The Recombinant DNA Debate

Stephen P. Stich

Philosophy and Public Affairs, Vol. 7, No. 3 (Spring, 1978), 187-205.

Stable URL:
http://links.jstor.org/sici?sici=0048-3915%28197821%297%3A3%3C187%3ATRDD%3E2.0.CO%3B2-9

Philosophy and Public Affairs is currently published by Princeton University Press.

STEPHEN P. STICH

The Recombinant DNA Debate

The debate over recombinant DNA research is a unique event, perhaps a turning point, in the history of science. For the first time in modern history there has been widespread public discussion about whether and how a promising though potentially dangerous line of research shall be pursued. At root the debate is a moral debate and, like most such debates, requires proper assessment of the facts at crucial stages in the argument. A good deal of the controversy over recombinant DNA research arises because some of the facts simply are not yet known. There are many empirical questions we would like to have answered before coming to a decision—questions about the reliability of proposed containment facilities, about the viability of enfeebled strains of *E. coli*, about the ways in which pathogenic organisms do their unwelcome work, and much more. But all decisions cannot wait until the facts are available; some must be made now. It is to be expected that people with different hunches about what the facts will turn out to be will urge different decisions on how recombinant DNA research should be regulated. However, differing expectations about the facts have not been the only fuel for controversy. A significant part of the current debate can be traced to differences over moral principles. Also, unfortunately, there has been much unnecessary debate generated by careless moral reasoning and a failure to attend to the logical structure of some of the moral arguments that have been advanced.

In order to help sharpen our perception of the moral issues under-

Philosophy & Public Affairs 7, no. 3 © 1978 by Princeton University Press
00048-3915/78/0703-0187$00.95/1

lying the controversy over recombinant DNA research, I shall start by clearing away some frivolous arguments that have deflected attention from more serious issues. We may then examine the problems involved in deciding whether the potential benefits of recombinant DNA research justify pursuing it despite the risks that it poses.

I. Three Bad Arguments

My focus in this section will be on three untenable arguments, each of which has surfaced with considerable frequency in the public debate over recombinant DNA research.

The first argument on my list concludes that recombinant DNA research should not be controlled or restricted. The central premise of the argument is that scientists should have full and unqualified freedom to pursue whatever inquiries they may choose to pursue. This claim was stated repeatedly in petitions and letters to the editor during the height of the public debate over recombinant DNA research in the University of Michigan community.[1] The general moral principle which is the central premise of the argument plainly does entail that investigators using recombinant DNA technology should be allowed to pursue their research as they see fit. However, we need only consider a few examples to see that the principle invoked in this "freedom of inquiry" argument is utterly indefensible. No matter how sincere a researcher's interest may be in investigating the conjugal

1. For example, from a widely circulated petition signed by both faculty and community people: "The most important challenge may be a confrontation with one of our ancient assumptions—that there must be an absolute and unqualified freedom to pursue scientific inquiries. We will soon begin to wonder what meaning this freedom has if it leads to the destruction or demoralization of human beings, the only life forms able to exercise it." And from a letter to the editor written by a Professor of Engineering Humanities: "Is science beyond social and human controls, so that freedom of inquiry implies the absence of usual social restrictions which we all, as citizens, obey, respecting the social contract?"

It is interesting to note that the "freedom of inquiry" argument is rarely proposed by defenders of recombinant DNA research. Rather, it is proposed, then attacked, by those who are opposed to research involving recombinant molecules. Their motivation, it would seem, is to discredit the proponents of recombinant DNA research by attributing a foolish argument to them, then demonstrating that it is indeed a foolish argument.

behavior of American university professors, few would be willing to grant him the right to pursue his research in my bedroom without my consent. No matter how interested a researcher may be in investigating the effects of massive doses of bomb-grade plutonium on preschool children, it is hard to imagine that anyone thinks he should be allowed to do so. Yet the "free inquiry" principle, if accepted, would allow both of these projects and countless other Dr. Strangelove projects as well. So plainly the simplistic "free inquiry" principle is indefensible. It would, however, be a mistake to conclude that freedom of inquiry ought not to be protected. A better conclusion is that the right of free inquiry is a qualified right and must sometimes yield to conflicting rights and to the demands of conflicting moral principles. Articulating an explicit and properly qualified principle of free inquiry is a task of no small difficulty. We will touch on this topic again toward the end of Section II.

The second argument I want to examine aims at establishing just the opposite conclusion from the first. The particular moral judgment being defended is that there should be a total ban on recombinant DNA research. The argument begins with the observation that even in so-called low-risk recombinant DNA experiments there is at least a possibility of catastrophic consequences. We are, after all, dealing with a relatively new and unexplored technology. Thus it is at least possible that a bacterial culture whose genetic makeup has been altered in the course of a recombinant DNA experiment may exhibit completely unexpected pathogenic characteristics. Indeed, it is not impossible that we could find ourselves confronted with a killer strain of, say, *E. coli* and, worse, a strain against which humans can marshal no natural defense. Now if this is possible—if we cannot say with assurance that the probability of it happening is zero—then, the argument continues, all recombinant DNA research should be halted. For the negative utility of the imagined catastrophe is so enormous, resulting as it would in the destruction of our society and perhaps even of our species, that no work which could possibly lead to this result would be worth the risk.

The argument just sketched, which might be called the "doomsday scenario" argument, begins with a premise which no informed person

would be inclined to deny. It is indeed *possible* that even a low-risk recombinant DNA experiment might lead to totally catastrophic results. No ironclad guarantee can be offered that this will not happen. And while the probability of such an unanticipated catastrophe is surely not large, there is no serious argument that the probability is zero. Still, I think the argument is a sophistry. To go from the undeniable premise that recombinant DNA research might possibly result in unthinkable catastrophe to the conclusion that such research should be banned requires a moral principle stating that *all* endeavors that might possibly result in such a catastrophe should be prohibited. Once the principle has been stated, it is hard to believe that anyone would take it at all seriously. For the principle entails that, along with recombinant DNA research, almost all scientific research and many other commonplace activities having little to do with science should be prohibited. It is, after all, at least logically possible that the next new compound synthesized in an ongoing chemical research program will turn out to be an uncontainable carcinogen many orders of magnitude more dangerous than aerosol plutonium. And, to vary the example, there is a non-zero probability that experiments in artificial pollination will produce a weed that will, a decade from now, ruin the world's food grain harvest.[2]

I cannot resist noting that the principle invoked in the doomsday scenario argument is not new. Pascal used an entirely parallel argument to show that it is in our own best interests to believe in God. For though the probability of God's existence may be very low, if He nonetheless should happen to exist, the disutility that would accrue to the disbeliever would be catastrophic—an eternity in hell. But, as introductory philosophy students should all know, Pascal's argument only looks persuasive if we take our options to be just two: Christianity

2. Unfortunately, the doomsday scenario argument is *not* a straw man conjured only by those who would refute it. Consider, for example, the remarks of Anthony Mazzocchi, spokesman for the Oil, Chemical and Atomic Workers International Union, reported in *Science News*, 19 March 1977, p. 181: "When scientists argue over safe or unsafe, we ought to be very prudent. . . . If critics are correct and the Andromeda scenario has *even the smallest possibility* of occurring, we must assume it will occur on the basis of our experience" (emphasis added).

or atheism. A third possibility is belief in a jealous non-Christian God who will see to our damnation if and only if we *are* Christians. The probability of such a deity existing is again very small, but non-zero. So Pascal's argument is of no help in deciding whether or not to accept Christianity. For we may be damned if we do and damned if we don't.

I mention Pascal's difficulty because there is a direct parallel in the doomsday scenario argument against recombinant DNA research. Just as there is a non-zero probability that unforeseen consequences of recombinant DNA research will lead to disaster, so there is a non-zero probability that unforeseen consequences of *failing* to pursue the research will lead to disaster. There may, for example, come a time when, because of natural or man-induced climatic change, the capacity to alter quickly the genetic constitution of agricultural plants will be necessary to forestall catastrophic famine. And if we fail to pursue recombinant DNA research now, our lack of knowledge in the future may have consequences as dire as any foreseen in the doomsday scenario argument.

The third argument I want to consider provides a striking illustration of how important it is, in normative thinking, to make clear the moral *principles* being invoked. The argument I have in mind begins with a factual claim about recombinant DNA research and concludes that stringent restrictions, perhaps even a moratorium, should be imposed. However, advocates of the argument are generally silent on the normative principle(s) linking premise and conclusion. The gap thus created can be filled in a variety of ways, resulting in very different arguments. The empirical observation that begins the argument is that recombinant DNA methods enable scientists to move genes back and forth across natural barriers, "particularly the most fundamental such barrier, that which divides prokaryotes from eukaryotes. The results will be essentially new organisms, self-perpetuating and hence permanent."[3] Because of this, it is concluded that severe restrictions are

3. The quotation is from George Wald, "The Case Against Genetic Engineering," *The Sciences*, September/October 1976; to be reprinted in David A. Jackson and Stephen P. Stich, eds., *The Recombinant DNA Debate*, forthcoming.

in order. Plainly this argument is an enthymeme; a central premise has been left unstated. What sort of moral principle is being tacitly assumed?

The principle that comes first to mind is simply that natural barriers should not be breached, or perhaps that "essentially new organisms" should not be created. The principle has an almost theological ring to it, and perhaps there are some people who would be prepared to defend it on theological grounds. But short of a theological argument, it is hard to see why anyone would hold the view that breaching natural barriers or creating new organisms is *intrinsically* wrong. For if a person were to advocate such a principle, he would have to condemn the creation of new bacterial strains capable of, say, synthesizing human clotting factor or insulin, *even if* creating the new organism generated *no unwelcome side effects.*

There is quite a different way of unraveling the "natural barriers" argument which avoids appeal to the dubious principles just discussed. As an alternative, this second reading of the argument ties premise to conclusion with a second factual claim and a quite different normative premise. The added factual claim is that at present our knowledge of the consequences of creating new forms of life is severely limited; thus we cannot know with any assurance that the probability of disastrous consequences is very low. The moral principle needed to mesh with the two factual premises would be something such as the following:

> If we do not know with considerable assurance that the probability of an activity leading to disastrous consequences is very low, then we should not allow the activity to continue.

Now this principle, unlike those marshaled in the first interpretation of the natural barriers argument, is not lightly dismissed. It is, to be sure, a conservative principle, and it has the odd feature of focusing entirely on the dangers an activity poses while ignoring its potential benefits.[4] Still, the principle may have a certain attraction in light of

4. It is important to note, however, that the principle is considerably less conservative, and correspondingly more plausible, than the principle invoked in the doomsday scenario argument. That latter principle would have us enjoin an activity if the probability of the activity leading to catastrophe is anything other than zero.

recent history, which has increasingly been marked by catastrophes attributable to technology's unanticipated side effects. I will not attempt a full scale evaluation of this principle just now. For the principle raises, albeit in a rather extreme way, the question of how risks and benefits are to be weighed against each other. In my opinion, that is the really crucial moral question raised by recombinant DNA research. It is a question which bristles with problems. In Section II I shall take a look at some of these problems and make a few tentative steps toward some solutions. While picking our way through the problems we will have another opportunity to examine the principle just cited.

II. Risks and Benefits

At first glance it might be thought that the issue of risks and benefits is quite straightforward, at least in principle. What we want to know is whether the potential benefits of recombinant DNA research justify the risks involved. To find out we need only determine the probabilities of the various dangers and benefits. And while some of the empirical facts—the probabilities—may require considerable ingenuity and effort to uncover, the assessment poses no particularly difficult normative or conceptual problems. Unfortunately, this sanguine view does not survive much more than a first glance. A closer look at the task of balancing the risks and benefits of recombinant DNA research reveals a quagmire of sticky conceptual problems and simmering moral disputes. In the next few pages I will try to catalogue and comment on some of these moral disputes. I wish I could also promise solutions to all of them, but to do so would be false advertising.

Problems about Probabilities

In trying to assess costs and benefits, a familiar first step is to set down a list of possible actions and possible outcomes. Next, we assign some measure of desirability to each possible outcome, and for each action we estimate the conditional probability of each outcome given that the action is performed. In attempting to apply this decision-making strategy to the case of recombinant DNA research, the assignment of probabilities poses some perplexing problems. Some of the outcomes whose probabilities we want to know can be ap-

proached using standard empirical techniques. Thus, for example, we may want to know what the probability is of a specific enfeebled host *E. coli* strain surviving passage through the human intestinal system, should it be accidentally ingested. Or we may want to know what the probability is that a host organism will escape from a P-4 laboratory. In such cases, while there may be technical difficulties to be overcome, we have a reasonably clear idea of the sort of data needed to estimate the required probabilities. But there are other possible outcomes whose probabilities cannot be determined by experiment. It is important, for example, to know what the probability is of recombinant DNA research leading to a method for developing nitrogen-fixing strains of corn and wheat. And it is important to know how likely it is that recombinant DNA research will lead to techniques for effectively treating or preventing various types of cancer. Yet there is no experiment we can perform nor any data we can gather that will enable us to *empirically* estimate these probabilities. Nor are these the most problematic probabilities we may want to know. A possibility that weighs heavily on the minds of many who are worried about recombinant DNA research is that this research may lead to negative consequences for human health or for the environment *which have not yet even been thought of*. The history of technology during the last half-century surely demonstrates that this is not a quixotic concern. Yet here again there would appear to be no data we can gather that would help much in estimating the probability of such potential outcomes.

It should be stressed that the problems just sketched are not to be traced simply to a paucity of data. Rather, they are conceptual problems; it is doubtful whether there is *any clear empirical sense* to be made of objective probability assignments to contingencies like those we are considering.

Theorists in the Bayesian tradition may be unmoved by the difficulties we have noted. On their view all probability claims are reports of subjective probabilities.[5] And, a Bayesian might quite properly note, there is no special problem about assigning *subjective* probabilities to

5. For an elaboration of the Bayesian position, see Leonard J. Savage, *The Foundations of Statistics* (New York: John Wiley & Sons, 1954); also cf. Leonard J. Savage, "The Shifting Foundations of Statistics," in Robert G. Colodny, ed., *Logic, Laws and Life* (Pittsburgh: University of Pittsburgh Press, 1977).

outcomes such as those that worried us. But even for the radical Bayesian, there remains the problem of *whose* subjective probabilities ought to be employed in making a *social* or *political* decision. The problem is a pressing one since the subjective probabilities assigned to potential dangers and benefits of recombinant DNA research would appear to vary considerably even among reasonably well informed members of the scientific community.

The difficulties we have been surveying are serious ones. Some might feel they are so serious that they render rational assessment of the risks and benefits of recombinant DNA research all but impossible. I am inclined to be rather more optimistic, however. Almost all of the perils posed by recombinant DNA research require the occurrence of a sequence of separate events. For a chimerical bacterial strain created in a recombinant DNA experiment to cause a serious epidemic, for example, at least the following events must occur:

(1) a pathogenic bacterium must be synthesized
(2) the chimerical bacteria must escape from the laboratory
(3) the strain must be viable in nature
(4) the strain must compete successfully with other micro-organisms which are themselves the product of intense natural selection.[6]

Since *all* of these must occur, the probability of the potential epidemic is the product of the probabilities of each individual contingency. And there are at least two items on the list, namely (2) and (3), whose probabilities are amenable to reasonably straightforward empirical assessment. Thus the product of these two individual probabilities places an upper limit on the probability of the epidemic. For the remaining two probabilities, we must rely on subjective probability assessments of informed scientists. No doubt there will be considerable variability. Yet even here the variability will be limited. In the case of (4), as an example, the available knowledge about microbial natural selection provides no precise way of estimating the probability that a chimerical strain of enfeebled *E. coli* will compete successfully

6. For an elaboration of this point, see Bernard D. Davis, "Evolution, Epidemiology, and Recombinant DNA," *The Recombinant DNA Debate*, forthcoming.

outside the laboratory. But no serious scientist would urge that the probability is *high*. We can then use the highest responsible subjective estimate of the probabilities of (1) and (4) in calculating the "worst case" estimate of the risk of epidemic. If in using this highest "worst case" estimate, our assessment yields the result that benefits outweigh risks, then lower estimates of the same probabilities will, of course, yield the same conclusion. Thus it may well be the case that the problems about probabilities we have reviewed will not pose insuperable obstacles to a rational assessment of risks and benefits.

Weighing Harms and Benefits

A second cluster of problems that confronts us in assessing the risks and benefits of recombinant DNA research turns on the assignment of a measure of desirability to the various possible outcomes. Suppose that we have a list of the various harms and benefits that might possibly result from pursuing recombinant DNA research. The list will include such "benefits" as development of an inexpensive way to synthesize human clotting factor and development of a strain of nitrogen-fixing wheat; and such "harms" as release of a new antibiotic-resistant strain of pathogenic bacteria and release of a strain of *E. coli* carrying tumor viruses capable of causing cancer in man.

Plainly, it is possible that pursuing a given policy will result in more than one benefit and in more than one harm. Now if we are to assess the potential impact of various policies or courses of action, we must assign some index of desirability to the possible *total outcomes* of each policy, outcomes which may well include a mix of benefits and harms. To do this we must confront a tangle of normative problems that are as vexing and difficult as any we are likely to face. We must *compare* the moral desirabilities of various harms and benefits. The task is particularly troublesome when the harms and benefits to be compared are of different kinds. Thus, for example, some of the attractive potential benefits of recombinant DNA research are economic: we may learn to recover small amounts of valuable metals in an economically feasible way, or we may be able to synthesize insulin and other drugs inexpensively. By contrast, many of the risks of recombinant DNA research are risks to human life or health. So if we are to

take the idea of cost-benefit analysis seriously, we must at some point decide how human lives are to be weighed against economic benefits.

There are those who contend that the need to make such decisions indicates the moral bankruptcy of attempting to employ risk-benefit analyses when human lives are at stake. On the critics' view, we cannot reckon the possible loss of a human life as just another negative outcome, albeit a grave and heavily weighted one. To do so, it is urged, is morally repugnant and reflects a callous lack of respect for the sacredness of human life.

On my view, this sort of critique of the very idea of using risk-benefit analyses is ultimately untenable. It is simply a fact about the human condition, lamentable as it is inescapable, that in many human activities we run the risk of inadvertently causing the death of a human being. We run such a risk each time we drive a car, allow a dam to be built, or allow a plane to take off. Moreover, in making social and individual decisions, we cannot escape weighing economic consequences against the risk to human life. A building code in the Midwest will typically mandate fewer precautions against earthquakes than a building code in certain parts of California. Yet earthquakes are not impossible in the Midwest. If we elect not to require precautions, then surely a major reason must be that it would simply be too expensive. In this judgment, as in countless others, there is no escaping the need to balance economic costs against possible loss of life. To deny that we must and do balance economic costs against risks to human life is to assume the posture of a moral ostrich.

I have been urging the point that it is not *morally objectionable* to try to balance economic concerns against risks to human life. But if such judgments are unobjectionable, indeed necessary, they also surely are among the most difficult any of us has to face. It is hard to imagine a morally sensitive person not feeling extremely uncomfortable when confronted with the need to put a dollar value on human lives. It might be thought that the moral dilemmas engendered by the need to balance such radically different costs and benefits pose insuperable practical obstacles for a rational resolution of the recombinant DNA debate. But here, as in the case of problems with probabilities, I am more sanguine. For while some of the risks and potential benefits of recom-

binant DNA research are all but morally incommensurable, the most salient risks and benefits are easier to compare. The major risks, as we have noted, are to human life and health. However, the major potential benefits are *also* to human life and health. The potential economic benefits of recombinant DNA research pale in significance when set against the potential for major breakthroughs in our understanding and ability to treat a broad range of conditions, from birth defects to cancer. Those of us, and I confess I am among them, who despair of deciding how lives and economic benefits are to be compared can nonetheless hope to settle our views about recombinant DNA research by comparing the potential risks to life and health with the potential benefits to life and health. Here we are comparing plainly commensurable outcomes. If the balance turns out to be favorable, then we need not worry about factoring in potential economic benefits.

There is a certain irony in the fact that we may well be able to ignore economic factors entirely in coming to a decision about recombinant DNA research. For I suspect that a good deal of the apprehension about recombinant DNA research on the part of the public at large is rooted in the fear that (once again) economic benefits will be weighed much too heavily and potential damage to health and the environment will be weighed much too lightly. The fear is hardly an irrational one. In case after well-publicized case, we have seen the squalid consequences of decisions in which private or corporate gain took precedence over clear and serious threats to health and to the environment. It is the profit motive that led a giant chemical firm to conceal the deadly consequences of the chemical which now threatens to poison the James River and perhaps all of Chesapeake Bay. For the same reason, the citizens of Duluth drank water laced with a known carcinogen. And the ozone layer that protects us all was eroded while regulatory agencies and legislators fussed over the loss of profits in the spray deodorant industry. Yet while public opinion about recombinant DNA research is colored by a growing awareness of these incidents and dozens of others, the case of recombinant DNA is fundamentally different in a crucial respect. The important projected benefits which must be set against the risks of recombinant DNA research are not economic at all, they are medical and environmental.

Problems about Principles

The third problem I want to consider focuses on the following question. Once we have assessed the potential harms and benefits of recombinant DNA research, how should we use this information in coming to a decision? It might be thought that the answer is trivially obvious. To assess the harms and benefits is, after all, just to compute, for each of the various policies that we are considering, what might be called its *expected utility*. The expected utility of a given policy is found by first multiplying the desirability of each possible total outcome by the probability that the policy in question will lead to that total outcome, and then adding the numbers obtained. As we have seen, finding the needed probabilities and assigning the required desirabilities will not be easy. But once we know the expected utility of each policy, is it not obvious that we should choose the policy with the highest expected utility? The answer, unfortunately, is no, it is not at all obvious.

Let us call the principle that we should adopt the policy with the highest expected utility the *utilitarian principle*. The following example should make it clear that, far from being trivial or tautological, the utilitarian principle is a substantive and controversial moral principle. Suppose that the decision which confronts us is whether or not to adopt policy *A*. What is more, suppose we know there is a probability close to 1 that 100,000 lives will be saved if we adopt *A*. However, we also know that there is a probability close to 1 that 1,000 will die as a direct result of our adopting policy *A*, and these people would survive if we did not adopt *A*. Finally, suppose that the other possible consequences of adopting *A* are relatively inconsequential and can be ignored. (For concreteness, we might take *A* to be the establishment of a mass vaccination program, using a relatively risky vaccine.) Now plainly if we take the moral desirability of saving a life to be exactly offset by the moral undesirability of causing a death, then the utilitarian principle dictates that we adopt policy *A*. But many people feel uncomfortable with this result, the discomfort increasing with the number of deaths that would result from *A*. If, to change the example, the choice that confronts us is saving 100,000 lives while causing the

deaths of 50,000 others, a significant number of people are inclined to think that the morally right thing to do is to refrain from doing *A*, and "let nature take its course."

If we reject policy *A*, the likely reason is that we also reject the utilitarian principle. Perhaps the most plausible reason for rejecting the utilitarian principle is the view that our obligation to *avoid doing harm* is stronger than our obligation to do good. There are many examples, some considerably more compelling than the one we have been discussing, which seem to illustrate that in a broad range of cases we do feel that our obligation to avoid doing harm is greater than our obligation to do good.[7] Suppose, to take but one example, that my neighbor requests my help in paying off his gambling debts. He owes $5,000 to a certain bookmaker with underworld connections. Unless the neighbor pays the debt immediately, he will be shot. Here, I think we are all inclined to say, I have no strong obligation to give my neighbor the money he needs, and if I were to do so it would be a supererogatory gesture. By contrast, suppose a representative of my neighbor's bookmaker approaches me and requests that I shoot my neighbor. If I refuse, he will see to it that my new car, which cost $5,000, will be destroyed by a bomb while it sits unattended at the curb. In this case, surely, I have a strong obligation not to harm my neighbor, although not shooting him will cost me $5,000.

Suppose that this example and others convince us that we cannot adopt the utilitarian principle, at least not in its most general form, where it purports to be applicable to all moral decisions. What are the alternatives? One cluster of alternative principles would urge that in some or all cases we weigh the harm a contemplated action will cause more heavily than we weigh the good it will do. The extreme form of such a principle would dictate that we ignore the benefits entirely and opt for the action or policy that produces the *least* expected harm. (It is this principle, or a close relation, which emerged in the second reading of the "natural barriers" argument discussed in the third part of Sec-

7. For an interesting discussion of these cases, see J. O. Urmson, "Saints and Heros," in A. I. Melden, ed., *Essays In Moral Philosophy* (Seattle: University of Washington Press, 1958). Also see the discussion of positive and negative duties in Philippa Foot, "The Problem of Abortion and the Doctrine of Double Effect," *Oxford Review* 5 (1967). Reprinted in James Rachels, ed., *Moral Problems* (New York: Harper & Row) 1971.

tion I above.) A more plausible variant would allow us to count both benefits and harms in our deliberations, but would specify how much more heavily harms were to count.

On my view, some moderate version of a "harm-weighted" principle is preferable to the utilitarian principle in a considerable range of cases. *However, the recombinant DNA issue is not one of these cases.* Indeed, when we try to apply a harm-weighted principle to the recombinant DNA case we run head on into a conceptual problem of considerable difficulty. The distinction between doing good and doing harm presupposes a notion of the normal or expectable course of events. Roughly, if my action causes you to be worse off than you would have been in the normal course of events, then I have harmed you; if my action causes you to be better off than in the normal course of events, then I have done you some good; and if my action leaves you just as you would be in the normal course of events, then I have done neither. In many cases, the normal course of events is intuitively quite obvious. Thus in the case of the neighbor and the bookmaker, in the expected course of events I would neither shoot my neighbor nor give him $5,000 to pay off his debts. Thus I am doing good if I give him the money and I am doing harm if I shoot him. But in other cases, including the recombinant DNA case, it is not at all obvious what constitutes the "expected course of events," and thus it is not at all obvious what to count as a harm. To see this, suppose that as a matter of fact many more deaths and illnesses will be prevented as a result of pursuing recombinant DNA research than will be caused by pursuing it. But suppose that there *will* be at least some people who become ill or die as a result of recombinant DNA research being pursued. If these are the facts, then who would be harmed by imposing a ban on recombinant DNA research? That depends on what we take to be the "normal course of events." Presumably, if we do not impose a ban, then the research will continue and the lives will be saved. If this is the normal course of events, then if we impose a ban we have *harmed* those people who would be saved. But it is equally natural to take as the normal course of events the situation in which recombinant DNA research is not pursued. And if *that* is the normal course of events, then those who would have been saved are not harmed by a ban, for they are no worse off than they would be in the normal course of

events. However, on this reading of "the normal course of events," if we *fail* to impose a ban, then we have harmed those people who will ultimately become ill or die as a result of recombinant DNA research, since as a result of not imposing a ban they are worse off than they would have been in the normal course of events. I conclude that, in the absence of a theory detailing how we are to recognize the normal course of events, harm-weighted principles have no clear application to the case of recombinant DNA research.

Harm-weighted principles are not the only alternatives to the utilitarian principle. There is another cluster of alternatives that take off in quite a different direction. These principles urge that in deciding which policy to pursue there is a strong presumption in favor of policies that adhere to certain formal moral principles (that is, principles which do not deal with the *consequences* of our policies). Thus, to take the example most directly relevant to the recombinant DNA case, it might be urged that there is a strong presumption in favor of a policy which preserves freedom of scientific inquiry. In its extreme form, this principle would protect freedom of inquiry *no matter what the consequences*; and as we saw in the first part of Section I, this extreme position is exceptionally implausible. A much more plausible principle would urge that freedom of inquiry be protected until the balance of negative over positive consequences reaches a certain specified amount, at which point we would revert to the utilitarian principle. On such a view, if the expected utility of banning recombinant DNA research is a bit higher than the expected utility of allowing it to continue, then we would nonetheless allow it to continue. But if the expected utility of a ban is enormously higher than the expected utility of continuation, banning is the policy to be preferred.[8]

III. Long Term Risks

Thus far in our discussion of risks and benefits, the risks that have occupied us have been what might be termed "short-term" risks, such as the release of a new pathogen. The negative effects of these

8. Carl Cohen defends this sort of limited protection of the formal free inquiry principle over a straight application of the utilitarian principle in his interesting essay, "When May Research Be Stopped?" *New England Journal of Medicine* 296 (1977). To be reprinted in *The Recombinant DNA Debate*.

events, though they might be long-lasting indeed, would be upon us relatively quickly. However, some of those who are concerned about recombinant DNA research think there are longer-term dangers that are at least as worrisome. The dangers they have in mind stem not from the accidental release of harmful substances in the course of recombinant DNA research, but rather from the unwise use of the *knowledge* we will likely gain in pursuing the research. The scenarios most often proposed are nightmarish variations on the theme of human genetic engineering. With the knowledge we acquire, it is conjectured, some future tyrant may have people built to order, perhaps creating a whole class of people who willingly and cheaply do the society's dirty or dangerous work, as in Huxley's *Brave New World*. Though the proposed scenarios clearly are science fiction, they are not to be lightly dismissed. For if the technology they conjure is not demonstrably achievable, neither is it demonstrably impossible. And if only a bit of the science fiction turns to fact, the dangers could be beyond reckoning.

Granting that potential misuse of the knowledge gained in recombinant DNA research is a legitimate topic of concern, how ought we to guard ourselves against this misuse? One common proposal is to try to prevent the acquisition of such knowledge by banning or curtailing recombinant DNA research now. Let us cast this proposal in the form of an explicit moral argument. The conclusion is that recombinant DNA research should be curtailed, and the reason given for the conclusion is that such research could possibly produce knowledge which might be misused with disastrous consequences. To complete the argument we need a moral principle, and the one which seems to be needed is something such as this:

> If a line of research can lead to the discovery of knowledge which might be disastrously misused, then that line of research should be curtailed.

Once it has been made explicit, I think relatively few people would be willing to endorse this principle. For recombinant DNA research is hardly alone in potentially leading to knowledge that might be disastrously abused. Indeed, it is hard to think of an area of scientific

research that could *not* lead to the discovery of potentially dangerous knowledge. So if the principle is accepted it would entail that almost all scientific research should be curtailed or abandoned.

It might be thought that we could avoid the extreme consequences just cited by retreating to a more moderate moral principle. The moderate principle would urge only that we should curtail those areas of research where the probability of producing dangerous knowledge is comparatively high. Unfortunately, this more moderate principle is of little help in avoiding the unwelcome consequences of the stronger principle. The problem is that the history of science is simply too unpredictable to enable us to say with any assurance which lines of research will produce which sorts of knowledge or technology. There is a convenient illustration of the point in the recent history of molecular genetics. The idea of recombining DNA molecules is one which has been around for some time. However, early efforts proved unsuccessful. As it happened, the crucial step in making recombinant DNA technology possible was provided by research on restriction enzymes, research that was undertaken with no thought of recombinant DNA technology. Indeed, until it was realized that restriction enzymes provided the key to recombining DNA molecules, the research on restriction enzymes was regarded as a rather unexciting (and certainly uncontroversial) scientific backwater.[9] In an entirely analogous way, crucial pieces of information that may one day enable us to manipulate the human genome may come from just about any branch of molecular biology. To guard against the discovery of that knowledge we should have to curtail not only recombinant DNA research but all of molecular biology.

Before concluding, we would do well to note that there is a profound pessimism reflected in the attitude of those who would stop recombinant DNA research because it might lead to knowledge that could be abused. It is, after all, granted on all sides that the knowledge resulting from recombinant DNA research will have both good and evil potential uses. So it would seem the sensible strategy would be to try to prevent the improper uses of this knowledge rather than trying to prevent the knowledge from ever being uncovered. Those who

9. I am indebted to Prof. Ethel Jackson for both the argument and the illustration.

would take the more extreme step of trying to stop the knowledge from being uncovered presumably feel that its improper use is all but inevitable, that our political and social institutions are incapable of preventing morally abhorrent applications of the knowledge while encouraging beneficial applications. On my view, this pessimism is unwarranted; indeed, it is all but inconsistent. The historical record gives us no reason to believe that what is technologically possible will be done, no matter what the moral price. Indeed, in the area of human genetic manipulation, the record points in quite the *opposite* direction. We have long known that the same techniques that work so successfully in animal breeding can be applied to humans as well. Yet there is no evidence of a "technological imperative" impelling our society to breed people as we breed dairy cattle, simply because we know that it can be done. Finally, it is odd that those who express no confidence in the ability of our institutions to forestall such monstrous applications of technology are not equally pessimistic about the ability of the same institutions to impose an effective ban on the uncovering of dangerous knowledge. If our institutions are incapable of restraining the application of technology when those applications would be plainly morally abhorrent, one would think they would be even more impotent in attempting to restrain a line of research which promises major gains in human welfare.

This essay is an abridged and somewhat modified version of my essay, "The Recombinant DNA Debate: Some Philosophical Considerations," which will appear in *The Recombinant DNA Debate* edited by David A. Jackson and Stephen P. Stich, to be published by Prentice-Hall. I am grateful to the editors of *Philosophy & Public Affairs* for their detailed and useful suggestions on modifying the essay to make it appropriate for use in this journal.