

Chapter 6

Justice as Efficiency

Pareto Efficiency, Justice, and the Free Market - A Pure Efficiency Conception of Justice

Introduction: Paretian concepts

The term 'efficiency' has, in this century, acquired a special use in social studies from the work of Vilfredo Pareto. A social state of affairs, according to this definition, is "efficient" or "optimal" if one cannot improve the situation of anyone in it without worsening the situation of someone else. A related (indeed, component) notion is that of Pareto-*improvements*: a social state of affairs, S2, is said to be such an improvement on another such state, S1, if and only if at least one person is better off and no person worse off in S2 than in S1. Such an alteration is said to be "Pareto-superior" to S1. An apparently still weaker concept says merely that no one is worse off, though possibly no one is better off either. Alterations in either sense are said to be Pareto-efficient, and that is the familiar sense of 'efficiency' in which I use the term here.

To lend a sense of verisimilitude to this discussion, let's first look at a good contemporary example of a professional economist regarding this matter. This author, Nicholas Barr,¹ defines 'economic efficiency' thus:

Economic efficiency [Barr's note 4: "Referred to synonymously as Pareto efficiency, Pareto Optimality, allocative efficiency, or external efficiency"] is about making the best use of limited resources given people's tastes. It involves the choice of an *output bundle*: $X^* = (X_1, X_2, \dots, X_n)$ (where X_i is the output of the i th good) with the property that any deviations from these

quantities will make at least one person worse off ... [T]he optimal quantity of any good, *ceteris paribus*, is that at which the value placed by society on the marginal unit equals its marginal social cost...²

Like the definition supplied in my opening paragraph, the definition just given is formal and indeterminate in certain crucial respects as it stands. Before it can be applied to real states of affairs, we must render the concept determinate, and this requires us to make decisions along two dimensions. In the first place, we must ask which conception of value we are to use. But this being for social application, this question might as well, for all practical purposes, be taken to be the question, *whose* conception, or criteria, of value we are to use, that is, whose judgments of value we are to accept as definitive for the purposes at hand. Professor Barr's characterization commits us to accepting that it is people's "tastes" that are in question. This is something many philosophers would quarrel with, but as we shall see below, I propose, essentially, to accept that view. And in the second place, we need a conception of that kind such that it is both logically and realistically possible for the criterion to be both applied and met: possible, that is, for it ever to be true, and capable of being known to be true, that someone *is* "better off" and no one else "worse off" in some situation or configuration, as compared with some other, for the purpose of assessing efficiency. Barr's discussion, referring us to "the value placed by society" on the marginal unit, clearly requires that we attach some definite significance to that idea if we are to apply the notion of efficiency. This question will occupy us shortly below.

Meanwhile, however, the reader may certainly wonder why these are said to be concepts of "efficiency" at all—it seems a bit of a stretch. But there is, I think, a good reason, and the idea is important. Efficiency in the most general ordinary sense has to do with the ratio between input and output in the context of production of some kind: the more output for a given input, the more efficient. It is, of course, most at home in mechanical devices, especially those powered by fuels such as gasoline, where the input is a continuously measurable quantity, such as gallons, while the output can be understood, at least in one major respect, as miles traveled, thus making it easy to get neat figures such as X Miles per Gallon. Other contexts are more difficult. A more efficient worker will be one who accomplishes more than the comparison worker of the desired task in the same time, perhaps, or with the same amount of effort. We can measure time easily enough, but effort is more difficult. Of course, as soon as there are multiple tasks and multiple products, it becomes very difficult indeed to quantify output. Here the role of money comes to the fore, along with market exchange. In other contexts, we carry on somehow, and eventually we might talk about aesthetic efficiency, legal efficiency, and so on. What enables us to do so in each case is some idea of inputs and outputs.

Thus Barr, in developing the notion of efficiency, proposes as the first of three conditions of social economic efficiency that of *Productive efficiency*:

Productive efficiency means that activity should be organized to

obtain the maximum output from given inputs. ... It is about building a hospital to a specified standard with as few workers as possible standing around waiting for something to do...³

But it could and presumably would also be about minimizing costs of production, which could, for example, include using vehicles that got more miles to the gallon. And where does it end? If a manager could improve the performance of his firm still further by working 13-hour days instead of 12-hour days, is it, then, “inefficient” if he foregoes this opportunity? Suppose he foregoes it at the request of his wife and family—does that count? One thing, at any rate, is clear: maximizing this kind of efficiency lies beyond the reach of economics. The best one can do with this aspect of the notion is to specify that each aspect of a productive process be “as efficient as possible” given the technology, management know-how, and so forth, of the day. And that is more like a wished-for ideal than a condition an economist, or anybody, could stipulate. Moreover, it is, as I shall point out below, dramatically contrary to a reasonable construal of the Pareto idea.

In any case, it is, as Barr notes, allocative efficiency that is more nearly what we would be after in talk of optimal social welfare. His next condition is *Efficiency in product mix*—“the optimal combination of goods should be produced given existing production technology and consumer tastes.” And what is that? As economists will do, Barr here says that “production is ... at the specific point ... at which the ratio of marginal production costs ... is equal to the ratio of marginal rates of substitution in consumption...”⁴ And how do we know when that point is reached? Here we go to the free activity of individual consumers, normally. And that is going to lead to a problem.

Efficiency in consumption, Barr’s third condition, “means that consumers should allocate their income in a way which maximizes their utility, given their incomes and the prices of the goods they buy—in formal terms, the marginal rate of substitution must be equal for all individuals.”⁵ Barr draws the familiar Edgeworth box⁶ for illustration, but we need not resort to such things for present purposes. We need only point out that there is a horrendous problem involved in applying the notion of utility-maximization for this purpose. Is it a fact about Sheila that she has “maximized her utility” when she buys Brand X rather than Brand Y, or for that matter when she decides not to bother with skin-cream at all, or when she decides to throw over her previous way of life and go into a convent? Will a professional economist be able to criticize her, in the latter event, for failing to do her duty and maximize her economic utility? The short answer is, of course, No. And the longer answer will be too, as I shall argue below. The concept that the rates of substitution for all consumers is equal is not verifiable.

Social Efficiency?

But what about societies—“social systems” as they might, rather misleadingly,

be called? What shall we say are “inputs” and what “outputs”? That question is closely enmeshed with fundamental questions of political philosophy, of course. However, on any view of the matter, there needs to be some sort of concept of social good, enabling us to say that some societies or some social institutions in a given society are better than others. And a notion of good will also give us some handle on an idea of costs. The costs may be described as “evils,” or less moralistically, simply referred to as costs. In any case, they are what you want to avoid, and if you can’t avoid them, then at least to minimize.

What makes the Pareto notions of special interest is that we are here focusing on the social aspect of society, on society *as such*. Each individual will have goals, purposes, ideals, and will have a corresponding notion of costs. But when we get to inter-individual relations, notoriously, the situation gets tricky. For one individual does not share another individual’s consciousness, his purposes, his capabilities, or his values. At best, he has a similar goal or purpose. But in general, a cost for Smith is not *necessarily* a cost for Jones.

The idea of Pareto efficiency can now be explained as follows. *Social* efficiency—the doing-well of Society as such—is a matter of how much interpersonal cost there is: to what extent do our actions impede each other, thus lowering our achievement of our various goals? In an inefficient engine, some fuel is expended merely in heating the engine rather than driving pistons; this impedes the operation of the engine somewhat, for it requires using still more energy to counteract the unfortunate effects of the heating. Fuel burnt with those effects is “wasted”; all we want out of it is forward motion of the vehicle. Somewhat similarly, a more Pareto-efficient society will be one in which people do better in relation to their expenditures of time and effort on trying to achieve whatever they want to achieve, insofar as they are less or not at all impeded in their pursuits *by* the impositions of others. And the presence of some who impede others creates a need for still others to spend time and energy combating the initial set of impederers. This idea fits somewhat, then, with ordinary efficiency. It is not an ideal fit, because a society could be perfectly efficient in the Pareto sense, and no machine, or indeed any individual human, can be that. (Let’s agree that in both cases it’s enormously unlikely anyway!) Moreover, it will surely be insisted, by many philosophers and probably by ordinary people, that it would be nice to have a corresponding notion about *benefit* to each other, and not just our non-harm to each other. We’ll say more about that below.

Why?

Why think Paretian? On the face of it, it sounds a plausible principle. That is because it looks to be a social application of a rudimentary principle of practical rationality. If we can have our choice between two bundles of goods, {G1} and {G1, G2}, the second bundle containing everything that the first one does plus one more good thing, and no bad side effects, then surely it is rational to choose the second? But that would be misleading. The status of the Pareto Principle is actually a matter of considerable controversy and some misunderstanding. What

is being advocated in these pages is by no means a truism, by no means self-evident, and certainly widely denied. Even so, I think it a powerful principle. To see why it is hardly self-evident, let us distinguish three quite distinguishable ideas about it.

1. Those who think the principle self-evident are perhaps thinking that it is a straight dominance principle: if x differs from y in being in some way better and in no way worse, then x is better than y . Whether we should say that that is self-evident or not I don't know. But whatever there is to be said for this plausible view, it is by no means enough to support the Pareto Principle, as will be made clear in the next two points.

2. As it stands, the principle means that we are not to do evil to person A in order to benefit person B. There is nothing logically self-evident about this, and it is not morally so either: to accept it, no matter how you understand the evaluations involved, is to go quite far out on a moral limb. It appears to deny, in particular, that the way to deal with A and B is to do what is "on the whole best" for them, just like that. For it says that if just one of the two, say A, is, on the whole, worsened *at all* by our doing x , then we are not justified in doing x even if it benefits B a great deal. As such, it behaves rather like what Nozick calls "side constraints."⁷ That is clearly controversial. Moreover, no one who has ever advocated the principle, I think, means it to be totally sharp-edged. If we can benefit B a lot by imposing a quite trivial cost on A, most of us will do it. But then, in doing so, most of us rely on A's being a benevolent enough person that he would be happy to pay the trifling cost if asked. And in the face of utter disasters, we will all suspend side constraints to greater or lesser extent. That is an important fact. Still, we have not yet arrived at the fully liberal version of the Pareto principle. As it stands, the principle still depends for its application on our idea of what is good, what bad, for individuals.

3. The liberal version of the principle says that we are to act in such a way that all affected persons *see themselves* to be no worse off, on the whole, in consequence of our having acted that way than in the status quo *ex ante*. Moreover, this is intended to imply that if consent, or approval, by A, of something involving A only, is forthcoming, then that is sufficient for presuming better-offness in the intended sense.

The requirement that at least one person's situation be actually bettered, and not merely not worsened, may be treated in either of two ways. We may eliminate it, and insist merely that no one be worse off; or we may simply *assume* that an agent always acts for something he takes to be good, so that the weaker requirement is superfluous. An agent, we assume, normally tries to achieve what she regards as a benefit, either for herself or for someone that *she* chooses to benefit in some way. She tries, in other words, to bring about what she takes to be some desirable effect on something or someone else; and so the requirement is pleonastically met by all rationally acting agents. At least, it is met in intention: we may fail to achieve the hoped-for benefit, but it is what we were *trying* to do, anyway. For this reason, then, we need not concern ourselves about a distinction between the weaker and the weakest formulation of Pareto improvements: everyone is always doing what betters things, in that agent's

view, relative to some other available choices.

Liberalism

In making decisions of the kind I have pointed to the need for, we inevitably engage in normative social philosophizing. It is too often said that talking in Paretian terms purports to be “value-free”. But in any context of policy decisions, that is obviously false, to the point that this brief denial really ought to be unnecessary. Though the Pareto criterion is used rather often for explanatory purposes in the social sciences, and there is still a perfectly reasonable, and much more than residual, reluctance to accept values as entering into explanation, my purpose here is not primarily explanatory, except in the sense of clarifying and explicating certain options in normative theory: I propose to identify a type, or rather, a portion, of normative theory and spell out some insufficiently recognized implications of what is widely, though not universally, agreed to be an important and plausible idea. And I propose to be quite forthright and self-conscious in making the value-judgments in question.

The two decisions I shall adopt for these purposes are, I think, in the first case straightforward and in the second case crucial to this or any coherent project that accepts the first. First, then, I shall identify the relevant assessments of better- or worse-offness with the *preferences of the individuals in question*—what Professor Barr counts as “tastes.” Alice will be declared to be “better off” for purposes of this theory if she prefers her new situation to the previous one, that is, if *she* believes that she is better off in it, whatever anyone else may think. And secondly, we will exclude *negatively tuistic* evaluations from our data-base of acceptable evaluations. This is a morally loaded move, not an arbitrary stipulation. It says that if a given policy, *p*, would benefit someone, *A*, by worsening the situation of some other person *B* who was the object of *A*’s hatred or resentment, say, then society may not account it a point in favor of policy *P* that it satisfies a desire of that kind. Suppose, for example, that Jane is envious of Alice, so that when Alice is better off in some non-other-regarding respect, Jane is *ipso facto*, in her own view, worse off: If we accept malevolent values as determinative of what is acceptable and unacceptable for social purposes, liberalism is impossible. For in that case, the occurrence of just *one* individual with a certain profile of values could make it impossible for *any* Pareto improvements to occur. Justice requires that we unhook individuals from each other, except in respects in which both such individuals agree to be hooked.

Suppose, for instance, that one individual is a “radical” egalitarian: She judges herself to be worse off if *anybody* is better off than *anybody else*; thus, only if that person’s situation is improved simultaneously and equally with the situation of everyone else in the society will she withhold her disapproval. If there are any persons with such profiles, that effectively renders it impossible for any Pareto-efficient changes to take place, since the condition is for all practical purposes impossible to meet. Liberalism, however, will account an

improvement in one person's situation in any respect that is *independent* of anyone else's as genuinely an improvement for the purposes at hand, without regard to the attitudes of others toward that improvement as such. Of course people may attach great weight to their nonindependent desires, including ones that are negatively tuistic. But the point of paretianism is that such desires are problematic in social contexts. As they stand, they can hardly provide a basis for the common good which, I hold—along with Aquinas and so many others—must be the basis of all reasonable law.

Between them, these two interpretations are, I believe, definitive of the normative outlook known as *Liberalism*.⁸ Injustice consists in worsening the situation of one or more others in respects in which those persons' levels of well-being or utility are intrinsically independent of those levels in others, or in which any other who are objects of the desires in question are not averse to being so; on the other hand, any situation that does not worsen anyone else's is acceptable on the score of justice, and in that sense, it is just. (As Hobbes says: "whatsoever is not unjust, is *just*."⁹)

What this means, specifically, is that we do not as such have any positive duties toward others in general—that is, we may not properly be forced to make positive contributions to the welfare of others. For to do so is to worsen one person's situation in order to improve another's.

Justice

Recall that our inquiry here is into justice in particular, and not into any and all normative assessments. What, then, is justice? The theory of justice is the theory of what society may do "in the way of compulsion and control," as Mill puts it.¹⁰ To say that a person's action is unjust is to say that those actions are of a kind that society should, or at least may, coercively intervene to prevent or punish, that is, to override the individual's own preferences in that instance. That someone else doesn't like the new situation so far as the situations of others are concerned has, on the liberal view of the matter, no bearing on its justice or injustice. For to allow that is to put control over the lives of the others into the hands of the person with those attitudes. We may say that such persons have an "attitude problem"—their attitudes are malevolent, or at least unvirtuous. But on the liberal view, we may not take them into account in making our assessments regarding Pareto improvements or the reverse.

Since we are here doing only what has come to be called "ideal theory," we ignore the necessary refinements for dealing with the need to rectify or prevent injustices, which would no doubt be very difficult in principle. However, there is no reason to expect a fundamental divergence from our results from that quarter. If force of law is used only to prevent or to rectify the previous use of force against innocents, then one can reasonably say that it does not really impose on person and property at all.

Rawls on the Supposed Inadequacy of Efficiency

It has become commonplace over the years to say, along with Rawls, that efficiency is inadequate to “account for” justice. Those who say this tell us, as if by way of mere reminder, that there are many different but equally efficient configurations or distributions of, say, a society’s assets which nevertheless differ, they claim, in the justness of those configurations. Consider this passage from Rawls:

It is not difficult to see, however, that ... this principle ... is an inadequate conception of justice ... There are presumably many arrangements of an institution and of the basic structure which are optimal in this sense. There may also be many arrangements which are optimal with respect to existing conditions, and so many reforms which would be improvements by this principle. If so, how is one to choose between them? It is impossible to say that the many optimal arrangements are equally just, and the choice between them a matter of indifference, since efficient institutions allow extremely wide variations in the pattern of distributive shares.¹¹

To be sure, it is not quite clear just which principle Rawls is there referring to. And, too, he talks of “Optimality”, which is related to but not identical to efficiency as discussed. An arrangement is optimal if no rearrangement *can* be made which is better for someone and worse for no one. Thus the situation in which one person is dictator to all others is the sort of situation that those who speak this way claim to be Pareto optimal. And if we assume that the dictator does not, in dictating, make others worse off, and that satisfying the dictator’s desire to dictate counts in favor of others doing what he dictates, then what he says is certainly right.

Rawls’s claim, then, is that the theory of justice in the contemporary sense of the term is underdetermined by Pareto Optimality. But it is not entirely clear how Rawls wishes to advance this criticism. Is he claiming that no conception of justice confining itself to Pareto efficiency can be coherently expounded? Or is he asserting a certain view about justice, as against other, rival views, claiming in effect that Paretian views are inadequate or implausible? There is some reason to think that it is the latter rather than the former. Be that as it may, theorists often hope to advance a view of that kind by showing that the alternative views really don’t make sense, don’t add up. My concern in this essay is to rebut that criticism, by arguing not only that a pure efficiency view does in fact make sense, but also that the claim of nonuniqueness which seems to be being advanced by Rawls in that passage is false.

His view, if interpreted in this latter way, involves, I shall argue, either a misunderstanding or a partisan view, or more likely both.

I shall now argue that there is exactly one view determined by taking Paretianism seriously, if advanced under the aegis of liberalism—that all alternatives require its violation in some respect or other. Whether the view in

question is also the right view is another matter, of course; though I am inclined to believe it is, that is not argued for here. For present purposes, it is enough to show that the claim that the pure efficiency view of justice is an “inadequate” view cannot be held on purely conceptual grounds, though it is surely reasonable to suggest that violations of Paretian efficiency are indeed *prima facie* objectionable, as Rawls himself implies.

It will be noted that Rawls talks of different specific “arrangements”, selection among which is not a matter of indifference. The reason, he says, for doubting that different optimal arrangements are equally just is that “efficient institutions allow extremely wide variations in the pattern of distributive shares”. Is this a good reason? It is, of course, *if* justice does indeed allow such wide variations. But what if it does not?

This question, in turn, leads to a related one: what does Rawls mean by an ‘arrangement’? From context, one gathers that he means by it, simply, a configuration or pattern of the components, which in this case is human actions, *no matter how it is brought about*. If that is what he means, though, then there is indeed reason to deny the implicit claim that the justness of a distribution is incompatible with “wide variations” in the resulting pattern of holdings of valued things.

To see this, consider a very rudimentary example, one that Rawls himself is unlikely to disavow, and that very few readers would disavow unless they are already in the grip of a quite remarkable theory. Suppose you have some good, say a certain amount of purchasing power. Any amount will do—we needn’t be talking about distributions of huge sums here: \$5 will do nicely. Now, suppose you elect to spend this money on some good, perhaps a box of breakfast cereal of a certain sort which you happen to like, or a movie at a modestly-priced theater. In doing this, you leave one firm with 100% of your \$5 and all the other people in the world with 0% of it. This is surely an unequal distribution pattern—all for one, zero for the rest. It is not clear how one could get a more unequal distributional pattern, if patterns are in question. But it is very commonly, indeed ubiquitously instanced. Whenever you spend money, you confer all of it on some very tiny selection of firms or persons, and none of it on the rest—repeating the distributionally unequal pattern. Nicholas Barr observes,

Natural-rights libertarians have little to say about the optimal distribution of goods. If the initial distribution is at *c*, then any point on the contract curve between *d* and *e* is optimal, provided that *c* accords with Nozick’s idea of justice in holdings, and that the movement from *c* to the contract curve is the result of individual utility maximization through voluntary trading in a competitive market system. More generally, depending on the initial distribution, *any* point on the contract curve can be an optimum.¹²

Well, does that result indeed depend on embracing some philosophical system, such as Nozick’s? Is it really true that whether a set of voluntary trades resulting

in some particular, assume unequal, distribution is just is a matter of serious dispute? If Paretianism is assumed, then the answer would seem to be that if it is, that can only be because there was something wrong with that “initial distribution.” Only if there was something wrong with it could there be said to be something wrong with a voluntary trade by persons in that system subsequent to the establishment of the initial distribution. We will return to that below.

Meanwhile, let’s also appreciate that one could take any number of noneconomic examples to the same effect. If you choose a spouse, you devote a great deal of love and affection to that one person, and none (of that particular kind) to anybody else, instead of distributing it equally among some very large number of potential recipients. Yet Rawls, I am sure, is not going to object to that, either.

Why not, though? His answer in the case of the monetary example could perhaps be that you have, after all, already paid your dues to society, your taxes, and so what is left is really yours, to spend as you like. This answer won’t do. If justice is allergic to inequitable patterns of distribution, then there is no inherent reason why it should tolerate them in the post-tax distribution as compared with a pretax one. Here was a good, capable of being distributed; here was a distribution that showed no regard whatsoever for patterns. But if justice is a matter of patterns, then why wasn’t it unjust, too? Most readers, along with this author and Rawls himself, will undoubtedly think that the example looks ridiculous—since, after all, it is ridiculous. But if it is, that establishes the point in question here: justice is perfectly compatible with wildly varying patterns of distribution of desirable things. No injustice is necessarily effected *merely* by distributing goods one way rather than another. Something has been left out.

This returns us to the question about ‘arrangements’. According to Rawls, “different arrangements” are compatible with Pareto Optimality. In the sense of ‘arrangements’ in which what results from distributing any good in one way rather than another counts as an “arrangement”, this is perfectly true. When you go forth with your \$5, or your \$N, there are innumerable possible recipients, *none* of whom will be made worse off by your employment of that money. You could, of course, use it to buy poison which you then put into someone’s coffee, or bullets which you then put into his brain. But if so, it won’t be the purchase, as such, that effects this violation of Pareto: it will, instead, be the use to which you put it. Obviously there are uses to which you can put resources that will cause violations of Pareto’s principle. (That, as Hobbes pointed out, is true right down to the minimal level of one’s bodily resources alone.) But we are here talking of distributions of *goods*, not of *evils*; voluntarily inflicted evils are, of course, violations of the principle, which, after all, says no more than that we are not to do evil to some person or persons in order to bring about what we suppose to be good for some other or others. Your expenditure of money, however, is, so far as it goes, the transfer of a *good*, from yourself to someone else—not of an evil.

Someone might, of course, regard worldly goods as inherently evil, to be

sure. You presumably couldn't give such a person anything without violating Pareto, it seems. However, Pareto's idea is irrelevant to such attitudes. Paretianism, coupled with the liberalism that economists and most of us assume here, takes the (self-regarding) interests of the persons in question as definitive of what is good or evil in their cases.

Otherwise, as we have already seen, a set of "arrangements" in which someone or some very few persons get all of a good, and everyone else gets none, is perfectly compatible with justice, on the face of it.

Suppose, on the other hand, that a social philosopher of the familiar Rawlsian stripe comes along and, noting his disaffection with some pattern of distribution resulting from your employment of a set of goods, proposes to "rearrange" them. His way of rearranging them, though, is to deprive you, forcibly, of some or all of those goods and effect some distribution of them different from the one you would have brought about had you had your choice in the matter. If so, you are now worse off. Someone has been made better off, indeed—but at the expense of making you worse off. Paretian efficiency has therefore been violated.

But Rawls, remember, was arguing that Pareto was "inadequate"—not that it was *wrong*. He talked, that is, as though a theory of justice needs to choose among the *alternative* distributions *compatible* with Pareto Optimality. Yet how is he going to defend a redistributive theory on such a basis? For after all, it is perfectly possible to hold that any of the patterns consistent with voluntary action by all concerned would be just. "Don't rearrange things by force" is the relevant principle; innumerable specific patterns are compatible with it. The reach of the Pareto principle is very deep indeed.

This forces us to reconsider the subject of "arrangements" or "distributions," insofar as they are examples of possible theories of justice. Characterizations of justice in Rawls are quite vague, to be sure. They are said to have to do with the "basic structure" of society, for instance.¹³ But this vagueness is also misleading. When does a society have one "structure," in the relevant respect, as opposed to another? Rawls would presumably not count a society almost all of whose members are chess addicts and never have picnics as having a different "structure" from a society in which no one at all plays chess but lots of people have picnics. For that matter, he also would not count one society as having a different structure from another if the first was considerably populated by coal-miners while in the other, not a single coal-miner was to be found. What, then, is the kind of difference that the term "structure" is meant to convey? In fact, it is quite easy to answer. Certain of its legal and political features are what we must point to.

But which? The answer here is also not so difficult, although not quite so quick. A law *imposes* an order. All sorts of different orders are compatible with a given law, but if there is a law, the order you have as compared with the one you would otherwise have differs precisely in that the one you have as a result of the law has been, or at least may be, enforced, imposed. (Of course, the libertarian theorist will claim that what is "imposed" is really only restoring or helping to ensure a natural order—what it imposes is, we might say, non-

imposition. True: but the apprehended thief has nevertheless been forced to restore the order he upsets.) More generally, in fact, justice is the rightful use of force, or more precisely, the principles concerning how it may rightfully be used.

If this is so, then let us go back to the choice among Pareto-optimal “arrangements”. As we have seen, these have nothing, as such, to do with justice. Justice comes on the scene when we do or should intervene, using force to bring about one distribution rather than another. Had a given distribution come about with no use of force at all, nor any threat of it, it would not differ in respects relevant to justice, though it might differ greatly in its patterns. The difference between one society and another, the first of which had many coal-miners and the second none, or the first many chess-players, the second devoid of them, is a great difference of pattern, but no difference at all in respect of justice. The social system, in the sense of the system determined by a given set of laws, could perfectly well be identical as between them. And the same could be true of two societies one of which displayed great variation in income, the other substantial equality in that respect.

Consider, then, the claim that Pareto efficiency is “inadequate” as a guide to justice. We now see that this claim is misleading. It is possible to have a system of justice in which Paretian efficiency is enforced.

The Uniqueness of Paretianism

Now we can make the point I am concerned with here. I wish to suggest that there is, actually, one and only one system of justice that is efficient in the relevant sense, one and only one system that is concerned exclusively with efficiency in the Pareto sense. That is a system in which force and fraud are the only things which force may legitimately be used to rectify. As for any other variations in “patterns,” so long as they result exclusively from the voluntary interaction of people, no forcible intervention would be permitted to alter them. They would all be reckoned to be just. This is, of course, a pure market system—indeed, that is the definition of a pure market system. (Some deny this; I will discuss their claim below.)

Rawls’ claim, then, that Pareto efficiency is “inadequate” is misleading in a quite fundamental way. For he makes it sound as though the further principles of justice that he proposes are meant to affect the choice of distributions compatible with Pareto efficiency. But that isn’t so. By definition, any other principle of justice would require people to do some things they do not want to do, or refrain from things that they want to do, *other* than actions taken to counter or correct Pareto suboptimal actions. For examples, the introduction of principles requiring equal opportunity, the forced social provision of minimal levels of income, and many other things, all advocated by Rawls and by most contemporary social philosophers, are incompatible with pure Paretianism. They are *not* selections of patterns of distribution, each compatible with that criterion, some of which are more just than others.

In other words: if we accept Paretianism as a principle of justice, and allow that we worsen people's situations when we force them to do what they do not want to do, unless what they want to do would in turn worsen the situation of others, then it crowds out all contrary views of that subject.

I am inclined, myself, to support the system in which Paretianism is enforced, in the sense discussed, as against all of its rivals. But my point here is not to do that; it is, rather, to correct a misconception inherent in, and very extensively fostered by, Rawls' extremely influential writings. It is disarming to begin in the way he does, for if indeed all of the various systems he discusses were literally compatible with Paretian efficiency, what would there be to complain of? None of the different systems would have been exacted at the expense of anybody! Of course this is false, and indeed quite spectacularly so. The different systems are systems in which some are made better off by making others worse off (the favored subjects being those "better off"). In only one of them is that not so, in principle: the pure free market system.

The "Ideal Market"

Mention of free markets makes it useful to add a brief note about one further possible misconception. There is a famous thesis to the effect that only a society with a so-called "ideal market" would result in a "Pareto Optimal equilibrium state." Barr, for example, adopts this view.¹⁴ Such so-called ideal markets have many extraordinary features: each participant has perfect information and is perfectly rational, transaction costs are nonexistent, and there is "perfect competition" in the sense that nobody is in a position to unilaterally affect the price of anything.¹⁵ While this proof is of some intrinsic interest, it has little to do with the argument of this paper, which in no way depends on the exotic conditions mentioned. One might suggest, indeed, that it has little to do with reality. But certainly the just society is not efficient in the sense that no conceivable trades could be made that would improve someone's condition without making someone else's worse. It is, rather, optimal in the sense that the only transactions that are *permitted*, apart from rectifications, are transactions that *impose no costs*, so far as any relevant persons can see, on other parties. How many and which of the infinity of mutually advantageous exchanges that could conceivably be made are actually made is irrelevant to justice. All permitted trades are presumptively Pareto superior to the status quo, because they are made between parties who, upon consulting their respective preferences and information bases, voluntarily engage in those trades. Liberalism takes that to be sufficient for supposing that those agents regard the aimed-at situation as superior to the one they are already in. But any proposed "improvement" brought about by force would automatically violate this condition. It is in that sense that the free market society is the only one that is efficient, despite not being at all likely to be in the optimal state meant by the theorem. For any other principles proposed as principles of justice would require imposing conditions on unwilling parties, thus effecting whatever improvements they

make to some people, if any, at the expense of others.

If you abandon liberalism, of course, you can help yourself to plenty of forcible interventions to make social improvements. You might even claim not to have abandoned it, insisting that when you take 25% of Smith's income to pay for Medicare for all, you're getting Smith what he really wanted all along, however much he protests. But Rawls and others who have accepted his claims about this matter claim to be liberals. Perhaps some who are persuaded by the foregoing sketch will also regard it as a *reductio* of that outlook. We'll have to leave that for another day.

Professor Barr holds that the imperfections of real-world markets are such as to justify a variety of interventions; and I presume that he speaks for a majority of his colleagues in the economics area these days. It is interesting to see how he proceeds on this matter. Take, as an example, rectification regarding imperfect information.

Simple theory assumes that consumers know what goods are available and their nature. The assumption can fail because economic agents may have imperfect knowledge of the *quality* of goods or their *prices*. The literature thus has two strands. The first analyses the effects of imperfect information about quality: consumers might be badly informed (e.g. about the quality of an automobile), so might producers (e.g. about the riskiness of an applicant for insurance). The resulting literature investigates such topics as 'lemons' and signaling. The second strand, imperfect information about prices, embraces search theory and reservations wages.¹⁶

Such failures, he supposes, could justify regulation in some cases, e.g. where "information is seriously deficient."¹⁷ Note that the question of *how* the consumer, or the producer, came to be in this condition of deficiency is not really addressed. But surely that makes all the difference. You can be poorly informed, for example, because I have deliberately misinformed you. In that case, I do you a hurt, and have violated the Pareto restriction, attempting to improve my own situation at the expense of you. But suppose your sorry state of information is in no way my fault, or perhaps anyone's. What right, then, would regulators have to intervene and "correct" the situation? Indeed, how did they know? And how can the criteria of efficiency Barr introduces enable us to judge that the situation in which some people buy products for bad reasons that are not the fault of others label the resulting situations as "inefficient"? On the contrary, it would seem that if I force you to do something that you think to be best, even though I think you are in error in so thinking it, I have done something wrong. Or if not, it might be a case in which your ignorance is indeed your fault, and I in friendship or sympathy act to correct it. But this correction can hardly be at the expense of third parties in no way involved.

It is hard to see, in short, that the consumer is entitled to be in some or other state of information. He is entitled not to be in a state of misinformation caused by the deliberate actions of those with whom he interacts, indeed. But

what else?

Barr does agree that sometimes misinformation can be corrected by market action. An example, I presume, is that people are willing to pay a good deal for education, with a view to improving their general state of information. But it is, to put it contentiously, not obvious that we all owe each other an education. To be sure, typical academics today *do* find that obvious. After all, we make our livings providing this great good. Not surprisingly, we are happy to have someone else pay for our efforts, thus improving the market for education. But this is no longer a free market we're talking about. Free markets, one would think, are based on freedom.

Real-World Paretianism

Two possibly important notes should be added here. The abstract idea that we are not to harm Peter in order to help Paul requires, as we saw, two things to complete it. First, we need to know what constitutes help and harm. Our answer to this is that it helps Peter to put him in a state that he prefers, and harms Paul to put him in a state that he prefers not to be in. But going ahead and applying the Pareto principle on those interpretations is subject to the important restriction that the people in question be relevantly innocent, that is, that they themselves are not in turn guilty of the sort of aggression that the Pareto principle prohibits. Those who are thus guilty are not automatically eligible for the treatment that the innocent are. On the other hand, being guilty of some small violation does not automatically deprive one of all the protections of morality. It means that you owe somebody something and need to pay it; it does not mean that you owe everybody everything.

Moreover, as we have seen, when we help Paula by doing something that in her view benefits another individual, say Johnny, we assume that Johnny is not unhappy with the arrangement. In the case where Johnny is a small child, applying this criterion may not be easy. It is, indeed, difficult enough that we will leave it for another essay.¹⁸

Restricting ourselves, then, to the somewhat easier case of ordinary people, we should note that, in the first place, the individuals we come in contact with are usually, if not always, "innocent" in the sense important for these purposes. In general, Paretianism applies to the *relevantly* innocent. However, I am claiming that we may in turn use Paretianism to identify who is innocent and who isn't. The criterion is intended to be reiterative in its application. The innocent are those who in turn have not worsened the situations of any other innocents. And where would this end? Presumably with those who have harmed no one at all. But why not?

Tricky questions can indeed be raised about this. Suppose that C. Milquetoast has been perfectly civil and nonviolent toward certain evil persons whom he might have been able to help bring to justice? I don't have a simple answer to that, but *prima facie* am doubtful that we in general have a duty to bring miscreants to justice, though it is certainly something we should do if we

can.

Second, and more important: presumably it is by no means unusual to inflict very tiny harms on other people. These may be of either of two sorts. First, we may raise by some vanishingly small amount the probability that someone will experience a certain sort of worsening. Or, second, there may have been actual worsenings, but again, vanishingly small.

In both cases, there is a fairly straightforward criterion to bring into play. It has two points. First: if the person allegedly harmed simply pronounces that he doesn't care, that's that. And second, if his situation is *unnoticeably* worsened, in that sense that our supposed victim can't even distinguish between the new situation and the old one, then we shall again ignore this particular supposed objection to our action. We deal, then, only with significant harms as seen by the persons affected.

There is a very interesting question, indeed, concerning applications at the other end of the spectrum—catastrophe, life-boat situations and the like. Hume argued that the restrictions of justice are simply inapplicable when no matter what we do, the greater number will perish.¹⁹ I do not attempt to deal with this problem here, but only to recognize that it is one, and that it is in real need of a solid theory. But certainly, for starters, Paretianism runs out of scope for application when it is impossible to avoid harm for all or most of those concerned, in particular for the agent himself. Generally speaking, morals is framed by the potential for all to benefit, or at least not be harmed, in the course of social interaction. That is the basis of the Social Contract approach to morals, one or another version of which is ubiquitous among those who think morality to have any foundation at all.

Meanwhile, there is also concern in many quarters with the supposed injustice in initial conditions, from reasons other than violence. In a sense this raises a large issue. But how large? It is not plausible to think that my fellow men are to blame for my lack of athletic ability or pianistic prowess, either of which I would be glad to be possessed of. It is not believable that I am allowed to invade my neighbor in order to rectify these disparities, even if that were possible. And it is obvious that innumerable differences among people will result from differing climates, available resources, and the particular interests, abilities, and proclivities of those around them. But in this there is no evident complaint of justice, which is all we are investigating here.

Conclusion

I conclude that it is a mistake to suppose that there are lots of different views about justice that are *equally efficient*. There are not. So far as I can see, there is only one such view, though the range of specific types of social states compatible with this view, e.g. specific “patterns” of enjoyments or holdings of goods, could vary wildly. All of them, however, would exemplify one and only one theory of justice, one principle of justice. That principle, in effect, is that justice consists in using force only on behalf of, literally, the common good,

liberally conceived in the Paretian spirit. Rival theories may differ in all sorts of interesting ways, but all of them, rather than exemplifying Pareto's principle plus further features, violate it in some area or other. Justice as pure efficiency is, I would contend, not only a viable conception, but an eminently plausible one; in any case, however, it is, contrary to the almost universally shared view of Rawls, a unique one.

Notes

Versions of this paper have been read at the University of Georgia in 1996, among others, and at the Canadian section of the International Society for Social and Legal Philosophy (IVR); it was to be published in a now-defunct On-line Journal of Social Philosophy, and has benefited from many comments, most of them adverse. I am indebted to Alistair MacLeod for much, in my view, erroneous but instructive criticism, and to Jan Lester, whose criticisms were not erroneous at all.

1. Nicholas Barr, *The Economics of the Welfare State* (Stanford, Calif.: Stanford University Press, 3rd ed., 1998).
2. Barr, *The Economics of the Welfare State*, 70.
3. Barr, *The Economics of the Welfare State*, 70.
4. Barr, *The Economics of the Welfare State*, 71.
5. Barr, *The Economics of the Welfare State*, 71.
6. Barr, *The Economics of the Welfare State*, 72.
7. Robert Nozick, *Anarchy, State, and Utopia* (NY: Basic Books, 1974), 28-30 and ensuing discussion.
8. The criterion is further discussed and argued for in Jan Narveson, "Liberal/Conservative: The Real Controversy." *Journal of Value Inquiry*, 2000, special issue on Liberalism, guest-edited by Jan Narveson and Susan Dimock vol. 34 nos. 2-3, 167-188.
9. Thomas Hobbes, *Leviathan* (New York: E. P. Dutton, Everyman Library, 1950) 119.
10. John Stuart Mill, *On Liberty*, chapter 1. Also *Utilitarianism*, Ch. 5.
11. Rawls, "Distributive Justice", in Peter Laslett and Runciman, W.G., *Philosophy, Politics, and Society*, Third Series (Oxford: Blackwell, 1967), 65.
12. Barr, *The Economics of the Welfare State*, 74.
13. John Rawls, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, 1971), 7.
14. Barr, *The Economics of the Welfare State*, 78.
15. I employ the description of the argument in Allen Buchanan's well-known *Ethics, Efficiency, and The Market* (Totowa, New Jersey: Rowman & Allanheld, 1985), 14-15.
16. Barr, *The Economics of the Welfare State*, 82.
17. Barr, *The Economics of the Welfare State*, 82.
18. See chapter 15 for an attempt in this direction.
19. David Hume, *An Inquiry Concerning the Principles of Morals* (Indianapolis, Ind.: Bobbs-Merrill, Library of Liberal Arts, 1957), 17.