CHAPTER 1

Outcomes of Scientific Research

LIKE MANY OTHER PHILOSOPHERS, I BELIEVE THAT IT IS important to be quite clear on how a question is to be interpreted, or "read," before any attempt is made to answer it. It is necessary to understand how pure science can lead to applications, and to other outcomes, if we are to address questions about responsibility in connection with pure science. So the vital clarification here concerns the role of pure scientific research—that is, research less directly concerned with outcomes than applied research is. Pure research is related to outcomes via applied research, so if we are able to determine how the former leads to outcomes, we will also understand how the latter does so.

Two kinds of outcomes will be considered in this chapter: those that have to do with *technology,* such as products and processes—the central concern of this book—and those forthcoming from science considered as a body of *ideas.* Put simply and crudely, science can affect both our material welfare and how we think about things, and therefore science is responsible for both sorts of outcome. However, it is not the case that "technological outcomes" are solely the province of applied research, while pure research exclusively gives rise to outcomes from science as a body of ideas. To clarify the relations among the two kinds of outcomes and the two kinds of research, it is necessary in the first place to carefully draw the distinction between pure and applied research and then to see how research leads to outcomes.

Outcomes

Suppose the work of a particular scientist leads to some outcome. While this seems to be the normal course of events, one might still ask whether that scientist is responsible for that outcome. More information, however, is necessary if this question is to be answered or even taken seriously. At the least, we need to know what the outcome is and how it is related to the scientist's work—to say that his work "led" to the outcome is not precise enough. If a scientist, for example, does his research at a university, then there is a good chance that it is undertaken without any particular application in mind beyond publication, which makes his results known to the scientific community for further discussion.[1] He might also see publication as a step toward a grant or promotion. The proximate result of pure scientific research is thus the scholarly scientific paper, which I will call the intended or "proper" product of such research. With respect to this aspect of science, the question posed above is not hard to answer: the fact that a scientist is named as an author of a paper should suffice to establish his responsibility for publication.

But other kinds of outcomes stem from pure research as well. Science has widespread application as a basis for technology, and technology in turn can have equally broad effects. If a scientist is an *applied* scientist who consciously seeks to design or develop some product, then again it initially appears obvious that he is responsible for that outcome if he did some of the work leading up to it. But there is a complicating factor: unlike scientific papers, technologies, products, and so on cannot be brought into the world by one person, at least not anymore. So while the single-author paper still exists, the "single-author technology" does not. Further, all different "kinds" of people, so to speak, are responsible for such outcomes: scientists, engineers, entrepreneurs, industrialists, even the military and politicians. If a scientist is an employee of a corporation, for instance, it might even be said that he has *no* responsibility for the outcome under consideration; he just sells his labor to others, who use it for their own ends.[2] Thus, answering the question now becomes much less straightforward than in the earlier case.

If these two cases are "combined," the following situation results: The scientist is aiming to publish a paper, the "proper" product of his pure research project. He does so, and the publication makes a contribution to some technology or some product thereof. In other words, "outcome" is here understood in accordance with the second case, but the scientist remains in the context of pure research, as in the first case. Now, it is possible that the scientist realizes that his work might well have such an external outcome—he might, for instance, have accepted a grant from a company in exchange for agreeing to provide it with his results before publication—although that is not his specific goal. Alternatively, the application might be unexpected—a possibility that is not uncommon, as pure research often becomes applied over time. This third case, then, raises new considerations, over and above those that must be taken into account for the other two. For example, can someone be responsible for an outcome that was not his aim? And what about an outcome that he did not have in mind or imagine might be possible?

To answer such questions, a sensible first step would be to gather some information about the research *group* to which our scientist belongs. Not all scientists work in groups, but most do, and thus most research is a cooperative effort. This is important from the present standpoint if there is something distinctive about group or *collective* responsibility. Suppose several individuals are needed to produce a given outcome. Does their collective activity add anything over and above each person's effort, as regards their responsibility for the outcome? Two views exist on this question: one holds that collective responsibility merely reduces to individual responsibility, and one holds that it does not. Those who maintain the latter position normally distinguish among different kinds of groups and attribute a special kind of collective responsibility to those with sufficient organization, such as corporations. Here the issue becomes whether any such thing as genuine collective responsibility exists, and, if so, whether research groups display sufficient organization to be charged with it. Of the scenarios just considered, the first and the third have the same institutional setting, or so I will assume—namely, the university or research institute. The second case, by contrast, is set in some industrial system, in some company or corporation that, I will assume, has its own research facility. How different, then, are these settings?

The Research System

Ascertaining the characteristics of these settings is key, because if they are radically different, then they may give rise to two altogether unconnected notions of scientific responsibility: that of the academic scientist on the one hand, and that of the industrial scientist on the other. But pure and applied research are allied activities. Taken together, they constitute the scientific research system, or what J-J. Salomon simply calls the "research system." I will argue that what distinguishes pure and applied research is not in fact the "content" of the work: the work of pure scientists is not of a different kind compared to the work of applied scientists. What then does distinguish pure from applied research, given that they are indeed different? I do not think it is possible to draw this distinction by means of philosophical analysis; rather, such fields as science policy studies and the sociology of science are most germane here.[3]

I want to return to the idea that the proper product of pure scientific research is the scientific paper, particularly the question of why scientists publish papers.[4] Publication might be a kind of end in itself, or it could have some further purpose. Scientists could simply be interested in making and communicating discoveries, which seems a more noble goal than getting promotions, swinging grants, and building a research empire. But browsing the specialist journals to see what has been discovered lately, I believe, would typically leave the questioner profoundly unexcited, unconvinced that science has much of interest to communicate, especially if one looks through journals in the physical sciences. Scientific papers on the whole record minute details of systems created in the laboratory, often the values of variables. Such values can provide the key to startling applications (see chapter 2), but surely they hold little interest in and of themselves. How, then, might the sociology of science help elucidate the purpose of publication?

Several different approaches exist in this field, some of which go beyond the traditional division of labor that would see the sociologist's task as discovering the nature of the scientist's activities and the structure of the scientific community and explaining away mistaken beliefs. Some sociologists, including Bruno Latour, have addressed issues that were earlier the exclusive province of philosophers, provoking sharp territorial reactions in some quarters. Philosophers' objections aside, Latour's view of what pure researchers do and why they do it is relevant. Specifically, what Latour discovered during his famous sojourn at the Salk Institute is, simply, that scientists publish papers: "The production of papers is acknowledged by participants as the main objective of their activity" (Latour and Woolgar 1979, 71). Much other evidence exists in favor of this claim, which is undisputed in the sociology of science. When Latour remarks that scientists take publication to be their main objective, he could mean that it is an end in itself, or that it is a common aim of all scientists, or that it is something that they must do if they are to do anything else. On any of these understandings, publishing papers is at least a candidate for the *criterion* necessary for distinguishing pure scientists in terms of what they do, even if Latour's claim is not anything like a definition of science.[5]

It is not, however, true that applied scientists never publish anything, so the criterion must be articulated a little further. It might then be suggested that what differentiates the products of applied research is that they are not the intellectual *property* of their authors, as are the products of pure research, and, further, that they are intellectual *capital* and, as such, that access to them is restricted. On the other hand, the products of pure research are freely available: anyone can buy a journal, access Web sites, and so on. The results of applied research, by contrast, are normally restricted and held in-house. While this may be the case on the whole, however, applied research or technical journals do exist, as do some examples of applied research results being made available immediately and for free. For instance, a U.S. biotechnology company sequenced the rice genome in 2001 and declared that it would make available any useful genetically modified variety free to Third World countries (one hopes for more success here than was achieved by the Green Revolution). While the original "wild-type" sequencing work seems to be typically pure research, the modified version is typically applied.

That pure scientists normally publish papers and applied scientists normally do not is thus not a sufficient distinguishing criterion. So it is appropriate to turn to another field: science policy studies. Science policy studies are focused on government and industry plans for science, including governmental funding. One issue here concerns the appropriate mix of pure and applied research for fulfilling national objectives: as taxpayers' money is being used to support science, it is only fair that taxpayers eventually benefit from the research. It appears, then, that applied research should get the lion's share of available funds, but strong arguments have also been made in favor of significant funding for pure research (see Bush 1945 for a classic statement). To carry on this debate and to implement policies, the two kinds of research must first be distinguished. This is done not with direct reference to the product of the research, but with reference to its purpose or aim. Thus, applied research, in the field of science policy studies, is research that aims at some *application,* while pure research is done *for its own sake* (ASTEC 1981, 3). The question now becomes not what counts as a scientific paper, but what counts as an application.

Applied science is *practical;* it often enables us to do things that we could not do otherwise. Applied research is therefore also connected to technology: in fact, applied research often gives rise to technology. To say that pure science is done "for its own sake" is to draw a contrast with this practical side of science. What is "applied" in applied research is "science itself," its theories and methods. Pure research, on the other hand—science done for its own sake—can also come to have applications later on. So again, there are not two distinct sorts of things, both called science; rather, there are two kinds of scientific projects, one that aims ultimately beyond science itself, to practical matters, and another that does not. If the "context" of a scientific research project is understood as what gives rise to or defines its aims, then in this sense, context will determine whether research is pure or applied. More could be said about the research system itself, from the perspectives of science policy, sociology of science, and other views, that could lead to an even greater understanding of the system's elements and how these condition the nature of research projects. Interesting though that might be, what is essential here is to recognize that different types of research projects represent different kinds of aims.

The Realist View of Pure and Applied Research

Suppose that the specific aim of an applied research project is to come up with a blueprint or design for a product or process. The project may have several stages, with various tasks and problems to be solved. When the design is achieved, in other words, the project moves on to the development of a prototype, testing, and so on. The design, then, eventually achieves some output or realizes some function—namely, what the product or process is supposed to accomplish. How is this possible? How is it possible, to take a simple example, to design a working clock or watch or chronometer? Is successful design the result of good luck or trial and error? Clearly not, at least in the case of sophisticated electronic or atomic clocks. A clock requires a reliable periodic system, such as a pendulum driven by gravity or a spring, or a vibrating crystal, or a radioactive element that decays at a constant rate. Clocks are possible because such systems exist, because we know about such systems, and because we are able to incorporate them into artifacts via design. But periodic systems are also studied "for their own sake." Radioactive decay, for instance, discovered at the beginning of the twentieth century, was investigated by Pierre and Marie Curie and by many others. But the radium studied by the Curies and the radium incorporated into an atomic clock are one and the same element, and the clock works as a consequence of the same "natural" properties of the element that the Curies investigated.

A design, then, makes use of the properties of natural things and the principles that describe these properties, which are the self-same things that are studied "for their own sake" in pure science. This is, in broad outline, the realist interpretation of science and technology.[6] This view, in the first place, explains the relationship between science and technology—how it is that science underpins technology. It also addresses the problem of the relevance of pure research to practical matters, because it holds that both pure and applied research make use of the same theories, ideas, and results. Thus, the maker of a grandfather clock uses the relation between the length and the period of the pendulum to design the mechanism—the clockwork—to fix how far the hands must move for each swing back and forth; it is the exact same relation studied, in another time, by Galileo. This, of course, is not a startling or radical conclusion: most will readily agree that science can be applied because artifacts, while not naturally occurring, are natural objects. As such, they conform to our theories: indeed, they must, because they are designed in light of these theories. It is possible, then, to conclude that there is no intrinsic difference between pure and applied research, in that the "content" of pure research is not different in nature or character from the "content" of applied research. While this conclusion clearly requires qualification, the view of science on which it rests might also be questioned: Is the realist account of science well supported? What are the alternatives?

The realist view of the relationship between science and technology is an extension of the realist account of science itself—of how science informs technology. Realism, of one stripe or another, has become the prevailing view in the philosophy of science, but this predominance in itself does not guarantee that it yields a good explanation of the relationship between science and technology. The core of all realist accounts of science, in other words, is still a realist *interpretation,* according to which scientific theories are literal statements about the natural world, in that their kind, property, and relation terms refer to corresponding items in the world and are true or false depending on what the world is like.[7] Realists, then, appeal to this interpretation to account for such things as experimental practice, explanatory and predictive roles, and the success of science.

Realism is compatible with several different scientific methodologies, but all acceptable methodologies must purport that theories are tested, or "applied" in some way, by means of experiment—experiment providing the means for scientists to confirm what theory says about the world. The fact that science is successful in this regard, with most predictions emerging in favor of the theory being tested, has a straightforward realist explanation: theories make predictions that are borne out in practice because theories describe the world, more or less accurately, and because experiment increases the store of information about the world, more or less reliably. This formulation is the so-called Ultimate Argument for realism.[8] This argument, first put forward in the early 1960s, was influential in eroding the (logical) empiricist view of science that held sway at the time. Part and parcel of the empiricist position was *instrumentalism,* according to which the characteristic terms of scientific theories do not refer literally to kind, properties, and relations but are shorthand for directly observable effects.

So, for instance, while realism takes "radium," "alpha particle," "gravity," and so on to refer to real things, instrumentalism denies the existence of such microscopic particles and forces. The instrumentalist understands "radium" as shorthand for talking about certain observable effects, such as those resulting when one uses particle counters and fluorescent screens in conjunction with radioactive substances. Theories in this view are "instruments" for making predictions, couched in the language of observation, and they are successful because they comprise descriptions of such observations. Both realists and instrumentalists would therefore predict that a Geiger counter will register if placed near a sample of radium, but only the realist could explain why it registers. The instrumentalist would instead perceive the "radium" as simply a compendium of observable effects, including effects on particle counters, so all the "explanans" could do is restate the explanandum. While the realist would appeal to the fact that radium emits charged particles that stimulate the counter to register, the instrumentalist can only observe that science succeeds when it employs good instruments but cannot explain why certain instruments are good.

The Ultimate Argument for realism (and the shortcomings of instrumentalism in this regard) is more pertinent here than any of the other arguments in favor of realism because it has an obvious extension to technology. Technology is successful, in the realist view, because the practical purposes for which it was devised are, on the whole, fulfilled—because it makes use of what we know about the world. Pendulums are reliable periodic systems because they are acted on by a constant gravitational force; radium can be used in an atomic clock because it emits alpha particles at a constant rate. The instrumentalist explanation of the success of technology, in contrast, is as unsatisfactory as the instrumentalist explanation of the success of science. "Radium" here is nothing but a shorthand description of observable effects, including the Geiger counter's clicking at a constant rate, which means that the instrumentalist has no basis on which to explain the constancy of this observable effect and hence why radium can be used in an atomic clock.[9]

Returning to my assertion that there is no intrinsic difference between pure and applied research, I do not mean that a particular piece of work could *always* appear either as a paper in a journal or as a classified technical report. For instance, suppose that a problem has already been solved for a given set of parameters (initial and boundary conditions), and the results published in a journal, but that another solution, for another set of parameters, is needed for an application, perhaps for "scaling up" an effect. An editor might tell the author submitting a paper addressing this application that it does not add to the scholarly literature on the subject, but not publishing the work on which some applied research is based would be only a contingent fact about the history of science.

There are also several examples of scientific research that was first done in an industrial setting but whose results clearly could have been, and in some instances eventually were, published in learned journals. One of the best-known examples is the discovery and subsequent use of the transistor by William Shockley and his coworkers in the Bell Laboratories in the 1950s. This discovery had an almost immediate application, but at the same time, it was important in the field of material science. Even better known are a series of discoveries associated with the making of the atomic bomb, in particular Enrico Fermi's realization of a self-sustaining nuclear chain reaction in 1942. This was simultaneously the first step in the manufacture of plutonium, the initiation of nuclear power, *and* confirmation of important theoretical work in nuclear physics. At another time, Fermi's work would have first been published in scholarly journals.[10] To clarify, the work of Shockley, Fermi, and others is classified as applied research because of the institutional setting or *context* in which it was done, not because of its "intrinsic" character. This context dictated that the research's aim was practical, that it was intended for some purpose beyond publication in a learned journal. It should be acknowledged here that some industrial research is classified as "basic" or "pure"—again, Shockley is an example—although recently the labels "strategic basic research" and "generic research" have become fashionable. Such research has no specific end in view and so differs from applied research, but it is conducted in areas that look promising from the point of view of future application and so differs from pure research in that its funding rationale is application.[11]

On the other hand, a paper on quantum cosmology written by Stephen Hawking would most likely not ever be reproduced in an industrial setting. Aside from this particular author's individual brilliance (and hence the unlikelihood of finding others like him), quantum cosmology is not a field that is likely to yield immediate (or any) applications.[12] So, just as some technical reports would not be suitable for publication in academic journals, some pure research would not be forthcoming from industrial laboratories. Indeed, highly theoretical research—such as quantum cosmology, the foundations of quantum and relativity theory, and so on—is not likely to be sponsored by those interested in applications. Nevertheless, such research *can* have application, as was famously demonstrated in the Manhattan Project, when Einstein's formula $E = mc^2$ was used to estimate the explosive power of a nuclear bomb and explain the source of its energy. This formula amounted to a prediction of the special theory of relativity in that it was a consequence of Einstein's use of Lorentz invariance. Einstein thought of this formula as a peculiarity and did not expect that it would ever be realized or even confirmed. Thus, although highly theoretical research is sometimes applied, this application is usually unanticipated or inadvertent and hence not something that can be planned.

If highly theoretical research does not typically lead to outcomes, or only does so well after publication—the Einstein formula was only confirmed some thirty-five years after special relativity was discovered—or only has applications that cannot be anticipated, then the issue of responsibility for outcomes would not appear to be a live one. It is worth asking here how common such research is, for if it is widespread, then pure research will not be of great relevance when it comes to outcomes, and hence the issue of responsibility for such outcomes will not really arise. In fact, however, highly theoretical research is not at all common, especially when considering the huge amount of scientific research undertaken worldwide. In Kuhnian language, most scientific research is puzzle-solving, work within a paradigm, whereas highly theoretical research would instead be directed toward formulating new paradigms, working on foundational problems within paradigms, addressing incompatibilities between paradigms, and so forth. Not many scientists are so engaged. Therefore, the question of the responsibility of pure researchers for outcomes is of more than merely marginal interest.

It will be necessary to return to the distinction between pure and applied research. But if the realist account of the relationship between science and technology is accurate, then we can at least generally see how pure science leads to outcomes—namely, it does so in the *same way* as does applied research. The difference, again, lies not in the research's intrinsic character but in its institutional setting, or whatever else determines a project's aims. Thus, if scientists working in industry, for the government, or for some other organization that aims to produce outcomes have different responsibilities than do pure scientists who work in universities or research institutes, then such differences will be a function and consequence of the institutional settings, not of the research's intrinsic character.

Pure Science as "Ideas"

I remarked earlier that the proper products of scientific research can often seem dull and uninteresting to the wider public. In the past, however, some scientific works—indeed, works full of numbers and calculations—were of profound interest and importance, and their influence has been felt far and wide: Copernicus's *De Revolutionibus Orbium Caelestum,* Kepler's *Astronomia Nova,* Galileo's *Dialogue on Two New Sciences* (Galileo published in the vernacular), and Newton's *Principia Mathematica Philosophia Naturalis* are examples from the Scientific Revolution. These books provided convincing evidence against the Greek worldview, which placed the earth at the center of the universe, and their authors, although devout Christians, thus seemed to challenge the authority of both the Bible and the Church.[13] It is clear that science as knowledge—as a body of ideas—has had a huge impact on our understanding of the world and our place in it. And of course, when we move beyond the Scientific Revolution to the late eighteenth and nineteenth centuries, we find geologists demonstrating that the earth is much older than previously believed, and Darwin bringing us to think of ourselves as having evolved naturally from animals, rather as than special divine creations.

What science does as a body of ideas is therefore perhaps of even greater significance than what it does as the foundation of technology—assuming that the comparison can be made in a meaningful way. If so, the notion of "outcome" must be broadened to include such things. Further, this broadening might serve to distinguish pure from applied science in a way not apparent when attention is directed exclusively toward technology; it may indeed reveal some "autonomous role" for the pure scientist with respect to outcomes. This impression is reinforced by the fact that it is not so much the process and reasoning that lead to such important discoveries but the discoveries themselves, detached from such details and expressed in simple terms, that have the greatest impact. Copernicus, for example, constructed a system of epicycles in which the earth moves around the sun (although the latter was not itself at rest) that was able to save the "quantitative phenomena," but it was the conclusion of his work, the heliocentric hypothesis, that carried the greatest force. Similarly, while Kepler's works are highly technical, the inclusion of three simple expressions that describe the motion of the planets widened its impact. What is important about pure science as a body of ideas therefore seems to differ from what is important about science as a basis for technology. Thus, beyond broadening the concept of outcome, taking account of science as a body of ideas might also lead to a reconsideration of the decision to assimilate the relation between pure scientific research and outcomes to that between applied research and outcomes.

However, while it would be a mistake to underestimate the importance of the work of such scientists as Copernicus, Galileo, and Darwin—to deny that responsibility for science as a body of ideas is as important as responsibility for science as the foundation for technology—the claims made in the previous paragraph overstate the case for the "autonomy" of pure research with respect to outcomes. First of all, does this aspect of pure science really distinguish it from applied science? Notice here that it is perfectly *possible* for an applied research project to produce novel and important ideas. It is also suggested above that the *conclusion* of research has the greatest impact here—that the earth goes around the sun, that the earth is many millions of years old, that we are descended from animals—rather than the painstaking reasoning supporting such propositions. But the same is of course true of the products of science-based technology: we simply want to use these products, not to understand the designs on which they are based or the physical principles in virtue of which they work. While some people might want to understand how, say, a computer works, or even build one, such knowledge is clearly not necessary to use a computer, nor is it the manufacturer's aim. Thus, the products of science-based technology and propositions rooted in science as a body of ideas are alike in that both are *outcomes* of science and can be used or grasped without understanding the underlying theory. By so broadening the concept of an outcome, certain conclusions or "propositions" are acknowledged as analogous to the products of technology in that both kinds of outcome stand in a *similar* relation to pure science.

The issue of the *relative* importance of the two kinds of outcomes becomes relevant here, but I believe in fact that science as a body of ideas is at the present time relatively *unimportant,* at least as regards the physical sciences, compared to science as the foundation for technology. One explanation for this situation is hinted at above—namely, that pure research in the physical sciences is aimed at uncovering minutiae; the paradigms under which it is conducted are themselves technical and as such difficult to understand. But Copernicus's epicycles were also hard to understand, so this cannot be the whole story. However, when propositions from modern physical science are "detached"—counterparts to "the earth goes round the sun"—the resulting statements are more likely to confound than to illuminate. For instance, it is said that quantum mechanics implies, among other things, that the world is "indefinite" until experienced by an "observer"; that while there is no action at a distance, there is "passion at a distance"; and that there exist two radically different equations of motion for measurement and nonmeasurement processes. Difficult books have been written on these matters, and popularizers of science have struggled to explain them, but their very reconditeness ensures that they will not have much of an impact outside the community of specialists.[14]

If this judgment about the physical sciences is accurate—and more would need to be said to back it up properly—the situation seems rather different in the biomedical sciences. Since the discovery of the structure of DNA, steady progress has been made toward uncovering the human genome, which is now essentially complete. This project seems to be a good example of strategic basic research, in that the rationale for all the resources that have been invested in it is the promise of future applications, although their precise form is not yet known. Advances are nonetheless anticipated in such areas as understanding the genetic basis of more diseases. Concurrently, however, these developments as a body of ideas could also have societal implications. The category of outcomes that was earlier characterized as significant propositions implied by science as a body of ideas is therefore not after all an empty one, and so it seems that it must not be overlooked. On the other hand, some argue that these outcomes can indeed be set aside when it comes to questions about the *responsibility* of the scientist. While this argument cannot be adequately supported, it is still worth considering, since it allows us to probe yet another perspective on the topic of scientists' responsibility for outcomes—the idea that science is a "mixed blessing."

Mixed Blessing?

The argument for science as a "mixed blessing" runs as follows: Science, as pure research, aims to discover the truth about the world, whether in cosmology, quantum physics, or genetics. Discovering the truth is *always* a good thing. Therefore, if we are asking questions about scientists' responsibility for outcomes because we are concerned about these outcomes—their potential to be used for harm, and so on—then *this* category of outcomes should not lead to concern, as it contains only true propositions. Three premises are assumed here. As to the first, since I accept the realist account of science, I am content with this description of the aim of pure research. The other two premises, however, are open to objection, especially the second, which holds that discovering the truth is always a good thing. The third premise is that we need only concern ourselves with the topic of the scientist's responsibility if we are worried about what scientists do.

It might be readily agreed that discovering or stating the truth, or trying to do so, is positive, if the alternative is fabrication or falsification *and* if our attention is restricted to the community of pure researchers. That is to say, it is surely better to publish findings that are correct rather than incorrect, for several reasons. Accurate findings are more probably based on careful and diligent experimentation, as opposed to careless and sloppy work, and the former have been recognized as virtues since at least the seventeenth century.[15] Moreover, correct findings provide a reliable basis for others who wish to build on original research, while incorrect findings mislead others and waste time. Finally, accurate research is independently verifiable by others, which safeguards the reputations of those doing the original work. So, when it is asserted that publishing the truth is always a good thing, and that no harm can come from it, one or more of these (or similar) reasons are, I believe, at base. But what is unwarranted is the claim that publishing the truth always secures good outcomes: this is most certainly false.[16] Even if the horrid experiments performed on human subjects during World War 2, at Auschwitz and Ping Fang, had been impeccably conducted, this would have weighed nothing in balance against the harm that was done to unwilling and innocent people.

It is not enough to reply that it is not the discovery of results but their publication that is wrong here. In order to produce results, it is necessary to embark on a specific research program, so the *intention* to discover certain things is in place at the start. It is not fanciful to suppose that in such a research environment, all such projects would be centrally approved and monitored. In any case, the idea that discovery is good in and of itself, independent of anything else—even publication—is at best a pun and at worst senseless. "Good," as I understand the word, refers to a state of affairs in which moral subjects are on the whole happy, have their interests looked after, and so on. "Good," then, is in this sense an evaluation of a state of affairs with reference to moral subjects and how things go for them. There can thus be no "good" that is abstracted from issues that affect moral subjects. The truth of this proposition, I think, explains the disgust we feel when we read about present-day experiments that cause pain and suffering, such as the animal experiments documented by Peter Singer (1993, 65–68). The only conceivable way that such experiments can be justified is if they produce some greater good for people or for other animals—good that outweighs the pain they cause. The suggestion that it is good in and of itself to discover, say, the circumstances that will bring an ape to kill its offspring is grotesque.[17]

Propositions derived from science as a body of ideas can be conceived of as a "mixed blessing" in two ways. First, different groups of people with different cultures, backgrounds, and traditions can have their most cherished beliefs confirmed or challenged by science—witness, for example, the implications of Darwinism for creationists and deists. Second, scientific propositions can give rise to different social and political policies and measures, such as eugenics and genetic counseling. It is this second sense that, I think, is more troublesome with regard to the scientist's responsibilities. If creationists are deluded, then it may be better to let them remain so, but not at the price of depriving others of the truth. But if scientific findings are going to support repressive policies, then the scientist is faced with a choice between suppressing what he believes to be true and contributing to the persecution of a minority. This representation of science as a mixed blessing becomes even more obvious in relation to the first category of outcomes, those of science-based technology.

Examples of such technology that can be considered "good things" are plentiful—for instance, our understanding of infectious diseases and the therapies and practices based on that understanding. Thus, Robert Koch's discovery of the bacterium responsible for tuberculosis was a great boon, as it led to health-care measures and eventually to drugs to fight the disease. On the other hand, some technological outcomes amount to methods for destroying life, not saving it, above all weapons of mass destruction, such as atomic and thermonuclear bombs. I classify these outcomes as "bad" and believe that we would be better off without them. However, the claim that science is a mixed blessing because it leads to both good and bad outcomes should not be understood to imply that some outcomes are *undisputedly* good and others are *undisputedly* bad. This would presuppose that we are in possession of a perfect value system that enables us to unerringly make such judgments. In practice, competing systems of values are informed by different ideas about what is right, and these can be expected to lead to conflicting judgments about which outcomes are good and which are bad.

If science were an unalloyed good, then discerning the nature of the scientist's responsibility would be relatively straightforward, and the scientist could simply press on with his work. Policy issues would remain, but these would be of technical interest only. Conversely, if science were thoroughly bad, the issue would be similarly clear: the scientist should stop work. However, the position contrary to both scenarios, that science is a mixed blessing, implies that scientists must acknowledge the issue of responsibility, and this is no longer straightforward.

Earlier I suggested that the question of whether a scientist is responsible for an outcome can be interpreted in two different ways. First, the question concerns how scientists *produce* outcomes, what conditions must be in place before it can be said that a scientist is responsible for an outcome. Second, the question concerns the responsibilities of the scientist, what he should do in relation to his work and what he should not do. Do these two interpretations divide the question neatly in two, leaving behind two distinct and separate topics? Or is there perhaps a more subtle connection to be made between these “backward-looking” and “forward-looking” aspects of responsibility? Some philosophers, such as Kurt Baier and most consequentialists, maintain that there is not much point in talking about responsibility unless we have an eye toward what will happen in the future. But I think that we cannot sufficiently address responsibility without taking both “directions” into account.