

Introduction to some basic ethical orientations

Rick Grush

[1] One of the major branches of philosophy is *ethics*, which involves, among other things, the attempt to articulate ethical principles that underlie our moral intuitions, as well as serve as a guide for morally relevant decisions we might face. We will follow convention in using 'moral' and its cognates to describe first-order situations and attitudes. Moral situations arise all the time, and moral behavior (or immoral behavior) is behavior that is assessable as good or bad. Ethics is *the study of* morally relevant situations and morally assessable behavior.

[2] The way I described the use of ethical principles just now had two components: a descriptive and a prescriptive component. A *descriptive* enterprise is one that simply observes and describes what happens. For example I might observe people who are speaking some language, and simply record what they say and try to determine the rules, if any, that determine what they say; perhaps they sometimes use double-negatives, like 'I don't want no eggs'. In such a case, I would simply observe what they do, and try to account for why it is they are doing that. This would be a descriptive grammar. In the context of ethics, I might undertake a descriptive ethical enquiry, trying to determine the principles that some person or group of people use when making their morally relevant decisions: perhaps they do whatever they feel like, or perhaps they do what they think God wants them to do, or perhaps they do what they think will lead to the greatest overall happiness of people. Whatever the details, the descriptive task would simply seek to determine what ethical principles, if any, are in fact being used by some person or group in their morally relevant behavior.

[3] By contrast, a *prescriptive* enterprise is one that attempts to articulate the correct principles that people ought to apply. For example we are all familiar with prescriptive grammars that articulate what somebody thinks is the 'correct' way to speak a language. Someone involved in a prescriptive enterprise might try to correct the speakers mentioned above, by telling them that double negatives are incorrect. In the domain of ethics, a prescriptive ethics tries to determine what ethical principles we ought to adopt and employ when making morally relevant decisions.

[4] By and large ethical theory involves both components. The problem can be set up more or less as follows: there are roughly two kinds of morally relevant situation: those in which moral intuitions are very clear and practically universally shared. For example 'If Smith offers Jones \$5 to kill Smith's family because Smith is tired of hanging around them, should Jones accept?'. In this sort of case, almost everyone's intuitions converge on the same answer. The second and more problematic area are cases where moral intuitions within a given individual are less clear, and opinions among different people may diverge: is abortion permissible? is euthanasia? should those convicted of violent crimes be given the same rights of access to donated organs as normal citizens when such organs are in short supply? is it right to steal money if you need it to do something that, considered in itself, is a good action (e.g. stealing money from a hard working single mother in order to buy Christmas gifts for orphans)? The typical tactic in ethical theory is to use the cases on which there is a high degree of agreement in order to try to figure out what the ethical principles are that we are implicitly employing to arrive at those judgments, and then using those descriptively arrived-at principles in order to make a prescriptive claim about the less obvious cases.

[5] This discussion has been a bit abstract, but as we go on it should become more clear. In the next few sections, we will cover a few of the most influential ethical orientations.

1. Utilitarian views

[6] The basic utilitarian view is that when presented with a morally relevant decision between a number of options, the right one, *the one you ought to do*, is the one that does the greatest good for the greatest number of people. Before we look more carefully at what this means, let's look first at what it does not say. First off, it does not give anyone, including the person making the decision and that person's friends or family, any special status; it is thus, like most of the views we will discuss, impartial. Second, it quite avoids a number of other ethical outlooks, such as divine command (which would maintain that you do what God tells you, regardless of what good or bad comes from it), or deontological views that prescribe certain duties (e.g. respect the dignity of all people) regardless of the results. Third, it makes no immediate distinction between acting and failure to prevent an action: and so accordingly pushing someone off a building is exactly as bad as failing to help someone from falling off (assuming that one is in a position to help without putting oneself at risk).

[7] The basic idea of utilitarianism is easy and plausible enough: the good thing to do is to make the world a better place. Thus what makes some act right or wrong is its utility in bringing about good consequences, and not, for example, the intentions behind the act. So when presented with some decision one calculates the expected utility of the different options, and then does the one that has the highest expected utility.

[8] So what is expected utility? The first idea is that utility is cashed out in terms of something like happiness, and so we suppose that we can assign various quantitative values to various kinds of happiness (hedons) and unhappiness (dolors). We then, for each possible action we might perform, we determine what all the possible outcomes are, and how likely that outcome is. We then multiply the outcome times the likelihood of that outcome, and add up the result for each outcome. The final result is the expected utility of that action.

[9] For example, suppose I have \$100 dollars and someone gives me the opportunity to make a bet. I roll a die, and if it comes up '6', I get \$1000, and if it comes up any other number, I lose my \$100. My two options are: take the bet, or don't take the bet. What is the expected utility of each? The expected utility of not taking the bet is easy. There is only one possible outcome (since I am not doing anything), and on that outcome, neither lose nor gain anything. So the expected utility of not taking the bet is \$0. (Note that I am now using dollars as the units of utility.)

[10] What about taking the bet? There are six possible outcomes, one for each number that could come up on the die. On a 1 through 5, I lose \$100; and on a 6 I win \$1000. So first I determine how likely each option is, and then I multiply that likelihood times the loss/gain, and then add them up. So:

$$1: 1/6 \times -\$100 = -\$16.67$$

$$2: 1/6 \times -\$100 = -\$16.67$$

$$3: 1/6 \times -\$100 = -\$16.67$$

$$4: 1/6 \times -\$100 = -\$16.67$$

$$5: 1/6 \times -\$100 = -\$16.67$$

$$6: 1/6 \times +\$1000 = +\$166.67$$

[11] Adding these up, we see that the expected utility of taking the bet is +\$83.34. This means that this is what I can expect, as a sort of average, if I take the bet: a gain of about \$83. So this is a good bet.

[12] In the context of ethics, we can either increase the number of students at UCSD, or keep the number of students constant. Each choice will have a number of certain and probable effects. Benefits of increasing the number of students is that a certain number of people who would otherwise not be able to attend UCSD would be able to attend. This would make them happy, and might also have beneficial effects for society in that it would make more people better educated, etc. On the other hand, it would lead to greater crowding on the campus and perhaps place larger demands on the facilities and staff. Taxes would have to be raised, which would make a great number of families slightly less well-off than they would have been, and it might also decrease the overall quality of education at UCSD if classes become overcrowded, etc. So what we would need to do is for each course of action, come up with some idea of what all the costs and benefits would be and then compare the totals for each course of action.

[13] Now that the basics are done, on to some more details: what is the 'good' that we are trying to maximize? Some historical options have been: pleasure, happiness, preferences, and values. Pleasure (Bentham's position) is easy enough to understand. Some things bring people pleasure, such as eating and having sex. And we can also understand physical suffering, such as having toothaches or being beaten up. On this conception, we look for what sorts of pleasures and displeasures might arise from various choices, and maximize the pleasure. While this is easy enough to understand, it has some disadvantages. Most obviously, it seems to place bodily pleasure and enjoyment at too fundamental a place in ethical theory. It might, for example, justify just hooking everyone up to morphine I.V.s if we could get robots to run everything else in the world. This seems wrong, it seems as though doing this would eliminate a number of things we value, and so perhaps mere pleasure is not the best thing to maximize.

[14] (Note that here we see the interplay between descriptive and prescriptive theories. On the basis of some moral intuitions about making the world a better place, etc., we come up with a tentative ethical proposal: the right thing to do is to maximize pleasure. If this were right, then it might be a principle we could appeal to in unclear cases. But we can see that the principle can't be right because it conflicts with some clear moral intuitions we have: most people would not think that a world run by robots full of morphine addicted humans experiencing great pleasure for their entire lives would be a good thing. So on descriptive grounds we know that this can't really be a principle that we in fact actually endorse, and perhaps nor should it be one we want to endorse. The hope is to find principles that accord with, and perhaps explain, our moral intuitions in all the clear cases, and can thus serve as guidelines in the unclear cases.)

[15] Another principle has been happiness (Mill's position), where happiness might include bodily pleasures, but also includes 'higher' pleasures such as appreciation of art, enjoying the company of family and friends; seeing one's children succeed; and so forth. While this seems better than pleasure, it has the problem that it is much more variable than pleasure: very different things make different people happy, and so determining what will lead to the greatest happiness may be no easy matter.

[16] A similar approach is to maximize preferences. We assume that each person has some set of preferences, and what we need to do is to maximize these preferences. An advantage of this is that in some contexts it might be possible to merge this with economic theorizing.

For example, assuming that people spend money as a function of their preferences, the free market will evolve to maximize preferences. Thus we don't have to figure out if people prefer McDonalds or Burger King: we just let the market evolve and it will, almost by definition, evolve to maximize preferences. (At least that is the hope.)

[17] A final option would be to maximize the prevalence of some sort of value, such as freedom or knowledge. This has obvious advantages, but it is subject to a number of problems, including the fact that different people and different cultures may have different ideas about what values are the right ones, and if this is the case, our ethical theory will be powerless to make any decisions.

[18] In addition to having different choices for the good that we might want to maximize, there are two versions of utilitarianism that differ over what we are calculating the utility of: acts or rules. *Act utilitarianism* says that for each individual action we try to assess its consequences and then perform the act with the best expected utility. On an act utilitarian view, it is actions rather than rules that we are assessing. According to *rule utilitarianism*, what we are interested in assessing are *rules*, not individual acts. The two approaches might give conflicting results. So for example, a rule utilitarian might come up with convincing data that show that a law prohibiting smoking in public buildings is a rule that will lead to the best overall outcome. And so we might adopt that as the rule, and on its basis kick someone out of a building for smoking on a given occasion. On the other hand, an act utilitarian might reason that on this particular occasion, since only the smoker is present and the ventilation system is such as to remove all smoke before anyone else arrives, the best thing would be for the smoker to light up.

[19] There are some common objections to utilitarianism. The first two are based on the problem of determining consequences. First, it is often very difficult if not impossible to determine what the expected utilities of a given set of acts or rules will be to any degree of certainty, even when we are using simple goods such as money and bodily harm. Things are only worse when we move to other more tenuous and variable things like happiness and values. To this the utilitarian might respond that in some cases we can either come up with a decent guess, and in others the relative goods and bads will be quite clear. The opponent will respond to this that this might be right, but we need ethical guidance exactly in a great many areas that are unclear, and it looks as though utilitarianism, either rule or act, may not be able to help us in many of these cases.

[20] A second objection is that because of the difficulty or impossibility of predicting all of the consequences of an act, we might end up doing acts that are, by utilitarian principles, quite bad, even though our intentions are good and we were as rational and as responsible as we could have been. The man who swerved his car in order to not hit the little 5 year old boy in the street might, in some sense, be morally culpable if that boy is Adolph Hitler, for example. In such a case, the 'right' thing to do would have been to run the boy down. The utilitarian might respond by saying that they are not in the business of assigning moral responsibility to anyone other than the responsibility to do the best they can with the utilitarian principles. While they might agree that, as it happens, running a five-year-old Adolph Hitler down in a car might have been a morally better act than swerving to miss him, they will refrain from saying that the agent is morally culpable, because the agent could not have known the future consequences. (But the utilitarian, especially the act utilitarian, would say that if the person knew that this boy would grow up to do the things he was going to do, and there was no less harmful way to prevent him from doing them, then the driver should have run the boy down.)

[21] Another set of objections has to do with determining whose values and goods we are to count in the utility calculation. This issue has a number of aspects. First, whose values and goods do we appeal to? Should we force another culture to change because we think it will increase happiness even if they do not think so? (Do we pressure cultures that maintain that women are subordinate to men to change, even in cases where all the people, including the women, claim that they like it that way?) Do we count the utility of ourselves, all humans, all adults, all animals? And do we include all future generations? (Arguably, if we take all future human generations into account, tremendous utility could be gained by subjecting a few hundred thousand people now to a wide range of medical experimentation and tests that we would think now are unacceptable; but looking forward to the future, which might have many thousands of generations over perhaps millions of years, the medical gains made now might well outweigh the current suffering.

2. Deontological views

[22] The deontological orientation in ethics had its first major proponent in Immanuel Kant. This approach is broad, but to a large degree can be seen as based on the concepts of *duties* and *rights*. What is meant by duties and rights here must be understood correctly. We often think of duties as things that are imposed upon us by external powers, and rights as things granted by external powers. And surely there are many such duties and rights. I have the duty to pay taxes and to serve jury duty when called upon; and my license gives me the right to drive a car.

[23] But these are not the kinds of duties and rights that deontology is concerned with. When we speak of human rights, for example, we speak of rights that human beings have not because of some external power, this or that government, granted them this right. In such a case this power could take that right away. Rather, we mean rights that a human being has simply in virtue of being a human being, not because of any external license was granted to them for the purpose, but simply because they are people. And the duty to respect human rights is not derived from some law that we have to follow, but rather, the deontologist argues, is derived from within, from the nature of reason and morality itself. To some extent systems of laws written, adopted and enforced by governments are designed, in part, to protect such rights and enforce these kinds of duties. But the deontologist claims that even in these cases, the law is not creating the rights and duties. These were there all along. The law is simply codifying these rights and duties in an enforceable way.

[24] For Kant, one of the primary duties we have is to treat people, including oneself, always as ends, and never as mere means. An end is the final purpose for which one does something, and the means are how one achieves that purpose. So Kant is saying that we have the duty to always treat people as ends – that is, always as worthwhile in themselves – and never as merely having a use for achieving some other purpose. This ethical regulation simultaneously defines a duty all people have, and a right all people have. We might, very roughly, put it this way: we all have the duty to treat all people with respect. Another categorical imperative for Kant was to always tell the truth.

[25] Exactly what the rights and duties are is a separate question. But pretty much all versions of deontological theories will recognize the duty to respect human dignity, and the right to such dignity. These core duties and rights are duties for all, and rights held by all. But the view also makes room for some special duties and rights that are not impartially distributed. For example, parents may have certain duties with respect to their children that their children do not have for their parents. Parents have the duty to provide for the

protection, health and education of their children, say, but children do not have the same duty towards their parents (at least not while they are very young).

[26] Thus while, as we saw, utilitarianism was impartial, treating everyone as on the same footing, deontological theories typically have an impartial, universal core (such as the right to human dignity), but can also accommodate special cases of rights and duties. (As a note, communitarian theories of ethics seek to show how one's position in a community, which might be a family, a team, a country, etc.) can impose various prima facie duties. We won't dwell on communitarianism in this course.)

[27] It is easy to see many cases in which utilitarian views and deontological principles give conflicting answers to ethical questions. For example in the movie *Minority Report*, the three siblings who could see the future were being used merely as means, their human dignity was being sacrificed in order to achieve some purpose. Now it is quite arguable that the overall benefit, in terms of happiness, received to society was greater than the harm caused, in terms of unhappiness suffered by the siblings. While a utilitarian does not like the fact that three people are being made to suffer in some sense, the fact that potentially many thousands of violent crimes were being prevented outweighed this fact. And so using them in this way would be the right thing to do by utilitarian standards. A deontologist, by contrast, is much less concerned with overall happiness, and would state that in fact doing this is wrong, because the rights of these three people are being violated. They are being treated as means, and not as ends.

[28] This leads to another difference between deontological and utilitarian views. Deontological views typically must claim that there is a very morally significant difference between causing something, and failing to prevent something. This is similar to the distinction between acts of commission and acts of omission. Whereas utilitarians do not recognize this as a significant difference. Note that in the *Minority Report* case, one could argue that in order to treat the siblings with dignity, one would have to allow thousands of people to be victims of violent crimes such as rape and murder. The utilitarian sees the two options as (a) imposing on the dignity of three people while respecting the dignity of thousands; or (b) respecting the dignity of three while effectively sentencing thousands to be the victims of violent crimes (which is also an affront to their dignity, of course). For the utilitarian, who is causing the crimes or indignities is not the present concern. But for the deontologist, this is a crucial concern, for if I am in the position of releasing the siblings or keeping them confined, then this is the issue that my duties concern. It is true that if I let them free, I will effectively be failing to prevent violent crimes. But on this view, what actions I perform carries more moral significance than what things I fail to prevent. And so my duty would be to treat these three siblings as ends, respect their dignity, and free them.

3. Contractarianism/contractualism

[29] Finally, a brief mention of contractarianism (my brief discussion will more or less conflate contractarianism and contractualism). The contractualist is concerned to find a justification for specific ethical norms, a justification that doesn't *necessarily* come from expected utility calculations or appeal to rights and duties. For the contractarian, ethical norms are those that are or would be agreed to by rational agents in some society. For example, Hobbes felt that the power of sovereign monarchs was justified because rational agents would realize that without some sort of power with various rights and duties, social life would dissolve back into a state of nature without order, and would be, as he put it, "solitary, poore, nasty, brutish and short." By agreeing to give up one's own autonomy on

many circumstances and give them to a sovereign, a group of people can thereby create an orderly, lawful and safe social setting. Notice that there are some similarities in this line to a utilitarian, but there are crucial differences. The similarity is that as Hobbes describes it, each agent reaches the conclusion that giving authority to a sovereign is a good choice based on utilitarian principles: doing that will lead to better consequences for me than not doing that. The difference is crucial, however, for the contractarian does not claim that the utilitarian calculus is what justifies the rules that are adopted. What justifies them is the explicit or implicit agreement of those in a society, their social contract. Why an agent would agree to such a contract is another matter, Hobbes felt it was by and large a matter of utilitarian reasoning.

[30] More recent versions of contractarianism are similar. Rawls, for example, asks us to imagine being behind what he calls 'the veil of ignorance'. Behind the veil, each agent knows more or less how the world is structured, but does not have any idea who one is, one's age, race, gender, cultural background, all particular information, is unknown. Rawls then asks what ethical principles would one agree to in those circumstances, knowing that once the agreement is made, one will then be placed in some particular situation in the world (particular gender, race, culture, etc.). Rawls claims that the ethical principles rational agents would agree to in such circumstances are then binding.

[31] Exactly what ethical norms would be agreed to by people in such a situation is not really clear. One might conclude that a system of rights and duties would be the centerpiece of a social contract one would want to endorse; or one might opt for a utilitarian ethics; or one might opt for some entirely different set of ethical principles. Various contractarians have of course not only defended the basic idea of what justifies a certain set of norms (the social contract), but have also argued about what specific set of norms rational agents in such a situation would agree to (as Hobbes argues that rational agents would agree to give power to a law-providing and law-enforcing monarch). But the main point is that contractarianism is not primarily an account of what specifically the correct ethical norms are, but it is rather an account of what gives some set of ethical rules their normative force. On this account, what gives a set of ethical rules normative force is the fact that a rational agent would agree to them in order to construct an acceptable social order.

GENETICS AND HUMAN MALLEABILITY

W French Anderson

From *Hastings Center Report*, vol. 20 (January/February 1990), pp. 21-24. © The Hastings Center.

[1] Just how much can, and should we change human nature by genetic engineering? Our response to that hinges on the answers to three further questions: (1) What can we do now? Or more precisely, what are we doing now in the area of human genetic engineering? (2) What will we be able to do? In other words, what technical advances are we likely to achieve over the next five to ten years? (3) What *should* we do? I will argue that a line can be drawn and should be drawn to use gene transfer only for the treatment of serious disease, and not for any other purpose. Gene transfer should never be undertaken in an attempt to enhance or "improve" human beings.

WHAT CAN WE DO?

[2] In 1980 John Fletcher and I published a paper in the *New England Journal of Medicine* in which we delineated what would be necessary before it would be ethical to carry out human gene therapy. As with any other new therapeutic procedure, the fundamental principle is that it should be determined in advance that the probable benefits outweigh the probable risks. We analyzed the risk/benefit determination for somatic cell gene therapy and proposed three questions that need to have been answered from prior animal experimentation: Can the new gene be inserted stably into the correct target cells? Will the new gene be expressed (that is, function) in the cells at an appropriate level? Will the new gene harm the cell or the animal? These criteria are very similar to those required before use of any new therapeutic procedure, surgical operation, or drug. They simply require that the new treatment should get to the area of disease, correct it, and do more good than harm.

[3] A great deal of scientific progress has occurred in the nine years since that paper was published. The technology does now exist for inserting genes into some types of target cells. The procedure being used is called "retroviral-mediated gene transfer." In brief, a disabled murine retrovirus serves as a delivery vehicle for transporting a gene into a population of cells that have been removed from a patient. The gene-engineered cells are then returned to the patient.

WHAT WILL WE BE ABLE TO DO?

[4] Many genetic diseases that are caused by a defect in a single gene should be treatable, such as ADA deficiency (a severe immune deficiency disease of children), sickle cell anemia, hemophilia, and Gaucher disease. Some types of cancer, viral diseases such as AIDS, and some forms of cardiovascular disease are targets for treatment by gene therapy. In addition, germ-line gene therapy, that is, the insertion of a gene into the reproductive cells of a patient, will probably be technically possible in the foreseeable future.

[5] But successful somatic cell gene therapy also opens the door for enhancement genetic engineering, that is, for supplying a specific characteristic that individuals might want for themselves (somatic cell engineering) or their children (germ-line engineering) which would not involve the treatment of a disease. The most obvious example at the moment would be the insertion of a growth hormone gene into a normal child in the hope that this would make the child grow larger. Should parents be allowed to choose (if the science should ever make it possible) whatever useful characteristics they wish for their children?

WHAT SHOULD WE DO?

[6] A line can and should be drawn between somatic cell gene therapy and enhancement genetic engineering. Our society has repeatedly demonstrated that it can draw a line in biomedical research when necessary. The Belmont Report illustrates how guidelines were formulated to delineate ethical from unethical clinical research and to distinguish clinical research from clinical practice. Our responsibility is to determine how and where to draw lines with respect to genetic engineering.

[7] Somatic cell gene therapy for the treatment of severe disease is considered ethical because it can be supported by the fundamental moral principle of beneficence: It would relieve human suffering. Gene therapy would be, therefore, a moral good. Under what circumstances would human genetic engineering not be a moral good? In the broadest sense, when it detracts from, rather than contributes to, the dignity of man. Whether viewed from a theological perspective or a secular humanist one, the justification for drawing a line is founded on the argument that, beyond the line, human values that our society considers important for the dignity of man would be significantly threatened.

[8] Somatic cell enhancement engineering would threaten important human values in two ways: It could be medically hazardous, in that the risks could exceed the potential benefits and the procedure therefore cause harm. And it would be morally precarious, in that it would require moral decisions our society is not now prepared to make, and it could lead to an increase in inequality and discriminatory practices.

[9] Medicine is a very inexact science. We understand roughly how a simple gene works and that there are many thousands of housekeeping genes, that is, genes that do the job of running a cell. We predict that there are genes which make regulatory messages that are involved in the overall control and regulation of the many housekeeping genes. Yet we have only limited understanding of how a body organ develops into the size and shape it does. We know many things about how the central nervous system works -- for example, we are beginning to comprehend how molecules are involved in electric circuits, in memory storage, in transmission of signals. But we are a long way from understanding thought and consciousness. And we are even further from understanding the spiritual side of our existence.

[10] Even though we do not understand how a thinking, loving, interacting organism can be derived from its molecules, we are approaching the time when we can change some of those molecules. Might there be genes that influence the brain's organization or structure or metabolism or circuitry in some way so as to allow abstract thinking, contemplation of good and evil, fear of death, awe of a 'God'? What if in our innocent attempts to improve our genetic make-up we alter one or more of those genes? Could we test for the alteration? Certainly not at present. If we caused a problem that would affect the individual or his or her offspring, could we repair the damage? Certainly not at present. Every parent who has several children knows that some babies accept and give more affection than others, in the same environment. Do genes control this? What if these genes were accidentally altered? How would we even know if such a gene were altered?

[11] My concern is that, at this point in the development of our culture's scientific expertise, we might be like the young boy who loves to take things apart. He is bright enough to disassemble a watch, and maybe even bright enough to get it back together again so that it works. But what if he tries to "improve" it? Maybe put on bigger hands so that the time can be read more easily. But if the hands are too heavy for the mechanism, the watch will run

slowly, erratically, or not at all. The boy can understand what is visible, but he cannot comprehend the precise engineering calculations that determined exactly how strong each spring should be, why the gears interact in the ways that they do, etc. Attempts on his part to improve the watch will probably only harm it. We are now able to provide a new gene so that a property involved in a human life would be changed, for example, a growth hormone gene. If we were to do so simply because we could, I fear we would be like that young boy who changed the watch's hands. We, too, do not really understand what makes the object we are tinkering with tick.

[12] In summary, it could be harmful to insert a gene into humans. In somatic cell gene therapy for an already existing disease the potential benefits could outweigh the risks. In enhancement engineering, however, the risks would be greater while the benefits would be considerably less clear.

[13] Yet even aside from the medical risks, somatic cell enhancement engineering should not be performed because it would be morally precarious. Let us assume that there were no medical risks at all from somatic cell enhancement engineering. There would still be reasons for objecting to this procedure. To illustrate, let us consider some examples. What if a human gene were cloned that could produce a brain chemical resulting in markedly increased memory capacity in monkeys after gene transfer? Should a person be allowed to receive such a gene on request? Should a pubescent adolescent whose parents are both five feet tall be provided with a growth hormone gene on request? Should a worker who is continually exposed to an industrial toxin receive a gene to give him resistance on his, or his employer's request?

[14] These scenarios suggest three problems that would be difficult to resolve: What genes should be provided; who should receive a gene; and, how to prevent discrimination against individuals who do or do not receive a gene.

[15] We allow that it would be ethically appropriate to use somatic cell gene therapy for treatment of serious disease. But what distinguishes a serious disease from a "minor" disease from cultural "discomfort"? What is suffering? What is significant suffering? Does the absence of growth hormone that results in a growth limitation to two feet in height represent a genetic disease? What about a limitation to a height of four feet, to five feet? Each observer might draw the lines between serious disease, minor disease, and genetic variation differently. But all can agree that there are extreme cases that produce significant suffering and premature death. Here then is where an initial line should be drawn for determining what genes should be provided: treatment of serious disease.

[16] If the position is established that only patients suffering from serious diseases are candidates for gene insertion, then the issues of patient selection are no different than in other medical situations: the determination is based on medical need within a supply and demand framework. But if the use of gene transfer extends to allow a normal individual to acquire, for example, a memory-enhancing gene, profound problems would result. On what basis is the decision made to allow one individual to receive the gene but not another: Should it go to those best able to benefit society (the smartest already)? To those most in need (those with low intelligence? But how low? Will enhancing memory help a mentally retarded child)? To those chosen by a lottery? To those who can afford to pay? As long as our society lacks a significant consensus about these answers, the best way to make equitable decisions in this case should be to base them on the seriousness of the objective medical need, rather than on the personal wishes or resources of an individual.

[17] Discrimination can occur in many forms. If individuals are carriers of a disease (for example, sickle cell anemia), would they be pressured to be treated? Would they have difficulty in obtaining health insurance unless they agreed to be treated? These are ethical issues raised also by genetic screening and by the Human Genome project. But the concerns would become even more troublesome if there were the possibility for "correction" by the use of human genetic engineering.

[17] Finally, we must face the issue of eugenics, the attempt to make hereditary "improvements." The abuse of power that societies have historically demonstrated in the pursuit of eugenic goals is well documented. Might we slide into a new age of eugenic thinking by starting with small "improvements"? It would be difficult, if not impossible, to determine where to draw a line once enhancement engineering had begun. Therefore, gene transfer should be used only for the treatment of serious disease and not for putative improvements.

[18] Our society is comfortable with the use of genetic engineering to treat individuals with serious disease. On medical and ethical grounds we should draw a line excluding any form of enhancement engineering. We should not step over the line that delineates treatment from enhancement.

Questions About Some Uses of Genetic Engineering

JONATHAN GLOVER

[1] There is a widespread view that any project for the genetic improvement of the human race ought to be ruled out: that there are fundamental objections of principle. The aim of this discussion is to sort out some of the main objections. It will be argued that our resistance is based on a complex of different values and reasons, none of which is, when examined, adequate to rule out in principle this use of genetic engineering. The debate on human genetic engineering should become like the debate on nuclear power: one in which large possible benefits have to be weighed against big problems and the risk of great disasters. The discussion has not reached this point, partly because the techniques have not yet been developed. But it is also partly because of the blurred vision which fuses together many separate risks and doubts into a fuzzy-outlined opposition in principle.

1. AVOIDING THE DEBATE ABOUT GENES AND THE ENVIRONMENT

[2] In discussing the question of genetic engineering, there is everything to be said for not muddling the issue up with the debate over the relative importance of genes and environment in the development of such characteristics as intelligence. One reason for avoiding that debate is that it arouses even stronger passions than genetic engineering, and so is filled with as much acrimony as argument. But, apart from this fastidiousness, there are other reasons.

[3] The nature-nurture dispute is generally seen as an argument about the relative weight the two factors have in causing differences within the human species: 'IQ is 80 percent hereditary and 20 per cent environmental' versus 'IQ is 80 percent environmental and 20 percent hereditary'. No doubt there is some approximate truth of this type to be found if we consider variations within a given population at a particular time. But it is highly unlikely that there is any such statement which is simply true of human nature regardless of context. To take the extreme case, if we could iron out all environmental differences, any residual variations would be 100 per cent genetic. It is only if we make the highly artificial assumption that different groups at different times all have an identical spread of relevant environmental differences that we can expect to find statements of this kind applying to human nature in general. To say this is not to argue that studies on the question should not be conducted, or are bound to fail. It may well be possible, and useful, to find out the relative weights of the two kinds of factor for a given characteristic among a certain group at a particular time. The point is that any such conclusions lose relevance, not only when environmental differences are stretched out or compressed, but also when genetic differences are. And this last case is what we are considering.

[4] We can avoid this dispute because of its irrelevance. Suppose the genetic engineering proposal were to try to make people less aggressive. On a superficial view, the proposal might be shown to be unrealistic if there were evidence to show that variation in aggressiveness is hardly genetic at all: that it is 95 per cent environmental. (Let us grant, most implausibly, that such a figure turned out to be true for the whole of humanity, regardless of social context.) But all this would show is that, within our species, the distribution of genes relevant to aggression is very uniform. It would show nothing about the likely effects on aggression if we use genetic engineering to give people a different set of genes from those they now have.

[5] In other words, to take genetic engineering seriously, we need take no stand on the relative importance or unimportance of genetic factors in the explanation of the present range of individual differences found in people. We need only the minimal assumption that different genes could give us different characteristics. To deny that assumption you need to be the sort of person who thinks it is only living in kennels which makes dogs different from cats.

2. METHODS OF CHANGING THE GENETIC COMPOSITION OF FUTURE GENERATIONS

[6] There are essentially three ways of altering the genetic composition of future generations. The first is by environmental changes. Discoveries in medicine, the institution of a National Health Service, schemes for poverty relief, agricultural changes, or alterations in the tax position of large families, all alter the selective pressure on genes. It is hard to think of any social change which does not make some difference to who survives or who is born.

[7] The second method is to use eugenic policies aimed at altering breeding patterns or patterns of survival of people with different genes. Eugenic methods are 'environmental' too: the difference is only that the genetic impact is intended. Possible strategies range from various kinds of compulsion (to have more children, fewer children, or no children, or even compulsion over the choice of sexual partner) to the completely voluntary (our present genetic counselling practice of giving prospective parents information about probabilities of their children having various abnormalities).

[8] The third method is genetic engineering: using enzymes to add to or subtract from a stretch of DNA.

[9] Most people are unworried by the fact that a side-effect of an environmental change is to alter the gene pool, at least where the alteration is not for the worse. And even in cases where environmental factors increase the proportion of undesirable genes in the pool, we often accept this. Few people oppose the National Health Service, although setting it up meant that some people with genetic defects, who would have died, have had treatment enabling them to survive and reproduce. On the whole, we accept without qualms that much of what we do has genetic impact. Controversy starts when we think of aiming deliberately at genetic changes, by eugenics or genetic engineering. I want to make some brief remarks about eugenic policies, before suggesting that policies of deliberate intervention are best considered in the context of genetic engineering.

[10] Scepticism has been expressed about whether eugenic policies have any practical chance of success. Medawar has pointed out the importance of genetic polymorphism: the persistence of genetically different types in a population. (Our different blood groups are a familiar example.) For many characteristics, people get a different gene from each parent. So children do not simply repeat parental characteristics. Any simple picture of producing an improved type of person, and then letting the improvement be passed on unchanged, collapses.

[11] But, although polymorphism is a problem for this crudely utopian form of eugenics, it does not show that more modest schemes of improvement must fail. Suppose the best individuals for some quality (say, colour vision) are heterozygous, so that they inherit a gene A from one parent, and a gene B from the other. These ABs will have AAs and BBs among their children, who will be less good than they are. But AAs and BBs may still be

better than ACs or ADs, and perhaps much better than CCs or CDs. If this were so, overall improvement could still be brought about by encouraging people whose genes included an A or a B to have more children than those who had only Cs or Ds. The point of taking a quality like colour vision is that it may be genetically fairly simple. Qualities like kindness or intelligence are more likely to depend on the interaction of many genes, but a similar point can be made at a higher level of complexity.

[12] Polymorphism raises a doubt about whether the offspring of the three 'exceptionally intelligent women' fertilized by Dr Shockley or other Nobel prize-winners will have the same IQ as the parents, even apart from environmental variation. But it does not show the inevitable failure of any large-scale attempts to alter human characteristics by varying the relative numbers of children different kinds of people have. Yet any attempt, say, to raise the level of intelligence, would be a very slow affair, taking many generations to make much of an impact. This is one reason for preferring to discuss genetic engineering. For the genetic engineering of human improvements, if it becomes possible, will have an immediate effect, so we will not be guessing which qualities will be desirable dozens of generations later.

[13] There is the view that the genetic engineering techniques required will not become a practical possibility. Sir MacFarlane Burnet, writing in 1971 about using genetic engineering to cure disorders in people already born, dismissed the possibility of using a virus to carry a new gene to replace a faulty one in cells throughout the body: 'I should be willing to state in any company that the chance of doing this will remain infinitely small to the last syllable of recorded time.' Unless engineering at the stage of sperm cell and egg is easier, this seems a confident dismissal of the topic to be discussed here. More recent work casts doubt on this confidence. So, having mentioned this scepticism, I shall disregard it. We will assume that genetic engineering of people may become possible, and that it is worth discussing. (Sir MacFarlane Burnet's view has not yet been falsified as totally as Rutherford's view about atomic energy. But I hope that the last syllable of recorded time is still some way off.)

[14] The main reason for casting the discussion in terms of genetic engineering rather than eugenics is not a practical one. Many eugenic policies are open to fairly straightforward moral objections, which hide the deeper theoretical issues. Such policies as compulsory sterilization, compulsory abortion, compelling people to pair off in certain ways, or compelling people to have more or fewer children than they would otherwise have, are all open to objection on grounds of overriding people's autonomy. Some are open to objection on grounds of damage to the institution of the family. And the use of discriminatory tax- and child-benefit policies is an intolerable step towards a society of different genetic castes.

[15] Genetic engineering need not involve overriding anyone's autonomy. It need not be forced on parents against their wishes, and the future person being engineered has no views to be overridden. (The view that despite this, it is still objectionable to have one's genetic characteristics decided by others, will be considered later.) Genetic engineering will not damage the family in the obvious ways that compulsory eugenic policies would. Nor need it be encouraged by incentives which create inequalities. Because it avoids these highly visible moral objections, genetic engineering allows us to focus more clearly on other values that are involved.

[16] (To avoid a possible misunderstanding, one point should be added before leaving the topic of eugenics. Saying that some eugenic policies are open to obvious moral objections does not commit me to disapproval of all eugenic policies. In particular, I do not want to be taken to be opposing two kinds of policy. One is genetic counselling: warning people of risks in having children, and perhaps advising them against having them. The other is the

introduction of screening-programmes to detect foetal abnormalities, followed by giving the mother the option of abortion where serious defects emerge.)

[17] Let us now turn to the question of what, if anything, we should do in the field of human genetic engineering.

3. THE POSITIVE-NEGATIVE DISTINCTION

[18] We are not yet able to cure disorders by genetic engineering. But we do sometimes respond to disorders by adopting eugenic policies, at least in voluntary form. Genetic counselling is one instance, as applied to those thought likely to have such disorders as Huntington's chorea. This is a particularly appalling inherited disorder, involving brain degeneration, leading to mental decline and lack of control over movement. It does not normally come on until middle age, by which time many of its victims would in the normal course of things have had children. Huntington's chorea is caused by a dominant gene, so those who find that one of their parents has it have themselves a 50 per cent chance of developing it. If they do have it, each of their children will in turn have a 50 per cent chance of the disease. The risks are so high and the disorder so bad that the potential parents often decide not to have children, and are often given advice to this effect by doctors and others.

[19] Another eugenic response to disorders is involved in screening-programmes for pregnant women. When tests pick up such defects as Down's syndrome (mongolism) or spina bifida, the mother is given the possibility of an abortion. The screening-programmes are eugenic because part of their point is to reduce the incidence of severe genetic abnormality in the population.

[20] These two eugenic policies come in at different stages: before conception and during pregnancy. For this reason the screening-programme is more controversial, because it raises the issue of abortion. Those who are sympathetic to abortion, and who think it would be good to eliminate these disorders will be sympathetic to the programme. Those who think abortion is no different from killing a fully developed human are obviously likely to oppose the programme. But they are likely to feel that elimination of the disorders would be a good thing, even if not an adequate justification for killing. Unless they also disapprove of contraception, they are likely to support the genetic-counselling policy in the case of Huntington's chorea.

[21] Few people object to the use of eugenic policies to eliminate disorders, unless those policies have additional features which are objectionable. Most of us are resistant to the use of compulsion, and those who oppose abortion will object to screening-programmes. But apart from these other moral objections, we do not object to the use of eugenic policies against disease. We do not object to advising those likely to have Huntington's chorea not to have children, as neither compulsion nor killing is involved. Those of us who take this view have no objection to altering the genetic composition of the next generation, where this alteration consists in reducing the incidence of defects.

[22] If it were possible to use genetic engineering to correct defects, say at the foetal stage, it is hard to see how those of us who are prepared to use the eugenic measures just mentioned could object. In both cases, it would be pure gain. The couple, one of whom may develop Huntington's chorea, can have a child if they want, knowing that any abnormality will be eliminated. Those sympathetic to abortion will agree that cure is preferable. And those opposed to abortion prefer babies to be born without handicap. It is hard to think of

any objection to using genetic engineering to eliminate defects, and there is a clear and strong case for its use.

[23] But accepting the case for eliminating genetic mistakes does not entail accepting other uses of genetic engineering. The elimination of defects is often called 'negative' genetic engineering. Going beyond this, to bring about improvements in normal people, is by contrast 'positive' engineering. (The same distinction can be made for eugenics.)

[24] The positive-negative distinction is not in all cases completely sharp. Some conditions are genetic disorders whose identification raises little problem. Huntington's chorea or spina bifida are genetic 'mistakes' in a way that cannot seriously be disputed. But with other conditions, the boundary between a defective state and normality may be more blurred. If there is a genetic disposition towards depressive illness, this seems a defect, whose elimination would be part of negative genetic engineering. Suppose the genetic disposition to depression involves the production of lower levels of an enzyme than are produced in normal people. The negative programme is to correct the genetic fault so that the enzyme level is within the range found in normal people. But suppose that within 'normal' people also, there are variations in the enzyme level, which correlate with ordinary differences in tendency to be cheerful or depressed. Is there a sharp boundary between 'clinical' depression and the depression sometimes felt by those diagnosed as 'normal'? Is it clear that a sharp distinction can be drawn between raising someone's enzyme level so that it falls within the normal range and raising someone else's level from the bottom of the normal range to the top?

[25] The positive-negative distinction is sometimes a blurred one, but often we can at least roughly see where it should be drawn. If there is a rough and ready distinction, the question is: how important is it? Should we go on from accepting negative engineering to accepting positive programmes, or should we say that the line between the two is the limit of what is morally acceptable?

[26] There is no doubt that positive programmes arouse the strongest feelings on both sides. On the one hand, many respond to positive genetic engineering or positive eugenics with Professor Tinbergen's thought: 'I find it morally reprehensible and presumptuous for anybody to put himself forward as a judge of the qualities for which we should breed.'

[27] But other people have held just as strongly that positive policies are the way to make the future of mankind better than the past. Many years ago H. J. Muller expressed this hope:

And so we foresee the history of life divided into three main phases. In the long preparatory phase it was the helpless creature of its environment, and natural selection gradually ground it into human shape. In the second -- our own short transitional phase -- it reaches out at the immediate environment, shaking, shaping and grinding to suit the form, the requirements, the wishes, and the whims of man. And in the long third phase, it will reach down into the secret places of the great universe of its own nature, and by aid of its ever growing intelligence and cooperation, shape itself into an increasingly sublime creation -- a being beside which the mythical divinities of the past will seem more and more ridiculous, and which setting its own marvellous inner powers against the brute Goliath of the suns and the planets, challenges them to contest.

[28] The case for positive engineering is not helped by adopting the tones of the mad scientist in a horror film. But behind the rhetoric is a serious point. If we decide on a positive programme to change our nature, this will be a central moment in our history, and

the transformation might be beneficial to a degree we can now scarcely imagine. The question is: how are we to weigh this possibility against Tinbergen's objection, and against other objections and doubts?

[29] For the rest of this discussion, I shall assume that, subject to adequate safeguards against things going wrong, negative genetic engineering is acceptable. The issue is positive engineering. I shall also assume that we can ignore problems about whether positive engineering will be technically possible. Suppose we have the power to choose people's genetic characteristics. Once we have eliminated genetic defects, what, if anything, should we do with this power?.

4. THE VIEW THAT OVERALL IMPROVEMENT IS UNLIKELY OR IMPOSSIBLE

[30] There is one doubt about the workability of schemes of genetic improvement which is so widespread that it would be perverse to ignore it. This is the view that, in any genetic alteration, there are no gains without compensating losses. On this view, if we bring about a genetically based improvement, such as higher intelligence, we are bound to pay a price somewhere else: perhaps the more intelligent people will have less resistance to disease, or will be less physically agile. If correct, this might so undermine the practicability of applying eugenics or genetic engineering that it would be hardly worth discussing the values involved in such programmes.

[31] This view perhaps depends on some idea that natural selection is so efficient that, in terms of gene survival, we must already be as efficient as it is possible to be. If it were possible to push up intelligence without weakening some other part of the system, natural selection would already have done so. But this is a naive version of evolutionary theory. In real evolutionary theory, far from the genetic status quo always being the best possible for a given environment, some mutations turn out to be advantageous, and this is the origin of evolutionary progress. If natural mutations can be beneficial without a compensating loss, why should artificially induced ones not be so too?

[32] It should also be noticed that there are two different ideas of what counts as a gain or a loss. From the point of view of evolutionary progress, gains and losses are simply advantages and disadvantages from the point of view of gene survival. But we are not compelled to take this view. If we could engineer a genetic change in some people which would have the effect of making them musical prodigies but also sterile, this would be a hopeless gene in terms of survival, but this need not force us, or the musical prodigies themselves, to think of the change as for the worse. It depends on how we rate musical ability as against having children, and evolutionary survival does not dictate priorities here.

[33] The view that gains and losses are tied up with each other need not depend on the dogma that natural selection must have created the best of all possible sets of genes. A more cautiously empirical version of the claim says there is a tendency for gains to be accompanied by losses. John Maynard Smith, in his paper on 'Eugenics and Utopia', takes this kind of 'broad balance' view and runs it the other way, suggesting, as an argument in defence of medicine, that any loss of genetic resistance to disease is likely to be a good thing: 'The reason for this is that in evolution, as in other fields, one seldom gets something for nothing. Genes which confer disease-resistance are likely to have harmful effects in other ways: this is certainly true of the gene for sickle-cell anaemia and may be a general rule. If so, absence of selection in favour of disease resistance may be eugenic.'

[34] It is important that different characteristics may turn out to be genetically linked in ways we do not yet realize. In our present state of knowledge, engineering for some improvement might easily bring some unpredicted but genetically linked disadvantage. But we do not have to accept that there will in general be a broad balance, so that there is a presumption that any gain will be accompanied by a compensating loss (or Maynard Smith's version that we can expect a compensating gain for any loss). The reason is that what counts as a gain or loss varies in different contexts. Take Maynard Smith's example of sickle-cell anaemia. The reason why sickle-cell anaemia is widespread in Africa is that it is genetically linked with resistance to malaria. Those who are heterozygous (who inherit one sickle-cell gene and one normal gene) are resistant to malaria, while those who are homozygous (whose genes are both sickle-cell) get sickle-cell anaemia. If we use genetic engineering to knock out sickle-cell anaemia where malaria is common, we will pay the price of having more malaria. But when we eradicate malaria, the gain will not involve this loss. Because losses are relative to context, any generalization about the impossibility of overall improvements is dubious.

5. THE FAMILY AND OUR DESCENDANTS

[35] Unlike various compulsory eugenic policies, genetic engineering need not involve any interference with decisions by couples to have children together, or with their decisions about how many children to have. And let us suppose that genetically engineered babies grow in the mother's womb in the normal way, so that her relationship to the child is not threatened in the way it might be if the laboratory or the hospital were substituted for the womb. The cruder threats to family relationships are eliminated.

[36] It may be suggested that there is a more subtle threat. Parents like to identify with their children. We are often pleased to see some of our own characteristics in our children. Perhaps this is partly a kind of vanity, and no doubt sometimes we project on to our children similarities that are not really there. But, when the similarities do exist, they help the parents and children to understand and sympathize with each other. If genetic engineering resulted in children fairly different from their parents, this might make their relationship have problems.

[37] There is something to this objection, but it is easy to exaggerate. Obviously, children who were like Midwich cuckoos, or comic-book Martians, would not be easy to identify with. But genetic engineering need not move in such sudden jerks. The changes would have to be detectable to be worth bringing about, but there seems no reason why large changes in appearance, or an unbridgeable psychological gulf, should be created in any one generation. We bring about environmental changes which make children different from their parents, as when the first generation of children in a remote place are given schooling and made literate. This may cause some problems in families, but it is not usually thought a decisive objection. It is not clear that genetically induced changes of similar magnitude are any more objectionable.

[38] A related objection concerns our attitude to our remoter descendants. We like to think of our descendants stretching on for many generations. Perhaps this is in part an immortality substitute. We hope they will to some extent be like us, and that, if they think of us, they will do so with sympathy and approval. Perhaps these hopes about the future of mankind are relatively unimportant to us. But, even if we mind about them a lot, they are unrealistic in the very long term. Genetic engineering would make our descendants less like us, but this would only speed up the natural rate of change. Natural mutations and selective pressures make it unlikely that in a few million years our descendants will be physically or

mentally much like us. So what genetic engineering threatens here is probably doomed anyway.

6. RISKS AND MISTAKES

[39] Although mixing different species and cloning are often prominent in people's thoughts about genetic engineering, they are relatively marginal issues. This is partly because there may be no strong reasons in favour of either. Our purposes might be realized more readily by improvements to a single species, whether another or our own, or by the creation of quite new types of organism, than by mixing different species. And it is not clear what advantage cloning batches of people might have, to outweigh the drawbacks. This is not to be dogmatic that species mixing and cloning could never be useful, but to say that the likelihood of other techniques being much more prominent makes it a pity to become fixated on the issues raised by these ones. And some of the most serious objections to positive genetic engineering have wider application than to these rather special cases. One of these wider objections is that serious risks may be involved.

[40] Some of the risks are already part of the public debate because of current work on recombinant DNA. The danger is of producing harmful organisms that would escape from our control. The work obviously should take place, if at all, only with adequate safe-guards against such a disaster. The problem is deciding what we should count as adequate safeguards. I have nothing to contribute to this problem here. If it can be dealt with satisfactorily, we will perhaps move on to genetic engineering of people. And this introduces another dimension of risk. We may produce unintended results, either because our techniques turn out to be less finely tuned than we thought, or because different characteristics are found to be genetically linked in unexpected ways.

[41] If we produce a group of people who turn out worse than expected, we will have to live with them. Perhaps we would aim for producing people who were especially imaginative and creative, and only too late find we had produced people who were also very violent and aggressive. This kind of mistake might not only be disastrous, but also very hard to 'correct' in subsequent generations. For when we suggested sterilization to the people we had produced, or else corrective genetic engineering for their offspring, we might find them hard to persuade. They might like the way they were, and reject, in characteristically violent fashion, our explanation that they were a mistake.

[42] The possibility of an irreversible disaster is a strong deterrent. It is enough to make some people think we should rule out genetic engineering altogether, and to make others think that, while negative engineering is perhaps acceptable, we should rule out positive engineering. The thought behind this second position is that the benefits from negative engineering are clearer, and that, because its aims are more modest, disastrous mistakes are less likely.

[43] The risk of disasters provides at least a reason for saying that, if we do adopt a policy of human genetic engineering, we ought to do so with extreme caution. We should alter genes only where we have strong reasons for thinking the risk of disaster is very small, and where the benefit is great enough to justify the risk. (The problems of deciding when this is so are familiar from the nuclear power debate.) This 'principle of caution' is less strong than one ruling out all positive engineering, and allows room for the possibility that the dangers may turn out to be very remote, or that greater risks of a different kind are involved in not using positive engineering. These possibilities correspond to one view of the facts in the nuclear power debate. Unless with genetic engineering we think we can already rule out

such possibilities, the argument from risk provides more justification for the principle of caution than for the stronger ban on all positive engineering.

DECISIONS

[44] Some of the strongest objections to positive engineering are not about specialized applications or about risks. They are about the decisions involved. The central line of thought is that we should not start playing God by redesigning the human race. The suggestion is that there is no group (such as scientists, doctors, public officials, or politicians) who can be entrusted with decisions about what sort of people there should be. And it is also doubted whether we could have any adequate grounds for basing such decisions on one set of values rather than another.

NOT PLAYING GOD

[45] Suppose we could use genetic engineering to raise the average IQ by fifteen points. (I mention, only to ignore, the boring objection that the average IQ is always by definition 100.) Should we do this? Objectors to positive engineering say we should not. This is not because the present average is preferable to a higher one. We do not think that, if it were naturally fifteen points higher, we ought to bring it down to the present level. The objection is to our playing God by deciding what the level should be.

[46] On one view of the world, the objection is relatively straightforward. On this view, there really is a God, who has a plan for the world which will be disrupted if we stray outside the boundaries assigned to us. (It is *relatively* straightforward: there would still be the problem of knowing where the boundaries came. If genetic engineering disrupts the programme, how do we know that medicine and education do not?)

[47] The objection to playing God has a much wider appeal than to those who literally believe in a divine plan. But, outside such a context, it is unclear what the objection comes to. If we have a Darwinian view, according to which features of our nature have been selected for their contribution to gene survival, it is not blasphemous, or obviously disastrous, to start to control the process in the light of our own values. We may value other qualities in people, in preference to those which have been most conducive to gene survival.

[48] The prohibition on playing God is obscure. If it tells us not to interfere with natural selection at all, this rules out medicine, and most other environmental and social changes. If it only forbids interference with natural selection by the direct alteration of genes, this rules out negative as well as positive genetic engineering. If these interpretations are too restrictive, the ban on positive engineering seems to need some explanation. If we can make positive changes at the environmental level, and negative changes at the genetic level, why should we not make positive changes at the genetic level? What makes this policy, but not the others, objectionably God-like?

[49] Perhaps the most plausible reply to these questions rests on a general objection to any group of people trying to plan too closely what human life should be like. Even if it is hard to distinguish in principle between the use of genetic and environmental means, genetic changes are likely to differ in degree from most environmental ones. Genetic alterations may be more drastic or less reversible, and so they can be seen as the extreme case of an objectionably God-like policy by which some people set out to plan the lives of others.

[50] This objection can be reinforced by imagining the possible results of a programme of positive engineering, where the decisions about the desired improvements were taken by scientists. Judging by the literature written by scientists on this topic, great prominence would be given to intelligence. But can we be sure that enough weight would be given to other desirable qualities? And do things seem better if for scientists we substitute doctors, politicians or civil servants? Or some committee containing businessmen, trade unionists, academics, lawyers and a clergyman?

[51] What seems worrying here is the circumscribing of potential human development. The present genetic lottery throws up a vast range of characteristics, good and bad, in all sorts of combinations. The group of people controlling a positive engineering policy would inevitably have limited horizons, and we are right to worry that the limitations of their outlook might become the boundaries of human variety. The drawbacks would be like those of town-planning or dog-breeding, but with more important consequences.

[52] When the objection to playing God is separated from the idea that intervening in this aspect of the natural world is a kind of blasphemy, it is a protest against a particular group of people, necessarily fallible and limited, taking decisions so important to our future. This protest may be on grounds of the bad consequences, such as loss of variety of people, that would come from the imaginative limits of those taking the decisions. Or it may be an expression of opposition to such concentration of power, perhaps with the thought: 'What right have they to decide what kinds of people there should be?' Can these problems be side-stepped?

[53] Robert Nozick is critical of the assumption that positive engineering has to involve any centralized decision about desirable qualities: 'Many biologists tend to think the problem is one of design, of specifying the best types of persons so that biologists can proceed to produce them. Thus they worry over what sort(s) of person there is to be and who will control this process. They do not tend to think, perhaps because it diminishes the importance of their role, of a system in which they run a "genetic supermarket", meeting the individual specifications (within certain moral limits) of prospective parents. Nor do they think of seeing what limited number of types of persons people's choices would converge upon, if indeed there would be any such convergence. This supermarket system has the great virtue that it involves no centralized decision fixing the future human type(s).'

[54] This idea of letting parents choose their children's characteristics is in many ways an improvement on decisions being taken by some centralized body. It seems less likely to reduce human variety, and could even increase it, if genetic engineering makes new combinations of characteristics available. (But we should be cautious here. Parental choice is not a guarantee of genetic variety, as the influence of fashion or of shared values might make for a small number of types on which choices would converge.)

[55] To those sympathetic to one kind of liberalism, Nozick's proposal will seem more attractive than centralized decisions. On this approach to politics, it is wrong for the authorities to institutionalize any religious or other outlook as the official one of the society. To a liberal of this kind, a good society is one which tolerates and encourages a wide diversity of ideals of the good life. Anyone with these sympathies will be suspicious of centralized decisions about what sort of people should form the next generation. But some parental decisions would be disturbing. If parents chose characteristics likely to make their children unhappy, or likely to reduce their abilities, we might feel that the children should be protected against this. (Imagine parents belonging to some extreme religious sect, who wanted their children to have a religious symbol as a physical mark on their face, and who wanted them to be unable to read, as a protection against their faith being corrupted.)

Those of us who support restrictions protecting children from parental harm after birth (laws against cruelty, and compulsion on parents to allow their children to be educated and to have necessary medical treatment) are likely to support protecting children from being harmed by their parents' genetic choices.

[56] No doubt the boundaries here will be difficult to draw. We already find it difficult to strike a satisfactory balance between protection of children and parental freedom to choose the kind of upbringing their children should have. But it is hard to accept that society should set no limits to the genetic choices parents can make for their children. Nozick recognizes this when he says the genetic supermarket should meet the specifications of parents 'within certain moral limits'. So, if the supermarket came into existence, some centralized policy, even if only the restrictive one of ruling out certain choices harmful to the children, should exist. It would be a political decision where the limits should be set.

[57] There may also be a case for other centralized restrictions on parental choice, as well as those aimed at preventing harm to the individual people being designed. The genetic supermarket might have more oblique bad effects. An imbalance in the ratio between the sexes could result. Or parents might think their children would be more successful if they were more thrusting, competitive and selfish. If enough parents acted on this thought, other parents with different values might feel forced into making similar choices to prevent their own children being too greatly disadvantaged. Unregulated individual decisions could lead to shifts of this kind, with outcomes unwanted by most of those who contribute to them. If a majority favour a roughly equal ratio between the sexes, or a population of relatively uncompetitive people, they may feel justified in supporting restrictions on what parents can choose. (This is an application to the case of genetic engineering of a point familiar in other contexts, that unrestricted individual choices can add up to a total outcome which most people think worse than what would result from some regulation.)

[58] Nozick recognizes that there may be cases of this sort. He considers the case of avoiding a sexual imbalance and says that 'a government could require that genetic manipulation be carried on so as to fit a certain ratio'. He clearly prefers to avoid governmental intervention of this kind, and, while admitting that the desired result would be harder to obtain in a purely libertarian system, suggests possible strategies for doing so. He says: 'Either parents would subscribe to an information service monitoring the recent births and so know which sex was in shorter supply (and hence would be more in demand in later life), thus adjusting their activities, or interested individuals would contribute to a charity that offers bonuses to maintain the ratios, or the ratio would leave 1:1, with new family and social patterns developing. The proposals for avoiding the sexual imbalance without central regulation are not reassuring. Information about likely prospects for marriage or sexual partnership might not be decisive for parents' choices. And, since those most likely to be 'interested individuals' would be in the age group being genetically engineered, it is not clear that the charity would be given donations adequate for its job.'

[59] If the libertarian methods failed, we would have the choice between allowing a sexual imbalance or imposing some system of social regulation. Those who dislike central decisions favouring one sort of person over others might accept regulation here, on the grounds that neither sex is being given preference: the aim is rough equality of numbers.

[60] But what about the other sort of case, where the working of the genetic supermarket leads to a general change unwelcome to those who contribute to it? Can we defend regulation to prevent a shift towards a more selfish and competitive population as merely being the preservation of a certain ratio between characteristics? Or have we crossed the boundary, and allowed a centralized decision favouring some characteristics over others?

The location of the boundary is obscure. One view would be that the sex-ratio case is acceptable because the desired ratio is equality of numbers. On another view, the acceptability derives from the fact that the present ratio is to be preserved. (In this second view, preserving altruism would be acceptable, so long as no attempt was made to raise the proportion of altruistic people in the population. But is this boundary an easy one to defend?)

[61] If positive genetic engineering does become a reality, we may be unable to avoid some of the decisions being taken at a social level. Or rather, we could avoid this, but only at what seems an unacceptable cost, either to the particular people being designed, or to their generation as a whole. And, even if the social decisions are only restrictive, it is implausible to claim that they are all quite free of any taint of preference for some characteristics over others. But, although this suggests that we should not be doctrinaire in our support of the liberal view, it does not show that the view has to be abandoned altogether. We may still think that social decisions in favour of one type of person rather than another should be few, even if the consequences of excluding them altogether are unacceptable. A genetic supermarket, modified by some central regulation, may still be better than a system of purely central decisions. The liberal value is not obliterated because it may sometimes be compromised for the sake of other things we care about.

A MIXED SYSTEM

[62] The genetic supermarket provides a partial answer to the objection about the limited outlook of those who would take the decisions. The choices need not be concentrated in the hands of a small number of people. The genetic supermarket should not operate in a completely unregulated way, and so some centralized decisions would have to be taken about the restrictions that should be imposed. One system that would answer many of the anxieties about centralized decision-making would be to limit the power of the decision-makers to one of veto. They would then only check departures from the natural genetic lottery, and so the power to bring about changes would not be given to them, but spread through the whole population of potential parents. Let us call this combination of parental initiative and central veto a 'mixed system'. If positive genetic engineering does come about, we can imagine the argument between supporters of a mixed system and supporters of other decision-making systems being central to the political theory of the twenty-first century, parallel to the place occupied in the nineteenth and twentieth centuries by the debate over control of the economy.

[63] My own sympathies are with the view that, if positive genetic engineering is introduced, this mixed system is in general likely to be the best one for taking decisions. I do not want to argue for an absolutely inviolable commitment to this, as it could be that some centralized decision for genetic change was the only way of securing a huge benefit or avoiding a great catastrophe. But, subject to this reservation, the dangers of concentrating the decision-making create a strong presumption in favour of a mixed system rather than one in which initiatives come from the centre. And, if a mixed system was introduced, there would have to be a great deal of political argument over what kinds of restrictions on the supermarket should be imposed. Twenty-first-century elections may be about issues rather deeper than economics.

[64] If this mixed system eliminates the anxiety about genetic changes being introduced by a few powerful people with limited horizons, there is a more general unease which it does not remove. May not the limitations of one generation of parents also prove disastrous? And, underlying this, is the problem of what values parents should appeal to in making their

choices. How can we be confident that it is better for one sort of person to be born than another?

VALUES

[65] The dangers of such decisions, even spread through all prospective parents, seem to me very real. We are swayed by fashion. We do not know the limitations of our own outlook. There are human qualities whose value we may not appreciate. A generation of parents might opt heavily for their children having physical or intellectual abilities and skills. We might leave out a sense of humour. Or we might not notice how important to us is some other quality, such as emotional warmth. So we might not be disturbed in advance by the possible impact of the genetic changes on such a quality. And, without really wanting to do so, we might stumble into producing people with a deep coldness. This possibility seems one of the worst imaginable. It is just one of the many horrors that could be blundered into by our lack of foresight in operating the mixed system. Because such disasters are a real danger, there is a case against positive genetic engineering, even when the changes do not result from centralized decisions. But this case, resting as it does on the risk of disaster, supports a principle of caution rather than a total ban. We have to ask the question whether there are benefits sufficiently great and sufficiently probable to outweigh the risks.

[66] But perhaps the deepest resistance, even to a mixed system, is not based on risks, but on a more general problem about values. Could the parents ever be justified in choosing, according to some set of values, to create one sort of person rather than another?

[67] Is it sometimes better for us to create one sort of person rather than another? We say 'yes' when it is a question of eliminating genetic defects. And we say 'yes' if we think that encouraging some qualities rather than others should be an aim of the upbringing and education we give our children. Any inclination to say 'no' in the context of positive genetic engineering must lay great stress on the two relevant boundaries.

[68] The positive-negative boundary is needed to mark off the supposedly unacceptable positive policies from the acceptable elimination of defects. And the genes-environment boundary is needed to mark off positive engineering from acceptable positive aims of educational policies. But it is not clear that confidence in the importance of these boundaries is justified.

CHANGING HUMAN NATURE

[69] Positive genetic engineering raises two issues. Could we be justified in trying to change human nature? And, if so, is genetic change an acceptable method? Most of us feel resistance to genetic engineering, and these two questions are often blurred together in our thinking. One aim of the discussion has been to separate the different sources of our resistance. Another has been to try to isolate the justifiable doubts. These have to do with risks of disasters, or with the drawbacks of imposed, centralized decisions. They need not justify total rejection of positive engineering. The risks are good reasons for extreme caution. The other drawbacks are good reasons for decentralized decisions, and for resisting positive genetic engineering in authoritarian societies. But these good reasons are quite separable from any opposition in principle to changing human nature.

[70] The idea of 'human nature' is a vague one, whose boundaries are not easy to draw. And, given our history, the idea that we must preserve all the characteristics that are natural to us is not obvious without argument. Some deep changes in human nature may

only be possible if we do accept positive genetic engineering. It is true that our nature is not determined entirely by our genes, but they do set limits to the sorts of people we can be. And the evolutionary competition to survive has set limits to the sorts of genes we have. Perhaps changes in society will transform our nature. But there is the pessimistic thought that perhaps they will not. Or, if they do, the resulting better people may lose to unreconstructed people in the evolutionary struggle. On either of these pessimistic views, to renounce positive genetic engineering would be to renounce any hope of fundamental improvement in what we are like. And we cannot yet be sure that these pessimistic views are both false.

[71] Given the risks that positive genetic engineering is likely to involve, many people will think that we should reject it, even if that means putting up with human nature as it is. And many others will think that, quite apart from risks and dangers, we ought not to tamper with our nature. I have some sympathy with the first view. The decision involves balancing risks and gains, and perhaps the dangers will outweigh the benefits. We can only tell when the details are clearer than they are now, both about the genetic techniques and about the sort of society that is in existence at the time.

[72] It is less easy to sympathize with opposition to the principle of changing our nature. Preserving the human race as it is will seem an acceptable option to all those who can watch the news on television and feel satisfied with the world. It will appeal to those who can talk to their children about the history of the twentieth century without wishing they could leave some things out.

What Kinds of People Should We Create?

JAMES HUDSON

Journal of Applied Philosophy, Vol. 17, No. 2, 2000

[1] ABSTRACT *The advent of genetic engineering will give people the power to design the members of the succeeding generation, and thus will render practical the issue how such a power ought to be exercised. Here I address that issue in a general way. I point out that the aim of making future people better adapted to the modern social environment is implicitly circular, since the natures of the people themselves will determine the nature of the social environment. I claim that the human property the enhancement of which would do most to enrich experience is intelligence; accordingly increased intelligence should be a primary aim of genetic designers. The tendency to feel pain should be attenuated, as positive motivation is substituted for negative (to some extent). People should be designed so as to be motivated more by reason than by any non-rational drives (though rational motivation may still involve pleasure and pain). The sex drive, having outlived its usefulness, will probably be replaced by some other source of pleasure. As a side effect of these changes in people, the arts and social sciences will be transformed beyond recognition.*

I

[2] The remarkable progress currently being made in molecular biology is providing us with techniques of genetic manipulation applicable to human beings as well as to other species. These techniques, which are constantly being improved, will become practically more important as we discover correlations between genes and gene combinations, on the one hand, and important properties of the human phenotype -- especially psychological properties -- on the other. With these techniques and the knowledge of these correlations, people will be able to engage in *designing* their own descendants; they will be able knowingly to form the natures of the people they are creating, insofar as their natures are genetically determined. This will enable them substantially to alter -- psychologically as well as physically -- the familiar, ordinary sort of human nature, exemplified by (for example) us present-day people.

[3] *Prima facie* this prospective ability can, and therefore should, be used for the good, with those and only those designs that will yield the *best possible people* being employed in the production of new people. This will require those choosing the designs to have a clear conception of what sorts of changes in human nature would constitute *improvements*. They will soon face the basic theoretical value question -- 'What constitutes a life's being *good*, or *one life's being better* than another? -- in a newly practical context. A firm understanding of value in general will become practically important as it has never been before.

[4] *Consequentialists* have all along considered the concept of value to be supremely important, at least in theory. Non-consequentialists accord it a less dominant theoretical role, since they believe that, besides considerations about value, there are "deontological" rules with independent force. But non-consequentialists, too, must depend entirely on the concept of value in answering the question *what sorts of people to produce*, for the usual non-consequentialist principles have little application in that context. The most plausible of such principles have a "person-affecting" form: they specify how *people who already exist* are to be treated (with respect, so that they receive equal distributive shares, according to their deserts, or however). Issues about the creation of new people -- "ontological" moral

issues, we might call them -- lie outside the scope of such rules. Even a moral theorist who, in ethics generally, regarded the production of value as only part of the basis for the distinction between right and wrong action should regard it as the whole story with respect to ontological issues.

[5] This special character of issues concerning future generations -- long-term issues about the rate of savings, procreation and population size, the environment, etc. -- has become widely recognized in recent years (thanks especially to Jan Narveson and Derek Parfit). But even in these ontological cases it has been standard until recently to regard the future possible people as having *innate natures* that were more-or-less given: they could be assumed to be like present people, the only question being how strongly they (given this presumed innate nature) would be affected by various degrees of crowding, resource depletion, pollution, etc. So the moral philosopher dealing with ontological questions tended not to consider *what sorts* of future people there ought to be: it was assumed that they would be of the same nature as us.

[6] What will be radically new in the practical situation we (or our near successors) will soon face, thanks to biotechnology, is that no longer will we be able to treat the future people affected by a proposed action as having fixed inherent natures, including the familiar basic drives, capacities, and other psychological characteristics. On the contrary, these will be subject to design, in a manner and to a degree without precedent; as a result, estimating the impact of a present action on people in the remote future will be subject to a new kind of uncertainty.

[7] Of particular importance is the procreative decision itself. Rather than a simple binary choice *whether* to produce offspring, this is in the process of becoming the more complex choice *what kind* of offspring (if any) to produce. One of the principal ways people will have of affecting the value of future lives will be by giving future beings improved versions of the innate combination of properties that we have inherited from our ancestors. Here I offer some general thoughts on how we (or our near successors) should approach the task of designing and creating future people.

II

[8] Let me start by observing that *homo sapiens* evolved under conditions considerably different from those in which we now live. One naturally expects to find our innate natures rather poorly adapted to present-day conditions, insofar as these differ from those under which our more remote ancestry developed. For example, we are nowadays faced with greater crowding, the need to communicate with many more other people, greater technological complexity and more extraordinary demands for skills and information, unprecedented prosperity and amounts of leisure, and so on. And one might hope that genetic tinkering could directly remedy some of the presumed mismatch between human nature and modern social conditions.

[9] But let us note that these conditions are to a great extent themselves the products of intelligent human design, and of semi-intelligent quasi-design of the "trial and error" sort, which should have made them rather well adapted to human nature. It would be a gross misrepresentation to portray our present social environment -- granting that it is quite different from that of the palaeolithic hunter-gatherer (to go no further back into prehistory) -- as alien to human nature. To the extent that our environment has been *designed* to accommodate our common nature, the mismatch will not be so great.

[10] Still, there must be *some* incongruity between modern living conditions and genotypes adapted mostly to more primitive ages, and genetic tinkering to secure better social adaptation retains some appeal. The chief difficulty about implementing it is that the social conditions to which adaptation is to be made will themselves become quite fluid, thanks to the new biotechnology. Suppose that in the year 2050 there were begun a programme of altering people's genetic endowments so as better to adapt them to the social conditions then prevailing. This very alteration would surely bring about changes in the nature of society -- that is, changes in the very conditions to which one was trying to promote adaptation. The nature of society at any given time is determined by the natures of the people then composing it; different sorts of participants would interact differently, giving rise to different social conditions. And so we may find that people designed in 2050 so as to be well adapted to the society of that time may not be so well adapted to the resulting society of 2075 or 2100. In short, the widespread attempt to produce beings well adapted to their social environment would run into vicious circularity, since the nature of that environment depends in considerable measure on what sorts of beings are produced.

[11] A similar difficulty would confront the attempt to make not so much *happy* individuals as *beneficent* ones -- not so much beings who would themselves thrive in their social environment as ones who would contribute maximally to the well-being of their fellows. Again, to pursue this objective we would need to know the natures of those others, which, however, would depend on the design decisions themselves.

[12] Why not broaden our perspective by considering how to improve *society as a whole* through genetically redesigning the individuals composing it? The answer is that any such programme would be quite unrealistic. First, there seems little prospect that the social sciences (political science being especially relevant) will ever provide the basis for such a "bottom-up" design of a whole society. Furthermore, co-ordinating the designs of the many different new people being brought into existence at any one time would probably be politically impossible. Thus our prospective ability to manipulate individual phenotypes is unlikely to be extended to an ability (by anyone) to manipulate the overall characteristics of *society* by genetic means. The shape of society as a whole is likely to remain largely the unforeseen, unintended consequence of a multitude of individual decisions aimed at local, limited ends.

[13] A special reason why our forecasts for the future of society must remain indefinite is that the current near-homogeneity of *homo sapiens* is unlikely to persist, once the genotypes of one generation have come under the conscious control of the previous one. The potential diversity of future people is enormous; and once the traditional methods of reproduction -- which largely tend toward homogeneity -- are abandoned, and parents are choosing the genotypes of their offspring independently of other parents, we should expect more of this potential diversity to be actualized. The implications for the workings of *society* are hard to fathom; they may well be quite revolutionary.

[14] We should expect future diversity not only because different people, in their role as designers, will have different ideas about how to improve upon pre-existing human nature (and perhaps different means for putting their ideas into effect), but because the projected environment to which adaptation is being attempted need not be the same (in prospect) for all those creatures being brought into existence. It will be possible to design different sorts of people for different niches in society, subject to the aforementioned reservation that it must always be somewhat uncertain just which niches will exist when the designed person reaches maturity. I am certainly neither defending nor predicting anything like Aldous Huxley's "Brave New World," but the general point that there is probably no one model on which all new people ought to be designed seems valid.

III

[15] These considerations incline me to de-emphasize social adaptation as a goal, and to adopt instead the rather atomistic approach of focusing on how the individual's life may be made better, neglecting broader social issues. In what general direction, then, should we wish to see human nature pushed by the prospective new techniques of biotechnological design? What sorts of changes in the genetic endowments of people would count as *improvements* from the individual point of view? As noted above, this is just the basic value question 'What is it that makes a life good, or *better* than some other life? in a newly practical form. The philosophical task is to answer this question in a way that can provide guidance for the users of the prospective biotechnology.

[16] This individualistic approach assumes not only that we can determine ways of making individuals better off while ignoring the intractable issue of social adaptation, but also that improving the representative individual's life is likely to be a good thing overall. This latter assumption appears at its most dubious in the particular case of individual life-extension.

[17] A simple and direct way to make an individual's life better already realized to a notable extent, and apparently destined for much greater realization is to make that life *longer*. Provided the value of each year of life added beyond the normal span is positive, even if it be less than the value of a year of life in his prime, the individual will be better off for the addition. But it may plausibly be suggested that *society* would be better off if the elderly would retire gracefully from the stage, making room for new people whose level of well-being, and whose contribution to the well-being of others, was greater. Indeed, once we are in the swing of making successive improvements in the design of new people, we might welcome the chance to trade in the older models for newer, better ones. Thus life-extension, though a benefit to the individual, might be detrimental to society, lowering the average level of well-being at any one time, and retarding the progressive improvement of the population through biotechnical engineering.

[18] True, the elderly may have valuable experience not attainable in a short span of time. If they can be made energetic enough even in their later years to contribute to society, rather than merely living in retirement, the extension of lifespans may after all bring a collective benefit. *Vita brevis, ars longa*; the artist whose working life has been extended will be able to exercise his mastery to fuller effect. So it may be that society should *welcome* the prospect of longer lifespans after all. But the issue is unclear, in spite of the fact that this development would obviously work to the individual's advantage.

[19] Nevertheless I shall here persist with the basic assumption that improving individual lives will make for a better world viewing *life-extension* as an atypically doubtful case.

IV

[20] Let us then consider direct methods of improving the lives of individual people through genetic redesign. In addressing this topic I shall have to rely upon some axiological assumptions, which I shall try to make appropriately weak and (therefore) plausible.

[21] First, even if we are unwilling to embrace hedonism we must, in my judgment, be inclined to accept some sort of experiential theory of value; at the very least, this will be a large part of the correct account. The difference between a human life and that of, say, an

elephant (to take a lower animal rather like us in many respects, including lifespan) must lie mainly in *the quality of the experience*. Admittedly it is difficult for us to appreciate just what an elephant's experience is like; but there is enough similarity, thanks to our common ancestry of less than 100 million years ago, that we can trust our intuitive empathy with elephants to a considerable extent. We engage in activities similar to many of those that make up an elephant's life; we can readily appreciate the pleasures of eating and drinking, nursing and being nursed, playing in water and wallowing in mud, playing simple games of mock-combat, socializing (i.e., agreeable low-intensity interacting) with others in one's group, engaging in sexual intercourse, battling a predator or fighting over a female (and winning), and so on. By contrast, as John Stuart Mill observed, a lower animal such as an elephant has no experience of many of the episodes and activities of greatest importance to us, those that involve our awareness of our own rational thinking and acting, as well as our arranging thoughts or sensory elements into complex wholes or our apprehending pre-existing complexes of such elements.

[22] These complex experiences, which make our lives more worth living than are those of any lower species, all involve *intelligence*. This is a main prerequisite of the excess value of our lives over those of lower animals, and presumably also of their lives over those of still lower ones. It should thus be a main focus of the biotechnical designer's attempts to improve upon nature in the design of future people. The designer should aim, *ceteris paribus*, at greater powers of perception, memory, and especially ratiocination, and at a closer integration of action with thought (at greater *rationality of action*). One of the most important bio-psychological research programmes for the near future will be to discover more about the biological basis for mental activity, and its genetic prerequisites. We have long had a theoretical interest in discovering whatever correlations there are between genes or gene-combinations and intelligence. We now have a practical spur to discovery the prospective ability to give our descendants combinations of genes that will greatly improve their mental functioning and, thereby, greatly enrich their lives.

[23] Of course, there are other, subsidiary concerns. Intelligence cannot function properly without health, and so genetic predispositions to disease ought certainly to be selected against. Any other opportunities to secure, by genetic engineering, the preconditions for value in life (better than by relying on unimproved nature) should be grasped. And possible trade-offs must always be considered: if a particular way of enhancing intelligence also had some deleterious side-effects, or ran some *risk* of producing such side-effects, it might not be worth using. Nevertheless, enhancing intelligence seems to be the main practical route to improving the lives of future people.

V

[24] Turning to pleasure and pain (both somatic and non-somatic): I shall adopt the traditional view, that pleasure has positive value just in itself, and pain negative. Arguably intelligence itself, aside from its obvious long-run instrumental value, has a more direct sort of value in that it makes available to us *kinds of pleasure* (kinds of pleasurable experiences) that are beyond the scope of lower animals.

[25] Should people accordingly be designed so as to have a great capacity for pleasure (perhaps from many different sources), and very little or no capacity for pain? My answer is a qualified affirmative, but I want to insist on the reasons for the qualification. They are obvious in the case of pain, only a little less so in the case of pleasure.

[26] The obvious rationale and justification for the capacity to feel pain is *instrumental*: though intrinsically bad, pain causes the experiencing organism to behave in ways that produce good results. Pain is a spur to beneficent action, and the good done generally outweighs the intrinsic bad of the pain itself. To design a being unable to feel pain (which I am taking broadly to include all sorts of unpleasant, negative feelings) would be unwise -- unless the motivational function now served by painful feelings were fulfilled in some other way.

[27] Is such a motivational substitution possible? As things now stand, a serious disorder in one's body, threatening the proper functioning of the organism, is usually felt as (somatically) painful, and this spurs the individual to take corrective action. (The non-somatic cases are somewhat similar.) Might there be a spur of a different kind that would be just as effective but less intrinsically bad? Human (and generally animal) motivation includes the carrot as well as the stick; we are motivated both by the appetite for good experiences and by the aversion to bad. Might one thrive with only the positive motivation? If so, it would clearly be advantageous to be so designed.

[28] However, I suspect that pain is and must be the strongest and surest motivator, and that we would be worse off if we were completely without the capacity to feel it, even if our positive motivations were adjusted optimally. If my suspicion is right, pain should not be eliminated from the psyches of our descendants. But I have offered no substantial argument, and I may be underestimating the wonders that future psychobiological designers will be able to perform. In any case, the individual's capacities for pain can very likely be *modified* through genetic engineering so that the average person, with no loss of instrumental value, felt less pain over the course of his life. This would be a boon to future people.

VI

[29] Regarding pleasure: it seems that we would make the individual's life better by increasing his capacity for pleasurable experiences, *ceteris paribus*; but other things will usually not be equal. Though pleasure is unlike pain in having *positive* intrinsic value, it shares with pain a considerable measure of instrumental value, in that it prompts the organism to actions that produce results that are good apart from the pleasure itself. This motivational aspect of pleasure is inherent in its nature. Modifying the details of the pleasure-mechanisms implanted in us through natural selection will entail modifying our motivations, and the behavioural consequences must be taken into account. In this light, what sorts of modifications should we engage in, once we acquire the technical skill to do so?

[30] Though pleasure is an intrinsic good, it seems beneath the dignity of a rational being or at least inefficient as a matter of design to have sources or channels of pleasure that fail to motivate behaviour that is instrumental to some further good. Such a condition would be rather like that of the drug addict, who behaves in a way that provides him with pleasurable sensations but who inspires pity and contempt rather than admiration and envy. So long as we are being motivated by the prospect of pleasure, the particular sources of pleasure for us ought to conduce to actions that have independently good consequences; thus may the designer kill two birds with one stone.

[31] Now, it is clear that genetic engineering should be used to modify our present sources of pleasure, since some of them incline us to actions that have predominantly harmful side-effects. The human taste for fat, salt, and sugar, brought about by evolutionary

pressures under nutritional conditions very different from those that currently prevail in the developed world, nowadays causes actions that are bad for the individual's health. President Bush would probably have been better off if at conception his genes had been modified so that he had a natural *liking* for broccoli and *dislike* for pork rinds (assuming his contrary tastes were genetically "determined"). In his day that was impossible, and so it is even now; but I am supposing that soon such tinkering with innate tastes will become practicable.

[32] Our tastes for food could be improved by genetic engineering; why not an even more powerful taste? *The sex drive*, with the associated capacity for sexual pleasure, is an enormously important part of our motivation. How, if at all, should genetic engineering modify this drive (assuming it will become technically possible to do so)? The extremist answer is that the sex drive should be eliminated from the psyches of future people, and there is a surprisingly strong case to be made for the extreme view.

[33] The main point is that sex is fast losing its instrumental justification as a spur to procreative activity. Sex used to be vitally important to the survival of the species; soon it will be merely distracting and pointless, except for the pleasure it affords.

[34] It has always been common to view sexual activity of the sort labelled "perverse" in this negative light. Homosexuality is the most salient example. Homosexual activity itself seems obviously pointless, apart from the pleasure it affords; and the typical homosexual might reasonably view with dismay the prominence of the desire for such activity in his motivational set. However, there is currently no point in dwelling upon such reflections, since no beneficial change is practically possible. We usually cannot alter deeply ingrained desires except by eliminating them without replacement, and in this case that would deprive the individual of an important source of pleasure.

[35] In the future it will be otherwise: any tendencies to develop homosexual desires assuming these have a genetic basis will be avoidable through genetic engineering; a designer will be able safely to channel his offspring's sex drive heterosexually. But I suggest that the basis for heterosexuality's claim to superiority is so weak that any such channeling must be regarded as an excessively timid half-measure, when we have the technical ability to do something more radical.

[36] What in popular thinking has seemed to raise heterosexual passion to a higher plane is its connection, however indirect, with procreation. In the past, the heterosexual has been able to some extent to rationalize his powerful desire for sexual intercourse by understanding the function of the sex drive in the system of nature. This would tend to reconcile him to his own sexual motivations, though he might still wish that the sex drive were not such a blunt instrument of procreation.

[37] But nowadays that rationalization is losing its basis, and it has become a cultural commonplace that heterosexual sex is coming more and more to resemble homosexual. Most heterosexual intercourse (at least in the advanced countries) is presently conducted with the intention *not* to produce children; at the same time, children can be produced *without* intercourse. This is true now, and seems on its way to becoming much more a matter of course in the near future. As the decoupling of heterosexual intercourse from procreation proceeds apace, the former will become as naked of even apparent meaningfulness as homosexual sex has always been. As regards sex, we will be very nearly in the position of the drug addict; sex will be a quasi-addiction a source of pleasure with no reference outside itself.

[38] When this has happened it will be reasonable to wish to replace the sex drive with a drive toward some activity that actually promotes people's survival and well-being. The activity promoted will, of course, produce immediate pleasure (once the drive is in place), but it must have some important value beyond that. Sexual intercourse will not be such an activity; so once the sex drive comes under the conscious control of rational designers it will be ripe for substantial modification most likely, for replacement by some other drive or source of pleasure.

[39] In defence of sex one might make the Freudian point -- that the sex drive is covertly implicated in an enormous range of human activities besides actual intercourse, and it serves as the source of many more kinds of motivation than the simple urge to copulate. Especially, the sex drive has an important social function, for it causes us to take a certain sort of interest in other people; it prevents us from being wholly self-absorbed. Even homosexual activity may in some cases be an element in friendship, which we may accept as undoubtedly good (for one reason, because it promotes survival, a precondition of any intrinsic good for the individual). And Freud suggests that the sex drive, "sublimated," is an essential factor in creative activity of all kinds, including artistic production.

[40] But equally good remote and secondary effects may well result from whatever drive is engineered to replace sex; and it may be possible to engineer drives directly for some of the supposed side effects of the sex drive, such as artistic creation. The fact is inescapable that the primary object of normal sexual desire is sexual congress, an activity which, if it plays no important role in reproduction, has no value at all apart from the pleasure it affords. Such a powerful drive ought directly to motivate some more useful activity. (I shall not here undertake to canvass the plausible candidates, though I do offer a tentative suggestion below.)

[41] Sex probably also serves to strengthen the family bond -- it may (however imperfectly) bind couples to each other so that they are more effective in raising their offspring; and the family bond will continue to be important for the foreseeable future, since in the short run genetic engineering will not eliminate the need to nurture children. But perhaps some substitute for sex might have a similar binding function; moreover, eventually there may be found a good substitute for the nuclear family in the nurturing process. (Already the advantage of the two-parent family seems to many observers less considerable than in the past.)

[42] In short, once we have the technical skill to replace it with some other drive, sex appears to be destined for elimination. But the discussion offered here must be viewed as merely tentative.

VI

[43] Motivation by pleasure or pain is potentially non-rational, since it produces actions that might not rationally have been done in the absence of the immediate feeling (or prospect thereof). Now, it is doubtful that we should accept *any* non-rational sort of motivation in the designed psyches of future people. Rational motivation, because of its greater accuracy and flexibility, is always to be preferred wherever it is achievable; therefore a designer of people should aim for as nearly *perfect* rationality as is technically achievable. Mechanisms of pleasure and pain can have at most the role of motivating actions that are on the whole desirable apart from the immediate pleasure or pain-avoidance that they produce; and would not a rational agent who knew of these good consequences find adventitious hedonistic/doloristic motivations superfluous?

[44] Granted, *immature* creatures will inevitably lack both knowledge and rationality; *they* will need non-rational motivation in order to function satisfactorily. But this does not justify carrying over such motivation to the mature agent. Even in maturity the agent will have deficiencies in knowledge, which might in theory be compensated for by non-rational motives. But in fact such compensation would be possible only if we could predict just what the ignorance would consist in; and if we could do that, we could remedy it directly. Other sorts of cognitive deficiency (besides ignorance, which is in part inevitable), such as various tendencies to misevaluate evidence in one's possession, or to fail to act rationally on the available evidence ("weakness of will"), should likewise be addressed directly. Non-rational motivation appears to be a second-best expedient to compensate for cognitive or motivational deficiencies; and should genetic designers not take as their goal to *eliminate* such deficiencies, so far as possible, rather than merely compensating for them?

[45] A partial reply to this line of thought is that pleasure and pain have a legitimate employment in *constituting* an agent as rational -- in *reinforcing* or even *producing* rationality in action. To make future people more rational it may be expedient or even necessary to give them, besides greater *powers* of rational thought and action, a greater *tendency to enjoy* the exercise of those powers. The greatest source of pleasure in the improved people of the future may lie in exercising their (greatly enhanced) rational capacities; thus may they be led to surpass us in rationality, actually as well as potentially. And -- to focus for a moment on the topic of the previous section -- if people can thus be made to act rationally in matters of *procreation*, they will have no use for a sex drive.

[46] In fact, it is hard to see why we should include *any* drive in the ideal person's innate motivational set other than a drive to *rational thought and action* in general. The tendency, over the whole span of our existence as a species, has been for reason to gain sway over innate drives in our motivation. It would be desirable, and should prove possible, greatly to accelerate this process through biotechnology.

[47] These remarks may remind the reader of Rousseau's paeans (in *On the Social Contract*) to the collectivist "citizen," or the Bolsheviks' to "New Soviet Man". But let us not overlook the contrast. Rousseau and the Bolsheviks were positing a transformation of human nature without an enabling transformation of the genetic underpinnings of personality. My imagined transformation is much more plausible, because it will have a solid biological basis.

[48] Let me add that the *altruism* that is a marked feature of those older collectivist personality constructs, and that is required by rationality according to both the Kantian and the utilitarian traditions, ought not to be construed as requiring "selfless" behaviour *regardless of the dispositions of those with whom one is dealing*. Agents may never be able in practice to be wholly confident that those with whom they are interacting are also selfless altruists at heart; accordingly their own behaviour will have to be rather guarded, lest they find themselves interacting with a ruthless egoist or psychopath. Rationality requires no more than a *cautious* altruism.

[49] I have not attempted to present here a complete theory of value, and the remarks I have offered on the topic have been disputable. Though not committing myself to axiological hedonism, I have mentioned no sort of intrinsic value besides pleasure and pain-avoidance. In fact, I can conceive neither of a good life wholly lacking in experiences that were "pleasurable" in a broad sense nor of a rival to pleasure/pain for the status of basic, intrinsic value. This hedonism suggests that designers should aim at producing "utility monsters" beings with great capacities for pleasurable experiences together with whatever

abilities and motivations would contribute to their obtaining those experiences in their conjectured environment. This conclusion is uncomfortable, but I see no attractive way to avoid it. It may, after all, be right to design our descendants simply for maximum pleasure, seeing to it that the means by which their pleasures are produced (and pains avoided) *immediately* are also generally beneficial in their remoter consequences.

[50] Perhaps a fuller and better treatment of the basic value-question would throw a different light on the problems of designing new kinds of people. If so, the time is ripe for such a treatment, since, I am predicting, we (or our near successors) will soon be confronted by these problems.

VII

[51] Genetic engineering promises to have revolutionary consequences that no one of a conservative temper can enjoy contemplating.

[52] The conservative's high valuation of the study of *history* has already come under increasing pressure in the last century or so; in the near future the relevance of history to our practical concerns will be further eroded. And (a related point) the value placed by conservatives on the *cultural treasures of the past* will be impossible to maintain. Indeed, as human nature is modified, people may lose the capacity to appreciate traditional art, in any direct way. This will surely be the case if, as I have suggested, the sex drive is altered beyond recognition: what will desexed people make of all the art -- a great proportion of the whole -- that rests on a sexual basis? Furthermore, as I have already mentioned in passing, the *human sciences* in their familiar form -- sociology, political science, psychology itself -- will be largely abandoned as no longer relevant to the new situation (one cannot say "the new human nature," since it is doubtful that any one uniform "nature" will emerge, or that the old biological term 'human' will still be appropriate). History will remain *true*, but with greatly diminished *relevance*. And the social sciences as presently constituted will have to be stripped of their scientific pretensions to general validity, and classified as departments of (largely irrelevant) history. (But it has all along been obvious that the "social sciences" did not possess the generality required to merit the label 'science', since they presupposed an unchanging substrate of human nature.) The content of present-day sociology, etc., tells us only how people have interacted up to now; it will not be valid for the genetically altered people of the future. Among the social sciences, only economics, which seems largely *a priori*, may survive intact.

[53] It would be quite inappropriate to lament these developments; rather, we should gladly embrace the proper conception of social "science," and we should accept the obsolescence of conservatism and of the older cultural values as an inevitable byproduct of progress. The changes I am predicting will be concomitants of the noble effort to produce better people, people capable of more exalted, more rewarding lives. Once we have mastered the biotechnology, we will be drawn irresistibly to making the sorts of enhancements of our inherited human nature that I have roughly sketched, guided by our conception of what would constitute an *improvement* -- of what sort of life would really be *better* than the sort most of us now live. There is every reason to be optimistic that this effort will succeed, and will do enormous good. Its negative aspects ought to be viewed -- even by those with an emotional attachment to older ways of thought as comparatively trivial.

[54] Philosophers have an opportunity to contribute to the coming effort to make better people, through their work in axiology, the practical importance of which will become much greater than ever before. In grasping the great opportunity to do good that biotechnology is

about to present to us, we must look to philosophy to clarify the conception of good that must guide our handling of the technical issues. In its combination of theoretical interest and practical urgency, this is the most important philosophical task now before us.

Genetically Based Handicap

ALAN HOLLAND

Journal of Applied Philosophy, Vol. 15, No. 2, 1998

ABSTRACT

[0] It is unclear what we should make of a policy designed to 'eradicate' genetically based handicap, and in particular whether it constitutes discrimination against people with a genetic handicap. After brief reference to the legal position, four arguments are examined which purport to justify differential treatment of handicapped lives either before conception or before birth: the argument from genetic 'error', the argument from parental responsibility, the argument from social consequences and the argument from impersonal harm. Weaknesses are detected in each of these arguments, and the conclusion is drawn that, although differential treatment of handicapped lives is sometimes justified, there are some circumstances in which it does amount to discrimination.

Introduction

[1] The problem addressed in this paper is whether we should do what we can to eradicate genetically based handicap, or whether such a policy would constitute discrimination against those with a genetically based handicap. The problem arises because, for the foreseeable future, the only realistic way of eradicating such handicaps is to eliminate those with the handicap, or prevent their coming into existence. And even for the unforeseeable future, these are always likely to be the simplest (and most 'cost-effective') ways of eradicating such handicaps. For this reason we must be aware that the 'crusade' against this class of handicap raises a special set of issues and gives off a wholly different cultural signal from that associated with campaigns against diseases such as AIDS, tuberculosis or cholera. The problem comes about for genetically based handicap in particular because in many cases the condition may be detected before, or immediately after, conception in other words, at a stage when intervention might seem to raise the fewest moral scruples.

[2] Genetically based handicaps include both those which are largely hereditary, such as Duchenne muscular dystrophy, Lesch-Nyhan syndrome and achondroplasia, and those which may be largely congenital, such as Down's syndrome. Some, but not all, of these conditions involve some 'clouding' of the mental faculties, a circumstance which, paradoxically, may render them more, rather than less supportable. By no means all of them preclude an adult life and the possibility of procreation; Huntington's chorea, phenylketonuria and achondroplasia, for example, are all compatible with something approaching normal life expectancy, and the possibility of parenthood.

I. The legal position

[3] So far as UK law is concerned, there is a provision in the Human Fertilisation and Embryology Act (1990) which bears on our problem. Section 37 of the Act permits abortion up to full term on three grounds:

- i) risk of grave injury to the physical or mental health of the mother

- ii) risk to the life of the mother, or
- iii) substantial risk that a child, if born, would suffer from such physical or mental abnormalities so as to be seriously handicapped.

[4] The first two grounds permit termination of the lives of both handicapped and normal foetuses under the conditions specified, and in this sense may be said to treat normal and handicapped foetuses 'equally'. But it is clear that in these cases there needs to be a strong and overriding consideration to justify abortion. This is not so in the third case. The third ground permits termination of the life of a seriously handicapped foetus whether or not there is any risk to the mother. Therefore, it is clear that the law permits differential treatment of the seriously handicapped, since it makes no such provision where the foetus is normal. We must assume that the law is based on a belief that it is morally permissible, in certain circumstances, to prevent the continuation of a life which is seriously handicapped, because it is seriously handicapped. Whether this belief is justified, and if so on what grounds, is the task of this paper to establish.

[5] One question which the provisions of the 1990 Act obviously raise is: who are the seriously handicapped? It is worth noting, to begin with, an ambiguity in the notion of a 'serious' handicap, which turns on the reference class that is chosen for the purpose of comparison. In the context of handicaps generally, you might judge a condition such as congenital deafness, Down's or achondroplasia to be not so serious 'as handicaps go'; whereas, if you take 'normal life' as your reference point, you might think them pretty serious, in the sense that anyone with such a condition is likely to lead a very different kind of life from the one they would lead if they were 'normal'. The ambiguity further reflects an underlying issue of some sensitivity (to which the Act appears insensitive) namely, whether it is in some way prejudicial to refer, not to the handicap as 'serious', but to the whole person as 'seriously handicapped', and further, whether it is indeed prejudicial to use the language of 'handicap' in the first place. Both the Clothier Committee and the Nuffield Council on Bioethics broach the question of what makes a handicap 'serious' with some diffidence, the latter partly explaining its stance by reference to cultural variations. Neither risks the clarity of judge Templeman's view that it should involve the tolerable certainty of an 'awful' life. The 1990 Act not unreasonably leaves the interpretation of the concept of 'serious handicap' to medical practice and, ultimately, to the courts; but there do not yet appear to have been any court cases in the UK subsequent to the Act which put this concept to the test. Nevertheless, there are two cases dating from before the Act, dealing with actions taken immediately after birth, from which one might extrapolate. The first, re B 1981, which is the context in which judge Templeman's comments occur, concerned the surgical treatment of a newborn baby with Down's, to remove an intestinal blockage. The courts ruled, contrary to the wishes of the parents, that the child was entitled to treatment. The second, re J 1990, concerned the giving of artificial ventilation to a baby with severe and irreparable brain damage, epileptic, potentially quadriplegic, blind and deaf, and with a limited life expectancy. In this case the courts authorised the withholding of treatment designed to prevent death from natural causes. Clearly, in both cases, if a 'normal' child had required the treatment in question, he or she would have been given it. So far as the UK courts are concerned, then, it seems to be generally true that in postnatal cases, a condition akin to Lesch-Nyhan is thought to justify differential treatment, whereas a condition such as Down's is not. A celebrated anomaly was the case of Dr Arthur, whom a jury acquitted of attempted murder after he had prescribed 'nursing care only' for a baby with Down's.

[6] A somewhat analogous situation is discernible in the United States, where the decision in Re Infant Doe (1982) to withhold relatively routine surgery from a baby with Down's

appears out of line with previous opinion such as that handed down in *Re McNulty* (1980) which said: 'If there is any life-saving treatment available, it must be given regardless of the quality of life that will result'. Subsequently, the 1984 Child Abuse Amendments of the 1974 Federal Child Abuse Prevention and Treatment Act have the effect of classing as 'medical neglect' any failure to correct a straightforward complication arising in conjunction with Down's syndrome.

[7] But, whatever the position of the courts regarding postnatal cases, one has the strong impression that society in general, and the medical profession in particular would, and perhaps does, take conditions such as Down's or achondroplasia as grounds for termination during pregnancy. Otherwise, it is hard to explain the excitement generated by the discovery of new ways of testing for these conditions during early pregnancy. The Nuffield Council on Bioethics acknowledges that prenatal screening for Down's is widespread in the UK, and does express some concerns about the practice, but raises no principled objections to it. If this is true, then it might seem to give rise to an anomaly. Either birth is viewed as a moral threshold, marking a difference of kind rather than a difference of degree between the moral judgements which can be made before and those which can be made after, or it is not. If it is, then we need some explanation of why it is not permissible to terminate any life before birth, rather than just handicapped ones. If it is not, then if differential treatment of Down's is unjustified after birth, it seems to follow that it is unjustified before birth also.

[8] However, it maybe that the anomaly is only apparent. If we focus again on the treatment of Down's before and after birth, it may not be so difficult to understand how a justification for termination could be found to exist before but not after. For example, life in general might be thought to be of less moral significance before, than after, birth. Hence, countervailing considerations that are thought sufficiently serious to override the claim to life before birth, might not do so afterwards. These considerations could be of two kinds: i) consideration for the handicapped life itself- for the sadness it is likely to entail; ii) consideration for the sadness and burden it might bring to others. Given the absence of such countervailing considerations in the case of the 'normal' life, we would then have a full explanation of how differential treatment of the normal and the handicapped life could be permissible before but not after birth. What still remains unclear, however, is whether the claim that differential treatment is justified before birth is compatible with the claim that people with Down's enjoy full equality of treatment thereafter. For example, the very fact that postnatal cases involving Down's can be brought to court at all would seem to belie the latter claim. There is also the more vague but palpable consideration that differential treatment of Down's before birth sends out a certain kind of signal to society at large.

II Arguments for the claim that differential treatment of genetically based handicap is justified

[9] The following discussion identifies four lines of argument that might be used to justify the differential treatment of genetically based handicap. These are: the argument from genetic 'error'; the argument from parental responsibility; the argument from social consequences; and the argument from impersonal harm.

II.1 The argument from genetic 'error'

[10] The first line of argument expresses what can be called a 'metaphysical' perspective, since it flows from the insinuation of certain basic categories into discourse about genes. Its outcome is the notion of genetic handicap as 'mistake' or 'defect'. The source of this

construal, or as I believe misconstrual, of genetic handicap, lies deep in the 'modern' understanding of evolutionary theory, and goes back at least to Schroedinger's notion of 'code-script' and, before that, to the work of Weismann. According to this way of understanding genes, we are invited to construe them 'literally' as instructions, or as vehicles for the transmission of information. The notion very obviously has official blessing. Thus, without ceremony, we have the opening paragraph of the 'Clothier Report' a report of the UK government Committee on the Ethics of Gene Therapy proclaiming: 'Genes are the essence of life: they carry the coded messages that are stored in every living cell, telling it how to function and multiply and when to do so'. The fateful result is that once the idea of genes as instructions, or carriers of coded messages, is in place, we have a ready use for the idea of 'faulty' instructions and 'mistaken' information, and hence for the notion of handicap as 'mistake' or 'defect'. This way of regarding handicap in turn can seem to make differential treatment legitimate. For the appropriate response to something which is perceived as defective is to discard it, send it back, or demand a replacement.

[11] But however useful this information theory model may be for giving genetic processes some semblance of intelligibility, at least in the present cultural context, it should not be assumed to be an adequate guide to the underlying realities. Perhaps it is no more than a debating point to observe that the 'defective' genes in question are somewhat good at the job of surviving; otherwise the condition which they help to engender would not be a hereditary one. Another 'debating' point might be to observe that the attitude of some of the leading protagonists of this point of view is distinctly fickle. For on the one hand they make free with the language of 'mistakes' and 'errors', yet at the same time, and a trifle sheepishly one feels, they acknowledge that major 'errors' of this kind are the very stuff of evolution. 'On rare occasions, such errors are harbingers of new evolutionary directions' writes S. J. Gould, while Richard Dawkins advises us to 'banish' from our minds 'all pejorative associations' of the word 'error', admitting that 'it is possible for an error to result in an improvement'. Thus, one generation's 'error' turns out to be the 'norm' for subsequent generations.

[12] But the authors cited fail to take the one further step of admitting that it is not simply the pejorative connotations of 'mistake', 'error' and 'defect' which are unwarranted, but the very concepts themselves. For the major objection to such ways of speaking must be that to speak of an organism as defective is to presuppose that there is a pattern or form with which it should comply, and that the organism is not as it should be. But to assume this is to reintroduce a teleological notion into the heart of the biological process, where the (purified) Darwinian theory of natural selection dictates that it has no place. Thus, my claim is that the construing of genetic handicap as 'mistake' or 'defect' is nothing less than a betrayal of the theory of natural selection. Furthermore, the construing of genes as 'instructions' is implicated in this betrayal. For to suppose that genes are 'instructions' or 'units of information' is to suppose that they are ways of ensuring that organisms are as they should be and that we can find out how they should be by reading the instructions. Nor is it only the language of 'instructions' and 'information' which carries the teleological implication, but also the accompanying and apparently innocent notions of 'copy', 'replica' and 'reproduction' (or 'self-replication'). For such notions imply that iteration the production of a 'copy' of the 'original' is the inherent purpose of procreation. By implying that iteration is the purpose of procreation, rather than the usually effective means by which natural selection secures it, such concepts ground the corresponding notions of 'error' and 'mistake' in cases when iteration fails. The reason why iteration is such a pervasive phenomenon, like the reason why organisms cluster around species norms rather than being strung out like beads on a string, is itself presumably to be explained by reference to the mechanism of natural selection. But the fact that iteration is the norm should be recognised for what it is, a de facto contingency resulting from the processes of natural selection rather than an

inherent feature of such processes. Regarding the notion of handicap as error, then, I conclude that no differential treatment of handicap can or should be grounded on such a baseless metaphysical importation.

II.2 The argument from parental responsibility

[13] A second consideration favouring the differential treatment of genetically based handicap might be based on the argument from parental responsibility, as outlined by Steinbock and McClamrock. They claim that a decision to create 'a child who is likely to have a life marked by pain and severe limitations' is not the act of a loving parent. Moreover, they claim not simply that it is wrong to have children under these circumstances, but that such children would be 'victims of their parents' decision to procreate': such an act would be unfair on the child. Thus, these authors argue that loving, concerned parents will want their children to have lives worth living, and that a principle of 'parental responsibility' means that it is unfair to children to bring them into the world 'with the deck stacked against them'.

[14] So far as their argument concerns a life that one has reason to believe would be judged by the person whose life it is to be 'not worth living', then they may have a basis for their view. However, Steinbock and McClamrock clearly mean it to include also cases where the person concerned would rather be alive than not. Now, they do indeed mount a successful argument to show that, in general, being satisfied with the situation one is in does not itself show that one has not been treated unfairly by being put in that situation. However, they do not show that the argument has application to the special case of being glad one is alive, which, unlike the examples they discuss, is distinguished by the fact that one would not otherwise have existed. Nor do they have an independent argument to show that a handicapped person would be treated unfairly by their parents' decision to procreate, except for the claim that it is unfair to bring a child into the world with 'the deck stacked against them'. But this claim is unconvincing. In the case of a child who comes into the world with the deck stacked against them, one is inclined to think that what is unfair is the fact that the deck is stacked against them, rather than the decision to procreate. Moreover the authors' principle of parental responsibility, which is said to imply that it is better not to have children unless they have a decent chance of a good life, is morally suspect. For it seems to imply that the very poor should not procreate and that those unfortunate enough to have been enslaved should choose to remain childless. If people in such a situation should decide not to procreate, they would no doubt have good reason, and their decision would be all too easy to understand; but this is not sufficient to licence the inference that in this situation they ought not to procreate.

[15] It may be felt that these responses to the appeal to parental responsibility do not do justice to the fact that the disadvantage faced by a handicapped person is intimately related to his or her handicap, and therefore differs significantly from the disadvantage faced by the person born into poverty or slavery. Although, as we shall see (section II.4), there are conceptual difficulties in supposing that a person can be unfairly treated or harmed simply by virtue of being born with a handicap, it is undeniable that being born with a handicap renders that person inherently vulnerable to unfair treatment and to harm. But there is a distinction, I suggest, between being harmed by a handicap and being harmed because of a handicap. Sometimes a handicapped person's needs, being greater, are less easy to meet; they may be less able to 'fend for themselves'; and sometimes, in some cases, their condition will attract cruelty and torment. Even so, I suggest, we should stand fast by the verdict already reached. The decision to bring a handicapped life into being should not be condemned as irresponsible simply because society may find it more difficult to adapt to.

Whether a different verdict may be justified on the basis of concern for the handicapped person themselves will be discussed in section II.4.

II.3 The argument from social consequences

[16] A third consideration thought to favour the differential treatment of genetically based handicap concerns the social implications of such handicap. These social implications might be divided into two kinds: (a) the implications for society at large; (b) the implications for the parents and the immediate family.

[17] Concerning the first kind of implication, the argument might be put that a handicapped life is likely to place a greater burden on society than a 'normal' life. This is not just a hypothetical point. In a context of limited health resources, calculations actually are being made concerning the relative costs of finding and terminating a Down's foetus, as against the outlay on care for Down's children. Against this argument it must be observed in the first place that the connection between burden and handicap is a purely contingent one. There are a number of genetically based handicaps, such as phenylketonuria and achondroplasia, which are unlikely to involve any extra burden on society whatever. In the second place, it can be argued, quite strongly, that this is a burden which society, in the name of both justice and humanity, should be encouraged to undertake. It would be morally objectionable, if not paradoxical, to discriminate against a life on the grounds that its needs were greater. It would also set a dubious precedent. Moreover, it totally ignores the many positive qualities which handicapped people can contribute, such as their capacity to touch and inspire more ordinary lives.

[18] Concerning the second kind of implication, the argument might be put that even a moderately handicapped life imposes additional burdens upon the parents and upon the immediate family. Early death, prolonged dependency or continual medical treatment is often involved. In varying degrees, proper regard for a handicapped person may require physical, financial and emotional reserves of an unusual kind. The legitimate aspirations of the parents may need to be foregone; the legitimate claims to attention of other members of the family may be neglected. In the light of these considerations it seems not unreasonable to claim as a general principle that differential treatment of a handicapped life is justified where it might result in significant harm to the parents and/or the immediate family. Differential treatment on the basis of anything less, however, would be discriminatory. Indeed, Jeffrey Botkin mounts the interesting argument that there is a limit on the parents' 'right to know', and that even to divulge information about the foetus to the parents when no significant harm is at stake is to violate the fetal right to privacy and confidentiality. However that may be, the chief point to make here is that what is being conceded is not the justification of differential treatment of the handicapped as such, but differential treatment of any life which significantly threatens the well-being of the prospective parents and their family. In some circumstances, any child at all, handicapped or not, will pose such a threat; in other circumstances, all it takes is for a perfectly 'normal' child to be just one too many. In effect, the argument from social consequences, like the argument from parental responsibility, provides no ground for the differential treatment of a handicapped life as such.

II.4 The argument from impersonal harm

[19] By far the most fundamental consideration pointing to the idea that differential treatment of genetically based handicap is justified invokes the notion of 'impersonal harm'.

A good place to begin our discussion of this idea is with Jonathan Glover's 'axiom' (more strictly 'a view which most people would take as axiomatic') that 'it is better, where possible, to bring into the world a child without handicap'. The 'axiom' is presented as the plausible conclusion of the following piece of reasoning: 'Consider the theoretical possibility of deliberately causing a child to be born handicapped. This would surely be a monstrous thing to do. And we think this because we do not believe it is just as good to be born handicapped'.

[20] Two interpretations of Glover's 'axiom' are possible:

- A. It is better for a child to be born normal rather than handicapped.
- B. It is better for a normal child to be born than a handicapped one.

[21] Let A be conceded. Yet even though a person may subscribe to B as well as to A, B in no way follows from A. In A, there is a single assignable individual of whom it is claimed that it is better that they be born normal. In B, there is no such single assignable individual. And returning to the 'monstrous thing', it may indeed be a monstrous thing to induce a handicap in an otherwise normal foetus, or deliberately to create a handicapped foetus by mixing the appropriate gametes in a test-tube; but it is far from clear that it is monstrous simply to procreate or to bring a handicapped child into the world, knowing or having reason to believe that it will be handicapped. Certainly, there is one kind of case where a decision to abort, or not to conceive, can readily be defended, namely, where the anticipated handicap is so severe as to warrant the judgement that if the person were born or conceived, they would wish that they had not been. In this case, if it is better for a normal child to be born than a handicapped one, it is simply because it is better for the handicapped child not to be born. (Even here, though, whether it would be 'monstrous' to go ahead with such a birth is quite another matter.) At first, and by way of objection, one might think: 'to allow the normal child to live, in circumstances where one would not allow the handicapped child to live, is still, surely, discriminatory; one is not allowing an equal right to life'. However, if, as is being hypothesised, it is actually in the interests of the handicapped foetus not to be born or not to have been conceived, then curtailing or forestalling the life of such a foetus, where one would not curtail or forestall the life of a normal foetus, ceases to be discriminatory because one would be considering the interests of both equally; it just happens to be in the interests of the one to live, and of the other to die, or not to have been conceived.

[22] The real difficulty arises over a handicap which is not so severe. Let us concede, what Glover claims, that it is not just as good to be born handicapped, and that this is a condition one would rather not have. In that event there is a choice between bringing into or continuing the existence of someone with a condition they would rather not have, and bringing into or continuing the existence of someone with no such condition. But even if the former person would have a condition they would rather not have, they would still let us suppose rather have a life than not have one. In this case it could not be argued to be in their interests not to have, or not to have had, a life. If, therefore, one were to give priority to the latter person over the former, or discontinue the former life under conditions in which one would not discontinue the latter, the conclusion seems inevitable that one would be discriminating against the former life.

[23] Glover's response is to invoke the notion of impersonal harm. Impersonal harm is explained as harm which is done without there being any identifiable person to whom it is done. This may be understood, I suggest, in a weak or a strong sense. In the weak sense, although there may be no currently identifiable recipient of the harm, either there will be in the future and in this sense it is possible to harm future generations or there was in the

past, so that we maybe said to harm someone if we blacken their reputation after they are dead. In the strong sense there neither was, is, nor will be any identifiable recipient of the harm. It is the concept of impersonal harm in this strong sense which Glover is out to defend. The suggestion is that differential treatment of the normal and the handicapped life is justified, because in failing to give priority to the normal life we should be doing an impersonal harm, even though there is no identifiable recipient of the harm. The reason why there is no identifiable recipient of the harm is that the only plausible recipient, the handicapped person, would rather be alive than not, and can scarcely be said to be harmed by the act which brings them into existence. In support of the suggestion that we might nonetheless be perpetrating impersonal harm Glover cites a hypothetical case of surrogacy in which the surrogate mother is to contribute the egg, and there is a choice to be made between a smoking and a non-smoking volunteer. If we chose the smoker, it is natural to think that harm would be done; but not to its only plausible recipient, the child of the smoker, because if the smoking mother had not been chosen, he or she would not have existed. Therefore the harm must be impersonal.

[24] Now the idea of a harm which is done, without there being any identifiable person to whom it is done, is already quite curious; but for the moment we shall let that pass. The chief flaw in this particular attempt to justify the idea is that the surrogacy case can be explained without invoking the notion of impersonal harm. There is someone who is harmed the child of the smoking mother. But he or she is not harmed (directly) by the person choosing the smoking mother, but by the mother herself. The fact that it is not the person choosing the surrogate mother who directly inflicts the harm, however, does not render that person's decision excusable. For it can be presumed to be unjustifiable, other things being equal, to bring about a situation in which one has reason to expect that a person will be harmed, if it is equally possible to bring about a situation in which this is not so. So we can say that the decision is wrong, and wrong because someone is harmed, although not because the person making the decision harms the child directly.

[25] But it might be suggested, on the basis of arguments originally propounded by Derek Parfit which centre around the so-called 'Non-Identity Problem', that if we do not countenance the notion of impersonal harm then we shall be unable to explain, for example, (a) why a very early pregnancy is undesirable, (b) how it is possible to do harm with respect to future generations say, by leaving behind radio-active waste, or (c) how it is possible to describe as harmful a chemical affecting the mechanism of conception so that it 'favours the conception of children with a moderately severe handicap'.

[26] In case (a) the problem arises because we cannot say that the early pregnancy harms the child of that pregnancy if he or she is at least glad to be alive and would not otherwise have existed. However, we do not need the notion of impersonal harm in order to explain why a very early pregnancy might be judged undesirable; we need only a non-relational notion of good and bad which, I argue, is different. The argument, briefly, is that ascriptions of good and bad need only to presuppose a point of view from which a thing is judged good or bad. The thing judged good or bad does not also need to be good or bad for someone or some thing. If we judge it bad for someone to suffer needlessly, we do not just mean that it is bad for the sufferer; we mean also that it is a bad thing that they should suffer. Harm, on the other hand, like suffering, does need to befall someone or some thing; it cannot be unattributed (I maintain). Hence we can say that a later pregnancy is better a better state of affairs, and not better for anyone in particular than an earlier pregnancy, because the child of an earlier pregnancy is more at risk of being (personally) harmed, as a result say- of inexperience or neglect; and because, where there is a choice, it is better to create a situation in which there is less risk of (personal) harm.

[27] In case (b) the problem arises if we suppose that had we not left behind this radio-active waste a different train of events would have ensued, and that people who come to be conceived and are affected by the waste would not otherwise have been conceived. Since such people would not have existed but for our actions, they cannot be said to have been harmed by those actions, since they are not worse off than they otherwise would have been. Yet surely we have done harm in leaving behind the waste impersonal harm. This conclusion does not follow. For although it is true that the people concerned would not have existed if we had not left behind the radio-active waste, this is a contingent fact only, from which it follows that they could have existed. That is to say, they are logically identifiable as individuals independently of the fact that they will be harmed. Since these very individuals could therefore have existed unharmed, then in leaving behind the nuclear waste we have in fact harmed them. This conclusion is compatible with the fact that they are not worse off than they otherwise would have been. Confirmation that someone can be harmed without being worse off than they otherwise would have been comes from the fact that a person who is punched to the ground is certainly harmed, even if it happens to be true that if they had not been punched to the ground they would have been run over by a bus. (I am assuming that they were not punched to the ground in order to prevent their being run over by a bus. In that event, we should say they were hurt, rather than harmed.) The case of harm with respect to future generations is not, therefore, an example of impersonal harm, and does not serve to legitimise the concept of impersonal harm.

[28] Case (c), a 'thought experiment', involves a factory which produces a chemical affecting the mechanism of conception so that it 'favours the conception of children with a moderately severe handicap'. The case is underdescribed (by Glover): quite how or when the mechanism is supposed to operate is not made clear. But if the mechanism affects the 24 hour process by which the zygote is formed (syngamy), or if it affects the ensuing zygote, then, contrary to what Glover says, it would be the child whom the zygote becomes who is harmed, and we have a straightforward case of personal harm. If on the other hand the chemical affects the gametes, the process (mitosis) by which they are formed, which gametes fertilise or are fertilised, or when, then it affects the reproductive processes of the parents, and it is the parents who are harmed. Again, this is another quite straightforward case of personal harm, and in no way legitimates the concept of impersonal harm.

[29] Does this mean, then, that in the case of bringing a handicapped life into existence, a personal rather than impersonal harm is inflicted? It does not. For in the first case, the 14 year old mother lays her earlier child open let us suppose to being harmed in ways that are less likely to afflict her later child, perhaps through being neglected or being put at risk. These are the reasons for which we judge that it would be better if she waited before starting a family. But, save for the considerations mentioned in section II.3, that is not what we are entitled to assume about the handicapped life. There can be no basis for saying that a handicapped person is harmed by their handicap. They may be harmed because of their handicap, e.g. because of the way people respond to their handicap, but that is a different matter. In the second case, no analogous distinction can be drawn, regarding the handicapped life, between what would have happened and what could have happened. One cannot in the same way speak of the possibility of the very same individual coming into existence without the handicap, since a genetically based handicap is constitutive of the individual. Nor, as in the third case, can we construe the conception of a moderately handicapped life as harmful either to the offspring or to the parents. We might well regard both offspring and parents as unfortunate; but that, once again, is a different matter. So, if the concept of impersonal harm is not available, and if there is no identifiable individual who can be spoken of as having been harmed when a moderately handicapped individual is brought into existence, we are forced to the conclusion that there are no reasons of an

intrinsic kind against bringing a moderately handicapped individual into existence, or against continuing such an existence.

[30] In fact, the same conclusion is suggested by an entirely different set of considerations. We can distinguish four kinds of case in which the question of bringing a handicapped life into existence might arise. First is the 'no alternative' case. The genetic history of a couple might be such that they have every reason to believe that any child of their own would be born with a moderate handicap. Second is the case where although a couple is at risk of conceiving or engendering a handicapped life, they can quite easily ensure that a normal rather than a handicapped life is born, for example by availing themselves of the screening process; or they might go ahead with the handicapped life and have a normal baby subsequently. Third might be a case where, for whatever reason, the couple has only one choice, but is able to choose between a normal and a handicapped life. Finally there is the possible case where a couple who already have ('normal') children decide to have just one more child, and then discover that the newly conceived life will be or is handicapped.

[31] Concerning the first case, Glover interestingly suggests that no impersonal harm would be done. The reason he gives is that in this case 'there is no alternative of having another child without the handicap'. Perhaps, too, it would seem a harsh judgement to make of a couple who want a child of their own, but can only hope for an offspring who is moderately handicapped, that they should not procreate. It would seem particularly harsh, and indeed self-serving from the point of view of this argument, in the case of a couple, one or both of whom has the handicap in question. For such a judgement would precisely seek to take away one of the joys of a life whose very joyfulness is in question. Nor is it satisfactory, as a way of retaining the hypothesis that harm would be done, while acknowledging the harshness of a judgement against procreation, to argue that harm would be done, but that given the parents' desire for a child of their own, this justifies their going ahead. The reason why it is not satisfactory is that it is not an argument that would work for the parents themselves. It is unlikely that such parents, if they were sensitive, would be satisfied with the notion that what justifies their going ahead is simply the strength of their own desire, while at the same time they have to admit that they would otherwise be perpetrating some harm. For this would be to suppose, implausibly, that their desire was unrelated to the consequences of its satisfaction. It seems more plausible to suppose that that they would want the reassurance of feeling that they were doing something worthwhile, or at least not harmful, over and above satisfying their own desire, before they would feel comfortable about going ahead with procreation.

[32] But whether one accepts Glover's suggestion for the sake of argument only, or whether one judges independently that no harm is done in the first of these cases, it is in fact possible to argue, from this admission, that no harm is done by bringing a handicapped life into the world in any of the other three cases either. For, consider the relation between case one and case four. In both cases a handicapped life is brought into the world; in neither case is a 'normal' life being displaced. There seems to be no other relevant consideration to warrant a difference of judgement in the two cases. So, if no impersonal harm is done in case one, then no impersonal harm is done in case four either. But if case four does not involve impersonal harm, then neither does either case two or case three. It is true that in both of cases two and three, if a handicapped life is brought into existence, this can be said to be 'at the expense of' a normal child: these are both cases where there is an 'alternative of having another child without the handicap', and therefore cases which Glover would diagnose as involving 'impersonal harm'. But there is no foundation for drawing any distinction here. For suppose, in case two, that the parents go ahead with the handicapped life and proceed to have a normal child subsequently. We might also suppose that if they had substituted a normal life for the handicapped one, they would have decided not to have

a subsequent child. In that event, case two differs in no significant particular from case four: there is simply one extra (handicapped) life than there would have been. Furthermore, case three differs from case two only in that no subsequent (normal) life is born. But this is merely to say that a 'normal' child will remain unconceived or born that might have been conceived or born. But this cannot constitute the perpetrating of any kind of harm, personal or impersonal. For if causing a normal child to remain unconceived were to count as the doing of an impersonal harm, then all contraceptive practice could be condemned for this reason. Although such a condemnation is not unknown, it could not be reconciled with the stance of anyone who was prepared to contemplate the termination of a handicapped life. And it seems clear that no personal harm can be done to a person who remains unconceived. If this argument is sound, the conclusion seems unavoidable that there is no argument against bringing a moderately handicapped life into existence, or continuing the existence of such a life, which rests on consideration for the handicapped life itself, or on some notion of 'impersonal' harm. Indeed, it seems to be emerging that this latter notion is little more than a label for an intuition which turns out to be unjustified.

[33] If this is true, does it then follow that any difference of treatment accorded to such a handicapped life on these grounds must be judged discriminatory, and therefore wrong? The answer, I suggest, depends on whether it is a matter of continuing the life, which is in question, or a matter of bringing it into existence. (It should be noted that if the answer does indeed lie here, then for practical purposes it will make a big difference when we judge a human life to begin whether at conception or at some later stage of fetal development. However, for the purpose of this argument we do not need to consider that question further here.) Prior to the beginning of a human life (by which I mean the life of an individual human being), whenever that is taken to be, differential treatment may be applied in the matter of screening and/or counselling. Preventing, failing to assist with, deciding or counselling against, the existence of a handicapped life because it is (to be) handicapped, where this would not apply in the case of a 'normal' life, will constitute differential treatment. But it would not be discriminatory because no interest would be adversely affected. No personal harm can be done to a handicapped life, any more than it can to a normal life, before it has begun, if it never will begin; and the notion of impersonal harm has been rejected. Nor is it clear how there can be a right to be brought into existence. On the other hand, the burden of our argument has been to suggest that it is not obligatory either, since a failure to screen or counsel against a moderately handicapped life does no intrinsic harm. Hence the favouring of a 'normal' over a handicapped life in these circumstances, simply because it is normal rather than handicapped, must be judged neither obligatory nor discriminatory.

[34] However, the case is different once the life is judged to have begun. Now there is a life to be adversely affected. In the absence of any sound reason for thinking that some harm is thereby prevented, discontinuing, or failing to assist the continuance of, a moderately handicapped life in circumstances in which one would not discontinue, or fail to assist the continuance of, a 'normal' life, just because it is handicapped, must be judged discriminatory, and therefore wrong, other things being equal. From the fact that it would be wrong, other things being equal, it does not follow that it would be wrong. For reasons mentioned in section II.3, other things may well not be equal: the social strains around even a moderately handicapped life can be considerable. But it is important to see that the legitimate reasons for differential treatment can only be these social ones, and that there are no legitimate reasons of any other kind; and to see, further, that if a supportive social fabric can be arranged, then there is no non-discriminatory reason whatever for discontinuing, or failing to assist the continuance of, a moderately handicapped life.

Active and Passive Euthanasia

JAMES RACHELS

The New England Journal of Medicine, Vol. 292, No. 2 (January 9, 1975), pp. 78-80.

[1] The distinction between active and passive euthanasia is thought to be crucial for medical ethics. The idea is that it is permissible, at least in some cases, to withhold treatment and allow a patient to die, but it is never permissible to take any direct action designed to kill the patient. This doctrine seems to be accepted by most doctors, and it was endorsed in a statement adopted by the House of Delegates of the American Medical Association on December 4, 1973:

The intentional termination of the life of one human being by another -- mercy killing -- is contrary to that for which the medical profession stands and is contrary to the policy of the American Medical Association.

The cessation of the employment of extraordinary means to prolong the life of the body when there is irrefutable evidence that biological death is imminent is the decision of the patient and/or his immediate family. The advice and judgment of the physician should be freely available to the patient and/or his immediate family.

[2] However, a strong case can be made against this doctrine. In what follows I will set out some of the relevant arguments, and urge doctors to reconsider their views on this matter.

[3] To begin with a familiar type of situation, a patient who is dying of incurable cancer of the throat is in terrible pain, which can no longer be satisfactorily alleviated. He is certain to die within a few days, even if present treatment is continued, but he does not want to go on living for those days since the pain is unbearable. So he asks the doctor for an end to it, and his family joins in the request.

[4] Suppose the doctor agrees to withhold treatment, as the conventional doctrine says he may. The justification for his doing so is that the patient is in terrible agony, and since he is going to die anyway, it would be wrong to prolong his suffering needlessly. But now notice this. If one simply withholds treatment, it may take the patient longer to die, and so he may suffer more than he would if more direct action were taken and a lethal injection given. This fact provides strong reason for thinking that, once the initial decision not to prolong his agony has been made, active euthanasia is actually preferable to passive euthanasia, rather than the reverse. To say otherwise is to endorse the option that leads to more suffering rather than less, and is contrary to the humanitarian impulse that prompts the decision not to prolong his life in the first place.

[5] Part of my point is that the process of being "allowed to die" can be relatively slow and painful, whereas being given a lethal injection is relatively quick and painless. Let me give a different sort of example. In the United States about one in 600 babies is born with Down's syndrome. Most of these babies are otherwise healthy -- that is, with only the usual pediatric care, they will proceed to an otherwise normal infancy. Some, however, are born with congenital defects such as intestinal obstructions that require operations if they are to live. Sometimes, the parents and the doctor will decide not to operate, and let the infant die. Anthony Shaw describes what happens then.

[6] When surgery is denied [the doctor] must try to keep the infant from suffering while natural forces sap the baby's life away. As a surgeon whose natural inclination

is to use the scalpel to fight off death, standing by and watching a salvageable baby die is the most emotionally exhausting experience I know. It is easy at a conference, in a theoretical discussion, to decide that such infants should be allowed to die. It is altogether different to stand by in the nursery and watch as dehydration and infection wither a tiny being over hours and days. This is a terrible ordeal for me and the hospital staff -- much more so than for the parents who never set foot in the nursery. [A. Shaw, "Doctor, Do We Have a Choice?" *The New York Times Magazine*, January 30, 1972, p. 54.]

[7] I can understand why some people are opposed to all euthanasia, and insist that such infants must be allowed to live. I think I can also understand why other people favor destroying these babies quickly and painlessly. But why should anyone favor letting "dehydration and infection wither a tiny being over hours and days"? The doctrine that says that a baby may be allowed to dehydrate and wither, but may not be given an injection that would end its life without suffering, seems so patently cruel as to require no further refutation. The strong language is not intended to offend, but only to put the point in the clearest possible way.

[8] My second argument is that the conventional doctrine leads to decisions concerning life and death made on irrelevant grounds.

[9] Consider again the case of the infants with Down's syndrome who need operations for congenital defects unrelated to the syndrome to live. Sometimes there is no operation, and the baby dies, but when there is no such defect, the baby lives on. Now, an operation such as that to remove an intestinal obstruction is not prohibitively difficult. The reason why such operations are not performed in these cases is, clearly, that the child has Down's syndrome and the parents and doctor judge that because of that fact it is better for the child to die.

[10] But notice that this situation is absurd, no matter what view one takes of the lives and potentials of such babies. If the life of such an infant is worth preserving, what does it matter if it needs a simple operation? Or, if one thinks it better that such a baby should not live on, what difference does it make that it happens to have an unobstructed intestinal tract? In either case, the matter of life and death is being decided on irrelevant grounds. It is the Down's syndrome, and not the intestines, that is the issue. The matter should be decided, if at all, on that basis, and not be allowed to depend on the essentially irrelevant question of whether the intestinal tract is blocked.

[11] What makes this situation possible, of course, is the idea that when there is an intestinal blockage, one can "let the baby die," but when there is no such defect there is nothing that can be done, for one must not "kill" it. The fact that this idea leads to such results as deciding life or death on irrelevant grounds is another good reason why the doctrine should be rejected.

[12] One reason why so many people think that there is an important moral difference between active and passive euthanasia is that they think killing someone is morally worse than letting someone die. But is it? Is killing, in itself, worse than letting die? To investigate this issue, two cases may be considered that are exactly alike except that one involves killing whereas the other involves letting someone die. Then, it can be asked whether this difference makes any difference to the moral assessments. It is important that the cases be exactly alike, except for this one difference, since otherwise one cannot be confident that it is this difference and not some other that accounts for any variation in the assessments of the two cases. So, let us consider this pair of cases:

[13] In the first, Smith stands to gain a large inheritance if anything should happen to his six-year-old cousin. One evening while the child is taking his bath, Smith sneaks into the bathroom and drowns the child, and then arranges things so that it will look like an accident.

[14] In the second, Jones also stands to gain if anything should happen to his six-year-old cousin. Like Smith, Jones sneaks in planning to drown the child in his bath. However, just as he enters the bathroom Jones sees the child slip and hit his head, and fall face down in the water. Jones is delighted; he stands by, ready to push the child's head back under if it is necessary, but it is not necessary. With only a little thrashing about, the child drowns all by himself, "accidentally," as Jones watches and does nothing.

[15] Now Smith killed the child, whereas Jones "merely" let the child die. That is the only difference between them. Did either man behave better, from a moral point of view? If the difference between killing and letting die were in itself a morally important matter, one should say that Jones's behavior was less reprehensible than Smith's. But does one really want to say that? I think not. In the first place, both men acted from the same motive, personal gain, and both had exactly the same end in view when they acted. It may be inferred from Smith's conduct that he is a bad man, although that judgment may be withdrawn or modified if certain further facts are learned about him -- for example, that he is mentally deranged. But would not the very same thing be inferred about Jones from his conduct? And would not the same further considerations also be relevant to any modification of this judgment? Moreover, suppose Jones pleaded, in his own defense, "After all, I didn't do anything except just stand there and watch the child drown. I didn't kill him; I only let him die." Again, if letting die were in itself less bad than killing, this defense should have at least some weight. But it does not. Such a "defense" can only be regarded as a grotesque perversion of moral reasoning. Morally speaking, it is no defense at all.

[16] Now it may be pointed out, quite properly, that the cases of euthanasia with which doctors are concerned are not like this at all. They do not involve personal gain or the destruction of normal healthy children. Doctors are concerned only with cases in which the patient's life is of no further use to him, or in which the patient's life has become or will soon become a terrible burden. However, the point is the same in these cases: the bare difference between killing and letting die does not, in itself, make a moral difference. If a doctor lets a patient die, for humane reasons, he is in the same moral position as if he had given the patient a lethal injection for humane reasons. If his decision was wrong -- if, for example, the patient's illness was in fact curable -- the decision would be equally regrettable no matter which method was used to carry it out. And if the doctor's decision was the right one, the method used is not in itself important.

[17] The AMA policy statement isolates the crucial issue very well; the crucial issue is "the intentional termination of the life of one human being by another." But after identifying this issue, and forbidding "mercy killing," the statement goes on to deny that the cessation of treatment is the intentional termination of a life. This is where the mistake comes in, for what is the cessation of treatment, in these circumstances, if it is not "the intentional termination of the life of one human being by another?" Of course it is exactly that, and if it were not, there would be no point to it.

[18] Many people will find this judgment hard to accept. One reason, I think, is that it is very easy to conflate the question of whether killing is, in itself, worse than letting die, with the very different question of whether most actual cases of killing are more reprehensible than most actual cases of letting die. Most actual cases of killing are clearly terrible (think, for example, of all the murders reported in the newspapers), and one hears of such cases

every day. On the other hand, one hardly ever hears of a case of letting die, except for the actions of doctors who are motivated by humanitarian reasons. So one learns to think of killing in a much worse light than of letting die. But this does not mean that there is something about killing that makes it in itself worse than letting die, for it is not the bare difference between killing and letting die that makes the difference in these cases. Rather, the other factors -- the murderer's motive of personal gain, for example, contrasted with the doctor's humanitarian motivation -- account for different reactions to the different cases.

[19] I have argued that killing is not in itself any worse than letting die; if my contention is right, it follows that active euthanasia is not any worse than passive euthanasia. What arguments can be given on the other side? The most common, I believe, is the following:

"The important difference between active and passive euthanasia is that, in passive euthanasia, the doctor does not do anything to bring about the patient's death. The doctor does nothing, and the patient dies of whatever ills already afflict him. In active euthanasia, however, the doctor does something to bring about the patient's death: he kills him. The doctor who gives the patient with cancer a lethal injection has himself caused his patient's death; whereas if he merely ceases treatment, the cancer is the cause of the death."

[20] A number of points need to be made here. The first is that it is not exactly correct to say that in passive euthanasia the doctor does nothing, for he does do one thing that is very important: he lets the patient die. "Letting someone die" is certainly different, in some respects, from other types of action -- mainly in that it is a kind of action that one may perform by way of not performing certain other actions. For example, one may let a patient die by way of not giving medication, just as one may insult someone by way of not shaking his hand. But for any purpose of moral assessment, it is a type of action nonetheless. The decision to let a patient die is subject to moral appraisal in the same way that a decision to kill him would be subject to moral appraisal: it may be assessed as wise or unwise, compassionate or sadistic, right or wrong. If a doctor deliberately let a patient die who was suffering from a routinely curable illness, the doctor would certainly be to blame for what he had done, just as he would be to blame if he had needlessly killed the patient. Charges against him would then be appropriate. If so, it would be no defense at all for him to insist that he didn't "do anything." He would have done something very serious indeed, for he let his patient die.

[21] Fixing the cause of death may be very important from a legal point of view, for it may determine whether criminal charges are brought against the doctor. But I do not think that this notion can be used to show a moral difference between active and passive euthanasia. The reason why it is considered bad to be the cause of someone's death is that death is regarded as a great evil -- and so it is. However, if it has been decided that euthanasia -- even passive euthanasia -- is desirable in a given case, it has also been decided that in this instance death is no greater an evil than the patient's continued existence. And if this is true, the usual reason for not wanting to be the cause of someone's death simply does not apply.

[22] Finally, doctors may think that all of this is only of academic interest -- the sort of thing that philosophers may worry about but that has no practical bearing on their own work. After all, doctors must be concerned about the legal consequences of what they do, and active euthanasia is clearly forbidden by the law. But even so, doctors should also be concerned with the fact that the law is forcing upon them a moral doctrine that may well be indefensible, and has a considerable effect on their practices. Of course, most doctors are

not now in the position of being coerced in this matter, for they do not regard themselves as merely going along with what the law requires. Rather, in statements such as the AMA policy statement that I have quoted, they are endorsing this doctrine as a central point of medical ethics. In that statement, active euthanasia is condemned not merely as illegal but as "contrary to that for which the medical profession stands," whereas passive euthanasia is approved. However, the preceding considerations suggest that there is really no moral difference between the two, considered in themselves (there may be important moral differences in some cases in their consequences, but, as I pointed out, these differences may make active euthanasia, and not passive euthanasia, the morally preferable option). So, whereas doctors may have to discriminate between active and passive euthanasia to satisfy the law, they should not do any more than that. In particular, they should not give the distinction any added authority and weight by writing it into official statements of medical ethics.

A Reply to Rachels on Active and Passive Euthanasia

TOM L. BEAUCHAMP

[1] James Rachels has recently argued that the distinction between active and passive euthanasia is neither appropriately used by the American Medical Association nor generally used for the resolution of moral problems of euthanasia. Indeed he believes this distinction - which he equates with the killing/letting die distinction -- does not in itself have any moral importance. The chief object of his attack is the following statement adopted by the House of Delegates of the American Medical Association in 1973:

The intentional termination of the life of one human being by another -- mercy killing -- is contrary to that for which the medical profession stands and is contrary to the policy of the American Medical Association.

The cessation of the employment of extraordinary means to prolong the life of the body when there is irrefutable evidence that biological death is imminent is the decision of the patient and/or his immediate family. The advice and judgment of the physician should be freely available to the patient and/or his immediate family.

[2] Rachels constructs a powerful and interesting set of arguments against this statement. In this paper I attempt the following: (1) to challenge his views on the grounds that he does not appreciate the moral reasons which give weight to the active/passive distinction; (2) to provide a constructive account of the moral relevance of the active/passive distinction; and (3) to offer reasons showing that Rachels may nonetheless be correct in urging that we ought to abandon the active/passive distinction for purposes of moral reasoning.

I

[3] I would concede that the active/passive distinction is sometimes morally irrelevant. Of this Rachels convinces me. But it does not follow that it is always morally irrelevant. What we need, then, is a case where the distinction is a morally relevant one and an explanation why it is so. Rachels himself uses the method of examining two cases which are exactly alike except that "one involves killing whereas the other involves letting die". We may profitably begin by comparing the kinds of cases governed by the AMA's doctrine with the kinds of cases adduced by Rachels in order to assess the adequacy and fairness of his cases.

[4] The second paragraph of the AMA statement is confined to a narrowly restricted range of passive euthanasia cases, viz., those (a) where the patients are on extraordinary means, (b) where irrefutable evidence of imminent death is available, and (c) where patient or family consent is available. Rachels' two cases involve conditions notably different from these:

[5] In the first, Smith stands to gain a large inheritance if anything should happen to his six-year-old cousin. One evening while the child is taking his bath, Smith sneaks into the bathroom and drowns the child, and then arranges things so that it will look like an accident.

[6] In the second, Jones also stands to gain if anything should happen to his six-year-old cousin. Like Smith, Jones sneaks in planning to drown the child in his bath. However, just as he enters the bathroom Jones sees the child slip and hit his head, and fall face down in the water. Jones is delighted; he stands by, ready to push the child's head back under if it is necessary, but it is not necessary. With only a little thrashing about, the child drowns all by himself, "accidentally," as Jones watches and does nothing.

[7] Now Smith killed the child, whereas Jones "merely" let the child die. That is the only difference between them.

[8] Rachels says there is no moral difference between the cases in terms of our moral assessments of Smith and Jones' behavior. This assessment seems fair enough, but what can Rachels' cases be said to prove, as they are so markedly disanalogous to the sorts of cases envisioned by the AMA proposal? Rachels concedes important disanalogies, but thinks them irrelevant:

[9] The point is the same in these cases: the bare difference between killing and letting die does not, in itself, make a moral difference. If a doctor lets a patient die, for humane reasons, he is in the same moral position as if he had given the patient a lethal injection for humane reasons.

[10] Three observations are immediately in order. First, Rachels seems to infer that from such cases we can conclude that the distinction between killing and letting die is always morally irrelevant. This conclusion is fallaciously derived. What the argument in fact shows, being an analogical argument, is only that in all relevantly similar cases the distinction does not in itself make a moral difference. Since Rachels concedes that other cases are disanalogous, he seems thereby to concede that his argument is as weak as the analogy itself. Second, Rachels' cases involve two unjustified actions, one of killing and the other of letting die. The AMA statement distinguishes one set of cases of unjustified killing and another of justified cases of allowing to die. Nowhere is it claimed by the AMA that what makes the difference in these cases is the active/passive distinction itself. It is only implied that one set of cases, the justified set, involves (passive) letting die while the unjustified set involves (active) killing. While it is said that justified euthanasia cases are passive ones and unjustified ones active, it is not said either that what makes some acts justified is the fact of their being passive or that what makes others unjustified is the fact of their being active. This fact will prove to be of vital importance.

[11] The third point is that in both of Rachels' cases the respective moral agents -- Smith and Jones -- are morally responsible for the death of the child and are morally blameworthy -- even though Jones is presumably not causally responsible. In the first case death is caused by the agent, while in the second it is not; yet the second agent is no less morally responsible. While the law might find only the first homicidal, morality condemns the motives in each case as equally wrong, and it holds that the duty to save life in such cases is as compelling as the duty not to take life. I suggest that it is largely because of this equal degree of moral responsibility that there is no morally relevant difference in Rachels' cases. In the cases envisioned by the AMA, however, an agent is held to be responsible for taking life by actively killing but is not held to be morally required to preserve life, and so not responsible for death, when removing the patient from extraordinary means (under conditions a-c above). I shall elaborate this latter point momentarily. My only conclusion thus far is the negative one that Rachels' arguments rest on weak foundations. His cases are not relevantly similar to euthanasia cases and do not support his apparent conclusion that the active/passive distinction is always morally irrelevant.

II

[12] I wish first to consider an argument that I believe has powerful intuitive appeal and probably is widely accepted as stating the main reason for rejecting Rachels' views. I will maintain that this argument fails, and so leaves Rachels' contentions untouched.

[13] I begin with an actual case, the celebrated Quinlan case. Karen Quinlan was in a coma, and was on a mechanical respirator which artificially sustained her vital processes and which her parents wished to cease. At least some physicians believed there was irrefutable evidence that biological death was imminent and the coma irreversible. This case, under this description, closely conforms to the passive cases envisioned by the AMA. During an interview the father, Mr. Quinlan, asserted that he did not wish to kill his daughter, but only to remove her from the machines in order to see whether she would live or would die a natural death. Suppose he had said -- to envision now a second and hypothetical, but parallel case -- that he wished only to see her die painlessly and therefore wished that the doctor could induce death by an over-dose of morphine. Most of us would think the second act, which involves active killing, morally unjustified in these circumstances, while many of us would think the first act morally justified. (This is not the place to consider whether in fact it is justified, and if so under what conditions.) What accounts for the apparent morally relevant difference?

[14] I have considered these two cases together in order to follow Rachels' method of entertaining parallel cases where the only difference is that the one case involves killing and the other letting die. However, there is a further difference, which crops up in the euthanasia context. The difference rests in our judgments of medical fallibility and moral responsibility. Mr. Quinlan seems to think that, after all, the doctors might be wrong. There is a remote possibility that she might live without the aid of a machine. But whether or not the medical prediction of death turns out to be accurate, if she dies then no one is morally responsible for directly bringing about or causing her death, as they would be if they caused her death by killing her. Rachels finds explanations which appeal to causal conditions unsatisfactory; but perhaps this is only because he fails to see the nature of the causal link. To bring about her death is by that act to preempt the possibility of life. To "allow her to die" by removing artificial equipment is to allow, for the possibility of wrong diagnosis or incorrect prediction and hence to absolve oneself of moral responsibility for the taking of life under false assumptions. There may, of course, be utterly no empirical possibility of recovery in some cases since recovery would violate a law of nature. However, judgments of empirical impossibility in medicine are notoriously problematic -- the reason for emphasizing medical fallibility. And in all the hard cases we do not know that recovery is empirically impossible, even if good evidence is available.

[15] The above reason for invoking the active/passive distinction can now be generalized: Active termination of life removes all possibility of life for the patient, while passively ceasing extraordinary means may not. This is not trivial since patients have survived in several celebrated cases where, in knowledgeable physicians' judgments, there was "irrefutable" evidence that death was imminent.

[16] One may, of course, be entirely responsible and culpable for another's death either by killing him or by letting him die. In such cases, of which Rachels' are examples, there is no morally significant difference between killing and letting die precisely because whatever one does, omits, or refrains from doing does not absolve one of responsibility. Either active or

passive involvement renders one responsible for the death of another, and both involvements are equally wrong for the same principled moral reason: it is (*prima facie*) morally wrong to bring about the death of an innocent person capable of living whenever the causal intervention or negligence is intentional. (I use causal terms here because causal involvement need not be active, as when by one's negligence one is nonetheless causally responsible.) But not all cases of killing and letting die fall under this same moral principle. One is sometimes culpable for killing, because morally responsible as the agent for death, as when one pulls the plug on a respirator sustaining a recovering patient (a murder). But one is sometimes not culpable for letting die because one is not morally responsible as agent, as when one pulls the plug on a respirator sustaining an irreversibly comatose and unrecoverable patient (a routine procedure, where one is *merely* causally responsible). Different degrees and means of involvement assess different degrees of responsibility, and our assessments of culpability can become intricately complex. The only point which now concerns us, however, is that because different moral principles may govern very similar circumstances, we are sometimes morally culpable for killing but not for letting die. And to many people it will seem that in passive cases we are not morally responsible for causing death, though we are responsible in active cases.

[17] This argument is powerfully attractive. Although I was once inclined to accept it in virtually the identical form just developed, I now think that, despite its intuitive appeal, it cannot be correct. It is true that different degrees and means of involvement entail different degrees of responsibility, but it does not follow that we are not responsible and therefore are absolved of possible culpability in any case of intentionally allowing to die. We are responsible and perhaps culpable in either active or passive cases. Here Rachels' argument is entirely to the point: It is not primarily a question of greater or lesser responsibility by an active or a passive means that should determine culpability. Rather, the question of culpability is decided by the moral justification for choosing either a passive or an active means. What the argument in the previous paragraph overlooks is that one might be unjustified in using an active means or unjustified in using a passive means, and hence be culpable in the use of either; yet one might be justified in using an active means or justified in using a passive means, and hence not be culpable in using either. Fallibility might just as well be present in a judgment to use one means as in a judgment to use another. (A judgment to allow to die is just as subject to being based on knowledge which is fallible as a judgment to kill.) Moreover, in either case, it is a matter of what one knows and believes, and not a matter of a particular kind of causal connection or causal chain. If we kill the patient, then we are certainly causally responsible for his death. But, similarly, if we cease treatment, and the patient dies, the patient might have recovered if treatment had been continued. The patient might have been saved in either case, and hence there is no morally relevant difference between the two cases. It is, therefore, simply beside the point that "one is sometimes culpable for killing . . . but one is sometimes not culpable for letting die" -- as the above argument concludes.

[18] Accordingly, despite its great intuitive appeal and frequent mention, this argument from responsibility fails.

III

[19] There may, however, be more compelling arguments against Rachels, and I wish now to provide what I believe is the most significant argument that can be adduced in defense of the active/passive distinction. I shall develop this argument by combining (A) so-called wedge or slippery slope arguments with (B) recent arguments in defense of rule

utilitarianism. I shall explain each in turn and show how in combination they may be used to defend the active-passive distinction.

[20] (A) Wedge arguments proceed as follows: if killing were allowed, even under the guise of a merciful extinction of life, a dangerous wedge would be introduced which places all "undesirable" or "unworthy" human life in a precarious condition. Proponents of wedge arguments believe the initial wedge places us on a slippery slope for at least one of two reasons: (i) It is said that our justifying principles leave us with no principled way to avoid the slide into saying that all sorts of killings would be justified under similar conditions. Here it is thought that once killing is allowed, a firm line between justified and unjustified killings cannot be securely drawn. It is thought best not to redraw the line in the first place, for redrawing it will inevitably lead to a downhill slide. It is then often pointed out that as a matter of historical record this is precisely what has occurred in the darker regions of human history, including the Nazi era, where euthanasia began with the best intentions for horribly ill, non-Jewish Germans and gradually spread to anyone deemed an enemy of the people. (ii) Second, it is said that our basic principles against killing will be gradually eroded once some form of killing is legitimated. For example, it is said that permitting voluntary euthanasia will lead to permitting involuntary euthanasia, which will in turn lead to permitting euthanasia for those who are a nuisance to society (idiots, recidivist criminals, defective newborns, and the insane, e.g.). Gradually other principles which instill respect for human life will be eroded or abandoned in the process.

[21] I am not inclined to accept the first reason (i). If our justifying principles are themselves justified, then any action they warrant would be justified. Accordingly, I shall only be concerned with the second approach (ii).

[22] (B) Rule utilitarianism is the position that a society ought to adopt a rule if its acceptance would have better consequences for the common good (greater social utility) than any comparable rule could have in that society. Any action is right if it conforms to a valid rule and wrong if it violates the rule. Sometimes it is said that alternative rules should be measured against one another, while it has also been suggested that whole moral codes (complete sets of rules) rather than individual rules should be compared. While I prefer the latter formulation (Brandt's), this internal dispute need not detain us here. The important point is that a particular rule or a particular code of rules is morally justified if and only if there is no other competing rule or moral code whose acceptance would have a higher utility value for society, and where a rule's acceptability is contingent upon the consequences which would result if the rule were made current.

[23] Wedge arguments, when conjoined with rule utilitarian arguments, may be applied to euthanasia issues in the following way. We presently subscribe to a no-active-euthanasia rule (which the AMA suggests we retain). Imagine now that in our society we make current a restricted-active-euthanasia rule (as Rachels seems to urge). Which of these two moral rules would, if enacted, have the consequence of maximizing social utility? Clearly a restricted-active-euthanasia rule would have some utility value, as Rachels notes, since some intense and uncontrollable suffering would be eliminated. However, it may not have the highest utility value in the structure of our present code or in any imaginable code which could be made current, and therefore may not be a component in the ideal code for our society. If wedge arguments raise any serious questions at all, as I think they do, they rest in this area of whether a code would be weakened or strengthened by the addition of active euthanasia principles. For the disutility of introducing legitimate killing into one's moral code (in the form of active euthanasia rules) may, in the long run, outweigh the utility of doing so, as a result of the eroding effect such a relaxation would have on rules in the code which demand respect for human life. If, for example, rules permitting active killing were

introduced, it is not implausible to suppose that destroying defective newborns (a form of involuntary euthanasia) would become an accepted and common practice, that as population increases occur the aged will be even more neglectable and neglected than they now are, that capital punishment for a wide variety of crimes would be increasingly tempting, that some doctors would have appreciably reduced fears of actively injecting fatal doses whenever it seemed to them propitious to do so, and that laws of war against killing civilians would erode in efficacy even beyond their already abysmal level.

[24] A hundred such possible consequences might easily be imagined. But these few are sufficient to make the larger point that such rules permitting killing could lead to a general reduction of respect for human life. Rules against killing in a moral code are not isolated moral principles; they are pieces of a web of rules against killing which forms the code. The more threads one removes, the weaker the fabric becomes. And if, as I believe, moral principles against active killing have the deep and continuously civilizing effect of promoting respect for life, and if principles which allow passively letting die (as envisioned in the AMA statement) do not themselves cut against this effect, then this seems an important reason for the maintenance of the active/passive distinction. (By the logic of the above argument, passively letting die would also have to be prohibited if a rule permitting it had the serious adverse consequence of eroding acceptance or rules protective of respect for life. While this prospect seems to me improbable, I can hardly claim to have refuted those conservatives who would claim that even rules that sanction letting die place us on a precarious slippery slope.)

[25] A troublesome problem, however, confronts my use of utilitarian and wedge arguments. Most all of us would agree that both killing and letting die are justified under some conditions. Killings in self-defense and in "just" wars are widely accepted as justified because the conditions excuse the killing. If society can withstand these exceptions to moral rules prohibiting killing, then why is it not plausible to suppose society can accept another excusing exception in the form of justified active euthanasia? This is an important and worthy objection, but not a decisive one. The defenseless and the dying are significantly different classes of persons from aggressors who attack individuals and/or nations. In the case of aggressors, one does not confront the question whether their lives are no longer worth living. Rather, we reach the judgment that the aggressors' morally blameworthy actions justify counteractions. But in the case of the dying and the otherwise ill, there is no morally blameworthy action to justify our own. Here we are required to accept the judgment that their lives are no longer worth living in order to believe that the termination of their lives is justified. It is the latter sort of judgment which is feared by those who take the wedge argument seriously. We do not now permit and never have permitted the taking of morally blameless lives. I think this is the key to understanding why recent cases of intentionally allowing the death of defective newborns have generated such protracted controversy. Even if such newborns could not have led meaningful lives (a matter of some controversy), it is the wedged foot in the door which creates the most intense worries. For if we once take a decision to allow a restricted infanticide justification or any justification at all on grounds that a life is not meaningful or not worth living, we have qualified our moral rules against killing. That this qualification is a matter of the utmost seriousness needs no argument. I mention it here only to show why the wedge argument may have moral force even though we already allow some very different conditions to justify intentional killing.

[26] There is one final utilitarian reason favoring the preservation of the active/passive distinction. Suppose we distinguish the following two types of cases of wrongly diagnosed patients:

1. Patients wrongly diagnosed as hopeless, and who will survive even if a treatment is ceased (in order to allow a natural death).
2. Patients wrongly diagnosed as hopeless, and who will survive only if the treatment is not ceased (in order to allow a natural death).

[27] If a social rule permitting only passive euthanasia were in effect, then doctors and families who "allowed death" would lose only patients in class 2, not those in class 1; whereas if active euthanasia were permitted, at least some patients in class 1 would be needlessly lost. Thus, the consequence of a no-active-euthanasia rule would be to save some lives which could not be saved if both forms of euthanasia were allowed. This reason is not a decisive reason for favoring a policy of passive euthanasia, since these classes (1 and 2) are likely to be very small and since there might be counterbalancing reasons (extreme pain, autonomous expression of the patient, etc.) in favor of active euthanasia. But certainly it is a reason favoring only passive euthanasia and one which is morally relevant and ought to be considered along with other moral reasons.

IV

[28] It may still be insisted that my case has not touched Rachels' leading claim, for I have not shown, as Rachels puts it, that it is "the bare difference between killing and letting die that makes the difference in these cases". True, I have not shown this and in my judgment it cannot be shown. But this concession does not require capitulation to Rachels' argument. I adduced a case which is at the center of our moral intuition that killing is morally different (in at least some cases) from letting die; and I then attempted to account for at least part of the grounds for this belief. The grounds turn out to be other than the bare difference, but nevertheless make the distinction morally relevant. The identical point can be made regarding the voluntary/involuntary distinction, as it is commonly applied to euthanasia. It is not the bare difference between voluntary euthanasia (i.e., euthanasia with patient consent) and involuntary euthanasia (i.e., without patient consent) that makes one justifiable and one not. Independent moral grounds based on, for example, respect for autonomy or beneficence, or perhaps justice will alone make the moral difference.

[29] In order to illustrate this general claim, let us presume that it is sometimes justified to kill another person and sometimes justified to allow another to die. Suppose, for example, that one may kill in self-defense and may allow to die when a promise has been made to someone that he would be allowed to die. Here conditions of self-defense and promising justify actions. But suppose now that someone A promises in exactly similar circumstances to kill someone B at B's request, and also that someone C allows someone D to die in an act of self-defense. Surely A is obliged equally to kill or to let die if he promised; and surely C is permitted to let D die if it is a matter of defending C's life. If this analysis is correct, then it follows that killing is sometimes right, sometimes wrong, depending on the circumstances, and the same is true of letting die. It is the justifying reasons which make the difference whether an action is right, not merely the kind of action it is.

[30] Now, *if* letting die led to disastrous conclusions but killing did not, then letting die but not killing would be wrong. Consider, for example, a possible world in which dying would be indefinitely prolongable even if all extraordinary therapy were removed and the patient were allowed to die. Suppose that it costs over one million dollars to let each patient die, that nurses consistently commit suicide from caring for those being "allowed to die," that physicians are constantly being successfully sued for malpractice for allowing death by cruel

and wrongful means, and that hospitals are uncontrollably overcrowded and their wards filled with communicable diseases which afflict only the dying. Now suppose further that killing in this possible world is quick, painless, and easily monitored. I submit that in this world we would believe that killing is morally acceptable but that allowing to die is morally unacceptable. The point of this example is again that it is the circumstances that make the difference, not the bare difference between killing and letting die.

[31] It is, however, worth noticing that there is nothing in the AMA statement which says that the bare difference between killing and letting die itself and alone makes the difference in our differing moral assessments of rightness and wrongness. Rachels forces this interpretation on the statement. Some philosophers may have thought bare difference makes the difference, but there is scant evidence that the AMA or any thoughtful ethicist must believe it in order to defend the relevance and importance of the active/passive distinction. When this conclusion is coupled with my earlier argument that from Rachels' paradigm cases it follows only that the active/passive distinction is sometimes, but not always, morally irrelevant, it would seem that his case against the AMA is rendered highly questionable.

V

[32] There remains, however, the important question as to whether we ought to accept the distinction between active and passive euthanasia, now that we are clear about (at least one way of drawing) the moral grounds for its invocation. That is, should we employ the distinction in order to judge some acts of euthanasia justified and others not justified? Here, as the hesitant previous paragraph indicates, I am uncertain. This problem is a substantive moral issue -- not merely a conceptual one -- and would require at a minimum a lengthy assessment of wedge arguments and related utilitarian considerations. In important respects empirical questions are involved in this assessment. We should like to know, and yet have hardly any evidence to indicate, what the consequences would be for our society if we were to allow the use of active means to produce death. The best hope for making such an assessment has seemed to some to rest in analogies to suicide and capital punishment statutes. Here it may reasonably be asked whether recent liberalizations of laws limiting these forms of killing have served as the thin end of a wedge leading to a breakdown of principles protecting life or to widespread violations of moral principles. Nonetheless, such analogies do not seem to me promising, since they are still fairly remote from the pertinent issue of the consequences of allowing active humanitarian killing of one person by another.

[33] It is interesting to notice the outcome of the Kamisar-Williams debate on euthanasia -- which is almost exclusively cast by both writers in a consequential, utilitarian framework. At one crucial point in the debate, where possible consequences of laws permitting euthanasia are under discussion, they exchange "perhaps" judgments:

[34] I [Williams] will return Kamisar the compliment and say: "Perhaps." We are certainly in an area where no solution is going to make things quite easy and happy for everybody, and all sorts of embarrassments may be conjectured. But these embarrassments are not avoided by keeping to the present law: we suffer from them already.

[35] Because of the grave difficulties which stand in the way of making accurate predictions about the impact of liberalized euthanasia laws -- especially those that would permit active killing -- it is not surprising that those who debate the subject would reach a point of

exchanging such "perhaps" judgments. And that is why, so it seems to me, we are uncertain whether to perpetuate or to abandon the active-passive distinction in our moral thinking about euthanasia. I think we do perpetuate it in medicine, law, and ethics because we are still somewhat uncertain about the conditions under which passive euthanasia should be permitted by law (which is one form of social rule). We are unsure about what the consequences will be of the California "Natural Death Act" and all those similar acts passed by other states which have followed in its path. If no untoward results occur, and the balance of the results seems favorable, then we will perhaps be less concerned about further liberalizations of euthanasia laws. If untoward results do occur (on a widespread scale), then we would be most reluctant to accept further liberalizations and might even abolish natural death acts.

[36] In short, I have argued in this section that euthanasia in its active and its passive forms presents us with a dilemma which can be developed by using powerful consequentialist arguments on each side, yet there is little clarity concerning the proper resolution of the dilemma precisely because of our uncertainty regarding proclaimed consequences.

VI

[37] I reach two conclusions at the end of these several arguments. First, I think Rachels is incorrect in arguing that the distinction between active and passive is (always) morally irrelevant. It may well be relevant, and for moral reasons -- the reasons adduced in section **III** above. Second, I think nonetheless that Rachels may ultimately be shown correct in his contention that we ought to dispense with the active-passive distinction -- for reasons adduced in sections **IV-V**. But if he is ultimately judged correct, it will be because we have come to see that some forms of active killing have generally acceptable social consequences, and not primarily because of the arguments he adduces in his paper -- even though something may be said for each of these arguments. Of course, in one respect I have conceded a great deal to Rachels. The bare difference argument is vital to his position, and I have fully agreed to it. On the other hand, I do not see that the bare difference argument does play or need play a major role in our moral thinking -- or in that of the AMA.

Right to Die or Duty to Live? The Problem of Euthanasia

WILLIAM GREY

Journal of Applied Philosophy, Vol. 16, No. 1, 1999

Of old, when folk were sick and sorely tried
The doctors gave them physic, and they died;
But here's a happier age, for now we know
Both how to make men sick and keep them so.
-Hilaire Belloc

1. Introduction

[1] The world's first legal euthanasia death occurred in the Australian city of Darwin on Sunday September 22, 1996 when Bob Dent ended his life under the Northern Territory's short-lived *Rights of the Terminally Ill Act 1995*. It is well known that euthanasia is openly practised in the Netherlands, even though the practice there remains technically in breach of the criminal law. The Northern Territory however was the first jurisdiction to boldly legislate where none had gone before.

[2] Dent's death intensified argument about euthanasia in Australia, transforming the debate from a textbook discussion in social ethics into a vigorous and divisive social dispute. Dent's death undermined the view that a *de facto* acceptance of the practice of euthanasia as part of palliative care might provide an avenue for defusing the debate. This may have been a rash suggestion anyway, even though there is evidence that euthanasia is covertly practised in Australia on a scale even greater than its more widely reported occurrence in the Netherlands.

2. The classical arguments

[3] The debate about the moral acceptability of the practice of euthanasia embodies a number of quintessentially philosophical elements. Several of the crucial concepts need to be articulated carefully to ensure that argument is not vitiated by vagueness and equivocation. Key terms which need careful delineation include 'death', the distinction between 'ordinary' and 'extraordinary' measures in medical treatment, and the distinction between active and passive killing the distinction between killing and allowing to die. This important elucidatory task has been ably and extensively addressed elsewhere and I have little to add to it here. I will however say something about the taxonomy of acts of dying which these categories help to delineate.

[4] Related to this clarificatory task we must consider several distinct *prima facie* categories of 'persons' (using the term widely to include 'pre-persons' and 'post-persons') for which the practice of euthanasia has been considered or proposed. These include, first, individuals who have descended into a persistent vegetative state, perhaps through injury or pharmacological mishap, who can be sustained by medical technology. Karen Ann Quinlan was a well known exemplar of this category.

[5] Secondly there are cases where persons have lapsed into a state of contented dementia, perhaps as a result of accident or the later stages of the 'long goodbye' of Alzheimer's disease. Such individuals may manifest little or no capacity for any significant personal interaction, let alone development or growth. Should someone in such a condition receive treatment for a life-threatening illness?

[6] Suppose that such an individual had earlier, when their faculties were intact, executed a living will requesting that treatment in such circumstances be withheld. Should that earlier decision be decisive? It is certainly not taken as decisive by a large number of physicians. Ruth Fischbach, Assistant Professor of Social Medicine at Harvard Medical School, reported at a recent medical conference that 66 percent of physicians interviewed felt that there was nothing wrong with overriding a patient's advance directive, even if that directive unambiguously stated the conditions for the withdrawal and withholding of medical treatment.

[7] Thirdly there are problematic cases of defective newborns. Cases of anencephalous, or Down's syndrome infants with duodenal atresia (intestinal blockage), are textbook clinical examples which raise vexing problems about permissible treatment, and permissible neglect.

[8] Fourthly and (for present purposes) finally there are cases where a life is one of agonising and unrelenting pain. This is perhaps the strongest category of claimants for the 'right to die' a right which was eloquently claimed by Bob Dent in a letter to Members of Australian Federal Parliament which Dent dictated to his wife the day before his physician-assisted death.

3. Types of euthanasia

[9] The taxonomy of euthanasia includes active and passive and voluntary, non-voluntary and involuntary. These can be depicted in a table, including *prima facie* candidates for the different categories.

[10] The non-voluntary and involuntary categories, set out in the two right-hand columns, are often conflated and this is an important source of confusion and misconception. Because they raise significantly different issues it is important to distinguish them carefully. Non-voluntary cases include individuals who are incapable of indicating a preference, or who are not even capable of having preferences at all. The permissible treatment and permissible neglect of those in this category is different from that of those who have indicated a preference about their treatment, as is the case in respect of both the voluntary and involuntary categories.

	VOLUNTARY	NON-VOLUNTARY	INVOLUNTARY
PASSIVE	Patient refusal of medical treatment	Persistent vegetative state (Karen Quinlan)	Medical rationing (drugs or treatment; organs recipients)
ACTIVE	Patient request for termination (Bob Dent)	Persistent vegetative state	Unlawful killing (Nazi 'euthanasia' programs)

[11] There is widespread consensus about some of the categories. Passive voluntary euthanasia, the right to refuse treatment, is almost universally accepted. Active involuntary

euthanasia, the infamous outcome of the so-called 'euthanasia' program of Nazi Germany, is unanimously condemned, and indeed has done much to raise alarm about permitting euthanasia in any form. One of the principal arguments marshalled by opponents is that sanctioning euthanasia in any of its forms will inexorably slide into an acceptance of acts in this inadmissible category.

[12] The category which is most vigorously contested is active voluntary euthanasia, the category into which Bob Dent fell. What are the arguments for and against this?

4. Arguments for active euthanasia

(a) Autonomy

[13] The strongest argument in favour of active voluntary euthanasia is based on respect for individual autonomy. The argument from autonomy is based on the claim that every person has the right to shape their own life through their choices and that must include, certainly under the unfavourable circumstances of pain or disability and some would argue more widely, the right to choose the time and circumstances of their death. On this view it is a matter of basic human dignity to be given the right to decide about the circumstances of our own lives and our deaths.

[14] The principle of autonomy is an expression of the essentially Kantian idea that what is of paramount importance for my life is that it consists of my own choices, for good or ill. For Kant, notoriously, it is never permissible to treat persons as a means rather than as ends-in-themselves, even if this involves attempting to use them as a means to their own well-being. Treating someone as a means to their own well-being would involve the dubious paternalistic presumption that someone other than that person might know better than they in what this well-being might consist.

[15] J.S. Mill was as resolute a defender as Kant of the right to determine one's own destiny autonomously, and to be free from heteronomous determination by others. This right is implicit in Mill's famous 'harm' principle:

the only purpose for which power can be rightfully exercised over any member of a civilised community, against their will, is to prevent harm to others. His own good, either physical or moral, is not a sufficient warrant.

... Mankind are greater gainers by suffering each other to live as seems good to themselves than by compelling each to live as seems good to the rest.

[16] Taking autonomy (literally 'self governance') seriously means acknowledging individual sovereignty over all purely self-regarding acts. Determining the circumstances of one's own death, according to this principle, should be allowed provided that it is a self-regarding act, and if so like other self-regarding acts it should be exempt from the interference of others. Mill, like Kant, thought that other people's taking control of one's life was an intolerable intrusion.

[17] In the spirit of Kant and Mill, Bob Dent clearly regarded Kevin Andrews' (ultimately successful) proposal to limit the freedom of individuals to seek medical assistance to end their intolerable life as an insufferable impertinence. In his last letter, written to the Australian Members of Parliament who were then debating the fate of the Northern Territory's *Rights of the Terminally III Act*, Dent wrote:

[18] I have no wish for further experimentation by the palliative care people in their efforts to control my pain. My current program involves taking 30 tablets a day! For months I have been on a roller-coaster of pain made worse by the unwanted side-effects of the drugs. Morphine causes constipation laxatives work erratically ... Other drugs given to enhance the pain-relieving effects of the morphine have caused me to feel suicidal to the point that I would have blown my head off if I had had a gun ... If I were to keep a pet animal in the same condition I am in, I would be prosecuted... What right has anyone, because of their own religious faith (to which I don't subscribe), to demand that I behave according to their rules until some omniscient doctor decides that I must have had enough and goes ahead and increases my morphine until I die? If you disagree with voluntary euthanasia, then don't use it, but don't deny me the right to use it if and when I want to. I am immensely grateful that I have had the opportunity to use the *Rights of the Terminally III Act* to ask my doctor, Philip Nitschke, to assist me to end this interminable suffering and to end my life in a dignified and compassionate manner.

(b) Beneficence and fairness

[19] There are at least three other important, and related, subsidiary arguments which are used to support a right to determine the circumstances of one's own death. First the golden rule 'do as you would be done by' requires that we provide aid to those in distress and in particular provide appropriate relief from suffering. In cases where palliation is ineffective, again illustrated in the case of Bob Dent, it would be unreasonable (many would say unconscionable) to deny anyone the right to make a dignified end to their suffering.

[20] Secondly, a corollary of the harm principle is that the denial of a right to die is unfair and cruel; no one should be obliged to endure unbearable suffering. The unreasonableness of those who would deny someone's rational choice to end their suffering is a principal concern of Bob Dent's letter. Thirdly, the denial of a right to die amounts to imposing a 'duty to live' no matter what the abject condition of that life might be. This imposition is presumptuous and intolerable.

5. Arguments against legalising euthanasia

[21] Not everyone opposed to the legalisation of euthanasia is opposed to the practice of euthanasia: we must therefore distinguish arguments against the practice of euthanasia from arguments directed solely against the legalisation of the practice. Those who oppose euthanasia in principle will of course oppose its legalisation; but there are some who, while supporting euthanasia in principle, have misgivings about its institutionalisation. Thus there are some who defend the right of people to choose the time and circumstances of their death but who find the requirement of satisfying a medical bureaucracy that their decision is sound both onerous and offensive. End-of-life decisions, on this view, should be a personal matter between patient and physician. Indeed one of the results of the survey of euthanasia in medical practice in Australia is that ending life by personal arrangement with one's physician is much more widespread than is commonly supposed. A significant number of physicians apparently are prepared to assist patients illegally to end their lives.

[22] According to a survey of doctors' attitudes towards, and practice of, euthanasia in New South Wales and the Australian Capital Territory, approximately half the respondents had been asked by a patient to hasten their death, and approximately a third had assisted

patients to end their lives. These figures were almost identical with the result of an earlier survey of doctors in Victoria. Medical practitioners who have included themselves in the substantial minority of doctors who have admitted to helping patients die sooner than they might have include Peter Baume, Professor of Community Medicine at the University of New South Wales and a former Federal Minister for Health, and the present (1997) Federal Minister for Health, Dr Michael Wooldridge.

[23] However, there are problems about tolerating euthanasia as a covert clinical procedure to be allowed to operate outside the law. It is arguable that it is social hypocrisy to advocate in private something which is rejected as a matter of public policy. Moreover it is an unjust policy because its benefits are not equally accessible to all, but depend to a large extent on individual resources, or resourcefulness, or luck.

6. Arguments against the practice of euthanasia

[24] Turning from arguments directed simply against the legalisation of euthanasia to objections in principle to the practice of euthanasia, there are three lines of objection which predominate, together with several auxiliary arguments. The three main objections are first the so-called 'slippery slope' or 'wedge' argument, second the worry that the overt practice of euthanasia would change the culture of medicine, and third an objection based on a denial of the claim that a decision to end one's life is purely self-regarding.

[25] The auxiliary arguments include the claim that palliation delivers adequate relief from indignity and pain, uncertainty arguments about the possibility of miraculous cures, the suggestion that people in the grip of a temporary physiological or psychological crisis may make an unwise and irrevocable choice to end their lives, the claim that all human life is sacred, and finally, an attempt to delineate which acts of killing are justified through the doctrine of 'double effect'.

[26] I will discuss each of these in turn and then briefly discuss a claim about power which was raised during the controversy generated by the case of Bob Dent.

(a) Thin end of the wedge

[27] The most commonly articulated worry is that if euthanasia is admitted for deserving cases, such as Dent's, then this will lead to the admission of less deserving or quite inappropriate cases. In time the principle will be progressively extended while the safeguards become diluted and weakened. Soon a policy motivated by compassion will become a vehicle for abuse and monstrous injustice. If we allow a case like that of Bob Dent, it is suggested, then we are on the slippery slope to Auschwitz and Buchenwald.

[28] This alarmist rhetoric is unpersuasive. Any overt practice of euthanasia will attract the closest scrutiny and for this reason it is very difficult to imagine safeguards being eroded without vigorous public opposition. The wedge argument is in any case in general unconvincing, and it fares no better here than in its other applications. It is worth quoting FM Cornford's droll assessment of 'wedge' arguments in general:

The Principle of the Wedge is that you should not act justly now for fear of raising expectations that you may act still more justly in the future -- expectations which you are afraid you will not have the courage to satisfy. A little reflection will make it evident that the Wedge argument implies the admission that the persons who use it

cannot prove that the action is not just. If they could that would be the sole and sufficient reason for not doing it, and this argument would be superfluous.

[29] The answer to wedge arguments, here as elsewhere, is not to forbid just practices, but to take careful steps to ensure that proper regulation is maintained so that standards are not eroded. In particular the line which most express fear about crossing is the line between voluntary and non-voluntary.

[30] If it is true that abuses have occurred where the practice is allowed, as has been alleged of the Netherlands, this does not indicate the need for complete prohibition. An obvious alternative is to ensure that adequate safeguards are established to prevent abuses. In particular the aim of safeguards is to ensure that the final responsibility for the decision and its execution rests squarely with the patient, not with the physician or with any other party. It would be crucial to ensure that regulations and safeguards were developed to guarantee that the line between voluntary and non-voluntary acts of euthanasia will not be transgressed. Considerable care was given to this important matter in developing the procedures approved in the Northern Territory.

(b) Culture of medicine

[31] A second argument used by opponents of euthanasia is that it changes the culture of medicine. Instead of sustaining and nurturing life, medical practice would come to include the deliberate killing of patients. Doctors it is said should be healers, not killers. One response to this is to point out first that compassionately-motivated killing has become a surprisingly common part of medical practice [211]. What needs to be done is to turn an unregulated procedure, whose benefits are now capriciously distributed, into a pattern of practice which is both safe and just.

[32] In an age of sophisticated life-sustaining technology some elements of the Hippocratic oath, such as the injunction to preserve life at all costs, have passed their use-by date.

[33] Contemporary medical circumstances demand that the physician's duty of care and beneficence is primary and sustaining life needs to be subordinated to this basic aim. It is no longer credible to regard the sustaining of life, regardless of its quality, as the primary and paramount concern of medical practice [231].

(c) Social limits of autonomy

[34] One of the strongest objections to euthanasia is that the autonomy which it is our duty to respect is not enjoyed by everyone. Even if it is granted that respect for individual autonomy is of paramount importance, it nonetheless applies only to socially empowered individuals or groups within society. There may be serious problems with the application of this principle to marginalised groups and especially to individuals who are, or can be, exploited. Legalising euthanasia, according to some, ignores the social reality of marginalised groups, and persons who might be exploited by unscrupulous relatives, or unscrupulous doctors. This is an essentially utilitarian argument drawing attention to grave social consequences of legalising the practice.

[35] A fundamental clash of intuitions arises here. On the one hand it is claimed by individualistically-inclined liberals that no one has a right to dictate to them a duty to live. Against this it is argued by those who reject this individualistic conception of society that

liberal individualists have no right to damage the social fabric of society by making choices which might lead to the degradation of the value of life within the community. The difference between these fundamentally opposed intuitions hinges on the extent to which acts of self-destruction are purely self-regarding or whether they are significantly social or political in character. We can distinguish individual acts, which primarily affect the agent and which affect others only minimally, from social acts, which significantly affect, and possibly harm, others.

[36] Mill's harm principle explicitly allows for restrictions on an individual's freedom to act in cases where their act harms others. Opponents of euthanasia can claim that individual acts of self-destruction, and the medical assistance for such acts, do in fact affect others and therefore are not purely self-regarding. They affect adversely the attitude of the community to the value of life and undermine the culture of medicine, to the detriment of all. Bob Dent's right to die diminishes the rights of others to live. Attitudes to life are degraded and disregard of its serious value is facilitated. Bob Dent's death, on this view, is not just Bob Dent's business, and we cannot appeal to Mill's principle to privilege his views of the matter.

[37] The dispute about euthanasia on this point resembles another which arises in discussions of pornography and prostitution. Some women claim a right to make commercial use of their bodies as a matter of individual liberty. To this it has been replied that such choices do not affect them alone but help to shape community attitudes about how women are perceived. In allowing themselves to be viewed or used as sex objects, they are fostering degrading attitudes towards all women.

[38] Whether pornography and prostitution involve choices which are purely self-regarding, and how to balance these competing views, are not matters to be addressed here. I want only to note the analogy and to point out its consequences: in particular that those who appeal to the claim that permitting the legalisation of killing has adverse social consequences are in effect insisting that Bob Dent had a duty to live for the greater social good, no matter how wretched his circumstances.

(d) It is unnecessary: palliative care provides all the comfort needed

[39] A subsidiary and supporting argument frequently used by opponents of euthanasia is the claim that the practice is unnecessary because the relief sought by patients can be adequately provided with palliative care. Perhaps in some cases the end-of-life quality can be improved by the skilful deployment of palliative care, but palliation is not effective in all cases certainly as far as Bob Dent was concerned it was not good enough.

(e) Ignorance and uncertainty objections

[40] Another objection to allowing euthanasia is that this irrevocable step should be avoided because a miraculous cure or spontaneous remission might reverse medical misfortune. Clearly this cannot happen for people who have chosen to end their lives prematurely. This suggestion is fairly unpersuasive when one examines real clinical cases like that of Dent. Nothing was more certain to Dent than the fact that his condition was very bad and would only get worse. He had a strong desire not to be remembered as a pain-wracked incontinent wreck. Dent, and the other three who died under the provisions of the *Rights of the Terminally Ill Act*, were all terminal cancer patients, and it is from this category that the majority of those seeking medically-assisted death could be expected to come.

(d) Sanctity of life

[41] The final objection to euthanasia I mention is based on the claim that all human life is sacred. As a simple unqualified claim this is much too strong to be taken seriously. If it were accepted then it would provide a powerful objection to passive as well as active acts of euthanasia, and that consequence is widely rejected. It is usually taken for granted that there is no serious moral objection to declining further medical treatment. In general, arguments about sanctity and sacredness need further bolstering with more substantial moral argument.

[42] The related suggestion that ending life is the prerogative of the deity is equally unhelpful precisely because divine prerogatives are conceived so disparately by different religious authorities. It is not helpful to be told that an action transgresses the prerogatives of the deity unless we know what these prerogatives are.

[43] Larue, for example, has recorded an immense diversity of religious opinion both between and within Christian and non-Christian denominations concerning the acceptability of euthanasia or physician-assisted suicide. There is in general no unified Christian, Jewish or other denominational position on this issue. The major problem with the claims based on the 'sacredness' of life or of 'playing God' is the danger that they operate as rhetorical devices to obfuscate rather than to illuminate discussion.

(g) Double effect

[44] Some opponents of euthanasia nevertheless allow that it is permissible to provide treatment to the terminally ill which hastens death. It is popularly, though not unanimously, believed that aggressive doses of morphine shortens life in many cases. When a patient's condition becomes sufficiently wretched it is permissible, on one view, for a doctor to end the misery with a lethal overdose of a pain killer, but this must be administered with the primary intention of alleviating pain even though there is also the foreseen but unintended consequence of killing the patient.

[45] The principle behind this view is the doctrine of double effect which draws a sharp moral line between what is intended and what is merely foreseen but which is not intended. Those who apply the principle usually have a robust confidence about its application. It is a confidence however which is not universally shared. It rests in particular on a dubious assumption that in analysing a complex act with multiple consequences (such as death and the alleviation of pain) we can identify which of these consequences are intended and which are not intended but merely foreseen.

[46] But how are we to distinguish, in any given act with complex or multiple consequences (relieving pain and killing a patient), what is intended from what is not intended but nevertheless foreseen? While the doctrine of double effect appears to provide a plausible reason for drawing a distinction in some cases there are many others where it delivers no unambiguous or unproblematic answer. Philippa Foot has suggested that the morally relevant distinction operating in cases in which the doctrine of double effect *appears* to provide guidance is actually the quite different distinction between negative duties not to harm others and the less stringent positive duties to give others aid. One does not in general have the same duty to help people as one does to refrain from injuring them.

[47] The distinction between what is really intended and what is merely foreseen, to which the doctrine of double effect appeals, was mercilessly lampooned by Pascal in his dazzling attack on some disreputable uses of casuistry. Pascal's attack on the methods of casuistry (moral reasoning through the use of cases) as an instrument of moral reasoning is an intellectual *tour de force* of such magnitude that it has taken more than three centuries for a serious attempt at its revival.

7. The issue of power

[48] A significant point raised by Bob Dent's physician, Philip Nitschke, is that a pivotal and underlying issue in the debate is the empowering of patients. The Northern Territory legislation, according to Nitschke, embodied a shift of power from the medical profession to the patient, and according to Nitschke the medical profession does not welcome any such shift of power. There is a clear insinuation in Dent's letter that physicians are prepared to help their patients put an end to their misery but the timing of the lethal overdose is to remain squarely in medical hands.

[49] There is evidence that unwillingness to give up power exercised over patients is a widespread feature of medical practice. Earlier we noted evidence from a recent medical conference of the propensity of physicians to disregard the advance directives of patients, indicating the reluctance of physicians to relinquish the power which they exercise over their patients. Another study which examined the views and experiences of physicians at a major Canadian teaching hospital regarding the use of advance directives in clinical care found that 40 per cent of the physicians questioned chose a level of care different from that requested in advance by patients who subsequently became incompetent. The physicians interviewed indicated that they would only follow a patient's advance directive if it was consistent with their own clinical judgment and indicated that they wanted to reserve the right to make clinical judgments about treatment regardless of a patient's request.

[50] Nevertheless, sometimes it seems not so much a case of an autonomous doctor exercising sovereignty over the disempowered patient as a doctor being as much a victim of the social and cultural circumstances in which the medical treatment of the patient takes place. This appears to be the case for example in the relief reported by a specialist upon learning of objections of family members to futile surgery which he had proposed, because he felt that in the absence of such objection he was ethically bound to proceed.

[51] This suggests that there may be cases in which physicians are acting against both their own wishes and against the wishes of patients and their families, because the parties are too reticent to express their preferences, and too prepared to submit to what they take to be conventional expectation. Perhaps such cases are instances of so-called 'defensive medicine' the overzealous or excessive provision of medical treatment to preclude any possible (moral or legal) charge of negligence. Whatever the reason, the complex and still unresolved issues surrounding the euthanasia debate point once again to the paramount importance of the effective communication of patients' hopes, values and aspirations to the health professionals who provide their medical care.

How to Argue Against Active Euthanasia

DAVID BOONIN

Journal of Applied Philosophy, Vol. 17, No. 2, 2000

[1] Let me begin with a confession: I am not really an opponent of active euthanasia. In fact, this paper began its life as a tentative defence of active euthanasia to be given at a public forum on the subject in the autumn of 1998 at the University of Colorado. But at the last minute, the one person on the panel who was planning to oppose the practice at the session had to cancel, and since that left the panel imbalanced, the chair of the session asked if I would instead present an argument against it. Since the chair of the session was also the chair of my department, and since I don't have tenure, I said, "Sure, that sounds like a great idea".

[2] In fact, it sounded like a very poor idea. But it turned out that being conscripted into the role of the critic of active euthanasia was valuable for two unexpected reasons. One was simply that it forced me to acknowledge that even though I do not really consider myself an opponent of the practice, I do have more qualms about it than I might initially have suspected. The other is that, in trying to locate an argument against decriminalizing active euthanasia that I felt I could recommend, even if only for the sake of the argument, I came to believe that the critic of active euthanasia is strongly disadvantaged by the fact that contemporary discussions in applied ethics generally take place in the context of what I would broadly call a modern as opposed to an ancient approach to moral theory. I came to believe, that is, that if there is a strong argument to be made against the decriminalisation of active euthanasia, it will be the sort of argument that is generally associated with classical rather than contemporary philosophy. It is this second belief that I want to examine in this paper. I will begin by saying a little more about the distinction that I have in mind between modern and ancient approaches to moral theory. I will then explain why I suspect that any argument against active euthanasia that is of the first sort is bound to fail. Finally, I will propose an argument of the second sort that I believe has a better chance for success. I am not myself prepared to endorse the argument outright, but I am concerned to defend the comparative claim that such an argument is more promising as a justification for policies forbidding active euthanasia than are the alternatives, and that this is how to argue against active euthanasia. If I am correct about this, then this may prove to be significant both as a means of making progress on the problem of euthanasia in particular, and as a way of showing that the neglect of the ancient approach to moral theory among practitioners of applied ethics in general can lead to importantly misleading results.

I

[3] The contemporary literature in moral theory is largely characterized by the debate between consequentialist, deontological and contractarian moral theories. In several respects, these are obviously importantly different moral theories. But in one respect, they are essentially the same. On any of these accounts of morality, one argues against the claim that society should sanction a particular practice by demonstrating that the practice involves actions that are forbidden by some principle justified by the theory itself. If the act itself is ruled out by the theory, then there is sufficient warrant for condemning the practice, but if the act itself is permitted by the theory, there are few resources remaining with which to construct a moral criticism of the practice. This is characteristic of what I will call the

modern approach to ethics. An opponent of capital punishment, for example, might argue that executing criminals has a desensitising effect on society, thus ultimately leading to more negative consequences than positive ones, that acts that produce more negative consequences than positive consequences are morally forbidden, and that therefore capital punishment is immoral. Or she might argue that executing criminals violates their right to life, that acts that violate the right to life are morally forbidden, and that therefore capital punishment is immoral. Or she might argue that self-interested agents agreeing to principles of behaviour in a suitably constructed social contract would choose to forbid capital punishment, that acts that violate principles that such agents would have agreed to under such circumstances are immoral, and that therefore capital punishment is immoral. Such arguments are substantively very different, but structurally the same. In each case, one argues that it would be morally wrong for society to practise capital punishment because (according to principles justified by the moral theory) society does not have the moral right to practise it.

[4] It is an unfortunate measure of the influence of this approach that some people seem to believe that this is the only kind of argument that can legitimately be aimed against a given practice. If a person can show that he has the moral right to do something, this is all too often taken as if it shows that there is nothing morally wrong with his doing it. This is the kind of thing that people like Jerry Springer say, for example, when people criticize his show: I have the right to broadcast it, you have the right to turn it off, and there is nothing more to be said. But, as anyone with even a cursory appreciation of ancient philosophy would recognize, this view is deeply mistaken. Morality can be concerned not simply with the moral status of the particular actions that people perform, but also (or, on some views, instead) with the moral status of the characters of the people who perform such actions. We might agree, for example, that people have the right to tell racist jokes, while still maintaining that morally decent people would refrain from exercising that right.

[5] A second and substantially different kind of argument that can be offered against a given practice, then, is representative of what I will broadly refer to as the ancient approach to ethics: it concedes (or, at least, does not deny) that society has the moral right to engage in the practice, but then tries to show that the sort of people we should strive to be would nonetheless refrain from exercising that right. A critic of capital punishment, for example, might agree that capital punishment is a cost-effective deterrent to crime, or that murderers have forfeited their right not to be killed, and she might therefore accept (or at least not deny) the claim that society has the moral right to execute them. But she might still argue, for example, that merciful people would refrain from exercising that right, and that merciful people are the sort of people we should strive to be.

[6] When applied to the question of whether or not the practice of active euthanasia should be decriminalised, this admittedly rough and very general distinction yields two distinct strategies that a critic of the practice might pursue: he can argue that society has no right to permit doctors to engage in active euthanasia, or he can argue that, even if society does have such a right, the sort of people we should strive to be would not exercise it. I want to argue that if there is a strong argument against society's permitting doctors to engage in active euthanasia, the argument will be of this second sort. I will begin by saying something about why I am sceptical about the possibility of any arguments of the first sort succeeding and will then present a version of the second kind of argument that may have some reasonable chance of success.

II

[7] So first: what sort of argument can a critic of active euthanasia construct if he limits himself to the theoretical tools of modern moral philosophy? One possibility would be to offer a consequentialist argument against the practice. But this seems to me very unpromising. When a doctor agrees to administer a fatal dose of morphine to a patient who is suffering interminably, the consequences seem to be plainly better than when she refrains from doing so. They are better for the patient, who is put out of his suffering, better for the patient's family, who are spared the torment of watching a loved one suffer, and better for other patients, who will have better access to their doctors when their doctors are not also being kept busy caring for patients who do not wish to go on living.

[8] Of course, a resourceful consequentialist might challenge this claim. He might appeal, for example, to the difficulties of calculating precisely which cases are such that killing the patient produces more overall good than does declining to kill him. A policy of permitting active euthanasia might therefore involve killing some patients when doing so would not produce more good overall, and thus the overall effects of the policy might be more harmful than beneficial. But this is a weak response. If such calculations are prohibitively difficult here, they are likely to be prohibitively difficult in almost any context. Indeed, if there is any situation in which it seems clear that one act would produce more overall good than any alternative, it is the act of honouring the request of a suffering patient who wishes his misery to end. So the appeal to the difficulty of doing the calculating in this instance seems to amount simply to abandoning consequentialism.

[9] A second consequentialist strategy that a critic of active euthanasia might employ involves appealing to a kind slippery slope argument: if we allow doctors to kill patients who want to die, it might be argued, this will eventually lead to doctors killing patients with comparable conditions even when they don't want to die. This gambit strikes me as unsatisfactory for two reasons. The first is that it rests on an empirical claim that is at best not clearly true, and at worst significantly implausible. We allow doctors to administer painful chemotherapy treatment to cancer patients who want them to, after all, but this does not lead us to allow doctors to impose such treatments on patients who decline them. We allow boxers to administer painful blows to one another when they consent to this, and this does not lead us to allow them to impose such treatment on people who decline to consent to them. There are, in general, all sorts of things that we allow people to do to one another when they consent to such treatment without then starting to impose comparable treatment on others involuntarily, and it is not at all clear why things should be any different in this case.

[10] The second reason for rejecting the slippery slope argument is that even if it could be shown that a policy of voluntary active euthanasia would inevitably lead to the practice of involuntary active euthanasia, it is not at all clear that the practice of involuntary active euthanasia would itself be objectionable on consequentialist grounds. The loss to those whose lives are already very impoverished might well be offset by the resulting gains to those for whom valuable resources, both medical and otherwise, might be made available as a result. Most of us, of course, will find the reasoning behind this second objection itself objectionably callous, but this is merely evidence that most of us are probably not consequentialists. And this provides a yet further reason to doubt that consequentialism can provide a sound basis for the opposition to active euthanasia.

[11] At first glance, a deontological approach to the issue may seem much more promising than a consequentialist one from the point of view of the critic of active euthanasia. The difference between passive euthanasia, which virtually everyone accepts as morally permissible, and active euthanasia, which is far more controversial, is that passive euthanasia involves letting the patient die while active euthanasia involves actively killing

him. And it is a central feature of the debate between consequentialists and their deontological critics that while consequentialists maintain that killing and letting die are morally on a par, deontologists maintain that killing is, all things equal, worse than letting die. Thus, it might seem, the resources of deontological moral theory can be used to show that there is a morally relevant difference between active euthanasia and passive euthanasia, and that even though passive euthanasia is morally acceptable, active euthanasia is not.

[12] But this appearance is the misleading result of ignoring two morally important distinctions. The first is the distinction between harm and benefit. The deontologist's claim that killing is morally worse than letting die is merely one instance of the more general deontological claim that causing harm is morally worse than allowing harm to happen. Just as it is worse to push a man into the water than to refrain from tossing him a life preserver should he fall, that is, it is worse to cut off his arm than to refrain from preventing it from being cut off, to make him sick than to refrain from curing him should he become sick, and so on. The disagreement between consequentialism and deontology, that is, is a disagreement about whether the fact that a negative result is caused rather than merely allowed is itself morally relevant. But nothing about the deontologist's claim that this distinction is, indeed, a morally relevant one implies that causing a result is worse than allowing it if the result is a good one. If being stuck in the arm by a sharp needle will harm you, say by introducing a poison into your body, then on the deontologist's account, it would be worse for me to stick you with the needle than to refrain from preventing you from being stuck with it. But if being stuck in the arm by the sharp needle will benefit you, say by providing you with a vaccination for some terrible disease, then nothing about the deontologist's position implies that sticking you with the needle would be wrong assuming that allowing you to be stuck by the needle in the first place would not be wrong. And if we are talking about a genuine case of euthanasia, then causing the patient to die will be causing a benefit to her, not a harm. So the basis for the deontologist's doctrine of killing versus letting die, the doctrine of causing harm versus allowing it, cannot be used to underwrite an argument against active euthanasia.

[13] A second reason that deontological moral reasoning cannot be used to underwrite opposition to active euthanasia arises from a second distinction, that between acts that treat people in ways that they consent to be treated and acts that treat them in ways that they do not consent to be treated. When the deontologist claims that causing harm is worse than allowing harm to occur, she has in mind cases in which the harm has not been consented to. When she claims that, for example, it is worse for me to burn your house down than to refrain from preventing it from being burned down, this is all with the tacit assumption that you have not hired me to burn your house down. If you have agreed to have me burn your house down, perhaps so that you can then use the land for something else, then nothing about the deontologist's position implies that my burning your house down would be morally worse than my standing by and refraining from preventing it from burning down. But the debate about active euthanasia is a debate about voluntary active euthanasia. No one is defending the claim that people should be euthanized against their will. If the critic of active euthanasia is to make a compelling argument, it must be an argument against voluntary active euthanasia. And, once again, the basis of the deontologist's position cannot be used to underwrite such an argument. I do recognize that the debate about active euthanasia versus passive euthanasia is often carried out in the context of the debate about killing versus letting die, but this seems to me to be unwarranted and an extremely unfortunate conflation of two distinct debates. Deontological moral theory seems to me to provide no more support for opposing active euthanasia than does consequentialism.

[14] What about contractarianism? In one respect, determining the practical implications of contractarianism is more difficult. Unlike consequentialist and deontological approaches, which are best understood as providing moral principles that are to be followed, contractarianism is best understood as providing a mechanism for generating moral principles which in turn are to be followed: the moral principles to be followed are those that would be agreed to by rational agents attempting to select principles for the regulation of behaviour that would be in their own interest, but making their selection under conditions that prevent them from unfairly exploiting certain sorts of information about their own particular circumstances. Since agreement about the conditions under which the agents must be understood as making their choices does not in itself ensure agreement about the nature of the principles that they would actually choose, there is room for a substantial divergence of views about what contractarianism would imply about a given moral issue. In the case of the problem of active euthanasia in particular, however, I believe that the verdict is nonetheless clear, and that it is also, once again, unlikely to provide support for critics of the practice.

[15] On what is perhaps the most typical account of the substantive output of moral contractarianism, an agent choosing principles of behaviour from behind a veil of ignorance would reason roughly as follows: if I knew which particular individual I would be, or even what racial group I would belong to, I would select principles that permitted me, or members of my group, to harm others for my own benefit. I would reasonably expect to do better under such a system of rules, than under one that treated all individuals equally. But since I do not know such facts about myself, selecting rules which permit such exploitation would be unacceptably risky. Agreeing to a rule that permits members of race A to enslave members of race B, for example, would turn out to be good for me if I turn out to be a member of race A, but disastrous for me if I prove to be a member of race B. And so I would instead select rules that protected each person from being exploited by others. On this account, the moral principles that are generated by contractarianism are essentially deontological in structure. And since I have already argued that such principles can provide little support for the critic of active euthanasia, this amounts to showing that on this account, moral contractarianism itself can offer little support for such a critic.

[16] Some people have argued, however, that the moral principles generated by the contractarian approach would be very different. On this account, an agent choosing from behind a veil of ignorance would reason more like this: I want to choose those principles of behaviour that will maximize my own expected utility. If I knew which particular person I was going to be, I would select principles that would direct everyone to maximize the wellbeing of that particular person. But since I don't know which particular person I will be, I should instead select those principles that will maximise the average level of wellbeing in society. That way, the amount of happiness that I can reasonably expect to enjoy will be greater than it would be on any other scheme. On this account, the moral principles that are generated by contractarianism are fundamentally different. Rather than providing restrictions on the right of one person to harm others, they instead direct agents to do those acts that would contribute to the greatest overall good even if they do involve the impositions of such harms. The substantive output of contractarianism, in short, would be essentially consequentialist on this account. And since I have also argued that consequentialism can provide little support for the critic of active euthanasia, the result is that on either way of cashing out the substantive principle of moral contractarianism, contractarianism itself can offer little support for such a critic. Consequentialist, deontological and contractarian moral theories represent the core approaches of modern moral philosophy, and I therefore conclude that modern moral philosophy can do little to ground opposition to the practice of active euthanasia.

III

[17] Let us now assume that I am correct about this, and that morally speaking, society does have the right to sanction the practice of active euthanasia. Is there any reason to think that an argument more along the lines of classical moral theory might better serve the position of the critic of active euthanasia? To ask this is to ask if there is any reason to believe that even though society has the right to sanction the practice of active euthanasia, the sort of people we should strive to be are the son of people who would refrain from exercising such a right. I want to suggest that there may well be such a reason, and to present it in the form of a two-part argument. The first part of the argument maintains that sanctioning the practice of actively helping to kill people who want to die involves endorsing a certain view about the value of the life of the individual being killed. And it claims that this view is not endorsed simply by sanctioning the practice of allowing such people to die by withholding or withdrawing life support. The second part of the argument then maintains that even if we have the right to endorse this view about the value of such lives, the son of people we should strive to be would refrain from doing so.

[18] Let me begin with the first claim. Suppose that you are a doctor with two patients, Larry and Moe, who have identical and very serious conditions and both of whom wish to die. Larry is on life support and asks you to pull the plug. He says he believes his life is no longer worth living, that it is worse than no life at all, that it has no value. Suppose you agree to pull the plug. Does this involve your endorsing the view that Larry's life is of no value? I think the answer is no for (at least) two distinct reasons. The first is this: to say that it is morally permissible for you to pull the plug on Larry is simply for you to say that you have no positive moral obligation to keep Larry alive given that he doesn't wish to remain alive. And you can certainly believe that you have no duty to force a person to remain alive against his will, even if you believe that his life has value in general, or value for him in particular. The claim that a person's having a particular good would be of value to him in particular or to society in general, after all, surely does not imply that you therefore have a duty to force him to acquire or maintain his possession of that good. You can agree that you have no positive moral duty to force a person to continue his education, for example, even if you believe that his education has value in general, or value for him in particular. The second reason for thinking that your decision to pull the plug on Larry does not involve endorsing Larry's claim that his life is no longer worth living is this: medical resources are not unlimited. Given this fact, there is nothing obviously objectionable about saying that priority in allocating resources should be given to those who wish to go on living. If we can do more good by helping others who want to live, then we can certainly redirect valuable resources away from Larry while still thinking that had he wanted to go on living, his life would have been one of value and one that we would have worked hard to extend. So in agreeing to honour Larry's request that we allow him to die, we need not agree to accept Larry's assessment of the value of his life.

[19] Now let's consider Moe. Moe also says that he believes his life is no longer worth living, that it is worse than no life at all, that it has no value. He wants you to give him a lethal injection of morphine to put him out of his misery. Suppose you agree to give him the injection. Does this involve your endorsing the view that Moe's life is of no value? I believe that the answer here is yes. In choosing to kill Moe, you cannot truthfully say that you are simply declining to provide him with scarce resources that someone else wants more. Nor can you truthfully say that you are merely declining to force him to go on living against his will. Instead, you must concede that you are deliberately destroying something that, at least under most circumstances, has tremendous value: the life of an innocent human being. And it is difficult to see how you could be willing to do this if you believed that the life

you were ending was a valuable one. If you were willing to do this, after all, then you would seem to be committed to the view that it is morally permissible for you to kill a person provided only that he gives his consent to your killing him. And this would imply that you would give a lethal injection to a perfectly healthy patient who asked for one if he came in with a minor headache complaining that his life was no longer worth living because his girlfriend had dumped him. But I find it highly unlikely that you will believe that this would be morally permissible. And if I am right about this, this provides strong evidence that you would be willing to help to kill someone only if both of two conditions were met: (1) they gave their consent to your killing them, and (2) you agreed that their life was not worth living.

[20] I have argued so far that sanctioning the practice of helping to kill Moe at his request involves endorsing the view that Moe's life is of no value, while sanctioning the practice of refraining from forcing Larry to remain alive against his will does not involve endorsing the view that Larry's life is of no value. Let us now assume that this is true. The question then is: is there a good moral reason for us not to endorse this view, even if we have the moral right to do the act that would amount to endorsing it? I want to suggest two distinct reasons that a critic of active euthanasia might appeal to here.

[21] The first reason concerns the fact that we live in a pluralistic society in which people disagree strongly, deeply, and sincerely about whether or not it is true that every human life is intrinsically valuable regardless of its quality. Some people believe that life is only of value to the extent that it contains more happiness than unhappiness. They might readily agree that Moe's life has no value, and so might find nothing objectionable in the thought that actively bringing about Moe's death involves endorsing this view. Others, however, believe that life itself (or, at least, human life itself) is of value even when it is a life that we would strongly hope never to have to lead. They believe very strongly that Moe's life is precious even if Moe no longer values it himself. They would thus be deeply offended by an act that amounted to endorsing the view that his life has no value.

[22] What we have here is thus a fundamental disagreement about basic values. And it seems very plausible to say that in a liberal, pluralistic society, toleration of such disagreement should be viewed as a virtue. Put in public policy terms, people should not use the power of the state to promote one value over the other unless there is a very strong reason to do so. Sanctioning the practice of active euthanasia amounts to declaring that one belief about the value of life is the correct one the belief that life is only of value when it contains more happiness than misery. But refraining from sanctioning the practice does not amount to endorsing either view. It simply amounts to remaining neutral about the question. Since this last claim may not seem clear at first, perhaps an analogy will be of some use. So consider the debate about the preservation of endangered species. Some people believe that each species in itself has a unique and irreplaceable value. Other people believe that this is not so. Suppose we wanted to formulate a policy that would be neutral with respect to this debate. Certainly we would not propose a policy that would allow the killing of members of endangered species, since that would clearly amount to endorsing the view that such animals have no such value. Nor would we select a policy that would require people to save animals from extinction, since that would just as clearly amount to endorsing the view that they do have such value. Instead, it would seem, the closest we could come to adopting a stance that was tolerant and neutral with respect to this dispute would be to adopt a policy on which people are obliged not to kill such animals but are not obliged to save such animals. That is the policy that tolerant people would endorse, unless there was a very strong reason not to. I am suggesting that the same would hold in the case of the question of the intrinsic value of each human life. In the absence of a very strong reason to the

contrary, tolerant people would endorse a policy on which people are obligated not to kill those who wish to die, but are not obligated to save those who wish to die. And this counts as one reason to believe that even if we have the right to sanction the practice of active euthanasia, the sort of people we should strive to be would refrain from exercising this right.

[23] The second reason for believing this is, I think, potentially even more powerful. Suppose that there are many other people who have precisely the same condition that Moe has and to precisely the same degree. And suppose that these people believe that their lives are worth living and they want to go on living as long as possible. By sanctioning the practice of helping to kill Moe, society is not only implicitly endorsing the view that Moe's life is no longer worth living. It is also, in effect, telling these people that it does not believe that their lives are worth living either. Remember, it must be saying that in order to avoid the implication that we would think it acceptable to kill anyone who asks us to kill him, no matter how clear it might seem to us that his life as a whole will be well worth living. There seems to me to be good reason to find this a powerful consideration in favour of the critic of active euthanasia, and the reason can be brought out most effectively by considering a more extreme analogy.

[24] So suppose that some blacks and some Jews felt that their lives were not worth living merely in virtue of the fact that they were black or Jewish. And they asked us to help them to die. Suppose that the state came up with some guidelines: we'll go ahead and help to kill you if you meet the standards, and if you don't, then while we won't force you to remain alive against your will, we won't assist in ending your life either. And suppose, in this case, that the result was this: we'll help you to die if you are black, but not if you are Jewish.

[25] One thing that should be very clear is that such a policy would be an enormous affront to black people. It would make them feel that they were living in a world in which society believed that their lives were of no value, or were of less value than those of other people. And notice that it would be of very little comfort to them for the state to say: don't worry, if you are black and you still want to go on living we aren't going to come around and kill you against your will. It's just that if you want to die, we'll help you do so, while we won't do that for Jews. I am not fully confident that I can identify some one virtue in particular that would definitively lead us to find the propagation of such a message toward black people repellent. But surely a virtuous person would find it so. Even if we believed that the adoption of such a policy did not, strictly speaking, violate anyone's moral rights, we could not bring ourselves to endorse it if we were compassionate, caring, decent human beings.

[26] Yet now substitute for blacks, people with advanced cases of AIDS, or various other serious diseases, or with extreme disabilities, who believe that their lives are worth living despite the grave obstacles they face. And conjoin with it a policy by which society permits doctors to actively kill people with such conditions when the patients consent, but which does not simply permit one person to kill another whenever such a request is made. The result here, just as surely as in the case of the policy regarding blacks and Jews, is a policy which must deeply offend those people with serious illnesses or extreme disabilities who believe that their lives are nonetheless worth living. If the sort of people that we should strive to be could not bring themselves to endorse the former policy in virtue of its giving offence to one group, then it should equally be the case that such people could not bring themselves to endorse the latter policy, and for the same reason. And since the latter policy just is the policy on which active euthanasia is permitted in some circumstances, it follows that virtuous people would not support such a policy.