[NOTE to the Jurisprudence Discussion Group: I have not had time to improve this paper by making it shorter. The result is that it is very long. For those who would prefer to read some but not all of it, I suggest secs. 3-5 (pp. 8-31), where I tell the four tales. For those who want to read something yet shorter, I suggest secs. 4.3 (pp. 19-21) and 6 (pp. 31-36), where I get to the philosophical heart of the matter. For the others, I'll give a quick summary that leaves nothing essential out.]

Four Tales About Fairness

Douglas MacLean

There is a well-known exercise called "The Ultimatum Game," which is used in business schools to show how a sense of fairness affects decisions in ways that cannot be explained by standard economic theory.[1] The game involves two players. One is given a sum of money and instructed to give some amount – however much or little he chooses – to the second player. The second player can decide only to accept or reject the offer. If she accepts, each player keeps the sum of money he or she has; if she rejects, then all the money is returned to the instructor, and the players keep nothing. Standard economic theory predicts that a rational second player will accept any amount she is given, since any positive gain is better than nothing. It follows that it would not be irrational (in the narrow sense familiar to decision theory) for the first player to give the second player as little as possible, keeping as much as possible for himself. It turns out, however, that unless the second player gets a percentage she deems to be fair – around 40 percent, the data suggest – she is likely to reject the offer. This exercise is used to demonstrate to

[1] See W. Guth, R. Schmittberger, and B. Schwarze, "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization* 3 (1982): 367-388; also R. Forsythe, J. Horowitz, N.E. Savin, and R. Sefton, "Fairness in Simple Bargaining Experiments," *Games and Economic Behavior* 6 (1994): 347-369; also G.E. Bolton, "A Comparative Model of Bargaining: Theory and Evidence," *American Economic Review* 81 (1995): 1096-1136.

business students that fairness matters, even when it involves some cost in wealth or welfare.

Some psychologists and behavioral economists have more recently conducted further studies to try to show in detail how our common intuitions about fairness affect our decisions.[2] The results of these studies are complex, but they confirm that most people express concerns about fairness that are incompatible with the maximizing assumptions of neoclassical economic theory.

1. Disagreements about Fairness

However interesting these results may be as a description of our attitudes and behavior, it isn't clear what normative implications they have. To address that issue, we would need to know not only what people's intuitions about fairness are, but also whether they are justified or true. We would need a normative or philosophical account of fairness. This shouldn't be a difficult task. After all, philosophers have been discussing fairness and justice for as long as philosophical ethics has existed. Here the problem is not that philosophy lacks a theory of fairness but that it has more than one theory, the leading contenders are incompatible with each other, and this situation has persisted for centuries. How can we explain this kind of persistent normative disagreement? My goal here is to answer this question.

There is another kind of normative theory of choice that must also be concerned about fairness, at least to the extent that fairness is important to people at a descriptive level. This kind of normative theory, which is closely tied to economics, is decision theory or policy analysis. The tendency here is to try to sidestep or avoid moral

---

[2] D. Kahneman, Knetsch, J., & Thaler, R., "Fairness as a Constraint on Profit Seeking: Entitlements in the Market" *The American Economic Review*, 76 (1986): 728-741; also D. Kahneman, Knetsch, J., & Thaler, R., "Fairness and the Assumptions of Economics," *Journal of Business*, 59 (1986): S285-S300.

controversy. One attempt to avoid these issues is to be pragmatic about fairness. One leading book on the subject, for example, claims at the outset that, "even granting that equity might not be entirely subjective, there is no sensible theory about it, and certainly none that is compatible with welfare economics."[3] Despite a lack of agreement in theory, however, people manage to carry on reasonable discussions about fairness. We discuss inequities in tax policies or in health care delivery; we debate the fairness of trade policies and affirmative action guidelines; we discuss alternative voting or representation schemes and how they can be procedurally fair or more equitable in their outcomes; we debate the importance of inequities in income and wealth, fairness in access to public goods and services, and so on. In many of these areas, moreover, we manage to find acceptable solutions, and the policy analyst can take these solutions as data to support some general principles that most people can agree to about the meaning of fairness and its implications.

Although I would not agree that there is no sensible theory about fairness or equity, the fact that we do not seem to agree on the correctness of any particular theory yet manage nevertheless to respond constructively to concerns about fairness in many important areas suggests that there is a lot to be said in favor of a pragmatic approach to fairness that does not try to settle the deeper normative disagreements. But some decision theorists also go beyond advocating pragmatism and argue that their theories of choice are in fact ethically neutral. For example, Ralph Keeney has argued that decision analysis can be made "consistent with any of the major ethical theories," and Ron

---

[3] Peyton Young, *Equity: In Theory and Practice* (Princeton: Princeton University Press, 1993), p. i.

Howard claims that, "since the formalism of decision analysis is amoral, like arithmetic," no moral commitments are made by using formal techniques to model fairness.[4]

Part of my aim in this paper is to challenge this assumption of neutrality. If philosophical differences about the meaning of fairness rely on different assumptions, and if these assumptions are incompatible with each other but jointly exhaust the space of logical possibility, then there may be no neutral ground on which to stand. Of course we can try to avoid this problem by attempting to measure different individuals' preferences or their willingness to pay for policies that strike some people as fair and others as unfair, but this is not a neutral way of making decisions. Nothing short of arguing for the truth or justifiability of some conception of fairness can resolve this kind of difference. And this remains true, even if philosophers after debating the issue for twenty-five centuries continue to disagree.

I will characterize what I see as the main differences among the contending philosophical perspectives about fairness. Although I will indicate briefly near the end where I think the truth most likely lies, my aim here is not to defend a philosophical theory of fairness but to show instead why different philosophers are drawn to one or another of the different and incompatible conceptions of what fairness means. I want to show where they agree and where they disagree on underlying principles.

2. Fairness, Equality, and Moral Philosophy

Philosophy is not content to find a few generalizations about fairness or justice to which all people can agree. Its ambition is not limited to enumerating the intuitions we all accept or describing some general principles that these intuitions may support.

---

[4] Ralph Keeney, "Ethics, Decision Analysis, and Public Risk," *Risk Analysis* 4 (1984): 117-129; Ron Howard, "An Assessment of Decision Analysis," *Operations Research* 28 (1980): 4-27.

Rather, philosophy's goal is to mine the foundations of the concept in the hope of uncovering something that might establish which of our different conceptions of fairness is correct. Philosophical theories aim to show how our intuitions can be defended or justified. Do they have a foundation in reason or human nature, or in anything more secure and basic than mere agreement in opinion based on experience?

The philosophical aim to reach foundations is not limited to ethics. For example, no philosopher genuinely doubts the existence of the external material world, but philosophers do wonder whether our belief in this world can be given a rational justification or has a basis in anything more secure than our mere psychological inability to doubt its existence for very long. Can we establish that our representations of the world correspond to anything outside our experience? David Hume famously argued that we cannot, but he also diagnosed the disposition that demands a non-skeptical solution to such doubts. He writes, "Philosophers, that may give themselves airs of superior wisdom and sufficiency, have a hard task when they encounter persons of inquisitive dispositions, who push them from every corner to which they retreat, and who are sure at last to bring them to some dangerous dilemma."[5] Hume was well aware that these persons of inquisitive dispositions are often other philosophers or simply oneself in a philosophical moment. The concern about foundations and justification – not any real doubt about the world's existence – that generates and sustains the problem of philosophical skepticism.

Similarly, philosophical disagreements about fairness have less to do with whether we share some set of intuitions than with how these intuitions are justified. Two main differences between moral or political philosophy, on the one hand, and

---

[5] David Hume, *An Inquiry Concerning Human Understanding*, Sect. IV, part ii.

epistemological skepticism, on the other, are these: First, we often disagree about the meaning of some basic moral concepts like fairness, whereas we pretty much agree on the meaning of the external world; and secondly, we are if anything less content to accept a "skeptical solution" about what constitutes fairness,which would license us to rely on our intuitions, because our intuitions about fairness differ, and the arguments surrounding them have important practical implications. It is not as easy with moral doubts as it is with doubts about the existence of the external world simply to let the fly out of the fly-bottle.

A philosophical account of fairness does not attempt to explain what some or most people believe about fairness, and it is not concerned with how people are in fact led to accept one or another view of the subject. A philosophical account of fairness is *normative*, not descriptive. It makes claims about what people *ought* to believe and how they *ought* to act. Its aim is to justify such normative claims. What distinguishes a philosophical account of fairness from descriptive theories is its focus on justification, i.e. the reasons or arguments for believing that some view is true. This is different from a causal account, which explains why people find some view attractive or psychologically irresistible.

I assume that everyone today accepts some idea of moral equality. There is disagreement, of course, about what equality means, but all normative conceptions of fairness presuppose some conception of moral equality. When some individual (or group) is accused of treating another individual or group unfairly, any attempt to explain or justify this action will either appeal to some idea of equality or else explain why equality in this particular situation is not required. The assumption of equality rules out

any attempt to respond to charges of unfairness by claiming that some people are simply not the moral equal of others or that some individuals are not full-fledged persons deserving equal moral consideration. These explanations are ruled out as incompatible with the idea of moral equality that we all share. This is why those who defend practices that many people regard as racist or sexist must attempt to show that such practices in fact show equal respect for all people or their rights. They cannot defend such practices by claiming that some people naturally deserve less than others.

It wasn't always so. The idea of equality is a modern idea, and one of the most significant differences between ancient and modern conceptions of justice is that the ancients tended to accept claims of natural inequality which justified fundamentally unequal treatment of different kinds of people. Some theories of justice were based on the idea that people should receive what is naturally their due, and that an individual's proper place in the social order, which include differences in rights and opportunities, should be determined in whole or in part by natural differences that were taken to be morally relevant. It is a mark of progress in morality that most people today regard natural differences as being, from the point of view of justice, morally irrelevant.[6] Some philosophers argue that the idea of moral equality should be extended beyond human beings to include at least the higher non-human animals, but there is not yet a consensus on this point. We do have consensus, however, on the idea that all human beings are in some sense morally equal. People understand moral equality in different ways, however, and different conceptions of fairness rely on these different interpretations of equality. This is what I want now to explain.

---

[6] See Stuart Hampshire's account of this difference in "A New Theory of the Just Society," *The New York Review of Books*, Feb. 24, 1972.

3. Equal Treatment Theories

We can begin by distinguishing between conceptions of fairness that depend on equal treatment alone, and those that do not. Equal treatment theories interpret fairness as requiring procedures for making decisions and policies that treat all the individuals involved in the process as equal in some respect. The fairness of the process thus determines the fairness of the outcome. A fair outcome is any outcome that results from procedures that are constrained by the relevant conception of equal treatment. Equal treatment theories, as I am characterizing them, see no role for fairness to operate on outcomes. They see no direct value in principles of distributive justice. Our first distinction, therefore, is between equal treatment conceptions of fairness and conceptions of distributive justice.

3.1. First Tale: Libertarianism

One equal treatment theory focuses on individual rights. The basic and essential moral property, which all normal people equally share, is the capacity to act freely to determine our own lives and control the world in which we live. Rights protect people in the exercise of their freedom over their natural entitlements, in particular their bodies and their property. The conception of fairness that follows from this view of equality is libertarianism, which is a view that gives priority to the protection of rights and insists that justice is historical.[7] This means that in order to determine whether some particular distribution of benefits and burdens in any situation is fair, we need only to know how it came about. If the situation resulted from actions that did not violate anyone's basic and equal rights, then the result is fair; otherwise, it is not fair and needs to be corrected.

---

[7] See Robert Nozick, *Anarchy, State and Utopia*, (New York: Basic Books, 1974).

The libertarian conception of fairness thus focuses entirely on how people are treated, i.e. on whether their rights have or have not been violated, and does not recognize as legitimate any principles of distributive justice that look only at the pattern of the distribution of benefits and burdens at a particular time without regard to how the pattern was produced. Unfairness on this account is not a matter of harm or disadvantage but only a failure equally to protect individual rights.

To illustrate this conception of fairness, suppose that Aziz invests everything he has and as much as he can borrow to open a restaurant, and a few months later another person opens a restaurant right across the street. Suppose that this new restaurant is in all respects better than his. It is decorated in ways that are more to customers' tastes, it has a better chef, and it is run more efficiently, so that the cost of a comparable meal is less than what Aziz can afford to charge. After a year of losing money, Aziz is forced to close his restaurant. He has lost his entire investment, is deeply in debt, and his life turns in every way for the worse. Assuming that the owner of the newer restaurant has operated within the law and obeyed all the regulations, then even though the superior restaurateur has caused Aziz to suffer great harm, he has done nothing unfair. In a capitalist system that protects everyone's rights equally, no injustice is committed. There may be harm, but on a libertarian conception of fairness there is no wrong. Wrongness, and more specifically unfairness, is determined solely by the violation of rights.

The simplicity of the structure of libertarianism makes it normatively appealing to some people, so one way to indicate why many philosophers reject this conception of fairness may be to mention some of the difficulties that such an account glosses over. It is unclear, for one thing, how a conception of fairness that is restricted to equality in the

initial allocation of rights can treat issues of fairness in risk-imposing activities.[8] If imposing a risk (let us say a small risk of death) on someone violates that person's rights, then it would seem that libertarianism requires that all such risk-imposing activities must either be prohibited, or they may be justified only with the explicit and informed consent of the victim, such as is needed by a surgeon prior to operating on a patient. The magnitude of the risk should make no difference. But this seems to be an absurd result, for the requirement of prior consent would mean the virtual elimination of many essential activities in modern societies. People could not drive automobiles without the consent of each pedestrian walking nearby, and power plants with any harmful emissions would never be built, because each person downwind of the pollution-emitting smokestack would effectively have a veto power over the construction of such a power plant.

In response to this objection, a libertarian might claim that risk-imposing activities do not by themselves violate any individual's rights, but the agent of such activities must be prepared to compensate victims if unallowable harms result. This alternative, though perhaps less absurd than the suggestion that risk imposition does violate a person's rights, nevertheless still seems unacceptable. Most of us would insist that at least *some* risk-imposing activities should be prohibited. Nobody should be allowed, for example, to fire weapons at well-attended sporting events, even if they are prepared to compensate anyone that happens to sustain be injured as a consequence. Robert Nozick indeed suggests a more interesting moderate solution to this problem, but it isn't clear that any solution, which would involve some detailed restrictions on rights

[8] Nozick has a fascinating discussion of this problem, one of the best to be found in contemporary political theory. See *Anarchy*, pp.

for the sake of achieving some highly desirable outcomes, is available by appealing only to the conceptual resources of libertarianism.

A second issue for libertarians is that entitlements alone fail adequately to address some important questions about fairness. There may be goods to allocate that are not privately owned, e.g. irreducibly public goods or a newly discovered good whose value is not a result of anyone's labor being mixed with it. Or there may be situations in which people have conflicting rights. Even libertarians need procedures for allocating scarce life-saving equipment, and so on. If two libertarians want to play chess, they must find a fair way to determine who shall play the white pieces. The solution to such problems may seem intuitively obvious enough, but it isn't clear how it can be expressed by appealing only to the structure of rights. If libertarianism is to explain our idea of fairness, therefore, then it needs to appeal to concepts other than rights, and it is not obvious that these additional concepts will not appeal in the end to some principles of distributive justice.

Finally, few people have the stomach for libertarianism in its purest form. Imagine a society in which all entitlements have been properly determined, people start out with equal opportunities, and everyone acts justly. Suppose that through luck, including luck in the genetic lottery that allocates the natural goods of intelligence, skill, and motivation, some people end up doing much better than others in increasing their share of wealth and other goods. Imagine further how these differences, once begun, might be magnified over time and passed on across generations. At some point the differences could become so great that the situation itself, regardless of the history that produced it, cries out morally for some intervention to change the pattern. Some sort of

redistribution, through taxation or land reform, will strike most people as a requirement of, well, justice. That is to say, I think that most of people feel some pull of egalitarianism, which extreme situations may make most salient to us. If this is true, then few philosophers will unreservedly embrace pure libertarianism, even though they may worry that striking a more moderate pose will make them inconsistent with their own emphasis on rights.[9]

3.2. Second Tale: Utilitarianism

A second equal treatment view, classical utilitarianism, is famously unconcerned with distributive justice. But it would be a mistake to think that this theory does not include a doctrine of fairness. According to utilitarianism, the morally right thing to do is determined by what maximizes well-being or net good, which is a measure of the value of the effects of any act or policy on all individuals, to the extent that these effects can be known or measured.[10] Utilitarianism is a consequentialist theory: what matters ultimately in determining right and wrong are the consequences of our actions. Right and wrong are determined by the aggregation of effects on all people, taken together. Utilitarianism is not directly concerned with distributive justice because it is not directly concerned with how these effects are distributed to different individuals.

Now, a theory with this structure, i.e. a possible consequentialist theory, could be radically inegalitarian in the way it measures the value of the consequences of alternative acts or policies on different individuals. It is possible, for example, to formulate a

---

[9] This point is developed in Thomas Nagel, "Libertarianism Without Foundations," *Yale Law Review* 85 (1975).

[10] J.S. Mill, *Utilitarianism* (London: Collins, 1861); Henry Sidgwick, *The Method of Ethics* (London: Macmillan, 1907). For a brief contemporary statement of the egalitarian assumption at the core of utilitarianism, see Peter Singer, *Practical Ethics* (New York: Cambridge University Press, rev. ed. 1993), ch. 1.

consequentialist theory in which the effects on women, or the aristocracy, or Shiite Moslems, were the only ones considered or in which the effects on members of these groups were weighed more heavily than the effects on others. Following the spirit of the U.S. Constitution, a consequentialist theory could count each white person as one individual and each African American as 3/5 of an individual. Such an inegalitarian moral theory would be outrageous, but it is not incoherent.

Classical utilitarianism rules out such theories by insisting that equal consideration is to be given to the interests of each person and, in the words of John Stuart Mill, "not to them only, but so far as the nature of things admits, to the whole sentient creation."[11] Each sentient being counts for one and none for more than one. Utilitarians insist that the experiences of different people can at least roughly be compared with respect to pleasure and pain. We need to weight these effects equally and take them all into account in determining what is the best or right thing to do. Utilitarianism thus makes a strong commitment to equal treatment, interpreted as equal consideration of the effects of any action or policy on the experiences of all people.

The distribution of effects *per se* has no value because the distribution does not show up in the experience of any individual, and what is best or right for the utilitarian is determined by aggregating these individual effects. But of course no plausible moral theory can fail altogether to take account of anything as morally basic as a concern for fairness in the distribution of benefits and burdens. So the utilitarian must give an indirect account of distributive justice in order to ward off criticisms of his theory which claim that equal consideration of interests is not a sufficient response to demands for

[11] *Utilitarianism*, Ch. II.

fairness. He typically has two stories to tell, and the adequacy of utilitarianism as a moral theory depends to a considerable extent on the plausibility of these stories.

The first story begins by acknowledging that people care about being treated in ways that they regard as fair and that perceptions about fairness contribute to a person's overall happiness or misery. Feeling that one is being treated unfairly or even that one's lot is much worse than that of others may foster resentment and other negative attitudes. A more equitable distribution of goods may lessen these feelings and thus lead to greater happiness overall. Now, the critic of utilitarianism will object to this reply that it responds only to the effects of different attitudes and not to their merit. The utilitarian wants only to reduce overall unhappiness, but this might be accomplished as effectively with drugs or brainwashing as with satisfying demands for receiving a more equal or just share of the benefits of social cooperation. The utilitarian has no means for assessing the merit of different attitudes and responding to those that are reasonable while rejecting others. He must treat the interests of the bigot and the spiteful as on a par with the interests of the reasonable or saintly. The critic of utilitarianism will point out that the utilitarian reason for claiming that rape is wrong and should be forbidden is that harm to the victims outweighs benefits to the rapist, but the critic will complain that benefits to rapists deserve no consideration at all.

The second story the utilitarian tells is about the decreasing marginal utility of extra resources. The assumption is that an amount of resources given to people who are worse off produces more happiness or well-being than the same amount of resources given to people who are better off. To the extent that decreasing marginal utility is a deep and pervasive fact of human life, utilitarianism will turn out to be egalitarian not

only in the way it considers each person's interests, but also in the resulting distribution of resources aimed at maximizing overall happiness. A more equal distribution of resources is likely to produce greater well-being overall. The critic of utilitarianism also has doubts about this story. One is that even if decreasing marginal utility is a pervasive fact of life, it is not a universal fact, and it may turn out that the next dollar will bring more happiness to the rich person who is then able to buy his dream sports car or vacation house than it will bring to the poor person for whom the difference of a dollar will affect what she eats. A more important but less remarked upon objection to this utilitarian story is that even if decreasing marginal utility were a universal fact of life, it is surely true that marginal utility decreases at different rates in different people. Individual utility functions will look different, even if they are all concave in the manner required for decreasing marginal utility to be true. In that case, however, overall utility is maximized not by equalizing the distribution of resources, but by distributing resources in a manner that equalizes the marginal rates of utility. So if marginal utility decreases at different rates, then even if it is constantly decreasing for all people, any distribution of resources that will maximize total well-being will still require some redistribution from those who are less well off to those who are better off.[12] This will strike the critics of utilitarianism as unfair.

4. Comparing Equal Treatment Theories

Of course just as there may be more complex libertarian theories than the one I described, there are more complex versions of utilitarianism that might give fuller or more adequate attention to the demand for fairness. My interest here is not in developing

---

[12] See John Broome, *Weighing Goods* (Oxford: Basil Blackwell, 1991), pp. 175-177.

more complex theories about fairness but in comparing the structure of theories that express very different conceptions of what fairness means. For this kind of comparison, highly simplified theories better illustrate the structural differences.

4.1. Similarities between Libertarianism and Utilitarianism

The similarities between libertarianism and utilitarianism are not often remarked upon, so the similarity in their regard for fairness or equity might seem striking. Both theories express equality at the beginning, so to speak, in their commitments to treating people equally. They are not directly sensitive to concerns for fairness in outcomes that might support arguments for a redistribution of resources, wealth, or income in response to claims of distributive justice.

Both libertarianism and utilitarianism also tend to favor market mechanisms for achieving the right allocation of benefits and burdens, although for different reasons. Libertarians like markets because of the freedom of action they allow to individuals, permitting them to exercise maximum control over their property and other entitlements. Utilitarians favor markets as a mechanism for making the preferences or values of different people comparable and for promoting efficiency in the allocation of resources.

4.2. Differences between Libertarianism and Utilitarianism

Libertarianism and utilitarianism are nevertheless importantly different in the way each expresses the requirement of equal treatment. This difference can be illustrated with an example that has been much discussed in recent philosophical literature. Suppose we must choose between saving one person's life and saving the lives of five other people at equal cost. (Assume that there are no other ethically relevant differences between them, and we have no obligations or special relationships to any of them.) A utilitarian would

say that in this situation it would be best to save the five lives, because each life has equal value and saving five lives is therefore five times better than saving one. The libertarian would deny this. He would claim that we have no better reason for saving five lives than we have for saving the one life.[13] The libertarian reasons as follows. None of the people has a claim that we should save *him* if we cannot save all of them. If we save the one, therefore, we violate nobody's rights, just as if we saved the five we do not violate the rights of the one. Moreover, any complaints that each of the five people who would die could make if we saved the one do not get added together to become a bigger complaint than the complaint of the one who is left to die if we save the other five.

Because the libertarian understands rights to be claims that must be satisfied to the fullest extent feasible and not balanced against other rights and interests in order to maximize well-being, a person's right to life cannot imply a right to be saved under all circumstances. If I can save only one life or five other lives, each person has an identical claim to be saved. These claims cannot all be satisfied, but neither can they be balanced and added up. This is why it would be no worse, according to the libertarian, to save the one instead of the five, and this is why it would be fairest to flip a coin to decide what to do. Each person would thus have an equal chance of being saved. The right thing to do in such a case is the same whether five or fifty lives are at stake. Each has a claim of equal strength as, but no stronger than, the claim of the one. For the libertarian, the numbers do not count.

If the libertarian argument here seems unconvincing, as it does to many people, then consider a different example. A libertarian would regard it as an unacceptable

---

[13] See John Taurek, "Should the Numbers Count?" *Philosophy & Public Affairs* 6 (1977): 295-316. For a reply to this article that defends an argument similar to the one I attribute here to the utilitarian, see Derek Parfit, "Innumerate Ethics," *Philosophy & Public Affairs* 7 (1978):

violation of rights to confiscate one person's property merely in order to use it to satisfy some important interests of five other people. He would consider it unfair and wrong to remove without consent one person's kidney in order to use it to save another person's life. To confiscate property in order to use it to increase benefits to others would be to treat someone's rights as if they were merely interests, which could be outweighed by the presence of enough other people who might have an even greater interest in his property. If someone needed a kidney in order to survive, and if you would likely survive the loss of one of your healthy kidneys, then arguably the other person has a greater interest in your kidney than you do. But to remove your kidney for this reason would violate the libertarian sense of fairness in the deepest way. A person's rights support claims she has over the use of her property that must be treated differently from concerns we have for the interests that others may have in her property. Interests may be compared and added together, but rights claims cannot be balanced in this way. They must each be satisfied, to the fullest extent feasible. This is how we explain the common thought that rights protect an individual against the interests of the majority. For the libertarian, an innocent person's right to his own body is the reason we forbid killing him in order to redistribute his organs to save the lives of some number – let us say five – of other people.

In sharp contrast to the libertarian, the utilitarian interprets fairness in terms of equal consideration of interests and insists that interests should be compared, so that the maximum good can be achieved. The utilitarian gives an answer that matches most people's intuitions in the example about saving the different number of lives. But the utilitarian has a harder time explaining why it would be wrong to kill an innocent person in order to use her organs to save the lives of five other innocent persons, just as the

utilitarian has a hard time explaining why it would not be permissible to harm another person any time that doing so would produce a greater benefit for others.

4.3. Different Conceptions of Moral Justification

These examples can help us comprehend the basis of the main difference between libertarian and utilitarian conceptions of fairness. Both theories interpret fairness to require equal treatment, but they have very different views about what equal treatment requires. This is because they have different views about how to reason about or justify moral claims. Utilitarians take equal treatment to mean that everyone's interests – all of them – should be counted together to determine the best action or policy. Every person's interests should count the same as every other person's similar interests in determining what is best. Fairness for the utilitarian requires impartiality with respect to whose interests they are. The utilitarian therefore reasons about what is best from an impersonal perspective which takes every person and all their interests into consideration together.

The libertarian, in contrast, emphasizes the separate importance of each person. She must figure out a way to reason about what to do that gives each person the respect that the libertarian conception of fair treatment requires. The libertarian thus rejects the idea that the interests of different and separate people should be combined and weighed from some impersonal perspective. For the libertarian, the collective entity represented by such an impersonal perspective does not share everyone's interests and experiences together. The collective entity does not share your happiness or feel your pain. Rather, the collective entity experiences nothing, and to reason from this perspective is to disrespect the separateness of each person, who is a locus of experience and a source of

choice. Individuals, not the collective entity, have basic moral rights that must be equally respected.

The libertarian does allow interests and experiences to be weighed and compared, but not across different people. Thus, if one person suffers for two hours, she suffers more than if she were to suffer for only one hour. But if two people each suffer for one hour, then although more people suffer than if only one had suffered, there is no greater amount of suffering than that experienced by either one of them. The experiences do not add together to become a worse suffering of the group that consists of the two of them. If a hundred people each enjoy watching a movie to an equal extent, this does not make the movie any more enjoyable than if only one person had watched it. The larger theater does not produce any more enjoyment; it simply allows for similar enjoyment to be experienced by more people. The utilitarian says that the larger theater is better for this reason. The libertarian says that this claim makes no sense. "Better for whom?" she asks. "Surely, not the theater!"

The libertarian reasons from the individual perspective of each person, and she tries to find principles of fairness that give equal consideration to each person as a separate entity, not as elements in some larger whole. To find such principles, she must try somehow to take into account the point of view of every person simultaneously without combining them. She might therefore test a candidate principle for acting by putting herself in the place of each person and testing whether the principle could not reasonably be rejected by any one of them.[14] She thus aims at unanimity of agreement to principles, where this is possible. If she cannot find some unanimously agreeable way of

[14] The most complete theory expressing this view of justification is worked out by T.M. Scanlon, *What We Owe to Each Other*, (Cambridge, MA: Harvard University Press, 1998).

distributing a good, then she reasons that nobody could claim that it is unfair to equalize the chances of getting the distribution that seems best to each individual. So she flips a coin to decide which life or lives will be saved.

5. Distributive Justice Theories

As the discussion so far makes clear, in much of our ordinary thought fairness seems to require sensitivity not only to the protection of individual rights or to an equal consideration of interests but also to the resulting distribution of resources or benefits and burdens. We should therefore consider some conceptions of fairness that defend principles that govern the distribution of valued outcomes directly.

5.1. Third Tale: Egalitarianism

One view of this sort holds that, other things being equal, a state of affairs is better when benefits and burdens are distributed more equally than when the same total is distributed less equally. We will call such views egalitarian. The essential feature of egalitarianism is that the *distribution* of outcomes has intrinsic value.

Different versions of egalitarianism may be distinguished both by the degree of emphasis given to the value of distributions and by the nature of the goods to be distributed. Regarding the latter kind of theories, some may be egalitarian with respect to welfare or happiness, while others may be egalitarian with respect to resources, opportunities, or some other good. I will not discuss these different versions of egalitarianism.

For the purpose of gaining a sense of the general structure of these theories, I will introduce a highly artificial device.[15] Let us suppose that the total of benefits, burdens, or whatever valued outcome will be distributed to each person affected by some action or policy can be added up and represented by a single number, such that a person with a higher number receives more of what is valued than a person with a lower number. Imagine a world that contains just two people, *A* and *B*.

Now, nobody would think that the world in which

A has 0
B has 0

is better than the world in which

A has 100
B has 99.

This example shows that any view that claimed that the distribution of goods takes absolute priority over all other considerations (including the total amount of good) is absurd. So egalitarianism must claim that there are (at least) two attributes that determine what is best, the total amount of good and its distribution. But to be an egalitarian as opposed to a utilitarian, one must also believe that the distribution of good has some independent weight, and this implies that it must be possible for some alternative or world with a lower (or equal) total of good to be better than an alternative or a world which contains a greater (or equal) total of good but in which this good is distributed in a worse or less equal way.

---

[15] My discussion of this and the following view owe much to unpublished writings of Derek Parfit, especially to his manuscript on Distributive Justice. My examples here are based on examples in his work.

An extreme egalitarian might thus favor a world in which benefits and burdens are more equally distributed even if nobody in that world is better off than in some alternative world. For example, the extreme egalitarian might judge the situation where

A has 99
B has 99

to be better than the situation where

A has 100
B has 105

because the good in the former world is more equally distributed than the good in the latter world. A more moderate egalitarian might want to resist such an extreme implication by insisting that a more equal distribution cannot make a world better overall unless at least one person in that world is made better off by the more equal distribution than she would be in an alternative world with a less equal distribution. The moderate egalitarian would thus reject the claim in the above example that the former world could be better than the latter world. But the moderate egalitarian might nevertheless claim that the world in which

A has 100
B has 99

is better than the world in which

A has 99
B has 105

even though the latter world contains more total good than the former. If the distribution of benefits and burdens is to have any real value for a moderate egalitarian, then it must carry some weight against less equal distributions that would otherwise add up to a higher or equal total value. This is the essence of the difference between egalitarianism and utilitarianism.

Most utilitarians believe that, as a matter of fact, less unequal distributions of well-being will increase the total, both because feelings of resentment will be reduced and because of the effects of the decreasing marginal utility of resources. An egalitarian believes that more equal distributions have value whether or not they increase total well-being.

A further difference between egalitarianism and utilitarianism is that for the egalitarian, the reason a situation with more equal distribution of good is better cannot be reduced to the negative feelings of people in that world who believe that their situation is unjust or unfair. Like the utilitarian, the egalitarian may take such feelings into account, but the egalitarian finds intrinsic value in the distribution itself. This is the essential feature of his theory.

With respect to the distinction between impersonal justification and justification to each person, egalitarianism obviously is included with utilitarianism in the class of theories that regard justification as impersonal. This is most clear in the case of radical egalitarianism. If a more equal distribution of good has intrinsic value, even in a world in which a more equal distribution is not better for any of the individuals than in an alternative world, then any justification for such a claim must come from an impersonal perspective. The same applies *mutatis mutandis* for the moderate egalitarian. For even if a more equal distribution is good only on condition that at least one person is made better off in the more equal situation than she would be in the less equal one, it is still the case that moderate egalitarianism allows that one situation or world might be better than another overall even if it has less total good. This implies that if at least one person is better off in the more equal situation, then at least one person is also worse off in that

situation. Unless there is a reason for giving more weight to one person's good than to another's good, the value of greater equality must be impersonal. If equality is an intrinsic good, its value can be appreciated only from an impersonal perspective.

Of course, not all theories of distributive justice are egalitarian. A merit-based theory would identify some morally relevant property – e.g. desert, contribution, or need – and argue that justice requires unequal distributions of resources that correspond to the morally relevant differences among people. We could, in an Aristotelian sense, regard these views of justice as calling for distributions that equalize some ratio, where the numerator is the value of the outcomes to be distributed, and the denominator is a measure of the merit-based property that is the criterion of justice. But this way of extending egalitarianism seems artificial, especially since these theories of distributive justice are likely to favor unequal distributions of benefits and burdens, if the inequalities are properly arranged.

In a less artificial sense, however, which differentiates theories according to the two distinctions we have introduced above, egalitarianism and merit-based theories of distributive justice are similar. Both theories regard fairness as an independent value placed on the distribution of benefits and burdens, such that it is possible that an outcome with less good overall may be better than one with more good overall, if the distribution more closely matches the criterion of distributive justice. The reasoning and justification for egalitarian views and for merit-based views will be from an impersonal standpoint. That is, both theories claim that some intrinsic value lies in the pattern of the distribution of benefits and burdens. This value cannot be seen from any individual's perspective.

5.2. Fourth Tale: Deteriorism

Our final tale about fairness considers an alternative account of distributive justice. It focuses especially on those who end up worst-off as a result of the natural lottery and social policies. This is a natural place to focus attention, because intuitively those who are worst-off seem most likely to have the strongest claims to have been treated unfairly. We will call a conception of fairness that emphasizes the plight of the worst-off deteriorism.

Now those who are worst-off might also be thought to be those who have lost the most under the terms of the social contract, but this is not always the case. For if a society decided that the best and fairest way to improve the situation of its worst-off members was by instituting a highly progressive tax on income or wealth, it is likely that those at the top would lose the most as a result of such a policy. Of course their wealth may owe much to begin with to the existence and enforcement of social cooperation. In any event, claims of unfairness tend to be taken more seriously when they are also focused on those whose situation is also in other respects bad. But the fact that those near the bottom may not be identical to those who have lost the most means that more has to be said in order to explain why deteriorism has a claim to be a plausible conception of fairness.

Deteriorism maintains that the worst-off have a reasonable complaint of unfairness to make if alternative social rules and arrangements would have made their lot better in a way that would not also result in some other group's situation being worse than theirs happens to be. The idea is that the hypothetical social contract rationalizes the terms of cooperation within a society and justifies the organization of the basic institutions that make cooperation possible. These institutions determine to a significant

degree how benefits and burdens are distributed. Therefore, the terms of the contract ought to be justifiable to each of the parties to it, i.e., the members of society.

Of course victims do not always have a reasonable complaint to make about their status. If someone starts out with an equal share of goods and opportunities, deliberately gambles and loses, then his outcome may be deserved. But this is often not the reason that some people are much worse off than others. The more typical situation is that outcomes are determined largely by luck – one's natural abilities, the resources one inherits, the situation into which one is born and raised, and the random ways that some people's efforts are rewarded while others are not, and so on – in ways that seem deeply unfair.

Now one possible way of focusing our attention and efforts on the worst-off would be to give greater weight to their interests in the aggregation of individual well-being. We could count improvements in the position of those near the bottom for more than we count improvements in the positions of others who are less badly off. This way of formulating a conception of fairness, however, is simply a variation of utilitarianism which rejects utilitarianism's principle of equal consideration of interests. When compared to utilitarianism, moreover, this way of focusing on the interests of the worst-off seems *ad hoc* and unfair. Why should we give the interests of the worst-off greater weight than the interests of others? Why not favor those with blue eyes, or the aristocracy?

The argument for the plausibility of deteriorism, therefore, must rest on an assumption of equality at some deeper level. Victim-based views thus assume that all people are equal in some fundamental moral sense, which means that the idea of what

people morally deserve must be consistent with this assumption of equality. To the extent that the distribution of benefits and burdens in a society results from natural differences between people or from luck in the social lottery, it cannot be said to be morally deserved. If a fair society is conceived in some sense to be a system of voluntary cooperation among equals, then unequal distributions of benefits and burdens need to be justified.

Deteriorism claims that the worst-off often have a legitimate complaint to make about their situation to the extent that their situation is not deserved, and fairness demands that we pay special attention to such complaints. The principle of distributive justice that emerges on this conception of fairness will be something like a maximin principle for distributing the benefits and burdens of social cooperation. This principle gives priority to improving the situation of the worst-off group, because their claims are most deserving and their complaints the strongest. But it gives priority in the choice of a maximin decision rule, rather than by giving greater weight to the interests of the worst off.[16]

Because we need to create economic incentives and sanction other measures as a means of increasing social wealth and net benefits, even complete satisfaction of a maximin principle of distributive justice is likely to sanction unequal outcomes and thus leave some groups worse off than others. Deteriorism thus rejects egalitarianism and

---

[16] Victim-based views are particularly popular these days in American moral philosophy, owing largely to the influence of John Rawls. See John Rawls, *A Theory of Justice*, Cambridge, MA: Harvard University Press, 1971); also Rawls, *Political Liberalism*, (New York: Columbia University Press, 1993). Also see Thomas Nagel, "Equality," in *Mortal Questions*, (Cambridge: Cambridge University Press, 1979, ch. 8; T.M. Scanlon, "Contractualism and Utilitarianism," in A. Sen and B. Williams, eds., *Utilitarianism and Beyond*, (Cambridge: Cambridge University Press, 1982); and Scanlon, *What We Owe to Each Other*, op. cit.

instead favors principles of distributive justice that would raise the level of the worst-off group higher than any other alternative.

Deteriorism is not an equal treatment theory. The moral basis of victims' complaints is determined by comparing their outcomes to the outcomes of others. In other words, for any procedure that affects the distribution of benefits and burdens, the justification of the procedure is determined in part by the justice of the outcomes.

Another implication of deteriorism is that it appeals to a personal conception of justification, which considers the claims of individuals (or representative members of a group) taken separately. It does not consider claims of unfairness from an impersonal perspective that aggregates all outcomes together. Thus deteriorism, like libertarianism, must aim to satisfy each person's claims and respect each individual's rights to the fullest extent feasible. Justification aims at unanimous consent from each perspective, not at what seems best or fairest from some impersonal perspective that any person is capable of adopting.

Deteriorism shares a further feature with libertarianism, which is the refusal to aggregate benefits and burdens across individuals. Thus, it is also an implication deteriorism that the numbers do not count. We can illustrate this implication by resorting once more to the highly artificial device of representing the goodness of a person's situation with a number. Thus, in a world that contains just two people, *A* and *B*, it is possible that the situation in which

A has 15
B has 7

is judged to be more fair or just than the situation in which

A has 5
B has 10.

The former situation has more aggregate good or well-being than the latter, but this fact is irrelevant to the deteriorist conception of justification. Rather, because justification appeals to each individual separately, a deteriorist theory might say that the former situation in this example is more just than the latter because the worst-off person in the former is better off than the worst-off person in the latter. Of course *B*'s outcome is worse in the former situation than in the latter, so we might ask: Why does *B* not have a complaint of unfairness to make, similar to the complaint that *A* can make which leads us to judge that the alternative in which her outcome is improved is the fairer alternative? The reason is that being worse off is not itself a ground for complaint on any plausible victim-based view. Any distribution that departs from equality will leave some people worse off than others. What gives *A* ground for complaint if the second situation is chosen is that she suffers a fate worse than it is necessary for anyone to suffer. This is an appeal to fairness that *B* should accept. *B*'s complaint, in contrast, is simply that he suffers a loss under one policy, whereas someone else would lose under the alternative policy. And while we can understand the reasonableness of *B*'s unhappiness, he does not have a complaint based on fairness. Even where justification is interpreted as individually-based, any conception of fairness will exclude the identities of the individuals differently affected as being irrelevant to the fairness of the outcomes.

Finally, we should also notice the way in which, by refusing to let the number of affected people count in determining what is fair, critics of deteriorism accounts can point to implications which strain the intuitive plausibility of these theories. For suppose we alter the situation above so that we are asked to choose between an alternative in which

A has 5
100 people have 10

and one in which

A has 15
100 people have 7.

The implication of a pure victim-based view is that the latter alternative is fairer than the former. This is because none of the 100 people who are made worse off in the latter situation has a stronger complaint than *A* has in the former situation. And complaints, like well-being, do not become stronger as the number of people increase.

Examples like this surely present a challenge to defenders of deteriorism. No matter how sympathetic we may be to the claims of the worst-off person or group, at some point the costs to the rest of society of satisfying these claims may become so great that we may be tempted to back off from believing that this is what we are morally required to do. Victim-based theorists might want to temper their views accordingly.

Now, one way to temper the view is to say that fairness is not everything. At some point the costs of satisfying claims of fairness can be too great, and we must give more weight to aggregate well-being. To take this path, however, threatens to collapse deteriorism into some version of egalitarianism. How much weight should we give to fairness in making a social decision or policy, and how much should we give to increasing aggregate well-being? This is a question that only makes sense from a detached and impersonal perspective and is thus not a question that can arise within a deteriorist theory as I have characterized it.

6. The Incompatibility of the Different Accounts

Perhaps we should say that fairness is not all that matters in deciding what to do. We should treat fairness as one relevant consideration which must somehow be compared and weighed against other considerations in figuring out what is best overall. If this is

what we think, then we may wish to adopt some neutral or pluralist framework that can incorporate fair treatment along with the justice of outcomes, or which can incorporate individual-based justification with impersonal justification. Having told four tales about fairness, I can now point out the difficulty of trying to satisfy such a hope.

Some years ago, Peter Diamond noticed that equity was a problem for classical utility theory, which is essentially a formalization of a consequentialist outlook.[17] Diamond argued that if only the aggregate good was valued in outcomes, then utility theory had no way to address the value of different ways of realizing those outcomes. There is nothing in the outcomes alone which, from an impersonal perspective, would make one process for realizing an outcome better than any other. In a situation where one of two persons must die, to use his example, utility theory does not distinguish between simply selecting one of the persons and using a procedure, like flipping a coin, which would give each person an equal chance of surviving. Diamond argued that, other things being equal, fairness in this situation demands distributing equally the *ex ante* risk of death. The demand for equity in this situation can be satisfied, he claimed, by equalizing expected utility, which could be achieved by a coin flip. This conclusion is now widely accepted by decision analysts.

What is philosophically most interesting about Diamond's argument, however, is that he accepts the decision analytic approach of assigning value impartially to outcomes but then interprets the demand for fairness as something that can be met by equalizing the value of *expected* outcomes. In his example, this means assigning value to each individual's *ex ante* risk or expectation of dying, and equalizing this value. In other

---

[17] Peter Diamond, "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility: Comment," *Journal of Political Economy* 75 (1984): 624-632.

words, Diamond thinks that expected good has value. And because we can equalize this value in many cases where we cannot equalize *ex post* the value of the outcomes to all parties, he regards the value of expected outcomes as the proper focus for concerns about fairness.

Now this assumption begs several questions. First, it isn't obvious that it makes sense to regard expectations, in the decision theoretical sense, as bearers of value.[18] There is a sense of expectation, of course, that can be good or bad. Someone who expects something good to happen may derive additional pleasure in anticipating the arrival of the good. But in decision theory, an expectation is simply a probability, and it isn't clear that this kind of expectation has any value at all. A 50-50 chance to win a dollar or nothing, for example, has an expected value of half a dollar. If it is rational for a person with no aversion to risk to be indifferent between taking the gamble and pocketing the half dollar, this is solely because the gamble represents a chance at getting the greater good of a dollar. The expectation itself is worth nothing. If I accept the gamble and lose, I have received nothing of value at all.

A second way in which Diamond's argument begs an important question about fairness can be seen by noticing that some people think that justice may require compensating people who, through no fault of their own, are the losers and suffer the burden of some social policy or action. Now compensation may be impossible where death is the outcome we are distributing, but in other cases it is not obvious that distributing equally the *ex ante* risk of losing satisfies the demand for fairness better than compensating people for the losses they in fact incur.

---

[18] See John Broome, "Uncertainty and Fairness," *Economic Journal* 94 (1984): 624-632.

A better explanation for satisfying the demands of fairness by equalizing the *ex ante* risk of dying may be available to someone who does not take an impersonal perspective on justification. A random procedure for selecting who will die, when death is unavoidable, is not something that has any intrinsic impersonal value, but it may be the best and often a fully adequate way of satisfying the equal claims of each individual that it would be unfair to select *her* to suffer the loss. A random procedure in such a case does not equalize the distribution of any good; rather, it equally satisfies each person's claim to be treated fairly.

Diamond's suggestion has more recently been developed in greater detail by other decision analysts as a way of incorporating demands for equity or fairness into an impersonal framework of justification that they assume to be morally neutral. One proposal for doing this is to develop an impersonal measure of the equity of different outcomes and adding the value of equity so considered as an attribute in an analytic framework that aims to maximize overall expected value.[19] I will ignore the interesting complexities of the axiomatic treatment of equity and instead make a simple observation. All these multi-attribute versions of utility theory turn out to be variations of one tale about fairness, viz., egalitarianism. While such theories might be able, with greater or less success, to model the results of applying other conceptions of fairness, they cannot capture the reasoning behind these other conceptions, so they cannot explain the justification of the results they seek to model. These multi-attribute utility models may nevertheless be useful information tools for policy makers, but the ability to model the

---

[19] See P. Fishburn, "Equity Axioms for Public Risk, *Operations Research* 32 (1984): 901-908; Ralph Keeney, "Equity and Public Risk," *Operations Research* 28 (1980): 527-534; and R. Keeney and R. Winkler, "Evaluating Decision Strategies for Equity in Public Risks," *Operations Research* 33 (1985): 955-970.

results of any of the rival conceptions of fairness is certainly not sufficient for claiming that the model is ethically neutral.

Recall our discussion earlier of the requirements of fairness in deciding whether to save one life or five other lives. The libertarian conception of fairness requires giving each person an equal chance of surviving, for example by flipping a coin. In his defense of the ethical neutrality of decision analysis, Ralph Keeney suggests that this libertarian condition of fairness can be captured in a consequentialist or egalitarian framework by using a more refined definition of outcomes. Thus, he suggests that if we accept the libertarian reasoning but also think, as he does, that saving more lives has greater value than saving fewer lives, even in cases where different lives are at stake, then the best or most highly valued outcome is one where we flip a coin to determine which group to save *and* as a result of the coin flip we end up saving the five lives. The least preferred outcome is one where we simply decide (without using any randomizing procedure) to save the one life. We would then rank the other two possible outcomes – no coin flip and save five lives; or flip a coin with the result that we save one life – according to the relative value we determine should be assigned to the randomizing procedure and to saving more lives.

Now this argument entirely misconstrues the reason why a libertarian "favors" a coin flip. She is not claiming that a randomizing procedure has some value, which can be weighed in an impersonal way against the impersonal value of saving more lives. Her view is that each person has a legitimate claim to equal treatment. Each has an equal right not to be sacrificed for the greater good of anyone else, taken individually or together. If each of them cannot be saved, then fairness demands respect for each of their

claims to an equal chance of being saved. Once these claims are met, the demands of fairness have been realized, each person is treated equally, and the right thing has been done. Keeney interprets such claims as if they were (merely) interests, which is not a neutral conception of fairness or equity.

The different conceptions of fairness are incompatible with respect to both of the dimensions that distinguish them. Equal treatment views are incompatible with distributive justice accounts. Libertarianism is a pure procedural view of fairness. This means that the just outcome is simply the one that is realized by treating all individuals in a way that equally respects their rights and legitimate claims. This is what Nozick means when he argues that justice is historical and cannot be respected by attempts to realize any pattern in the outcomes. Similarly, utilitarianism expresses fairness in the equal consideration of every person's (or sentient being's) interests. To assign any intrinsic weight to more equal outcomes is essentially to undermine the equal consideration of interests.

The different perspectives on justification are likewise incompatible. The impersonal perspective of justification implies that benefits and burdens can be aggregated across persons. But the essential element of the conception of fairness from an individual perspective of justification is that it is unjust and therefore wrong to sacrifice the interests or legitimate claims of any person for the sake of realizing greater benefits for others. Unfairness from this perspective is a failure of respect for persons.

7. Concluding Remarks

In order to understand why philosophical disagreements about the meaning of fairness or justice have such a long history and remain unresolved, it is necessary to

comprehend the underlying bases of this disagreement. I have argued that philosophical views in this area diverge along two different dimensions, each of which has conflicting poles of strong attraction. I believe that any prima facie plausible account of fairness can be located in one of the quadrants that these two dimensions define. Disagreement persists because the poles on each dimension are incompatible, and yet, as I have tried to suggest, each pole on both dimension has features that make it attractive.

We should not conclude, however, that differences of this sort simply show that moral philosophy, or even the part of it that concerns justice or fairness, is incapable of progress or the ultimate resolution of disagreements. The concept of equality was not a central part of ancient moral thought, but it is central to all conceptions of fairness today. This is one example of progress in moral philosophy, and it has changed the world in ways as profound as any scientific discovery.

Nor do I think we are permanently at an impasse in resolving differences in our conceptions of fairness. It seems to me that the individual perspective of justification expresses our evolving sense of fairness more adequately than the impersonal perspective of justification, and it seems to me also that our evolving sense of fairness cannot avoid taking seriously the claims of distributive justice and the demand to focus on those who through no fault of their own are worse off than others as a result of social policies and decisions. Someday, I believe, we will find that we agree that the true or most plausible conception of fairness lies in the deteriorist quadrant. But since the goal of this paper has been to explain why disagreements about the meaning of fairness persist and to describe the roots of that disagreement, it would be foolish for me here to try to defend one of

these views over the others. That would require a much fuller argument than I can give, and it probably would not move many critics in any case.