responsibility. In general there should be broad symmetries between what warrants praise and what warrants blame. In other words, if there are circumstances where we are willing to praise someone for some particular action or result, for those same general kinds of circumstances we should also attribute blame. The crucial point here is that the attribution of praise or blame is not equivalent to the attribution of cause. While some kind of causal connection is necessary, we are not held morally responsible for all the things in which we play a causal role.

When does a causal responsibility turn into a moral responsibility? Minimally, we are morally responsible for those things we intend to bring about. In these cases, a person chooses to do something deliberately, either because they think it is inherently the right (or wrong) thing to do or because of a particular sought consequence. The deliberate choice brings on the moral responsibility. Thus, if I intend to help or harm someone and I succeed, I am morally responsible in both cases, and usually praiseworthy in the former, blameworthy in the latter.

However, we are not morally responsible merely for those things we intended to bring about. We are also morally responsible to some extent for side effects of our actions. While this is widely accepted, it is a difficult question under which circumstances and to what extent we should be responsible for unintended consequences. Two general categories cover unintended consequences: recklessness and negligence. In discussing these terms, I follow Joel Feinberg's work and general legal usage. Feinberg (1970) writes, "When one knowingly creates an unreasonable risk to self or others, one is reckless; when one unknowingly but faultily creates such a risk, one is negligent" (193). When one is *reckless*, one is fully aware of the risks one is taking or imposing on others, and those risks are unjustified. What justifies certain risks can be contentious, particularly when the risks involve other people not given decisionmaking authority or choice in a situation. Nevertheless, there are clear examples of justified risk (speeding on city streets to get a seriously injured person to the hospital) and unjustified risk (speeding on city streets for the fun of it). The key point is that we expect moral agents to carefully weigh such risks and to determine whether they are, in fact, justified.

If, on the other hand, one is not aware that one is risking harm, but one *should* be aware, then one is being *negligent*. When being negligent, one does not bother to evaluate obvious risks of harm, or one does not think about potential consequences of one's actions. As Feinberg notes, there are many ways in which to be negligent:

> One can consciously weigh the risk but misassess it, either because of hasty or otherwise insufficient scrutiny (rashness), or through willful blindness to the magnitude of the risk, or through the conscientious exercise of inherently bad judgment. Or one can unintentionally create an unreasonable risk by failing altogether to attend either to what one is doing (the manner of execution) or to the very possibility that harmful consequences might ensue. (Feinberg 1970, 193–94)

The difficulty with negligence, in addition to determining whether a risk is justified, is to determine what should be expected of the agent. How much foresight and careful deliberation should we expect the individual to have? Often, this question is answered through an examination of community standards, couched in terms of what a *reasonable person* would have done in like circumstances.

Through recklessness and negligence, one can be held morally responsible for unintended consequences both when things go the way one expects them and when things go awry. Through negligence, things may go exactly as planned (as far as you planned them), and still harmful and clearly foreseeable consequences would be your fault. You would be responsible because you should have foreseen the problems and planned further. For example, suppose you set fire to a field one dry summer to clear it of brush. You didn't bother to think about how to control the fire, not recognizing the obvious risk. Because of your negligence, harm caused by the fire raging beyond your property is your moral responsibility. If, on the other hand, you are aware of the risks in setting the fire, decide not to care and proceed anyway, then you are morally responsible for the damage when the fire escapes because of recklessness. The distinction between recklessness and negligence thus rests on the thought processes of the agent, on whether they reflect on potential consequences (either if events go as planned or unexpected "errors" occur), and on whether there was any attempt to prevent or mitigate possible harms that could arise from the chosen action. Recklessness is proceeding in the face of unreasonable risk; negligence is the failure to foresee and mitigate such risk.

This discussion of moral responsibility for choices might seem far removed from the context of science, where one is primarily concerned with coming to the correct empirical beliefs (and thus the province of theoretical instead of practical reason). However, science is a social enterprise, involving not just isolated individuals coming to have beliefs, but social communities working together to develop empirical work (Longino 1990, 2002; Solomon 2001). Even more pertinent, scientific work is developed and discussed within a society that takes the claims made in the name of science with special authority—hence the importance of science advising. Making empirical claims should be considered as a kind of action, with often identifiable consequences to be considered, and as a kind of belief formation process. We can consider everyday nonscientific examples to show that the responsibility to consider consequences of error embodied in concerns over recklessness and negligence applies not just to direct interventions in the world, but also to the making of claims, including descriptive claims.

For example, suppose one sees an unattended briefcase. Should one report the possible presence of a bomb? There are clear risks of error for making the descriptive claim that a bomb may be present. If one does not report it and it is a bomb, death and destruction may result. If one does report it and it is not, disruption of people's daily lives and distraction of resources away from more serious problems may result. Obviously, the former is a more serious consequence than the latter, but one should also weigh the uncertainties involved. Suppose, for example, that the briefcase is spotted in a busy subway station. In that context, where leaving a briefcase seems ominous, one should report the unattended briefcase. This is the reasonable weighing of risk and uncertainty in this context. Consider, however, a briefcase left in a college classroom, where the classroom is known to be used by a particularly absentminded colleague. In that context, while the consequences are similar, the uncertainties shift, and it is far more likely it is the colleague's briefcase than a bomb. Checking with the colleague first is the more prudent measure. In both cases, we expect each other to reflect upon the risks of making a claim, particularly the consequences of error and the uncertainties and likelihoods involved. Thus, we can be negligent or reckless in the making of descriptive or empirical claims.

This brief sketch of moral responsibility holds several crucial insights. Moral responsibility involves the attribution of praise or blame (and sometimes both). We are held morally responsible for our intentional choices and the intended consequences of those choices, but also for some (not all) of the unintended consequences. Those consequences that are reasonably foreseeable (more on this below), even if not foreseen by the individual, due to negligence, or even if ignored, due to recklessness, can be evaluated and the person held morally responsible for them. And finally, we are morally responsible for reflecting upon the possible negligence or recklessness of making claims, particularly when the making of a claim will have clear consequences for others. These are the standards we all live with in our daily lives, that help to shape our choices. The question is whether these standards should apply to scientists, and if so how, particularly in their role as advisors and public authorities.

Scientists and Moral Responsibility

What are the moral responsibilities of scientists? If scientists have the same responsibilities as the rest of us, they have the basic responsibilities we all share for the intended consequences of their choices, as well as for some of the unintended consequences. Specifically, they are responsible for the foreseeable consequences of their choices, whether intended or not. Thus, scientists have the responsibility to be neither reckless nor negligent in their choices. Because the impacts of science have such substantial, and often foreseeable, reach, meeting these basic responsibilities could profoundly influence the practice and direction of science.

We can set aside here the rare cases of scientists with genuinely malicious intent, and thus the problem of failing to meet the general moral responsibility of having a good aim, of one's intended consequences being morally acceptable. (Scientists can be as mistaken about the goodness of an aim as anyone else, but this poses no special problems for us here.) Even scientists who solely concern themselves with the pursuit of knowledge can be said to pursue a good, as it seems clear knowledge in general is thought to be a good. Unintended consequences are the more important—and murkier—ground here, even if we restrict our discussion to foreseeable consequences.

There are two kinds of unintended foreseeable consequences that may be of concern to the scientist. The first is the consequence that will likely result as a side effect even if the knowledge produced is perfectly reliable and accurate. Consequences to human subjects of research have been a focus of research ethics for over half a century, and they place restrictions on what kinds of research projects are morally acceptable. Even if the research goes exactly as planned, and the scientists intend the best, moral concerns may trump epistemic drives. This kind of concern takes on trickier overtones when one considers the impact of knowledge itself on society. Recent concerns have been raised over knowledge that, even if produced ethically and even if true, may harm society. This problem, what has been called the problem of forbidden knowledge, has received some attention lately, but discussions have stopped short of arguing that the potential harms to society are sufficient grounds for forbidding the pursuit of such knowledge (Johnson 1996, 1999; Kitcher 1997). Because these discussions center more on policy for science and democratic input into research agendas than on science for policy and its implications for scientific reasoning, this issue will be set aside here. More central to our purposes are the potential unintended consequences of making inaccurate or unreliable empirical claims, the second kind of unintended consequence. Given the public authority scientists wield, and the understanding of making empirical claims as a public action, there can be clear consequences for making a well-intended but still ultimately incorrect claim, just as with the briefcase example above. Thus, it seems that scientists can be reckless or negligent if they improperly consider the potential consequences of error based on their basic moral responsibilities.

These basic or general responsibilities are not the only ones relevant for scientists, however. There are also the responsibilities that scientists must meet because they are scientists. Responsibilities special to science can be considered *role responsibilities*, which are those that assist scientists in achieving the central goals of science: improved explanations and predictions about the world around us. They include the precepts of basic research ethics, such as the honest reporting of data, the open discussion of scientific results, and the fair consideration and evaluation of the work of others. There is little or no debate about whether these kinds of role responsibilities should hold for science. For example, all fraudulent reporting of data is roundly condemned. Debate centers, rather, on the extent of fraud in science (see, for example, Callahan 1994). Concern has also been raised about whether the rise of propriety research and the importance of business interests in science pose a threat to scientific integrity (Maker 1994). Few dispute that basic research ethics are central to the good functioning of science.

How should we understand the relationship between general or basic responsibilities and role responsibilities? Role responsibilities arise when we take on a particular role in society, and thus have additional obligations over and above the general responsibilities we all share. For example, when one becomes a parent, one takes on the additional responsibility of caring for one's child. This additional responsibility does not excuse one from the general responsibilities we all have, but must be met in conjunction with one's general responsibilities.[4] Thus, role responsibilities usually expand one's set of responsibilities. Examples range from family relationships to professional obligations. For scientists, the role responsibilities add the extra burden of care in reporting results and properly dealing with colleagues and students, for example. It is possible, but rare, that role responsibilities call for a contraction of general responsibilities. For example, consider a defense lawyer who has learned of past criminal activity from a client. The lawyer, unlike the rest of us, is under no obligation to report such activity because of the need to protect lawyer-client confidentiality. This exemption for the lawyer is possible only because of the clear structure of our criminal justice system. It is the responsibility of others to discover and investigate the past criminal activity. In the rigid adversarial system of our criminal justice system, the defense lawyer has the responsibility of being a client advocate, while others must discover and prosecute crimes, although even defense lawyers must report knowledge of crimes ongoing or not yet committed. Only within this kind of rigid system with clearly defined roles covering all the important responsibilities can a role responsibility lead to a reduction of general responsibilities.

Scientists do not currently work within such a rigid system with clearly defined roles covering all important responsibilities. Would scientists want to consider placing their work within such as system? In other words, would it be possible and desirable for someone other than a scientist to shoulder the burden of the general responsibility to consider the consequences of error in the scientist's work, so that the scientist does not have to worry about being negligent? There are several reasons such a burden-shifting is neither possible nor desirable. The most important reason is that it is doubtful anyone could fully take over this function for scientists. Because science's primary goal is to develop knowledge, scientists invariably find themselves in uncharted territory. While the science is being done, presumably only the scientist can *fully* appreciate the potential implications of the work, and, equally important, the potential errors and uncertainties in the work. And it is precisely these potential sources of error, and the consequences that could result from them, that someone must think about. The scientists are usually the most qualified to do so.

Despite the implausibility of scientists turning this aspect of general responsibility over to someone else completely, there is the possibility for some partial sharing of these general responsibilities. For example, in order to assure proper weighing of knowledge goals against the possible (even unintended) harms to humans in the pursuit of that knowledge, it is now standard practice for scientists using human subjects to submit their proposed methodologies to ethical review boards before proceeding with the experiment. Similar practices are becoming common for animal subjects as well. And I will suggest in chapter 8 that public involvement in some kinds of research can help to remove part of the burden from scientists as well. However, such sharing of the burden is the most that can be accomplished because scientists often encounter the unexpected and are the only ones aware of its presence, nature, and novelty. For scientists not to bear a general responsibility to consider potential unintended consequences of mistaken empirical judgment would require constant ethical oversight of all scientific practice. It is highly doubtful this could be accomplished (although moving toward such an option would surely mean full employment for applied ethicists), and it is even more doubtful scientists would prefer this option to shouldering most of the general responsibilities themselves. To abandon their general responsibilities would be to simultaneously relinquish most of their autonomy.

In order to see why, consider what would be required if we implemented a system that provided ethical oversight of all scientific decisions in order to remove the burden of these general responsibilities from scientists. The consideration of nonepistemic consequences could be neither an afterthought to the research project nor a process merely at the start of the project if the general responsibilities are to be properly fulfilled. Instead, such consideration would have to be an integral part of it and involved *throughout* the research project. Those shouldering the general responsibilities to consider social and ethical consequences of research (and in particular of errors in research) would have to have decisionmaking authority with the scientists in the same way that research review boards now have the authority to shape scientists' methodological approaches 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 all of the nonepistemic consequences of scientific choices would have to be kept abreast of the research program at every stage (where choices are being made), and would have to have the authority to alter those choices when necessary. Otherwise, the responsibility would not be properly fulfilled and would not be able to keep pace with the developments accompanying discovery. Such intensive interference in scientific practice is anathema to most scientists.

In sum, there is no one currently available to shoulder the general moral responsibility to consider the consequences of error for scientists. Scientists themselves are the best qualified to do so, and to develop a body of people to take this burden from them would probably not be acceptable to scientists. This last point could use some emphasis. A cadre of ethical overseers would certainly draw the resentment of scientists. In order to thwart such oversight, scientists would likely become less thoughtful about both the presence of uncertainty in their work and the potential consequences of error. After all, if they do not identify the possibility of error, then there would be less need for oversight. Thus, not only would an oversight system be undesirable, it would likely be self-defeating.

On the face of things, it seems that scientists should meet both their role responsibilities and their general responsibilities. Despite the argument that no one can effectively take over the general responsibility to consider the consequences of error for scientists, some have suggested that scientists should nevertheless ignore or abandon such responsibilities. So, knowing that no one else can do the job, should scientists be shielded from the general responsibility to consider the consequences of error in their choices? When deciding how to proceed in a research project or considering which empirical claims to make, should scientists meet the general responsibilities discussed above? Some have thought that they should not.[5] As noted in chapter 3, many philosophers of science have assumed that scientists should be insulated (or act as if they were insulated) from such considerations. Others have made the claim more explicit. For example, Lübbe (1986) suggests that scientists, by virtue of being scientists, enjoy "a morally unencumbered freedom from permanent pressure to moral self-reflection" (82). However, role responsibilities of a profession do not generally provide an exemption from general responsibilities. In the years after World War II, Percy Bridgman attempted to provide a more considered argument for such an exemption as controversy swirled over the role of scientists in the construction of the atomic bomb. Bridgman, a Harvard physicist and Nobel laureate, argued, "The challenge to the understanding of nature is a challenge to the utmost capacity in us. In accepting the challenge, man can dare to accept no handicaps. That is the reason that scientific freedom is essential and that the artificial limitations of tools or subject matter are unthinkable" (Bridgman 1947, 153).

The knowledge that scientists produce is so valuable to society, Bridgman suggests, that we must relinquish other claims of social or moral responsibility on scientists so that they can produce this valued end. Scientists, under this view, not only have a need for autonomy (that is, the ability to be the primary decisionmakers in determining the direction of their work), but also have a need to be free from considering the potential consequences of their work beyond the realm of science. The ideas latent in Bridgman's arguments can be cast in two ways, a stronger and a weaker version. The stronger version is that scientific knowledge is valuable beyond price, and thus any sacrifice is worth its achievement. The weaker version is that the price of accepting the burden of moral reflection is too high compared to the value of science.

Is the knowledge produced by scientists so valuable that it is worth the price of scientists' moral exemption from the basic responsibilities articulated above? One way to fashion this claim is to argue that epistemic concerns trump all other values, that is, that the search for truth (or knowledge) is 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 concerns alone would be sufficient for considering the consequences of research. The search for truth would overshadow other values, and there would be no need to weigh epistemic concerns against other values. However, there is substantial evidence that we do not accord epistemic concerns such a high status. We place limits on the use of human (and now animal) subjects in research, which indicates we are not willing to sacrifice all for the search for truth.

Such considerations came strongly to the fore after Bridgman wrote his 1947 essay, with the Nuremburg trials and the subsequent concerns over the use and protection of human subjects in scientific research. In addition, our society has struggled to define an appropriate budget for federally funded research, and some high-profile projects (such as the Mohole project in the 1960s[6] and the superconducting supercollider project in the 1990s) have been cut altogether. This suggests that in fact we do weigh epistemic goals against other considerations. That knowledge is important to our society is laudable, but so too is the fact that it is not held transcendently important when compared to social or ethical values. Thus, the knowledge produced by scientists should not be and is not considered priceless. The stronger version of the argument for the moral exemption of scientists fails.

The weaker version of Bridgman's argument is that requiring scientists to consider the consequences of their work (including the consequences of error) is a burden on science and would thus unduly hamper science. In other words, the price of fully morally responsible science is too high. Unfortunately for this line of argument, what that price is has not been articulated or developed. Yet the opposite *can* be clearly argued. One can think of cases where a failure of scientists to consider the unintended consequences, or the implications of error, would be catastrophic. Consider two examples where scientists happily did not view their work as subject to a moral exemption, one from the development of atomic weapons and another from cellular biology. In July 1945, the first atomic bomb was tested in New Mexico. Whatever one may think of the morality of building such weapons, this test of the first plutonium bomb, the Trinity test, was not just a test of a new technology. It was also a decisive test of some of the physical principles that went into the development of the bomb, from the fissibility of plutonium to the calculations behind the implosion device developed by George Kistiakowsky. It was also an experiment about what happens when you produce an explosive chain reaction in the atmosphere. No one had done this before, and there were some worries. One worry that was considered well before the test, and worked on by Hans Bethe, was that the energy in the explosion might produce an explosive chain reaction in the constituents of the earth's atmosphere itself, thus obliterating human life on earth. Happily, the scientists not only thought of this potential outcome, but Bethe pursued the possibility and determined it was scientifically impossible (Rhodes 1986, 419). Only the scientists immersed in the project could have foreseen this and determined the risk was extremely negligible. For a second example, consider the concern scientists raised over recombinant DNA techniques and the resulting Asilomar conference (Culliton 1979, 150–51; Krimsky 1982). Scientists in the midst of exciting research realized that there were risks associated with a new line of research, that serious consequences for public health could occur. They moved to mitigate those risks, even accepting a moratorium on that line of research, thus reducing the likelihood of a damaging unintended result.

In both these cases, scientists, while doing science, reflected on the potential unintended consequences and found the risks unacceptable. Before proceeding with the development of science, they paused and either made sure that the harmful consequences were nearly impossible or figured out ways to make them so. When looking back on these cases, we should be relieved that scientists did not view themselves as morally exempt from considering the risks and consequences of error. The price of morally exempting scientists from the general responsibility to consider the consequences of error looks much higher than the price of having scientists shoulder this burden.

While both of these examples involve scientists considering which testing procedures to follow, or more generally, which methodologies, they do exemplify the difficulties of producing a blanket moral exemption for scientists. The methodological difficulties that might be encountered at the Trinity test were certainly not foreseen at the start of the Manhattan Project. Instead, the concerns with the atmosphere came up in the middle of the project. Given this example, it is clear that we want scientists to think about the potential consequences of error (that is, the potential harms that occur when things do not go as planned) throughout a particular project. Similarly, the recombinant DNA concerns grew out of a set of research projects. Scientists thinking about possible consequences of error in the midst of doing science seems a clearly desirable rather than undesirable aspect of scientific work.

Perhaps those who wish to define a moral exemption for scientists might say that scientists should not consider the consequences of error when they are making general empirical claims, but only when they are considering certain actions they might take that may harm others. However, the distinction between making a general empirical claim and making a judgment about the safety of a scientific process seems dubious at best, showing how difficult it is to keep practical and theoretical reason distinct in this border area. No discussion of the recombinant DNA controversy or the risks of the Trinity test could go forth without making claims about potential risks and the benefits of proceeding. One would sincerely hope that in a case like the Trinity test, scientists would demand a greater burden of proof that an atmospheric chain reaction would not occur than if they were considering the risks of accepting some more mundane claim. Less uncertainty is tolerable when the consequences of error are so high, and in making the claim that the risk of catastrophe is low, we would want a high degree of surety. This *is* the weighing of uncertainty around an empirical claim against the ethical consequences of error. In making crucial empirical claims about the likelihood of error, scientists should consider the consequences of error, for such consideration is what requires different burdens of proof in different contexts. No tenable grounds for a general moral exemption for scientists may be found in a distinction between a judgment regarding the safety of an action and a judgment about an empirical claim.

Why go through these rather desperate attempts to articulate some kind of moral exemption for scientists from a general responsibility we all share? The fear of many proponents of such a moral exemption for scientists seems complex, embedded in a host of considerations. Some are concerned that science will lose its general public authority if a role for social or ethical values is admitted. Yet as chapter 1 demonstrates, no particular piece of science is obviously authoritative anymore, as debates over sound science and junk science rage. The value-free ideal has done nothing to help this debate. Some defenders are concerned that if the current value-free ideal fails, we will revert back to the horrors of science history such as Lysenkoism and Galileo's trial. As I will show in the next chapter, this leap is not warranted. We can reject the current value-free ideal while still holding a constrained role for values in science. Some are concerned that there will be no objectivity left in science if the value-free ideal is gone. Chapter 6 shows that this is not the case, that science can be objective even as it is value-laden. Some are concerned that scientists will be usurping undue authority in our society if the value-free ideal is relinquished. In chapters 7 and 8, I will argue that this need not be the case, particularly as scientists are encouraged to make their judgments explicit. Indeed, the possibility for genuine public input into science increases when the value-free ideal is relinquished.

In sum, there are two important arguments against a moral exemption for scientists. First, there are many cases where such a moral exemption would be very harmful, and no clear boundary can be located that would require moral reflection in some cases and not others. Thus, no blanket exemption is tenable. Second, no convincing argument has been articulated to give scientists even an occasional moral exemption from the consideration of the consequences of their work. With no clear argument for the exemption and a host of considerations against the exemption, I will reject the idea of a general exemption for scientists for the remainder of this work. What are the implications of this rejection for scientists? In the remainder of this chapter, I will discuss the implications for scientists in the advising process. I will also articulate the limits placed on the burden of this responsibility and how the standards of reasonable foresight might work for scientists in practice.

I have argued that scientists can and should bear the burden of the general responsibility to consider the consequences of error. What does the bearing of this burden mean for the role of science in policymaking? As noted in chapter 2, the importance of scientific advice over the past century has been increasing. This advice is used to guide decisions of major social importance. We certainly need scientists in this advisory role, helping to shape and guide decisions with significant technical aspects. We need scientific advice on such issues as which health problems may be posed by air and water pollutants, what options we have for nuclear waste disposal and what risks are associated with them, which drugs should be released on the market and with what labeling, and so forth. What are the moral responsibilities of scientists in this crucial role?

Certainly we should expect honesty and forthrightness from our scientists. To deliberately deceive decisionmakers or the public in an attempt to steer decisions in a particular direction for self-interested reasons is not morally acceptable. Not only would such a course violate the ideals of honesty central to basic science, but it would violate the trust placed in scientists to provide advice, and it would violate the basic ideal of democracy, that an elite few should not subvert the will of the many for their own gain.

But whether scientists should be honest is not the dispute at the heart of the value-free ideal for science. The issue is whether scientists should consider the consequences of error in their advising, including errors that may lead to harm. Should scientists worry about being reckless or negligent when giving advice; that is, should they consider the consequences of error when deciding how much evidence is enough to support making an empirical claim? As noted in chapter 3, Rudner developed this line of argument in the 1950s, suggesting that scientists should consider not just the extent of the uncertainty inherent in any scientific statement, but should also weigh the importance of the uncertainty by considering the consequences of error. Such a weighing would require the use of social and ethical values in scientific judgments. So, the issue is whether or not scientists, when placed in official advisory positions or simply providing the general public with authoritative advice, should consider the consequences of possible errors in their advice.

If my arguments in the preceding section are convincing, we certainly would want scientists to consider the potential consequences of error when giving advice. Scientists have the same obligations as the rest of us not to be reckless or negligent, and this obligation also holds when a scientist is making an empirical claim (the basic component of advising). This means that when a scientist makes an empirical claim in the process of advising, they should consider the potential consequences if that claim is incorrect. In the advising context, this includes possible social and ethical consequences of policymakers acting on the basis of the empirical claim. The scientist acting as an advisor should consider the extent of uncertainties around the claim and the possible consequences of incorrectly accepting or rejecting the claim, and they should weigh the importance of the uncertainties accordingly. Thus, science advising should not be value free.

Note, however, that social and ethical values can legitimately enter into the advice only through the weighing of uncertainty. The scientist should not think about the potential consequences of making an accurate empirical claim and slant their advice accordingly. Only in the weighing of uncertainty do social and ethical values have a legitimate role to play when deciding, based on the available evidence, which empirical claims to make.

An example will help to clarify this ideal for fully responsible science advice. Suppose a scientist is examining epidemiological records in conjunction with air quality standards and the scientist notices that a particular pollutant is always conjoined with a spike in respiratory deaths. Suppose that this pollutant is cheap to control or eliminate (a new and simple technology has just been developed). Should the scientist make the empirical claim (or, if on a science advisory panel reviewing this evidence, support the claim) that this pollutant is a public health threat? Certainly, there is uncertainty in the empirical evidence here. Epidemiological records are always fraught with problems of reliability, and indeed, we have only a correlation between the pollutant and the health effect. The scientist, in being honest, should undoubtedly acknowledge these uncertainties. To pretend to certainty on such evidence would be dishonest and deceptive. But the scientist can also choose whether or not to emphasize the importance of the uncertainties. And this is where the weighing of the consequences of error comes in. If the scientist accepts that claim as sufficiently reliable (not certain) and is wrong, little expense will have been accrued as policymakers act on the scientist's advice. If the scientist rejects the claim as insufficiently reliable or well supported and is wrong, public health will be damaged substantially. In this admittedly easy case, the fully responsible advice would be to note the uncertainties but, on the basis of the consequences of error, suggest that the evidence available sufficiently supports the claim that the pollutant contributes to respiratory failure. Such advice would be fully morally responsible, and not value free.

One might argue that considering the consequences of error when giving advice on an advisory panel or some other formal advising mechanism is acceptable, and indeed a responsibility of the participants in such a mechanism. However, one may still want to claim that scientists in general should eschew the considerations of such consequences in other forums when making empirical claims. In other words, one may want to hold to a distinction between the scientist qua scientist and the scientist qua advisor, accepting the need for values in the latter role while rejecting them in the former.

The attempt to draw a distinction in the practices of scientists between the two roles, scientist and advisor, is dubious at best. The scientist is called on to be an advisor because she is a scientist, and the advice is to be based on her expertise as such. But the problem runs even deeper. Scientists hold a broadly authoritative position in our society, regardless of whether they are functioning in a formal advising role or not. Thus, when scientists make empirical claims, whether in scientific conferences, in science journals, or on an advisory panel, those empirical claims carry with them a prima facie authority. This is why science journalists are interested in science conferences, why scientific journals are covered in the general press, and why scientists are asked for their views in areas related to their expertise. This basic authority is what turns a general responsibility to consider the consequences of error into a particularly important responsibility for scientists. Because their empirical claims carry with them this prima facie authority, the potential consequences of error can be more far-reaching than for a nonscientist. If I make a claim that a plant is a dangerous invasive species that can be eradicated, no one is likely to listen or act on the basis of my claim. There are few consequences of error for me to consider. But an ecologist making such a claim is far more likely to be believed, and to have their claim acted upon. The authority of science in society makes a distinction between scientist qua scientist and scientist qua advisor untenable. It also can place a heavy burden on scientists to meet their moral responsibilities fully. It is time to reflect upon the limits of this burden, not just the full extent of its reach.

Limits to the Moral Responsibilities of Scientists

The positions of Bridgman and Lübbe are untenable. Simply because scientists provide us with important knowledge cannot and does not exempt

them from basic moral responsibilities, which include reflecting on the implications of their work and the possible consequences of error. Yet scientists cannot be responsible for every use or misuse of their work. Fortuitously, the basic understanding of moral responsibility articulated in the beginning of this chapter provides guidelines for when and how much moral responsibility scientists must bear for their work and their advice.

As noted above, when considering the consequences of error, a person is not responsible for every consequence that follows from their action. Rather, a person is held responsible only for those consequences that are reasonably foreseeable. Thus, in our earlier example of setting a field ablaze to clear brush, it is easily foreseeable that on a hot, windy day the fire could get out of control and burn a neighbor's property. To proceed with a burn under such conditions is reckless (if one sees the risk but does not care) or negligent (if one fails to foresee the risk). We can demand such foresight in these cases because any reasonable person would be able to foresee the risks from such an action.

Scientists should be held to a similar standard of foresight, but indexed to the scientific community rather than the general public. Because scientists work in such communities, in near constant communication and competition with other scientists, what is foreseeable and what is not can be readily determined. As with other ideas in science, potential consequences of error spread quickly, and scientists discuss pitfalls, dangers, and uncertainties readily. Another example from nuclear physics shows the ready benchmark of foreseeability that can exist in science. Throughout the 1930s, after the discovery of the neutron by James Chadwick, nuclear physics blossomed. The neutron provided a ready tool for the probing of the nucleus, and many discoveries followed. However, none of these discoveries seemed to have any important implications outside of nuclear physics. The idea that usable energy may be derived from nuclear processes was thought to be unfounded speculation, or "moonshine," as Ernest Rutherford put it. All this changed in December 1938 with the discovery of fission. No one had thought that an atom might actually split into two large chunks before Lise Meitner's insight into Otto Hahn and Fritz Strassmann's experiments. There was no foreseeability of the usefulness of nuclear processes for a bomb or useful energy production until this point.[7] As word of fission crossed the Atlantic in January 1939, it was clear to all what fission meant: the possibility for useful nuclear energy, either as a power source or a bomb.[8] In the political climate of the time, this worried many, and debate soon followed on how to proceed. But the foreseeability of nuclear weapons can be pinpointed to this moment in time. It would be absurd to have expected Chadwick to have foreseen this when he discovered the neutron in 1932. By the spring of 1939, few nuclear physicists had not foreseen the disturbing potential of fission.[9] The nuclear physics community provides a ready benchmark for what a reasonable scientist could foresee in this case.

The example of fission and the sudden usefulness of nuclear physics is an example of the foreseeability of the direct consequences of a line of research. Much of this chapter is about the foreseeability of error and its consequences. Here, too, scientific communities provide ready benchmarks. The concern of scientists working with recombinant DNA over the potential consequences of laboratory error, particularly the accidental generation and release of new biohazards, led directly to the moratorium on such research by 1974. Scientists in the field could readily foresee the dangers and worked together to forestall them. And the consequences of error in making empirical claims were also generally agreed upon and foreseeable. Today, it is precisely the concern over these consequences that drives so many of our more contentious technical debates. For example, if a chemical is known to cause cancer in humans, a regulatory response can be predicted. So making the empirical claim that a chemical causes cancer brings with it clear consequences of error, namely unnecessary regulation (unnecessary if the claim is erroneous). On the other hand, not making the empirical claim (perhaps because one suspects the chemical but does not think the evidence sufficient) also carries with it clear consequences of error, namely the cancer will continue to be caused. The consequences of error are readily foreseeable by all, and are often a central engine of the debate.

Requiring that scientists consider the consequences of their work does not mean requiring that they have perfect foresight. The unexpected and unforeseen can and does happen. Holding scientists responsible for unforeseen consequences is unreasonable. What is reasonable is to expect scientists to meet basic standards of consideration and foresight that any person would share, with the reasonable expectations of foresight judged against the scientist's peers in the scientific community. Thus, the moral burdens on scientists are not unlimited. They are held to only what can be foreseen, and thus discussed and considered.

It must also be noted here that scientists need not carry this burden alone. As mentioned above, scientists already use the assistance of internal review boards to help them meet their responsibilities with respect to methodologies involving human and (some) animal subjects. However, if

scientists found it useful, they might consider convening similar kinds of bodies, either permanent or temporary, that could help them make difficult choices when they arise. I will discuss some of these possible mechanisms in the final chapter. There are many interesting ways in which scientists can shift the burden of reflecting on these issues to others. However, the scientist can never abdicate the responsibility completely. Often only the scientists on the cutting edge will fully understand the implications and the risks of their work. We will always need them to reflect on those as they proceed into the unknown.

Conclusion

We all share a general responsibility to consider the consequences of our choices, including the consequences of error. This responsibility extends to the making of empirical claims, an activity central to any examination of science in public policy. Because scientists have no good reason to be exempt from this general responsibility, and indeed we have good reason to want them to shoulder (at least in part) the responsibility, scientists must weigh the consequences of error in their work. The values needed to weigh those consequences, and thus determine the importance of uncertainties in science, become a required part of scientific reasoning. The value-free ideal for science can no longer be held up as an ideal for scientists.

Some might argue at this point that scientists should just be clear about uncertainties and all this need for moral judgment will go away, thus preserving the value-free ideal. It is worth recalling Rudner's response to a similar argument from Jeffrey discussed in the previous chapter. Even a statement of uncertainty surrounding an empirical claim contains a weighing of second-order uncertainty, that is, whether the assessment of uncertainty is sufficiently accurate. It might seem that the uncertainty about the uncertainty estimate is not important. But we must keep in mind that the judgment that some uncertainty is not important is always a moral judgment. It is a judgment that there are no important consequences of error, or that the uncertainty is so small that even important consequences of error are not worth worrying about.[10] Having clear assessments of uncertainty is always helpful, but the scientist must still decide that the assessment is sufficiently accurate, and thus the need for values is not eliminable.

The demise of the value-free ideal may be disturbing to some. What is the proper role for values in science, if science is not value free? Can values, any values, play any role whatsoever in science? Can they dictate a scientific result? Is the objectivity of science, the source of its authority, doomed if the value-free ideal is rejected? I will take up these difficult questions in the next two chapters and argue that values should play only a constrained role in scientific reasoning. Thus, the demise of the value-free ideal does not mean values can run roughshod over evidence and reasoning. Science can be objective while remaining value saturated.

CHAPTER 5

THE STRUCTURE OF VALUES IN SCIENCE

EVEN WHEN MAKING EMPIRICAL CLAIMS, scientists have the same moral responsibilities as the general population to consider the consequences of error. This apparently unremarkable statement has some remarkable implications. It means that scientists should consider the potential social and ethical consequences of error in their work, that they should weigh the importance of those consequences, and that they should set burdens of proof accordingly. Social and ethical values are needed to make these judgments, not just as a matter of an accurate description of scientific practice, but as part of an ideal for scientific reasoning. Thus, the value-free ideal for science is a bad ideal. However, simply discarding the ideal is insufficient. Although scientists need to consider values when doing science, there must be constraints on how values are considered, on what role they play in the reasoning process. For example, simply because a scientist values (or would prefer) a particular outcome of a study does not mean the scientist's preference should be taken as a reason in itself to accept the outcome. Values are not evidence; wishing does not make it so. There must be some important limits to the roles values play in science.

To find these limits, it is time to explore and map the territory of values in science. This will allow me to articulate a new ideal for values in science, a revised understanding of how values *should* play a role in science and of what the structure of values in science *should* be. I will argue that in general there are two roles for values in scientific reasoning: a direct role and