

WILEY

Distributive Justice, Welfare Economics, and the Theory of Fairness

Author(s): Hal R. Varian

Source: *Philosophy & Public Affairs*, Spring, 1975, Vol. 4, No. 3 (Spring, 1975), pp. 223-247

Published by: Wiley

Stable URL: <https://www.jstor.org/stable/2265084>

REFERENCES

Linked references are available on JSTOR for this article:

https://www.jstor.org/stable/2265084?seq=1&cid=pdf-reference#references_tab_contents

You may need to log in to JSTOR to access the linked references.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



Wiley is collaborating with JSTOR to digitize, preserve and extend access to *Philosophy & Public Affairs*

JSTOR

HAL R. VARIAN

Distributive Justice,
Welfare Economics,
and the Theory of Fairness

Robert Nozick's recent article "Distributive Justice" contains a provocative discussion of the meaning of justice.¹ A central point is his distinction between "historical" and "end-state" principles, the idea being that certain approaches to distributive justice, such as Rawls' approach and the approach of welfare economics (what Arrow calls ordinal utilitarianism), have concentrated on evaluating only "current time-slices" of a distribution and have not focused on the procedural aspects of distributive justice.

Nozick has done us a great service by describing an alternative to these "end-state" theories of justice; for certainly *some* of our moral intuitions about justice depend not only on the current description of a distribution but also on how the distribution came to be. Thus discussions of the issues of procedural justice should certainly spread some light on the questions of distributive justice. Indeed they do; Nozick's article has—at least for me—clarified some of the important interrelationships among the concepts of justice, individual rights, liberty, and so on. Nevertheless, I do not entirely agree with his analysis.

My objective here is threefold: (1) to question some of Nozick's arguments and assumptions; (2) to clarify some of the results of

I would like to acknowledge the helpful comments made on this paper by Lester Thurow, Robert Nozick, and Susan Foster. I also wish to thank the editors of *Philosophy & Public Affairs* for aid in clarifying the exposition of the economic issues and for valuable editorial assistance.

1. *Philosophy & Public Affairs* 3, no. 1 (Fall 1973): 45–126.

welfare economics to which Nozick refers; and (3) to present yet another theory of distributive justice, which I believe is novel, interesting, and distinct from previous theories.

THE ENTITLEMENT THEORY

Nozick's meta-ethical remarks imply that a valid theory of justice must include three parts: a description of how people legitimately acquire holdings, a description of how people legitimately transfer holdings, and a description of how past injustices should be rectified (pp. 46–49). He makes the following specification: "A distribution is just if it arises from another (just) distribution by legitimate means" (p. 47). Nozick's theory is a procedural theory; the justice of a distribution is entirely dependent on the path used to reach it: "Justice in holdings is historical; it depends on what actually has happened" (p. 48). Most of Nozick's article is devoted to justifying this *form* of a theory of justice; he sees this form in opposition to that of welfare economics and Rawls' theory, which, he argues, are basically ahistorical in that they examine only a current time-slice of a distribution.

Nozick hesitates to describe a specific theory:

To turn these general outlines into a specific theory we would have to specify the details of each of the three principles of justice in holdings: the principle of acquisition of holdings, the principle of transfer of holdings, and the principle of rectification of violations of the first two principles. I shall not attempt that task here [p. 49].

But despite this disclaimer, Nozick *does* attempt the task of describing what he believes are reasonable principles of acquisition, transfer, and rectification, as indeed he must. For we can always put any end-state principle into Nozick's form: we determine the desired distribution according to an end-state criterion and then rule that agents can only acquire or transfer holdings if doing so leads to the desired distribution. But Nozick rightly rejects this trick (p. 53). He devotes a considerable amount of space to examples of what he believes would be acceptable rules of acquisition, transfer, and rectification. The primary principle he uses in choosing acceptable rules

is that the chosen rules should not violate agents' rights. Nozick's theory of rights is discussed at length elsewhere.²

Let us begin with rectification. How *are* we to rectify an unjust distribution? Nozick's answer goes something like this: if an injustice has occurred, we must determine what the expected outcome would be if that injustice had not occurred; in case of ties we can then perhaps use end-state principles to resolve which of these expected outcomes to choose.

As an example, consider Table 1. Let us suppose that some injustice occurred in the past and that if the injustice had not occurred, the three listed allocations had the various listed probabilities of occurring today. (All other allocations have zero probability of occurring.) But because of the effect of the injustice, the actual distribution today is something else. Nozick suggests that we rectify the past injustice by redistributing today's allocation according to a weighted average of the allocations in the other possible states. This distribution is also listed in Table 1; it is constructed by multiplying each possible allocation by the probability of its occurrence and then adding all these numbers together.

TABLE 1

	Possible allocations			Weighted average
	1	2	3	
	p = .2	p = .3	p = .5	
Agent 1	5	1	6	4.3
Agent 2	0	3	0	.9
Agent 3	5	6	4	4.8

I find the idea of this weighted average somewhat at odds with Nozick's explicitly historical theory. If the justice of an end-state is really so dependent on the historical process used to reach it, isn't a more reasonable choice the distribution that has the greatest probability of actually having occurred? In the example described above,

2. Robert Nozick, *Anarchy, State, and Utopia* (New York, 1974).

the weighted-average distribution actually has zero probability of having arisen from some historical process, while distribution three is likely to have actually arisen. In any probability distribution where extreme distributions of wealth or income are likely, the weighted average will often be very different from the most likely allocation.³ One can well argue that the private ownership, laissez-faire society Nozick describes does indeed tend to generate extreme distributions of wealth; thus the problem of choosing the most likely allocation or the weighted-average allocation in rectification becomes rather serious.

It seems to me that the weighted-average allocation *really is* a more reasonable way to rectify, the primary reason being that it takes into account the whole range of possible outcomes in a way that the most likely allocation does not. The expected value—i.e. the weighted average—is after all a measure of central tendency. The case depicted in Table 1 has two rather extreme allocations as possible outcomes, and the expected value tends to choose an intermediate allocation. But it also seems that the most likely allocation is more consistent with the historical approach Nozick describes. If the justice of a particular allocation depends on what actually happened, then one allocation should be more just than another if it is more likely to occur. One cannot say this is a logical inconsistency in Nozick's theory. After all, the principles of rectification are presumably independent of the other principles. But if the historical process is so important to the justice of an end-state, it seems peculiar to have a rectification procedure that can pick a rectifying allocation which is actually unlikely to occur.

I think that Nozick is led to this position precisely because the extreme historical approach he advocates is incompatible with some of our moral intuitions of justice. Of course, how an allocation was reached has *something* to do with justice; but it is just not the whole story. This becomes clearer when we consider processes that include large random elements. Here we may start the process with an equal distribution and let legitimate transfers lead us to new allocations; but the large random component—acts of God, depressions, accumulation of fortunes that may be transferred to new generations—may

3. Only in the case of a symmetric, unimodal probability distribution will the expected value and the mode (maximum probability point) coincide.

very well move the society to an unequal final allocation. One might well argue that the role of the state in such a situation is to avoid such distortions. Nozick chides Rawls for worrying about the fact that natural assets are distributed randomly; however, he himself proposes no mechanism to correct for any kind of randomness. If a process can be radically affected by turns of fortune that are arbitrary from a moral point of view, it seems unreasonable to attach great moral significance to the outcome of such a process.

The impression one gets from reading Nozick is that the problem of rectification is somehow minor. It seems to me that the reverse is the case: the problem of rectification is central to the issue of justice. We are interested in the question of justice precisely because we live in an unjust world; injustices have occurred in the past and are occurring now. The question is what should we do about them. Proponents of end-state principles, such as Rawls, are attempting to answer this question: we decide what a perfectly just state is and then try to move toward it.

HISTORICAL PRINCIPLES AND END-STATE PRINCIPLES

As a justification for his consideration of historical facts in evaluating distributional equity, Nozick points out that

most persons do not accept current time-slice principles as constituting the whole story about distributive shares. They think it relevant in assessing the justice of a situation to consider not only the distribution it embodies, but also how that distribution came about. . . . We think it relevant to ask whether someone did something so that he . . . deserved to have a lower share [p. 50].

Nozick is perfectly correct here; naive theories of justice that provide for a fixed social product to be divided without regard for those who contributed to the formation of that product ignore the most difficult and important problem of formalizing our notions of justice.

At some points Nozick interprets “how the distribution came about” in the explicit sense of how the patterns of production were arranged. Thus he criticizes Rawls’ theory and welfare economics for ignoring the interaction of production and distribution:

To think that the task of a theory of distributive justice is to fill the blank in “to each according to his ——” is to be predisposed to search for a pattern; and the separate treatment of “from each according to his ——,” treats production and distribution as two separate and independent issues. On the entitlement view these are *not* two separate questions [p. 56].

This is a peculiar criticism; both Rawls and welfare economics devote very careful consideration to the interaction of production and distribution. The whole point of the difference principle in Rawls is to take account of production; welfare economics deals explicitly with production, as I will next discuss.

WELFARE ECONOMICS

Because several points I want to make subsequently depend on some of the basic concepts and results of welfare economics, a summary of a few of these ideas will be of use here. It also, I hope, will be of independent interest, since I believe these concepts and results to be crucial to any discussion of distributional issues.

Suppose that we have a set of economic agents—producers and consumers—and that the consumers have some *initial endowments* of goods. The definition of goods considered here is very broad—we can think of goods as being indexed by time, location, or state of the world, and in particular we can consider the consumer’s original endowment of his own labor time as being one of his goods. Since one agent’s labor is different from another’s, we often think of each agent’s labor as being a separate good. If we want to be very simple-minded about it, we can think instead that there is only one kind of “labor power”: “able” agents are endowed with a lot of this labor power, and “unable” agents with only a little. Given the possible models of production available, there will be some set of *feasible allocations*, that is, a set of descriptions of how much of each good each agent consumes, how much each agent works, and how much each producer produces. The set of feasible allocations is just the set of all *possible* allocations; this of course includes all possible redistributions of the initial endowments.

A simple story may clarify the concepts here. We consider a group

of agents who each morning take to the town market their bundles of goods, which we can think of as farm produce—carrots, eggs, tomatoes, and so on. In particular, one element of each agent's bundle is his available labor time for that day. During the day agents may produce new goods and transfer these and the old goods among themselves in various ways, not necessarily voluntarily. For example, one agent could appropriate all of the other agents' goods, or the agents could vote on a way to divide up the goods, or whatever. However they decide to divide the goods, it must be a feasible way in that the total amount of each good must be equal to the total amount started with plus the amount produced minus the amount used up in production. The set of all such allocations—the set of all possible ways of dividing the goods—is called the set of *feasible allocations*. Thus a particular feasible allocation is simply a description of what each agent takes with him in the evening, when he returns home from the town market.

The basic problem of welfare economics—and of course the basic problem of distributive justice—is to determine at which feasible allocation the economy *should* operate, i.e. what bundles each agent should be allowed to take home. One very weak criterion that has been proposed is that of pareto efficiency: an allocation is defined as being pareto efficient if and only if there is no feasible allocation where *all* agents are at least as well off and some agents are strictly better off. (I will often abbreviate this clumsy condition and say “. . . where all agents are better off.”) An equivalent way of stating the definition is that at a pareto efficient allocation there is no way to make some one agent better off without making some other agent worse off.

As it stands, the criterion of pareto efficiency is certainly reasonable: if there is some way to make everyone better off, why not do it? Unfortunately, it is a weak criterion. For example, the allocation where one agent gets everything is pareto efficient. Why? Because the only way we can make one of the “slaves” better off is to take something away from the “privileged” agent—there is no way to make *everyone* better off.

Even though we may choose to limit ourselves to pareto efficient allocations, there is still a large set of allocations to choose from. The basic idea of welfare economics is to assume that there is a welfare

function which evaluates the “goodness” of the social states as a function of the utility evaluations of those states by the agents in the society. Thus every welfare function is of the form $W(u_1(x_1), u_2(x_2), \dots, u_n(x_n))$, where u_i is the utility function of the i th agent and x_i is a description of the consumption-labor bundle of the i th agent. The important restriction is that W depends only on the utility evaluations of each individual’s bundle and not directly on the bundle itself. Of course, we also require that a welfare function depends on these utilities in a positive way: if the utility of any agent increases and no agent’s utility declines, we want the value of the welfare function to increase. The classic utilitarian welfare function, $\sum U_i(x_i)$ is of course a very special case of a function of this form. Nevertheless, for any function of this form, one can show that if we choose a feasible allocation of maximum welfare, it must necessarily be a pareto efficient allocation. Hence, the choice of a welfare function “solves” the problem of choosing a best pareto efficient point. The proof is very simple. Suppose x is a feasible allocation that maximizes some *specific* welfare function $W(u_1(x_1) \dots u_n(x_n))$. Suppose x were not pareto efficient; then by definition there is some feasible allocation y such that $u_i(y_i)$ is greater than $u_i(x_i)$ for all agents i . Since a welfare function is required to be increasing in all of its arguments, this implies $W(u_1(y_1) \dots u_n(y_n))$ must be greater than $W(u_1(x_1) \dots u_n(x_n))$ which is a contradiction.

There is a completely different way to choose an allocation at which the economy can operate; namely to use the market mechanism. That is, we assume that a market forms with well-defined prices and each agent takes these prices as being outside of his control and does the best he can under this constraint—that is, each consumer tries to purchase the best bundle he can afford given the value of his initial endowment, and each producer attempts to maximize profits at the going prices. If these prices are such that supply equals demand in all markets, this mechanism determines a well-defined allocation called a *market equilibrium*.⁴

The prices paid for the various goods are not arbitrary; one can show that the hypotheses presented above imply, loosely speaking,

4. Such a market-clearing price-system will generally exist; see K. Arrow and F. Hahn, *General Competitive Analysis* (San Francisco, 1972), chap. 2.

that the “factors of production” will be paid their “marginal product.” That is, the wage rate of one unit of a certain kind of labor will be equal to the difference between the value of the output produced using the total labor used by a profit-maximizing firm minus the value of output produced using one less unit of that kind of labor. Similarly, the owner of a particular plot of land or the owners of capital equipment will be paid according to the marginal contribution of their product to the production process. Hence, each factor will be paid according to its *marginal contribution* to the value of the output.

There are two important facts about this concept of marginal product that are often misunderstood. First, only the *price* of a factor of production is determined by its marginal product; the rewards reaped by an owner of this factor depend on how much of that factor he owns, i.e. on his *initial endowment* of that factor. An agent’s *net* reward, his income, is the product of the price of the goods and labor he holds and the *amount* of his endowment of those goods and labor.

Second, the marginal product of a factor itself will in general depend on the level of output and on the initial distribution of endowments. Thus to think of the market mechanism as “solving” the distribution problem by itself is unreasonable. The market distribution depends completely on the initial distribution of resources in the economy. How is *that* distribution to be determined?

Let us restate this crucial point in terms of the simple story we discussed earlier. We now imagine that the agents use a *particular* method of determining a final allocation; namely, the market mechanism described above. The prices of the items they sell on the market will depend in general on the available supplies and demands for the goods, and the value of each agent’s final bundle of goods will depend, of course, on the value of his initial bundle of goods. The market mechanism will determine a *particular* allocation that depends very much on the pattern of initial endowments.

Thus when Nozick describes a *laissez-faire* world, in which each agent is paid his “marginal product,” and asks what is the role of a theory of justice here (pp. 82–83) we can well answer: in the determination of the initial endowment—for the market equilibrium is completely indeterminate until it has been specified who owns what in the beginning.

Now that the appropriate concepts have been defined I can present the two main results of welfare economics.⁵

Proposition A

The outcome of the market mechanism as previously described—the market equilibrium—will, under very general assumptions, be a pareto efficient allocation.

Thus the laissez-faire economy described above generally will operate in an efficient manner; however, there is no other ethical content to this result. If, for example, the initial endowment gives one agent everything, the market equilibrium from this endowment will also give this agent everything. There is nothing *just* about that.

Proposition B

Under somewhat more restrictive assumptions, one can show that every pareto efficient allocation is a market equilibrium for some initial endowment of goods.

Since the pareto efficient allocations are precisely those that maximize *some* welfare function of the form discussed earlier, we can restate this result as:

Proposition B'

If an allocation is a point of maximum welfare of some particular welfare function, it can be achieved by a suitable reallocation of endowments followed by trading to a market equilibrium.

This theorem is considerably deeper; the point is that the market mechanism can be used to support whatever efficient allocation society wishes. It can therefore serve as a “self-correcting” distribution mechanism, once the initial endowments are determined. This result inspired the school of market socialists such as Lange and Lerner, who argued that the appropriate way for a socialist society to handle the immense task of allocating efficiently all goods to both consumption and production activities was to allow the laissez-faire market to do the allocation, while the government worried only about

5. For a detailed discussion, see A. Bergson, “A Reformulation of Certain Aspects of Welfare Economics,” *Readings in Welfare Economics*, ed. K. Arrow and T. Scitovsky (Homewood, Ill., 1969); A. Lerner, *Economics of Control* (New York, 1970); and Arrow and Hahn, *General Competitive Analysis*.

ensuring a just distribution of initial endowments (wealth). In this way the socialist state could be assured of efficient operation no matter what the desired distribution was.

When Nozick remarks that “the socialist society would have to forbid capitalist acts between consenting adults” (p. 59), he is implicitly assuming that the socialist society must be operating in a pareto inefficient manner. For if the society were pareto efficient, no such trades would be possible because they would contradict the assumption of pareto efficiency. If the goal of the socialist society were pareto efficient use of resources, there would be no need to forbid agents to use the resources in any way they saw fit. As Lerner puts it. “The fundamental aim of socialism is not the abolition of private property, but the extension of democracy” (p. 1).

At first glance this kind of market socialism seems quite feasible; however, some further thought shows some problems in the concept of redistribution of endowments. Recall that a description of an agent’s endowment includes a description of how much labor he has; a reallocation of endowments may well involve giving some agent control over some other agent’s labor—in a word, reallocation of endowments may well involve slavery. (Compare to Nozick: “This process . . . makes them a *part owner* of you; it gives them a property right in you” [pp. 68–69]).

Now the situation is not quite as bad as all that; one could also reallocate money just as well, i.e. impose a lump-sum tax on initial endowments and redistribute the proceeds. After all, we are assuming that an agent’s choice set is described by the set of bundles he can *afford*, the set of bundles whose value is less than or equal to the value of his endowment. We can change this choice set by changing *what* he has to sell, or merely by transferring money to begin with.

This lump-sum tax must be a tax on *endowments*—not a tax on what an agent actually sells. In particular, when we tax an agent’s labor, we tax his whole endowment of labor. Thus agents may find a tax bill waiting for them before they even begin to trade; they will then be forced to work to get the money necessary to pay the bill. In the market equilibrium it matters not at all whether an agent works

to pay his taxes, which are then transferred to another agent, or whether at the outset the other agent owns some part of the first agent's labor power.

Of course the lump-sum taxation scheme is almost as impractical as the slavery scheme. The main problem is one of information: different agents have different kinds of labor. To be rather one-dimensional: different agents have different abilities. Thus to determine the lump-sum tax, one needs to determine with how much labor power each agent is endowed. The lump-sum tax would therefore be based on *ability*, as opposed to the normal sort of tax based on value of labor sold, i.e. an income tax.

To determine an agent's ability may be difficult, so difficult that no one takes this lump-sum tax very seriously. The common practice of redistribution via an income tax of course distorts relative prices and produces an inefficient outcome—by placing a tax on labor sold, one encourages people to sell less labor and thus to produce less total output. It should be emphasized that even though the result of an income tax may be *inefficient*—in that there is some way to make everyone better off—it still may be better in terms of general social welfare than no tax. Even though there may be some way that everyone can be made better off, the information is not available to determine how to do it.

Nevertheless, there is a large part of agents' endowments that is not labor but consists instead of physical wealth—primarily land and capital. After all, in a capitalist society, every factor of production is owned by someone and the payments to each factor contribute to the incomes of agents. Even if we allow only the reallocation of physical wealth, we could still achieve a sizable number of different pareto efficient allocations. Such a scheme is discussed later under the name "people's capitalism"; it could just as well be called "market socialism."

It should be emphasized that both of the propositions of welfare economics follow from a simple and unrealistic model of economic activity. If the assumptions of the model are violated, the propositions need not be true. In particular, if agents do not take prices as given, but instead take into account their influence on the market price—that is, act as monopolists—the resulting equilibrium need not be pareto efficient. In other words, there will be in general some way

to make both the consumers of the monopolist's product *and* the monopolist better off. Thus when Nozick says, "But an entitlement theorist would find acceptable whatever distribution resulted from the party's voluntary exchanges" (p. 84), he leaves himself open to the charge that without some restrictions on the way agents behave, i.e. through antitrust legislation or whatever, his system of voluntary exchanges may result in an allocation that is pareto *inefficient*.

The outcome of this discussion of welfare economics is as follows: yes, the market (if it works) leads to efficient allocations where each agent is paid his marginal product; however, each agent's total earnings depends crucially on his initial endowment of factors. Thus the resulting allocation need have no particular significance with respect to a measure of welfare. However, one can further assert that any welfare-maximizing (and thus pareto efficient) allocation can be supported by a market mechanism, after endowments have been suitably redistributed.

THE QUESTION OF ENDOWMENTS

The above discussion shows that the naive notions of laissez-faire economics do not take us very far in determining what a just allocation of goods (and labor) should be. Thus when Nozick says, "From each as they choose, to each as they are chosen," we might well reply, "Fine, but how are the initial endowments of agents to be determined?"

As far as I can tell, Nozick's discussion of this issue is contained almost entirely in the sections entitled "Locke's Theory of Acquisition" and "The Proviso." Our previous discussion of the workings of the free market and welfare economics indicates how crucial this question of acquiring initial endowments is. For the theorems of welfare economics show that a certain *end-state principle*—that of maximum social welfare—can be supported by a certain *procedure*—the use of the market mechanism. In this sense the "historical" and the "end-state" need not be in contradiction. The question that is left open, of course, is that of the original endowments of consumption goods and of the factors of production.

Nozick begins the discussion by considering Locke's view that one

acquires property rights in an unowned object by “mixing his labor with it”; however, he soon rejects this theory: “No workable or coherent value-added property scheme has yet been devised, and any such scheme presumably would fall to objections . . .” (pp. 71–72). He then goes on to present his own view that nearly *any* appropriation is legitimate so long as it “does not worsen the situation of others” in the sense that they can still use freely the same things that they previously could (p. 75).

Nozick then considers an immediate objection to this: eventually nearly all *valuable* things become owned (almost by definition) and the unfortunate people who are born at this late date have nothing left to appropriate while the descendants of the original appropriators live in unearned wealth. The system of appropriation of anything that is unowned seems rather vacuous when virtually nothing is unowned. When we couple this with the historical fact that little has been unowned for the last few hundred years (at least) and that most “property rights” held today can trace their lineage back to forceful (presumably illegitimate) appropriation, one wonders again at the workability of Nozick’s theory. Are we supposed to trace back ownership of American land to the Indians and then try to rectify the wrongful appropriation of it by the methods Nozick discusses?

Furthermore, Nozick’s theory that the current distribution of wealth should depend on initial endowments of wealth randomly determined centuries ago seems totally unacceptable. For if the initial endowments really are random from the moral point of view, depending as they do on historical accidents, how can one base a theory of justice on such a foundation? (The similarity with my remarks on rectification should be apparent.) It is not enough even to distribute wealth equally in the first generation; for if we adopt a principle of equal distribution, why should we discriminate among generations? If we desire equality of endowments it seems that we should demand an equal distribution of wealth *each generation*.

Of course, Nozick does not necessarily want equality of endowments; his suggestions for a principle of acquisition allows virtually any kind of acquisition that “does not worsen the situation of others.” The problem, as Nozick well realizes, lies in the choice of the baseline

against which we make our comparison: “Lockean appropriation makes people no worse off than they would be *how?*” (p. 74).

Am I worse off now because of the existence of private property? Does private property itself fall afoul of the Lockean proviso? “Is the situation of persons who are unable to appropriate (there being no more accessible and useful unowned objects) worsened by a system allowing appropriation and permanent property?” (p. 73).

Nozick’s answer to this question is a brief defense of the market mechanism: “Here enter the various familiar social considerations favoring private property: it increases the social product by putting means of production in the hands of those who can use them most efficiently (profitably); experimentation is encouraged, because with separate persons controlling resources, there is no one person, or small group whom someone with a new idea must convince to try it out; . . .”

But what does this have to do with the question of *permanent* property? Suppose we considered a scheme of the type where the ownership of the factors of production was indeed private, but this ownership was nontransferable except through the market and reverted to the state upon death to be redistributed equally to new generations. This “people’s capitalism” still has the above desirable characteristics of encouraging efficiency and innovation, without any notion of *permanent*, inheritable, property.

“. . . Private property enables people to decide on the patterns and types of risks they wish to bear, leading to specialized types of risk bearing; . . .”

Arguments concerning risk bearing are often used to show that production should be centralized: by the virtue of the Law of Large Numbers, the larger an organization is, the better equipped it is to pool risks. Several economists have argued that, for example, government-owned electrical power plants are more innovative technologically than privately owned plants because mistakes made in one plant are outweighed by successes in another plant. A private firm can af-

ford only one shot at success and therefore must be more conservative in its approach.

“. . . Private property protects future persons by leading some to hold back resources from current consumption for future markets; . . .”

Again, the reverse may often be the case: free-market capitalism may overuse resources, resulting in an inefficient social state. If a number of independent, competitive wildcatters are all drilling on the same pool of oil, they will tend to sink *too many* wells, because each producer views his actions as being independent from those of the rest and ignores the total effect of the feedback from his decisions. (Another good practical example is the situation of the many small fishing companies which, without cooperation, tend to overfish common waters.)

Private property “provides alternate sources of employment for unpopular persons who don’t have to convince any one person or small group to hire them, and so on.”

But “people’s capitalism” does the same. Nozick’s arguments here are, as he says, familiar. One can quarrel with them in many respects—I have mentioned a few caveats above—but overall they do provide several good reasons for the use of the market mechanism as a means of allocating resources. *However, they are not arguments for the existence of permanent private property.*

Nozick himself brings up these reasons to show that private property does not violate the Lockean proviso, not as a utilitarian justification for property (p. 73). However, such reasons *are* often used by others to support arguments for permanent private property. It is important to realize that the market mechanism, which does have many desirable features, can work perfectly well without such a notion of property. The fundamental feature of the market mechanism is not private property but the price system. Within the market mechanism, prices serve two roles: an allocative role and a distributive role. The allocative function of prices is to indicate the scarcity value of goods and thereby reward efficient use of resources. This is quite distinct from

the distributive function which simply provides one way of distributing wealth among agents—namely, via permanent private ownership of the factor payments to property. It is perfectly possible to use prices for *allocation*, while basing *distribution* on factors other than the blind-chance assignment of initial endowments.⁶

Some of Nozick's arguments for the market mechanism are clear, others are questionable. It is a technical question of economics—perhaps the *central* question of economics in a mixed economy such as ours—what should be done by the market and what should be done by the government. However, this issue in itself has very little to do with the question of whether society should have permanent property rights. Nozick presents arguments for the market as a process of allocating resources; but, of course, the description of the market must include a description of the *initial* allocation or endowments, and Nozick's (and Locke's) theories shed little light on how these are to be determined.

Nozick's meta-ethics imply that a correct theory of justice requires three parts: a description of how people legitimately acquire holdings, a description of how people legitimately transfer holdings, and a description of how past injustices should be rectified.

As we have seen, Nozick believes that a free-market economy is a reasonable way to achieve the second part and that the third part should be achieved by asking what would have happened if the injustice had not occurred. But he gives no acceptable analysis of how the initial endowments of the agents are to be determined. And yet, as we have seen, this question is absolutely crucial to the whole analysis. The description of a process is incomplete until a description of its starting position is given, and to determine the justice of a given starting position we must use end-state criteria.

The interesting result of welfare economics is that we can relate an end-state principle of justice—maximum “social welfare”—to an allocative procedure—the market mechanism. Nozick's own theory is most deficient in failing to provide such a relationship: the first

6. See John Rawls, *A Theory of Justice* (Cambridge, Mass., 1971), sec. 42, p. 273.

part—how agents come to acquire legitimate holdings—seems to require some sort of end-state principle and is crucial in determining the entire outcome.

Unfortunately, welfare economics is itself too arbitrary in that it leaves unanalyzed the basic normative question of the choice of the social welfare function. In the next section I shall consider an alternative to classical welfare theory. This alternative leads to a more determinate answer to the distributional question as well as relating end-state criteria of justice to a procedure for achieving just allocations.

THE THEORY OF FAIRNESS

I now wish to discuss another theory of distributional justice, which, I believe, can serve as a viable alternative to the basically utilitarian theory of welfare economics, the contractual theory of Rawls, or the entitlement theory of Nozick. The theory of fairness, as I shall call it, is founded in the notion of “extended sympathy”⁷ and in the ideas of “symmetry” in the treatment of agents.

Let us consider the simplest situation of distributive justice: a group of agents has some bundle of goods to be divided among them in a “fair” way. No production is possible, and each agent’s bundle consists only of goods; no labor is present. What criteria can we use to choose a fair division? First of all, our discussion of the issues of welfare economics leads us to limit ourselves to pareto efficient allocations—for if all agents are to be made better off, we might as well do it. But then which of the efficient allocations shall we choose? We notice that the division problem is a symmetric one—no agent is privileged over any other agent—hence we wish the solution to be symmetric. But symmetric in what sense? I submit that we want the solution to be symmetric in the sense that no agent wishes to hold any other agent’s final bundle. I shall *define* an allocation that has this property as an *equitable* allocation. This is a formal definition of an abstract concept and is not necessarily meant to reflect ordinary usage. Of course, I hope to show that this definition is of interest in formalizing certain ordinary concepts of equity, for example, the con-

7. K. Arrow, *Social Choice and Individual Values* (New Haven, 1963), p. 114.

cepts of equal distribution of wealth, equal distribution of income, and so on.

Thus to determine whether an allocation is equitable, we have only to present each agent with the consumption bundle held by each of the other agents to see whether any agent would wish to exchange his bundle for another agent's bundle. If not, the allocation is equitable. It is clear that equitable allocations exist; for example, the allocation where everyone gets the same bundle—the even-division allocation—is of course equitable. However, there is no reason for this allocation to be pareto efficient; in general it will not be.

The properties of equity and efficiency are both desirable. Is it possible to find an allocation that has both of these properties? Such an allocation will be called a *fair* allocation. Notice that the concept of fairness is quite operational; we have not postulated some hypothetical welfare function or some hypothetical original position. Instead we have given a simple criterion based on the preferences of the individual agents which can be used to determine a just solution to the problem of fair division. In effect, we are asking each agent to put himself in the position of each of the other agents to determine if that is a better or a worse position than the one he is now in.⁸

The first question is, do fair allocations exist? Is it always possible to find an allocation which is equitable and at the same time efficient? The answer is yes. We can demonstrate this in the following manner: first, make an equal division of the socially owned bundle. This is certainly equitable, but if the agents have different preferences it will generally not be efficient. Next, allow the agents to trade to a market equilibrium by the use of a price system. It is important to ensure that this particular method of trading is used. Now, by the basic theorems

8. As far as I know, the original definition of equity is due to D. Foley, "Resource Allocation and the Public Sector," *Yale Economic Essays* 7 (Spring 1967). The idea of combining the concepts of equity and efficiency into the concept of fairness is due to D. Schmeidler and M. Yaari. Others who have treated this and similar topics include S. C. Kolm, *Justice et Équité*, Editions du Centre National de la Recherche Scientifique (Paris, 1972); K. Vind; D. Schmeidler and K. Vind, "Fair Net Trades," *Econometrica*, 1974; A. Feldman and A. Kirman, "Fairness and Envy," *American Economic Review*, 1974; and H. Varian, "Equity, Envy, and Efficiency," *Journal of Economic Theory*, September 1974, pp. 63–91, in which all of the results I describe below are proved.

of welfare economics, the resulting allocation is efficient; the question is, is it still equitable? Let us suppose not and derive a contradiction.

Recall the definition of a market equilibrium: each agent has chosen the bundle he prefers most out of all the bundles he can afford at the current prices. Thus if some agent *I* envies some other agent *J*—that is, if agent *I* prefers the bundle of agent *J* to his own bundle—it must be that agent *J*'s bundle costs more than the value of agent *I*'s endowment. Otherwise *I* could have bought it himself. But agent *J* and agent *I* both had the *same* endowment since the original bundle was divided equally. Thus agent *J* is holding a bundle which costs more than the value of his initial endowment, a contradiction of the definition of market equilibrium.

This approach to the question of a just distribution has the very appealing feature that one can draw on the methods and techniques of economic analysis to *prove* that fair allocations will in general exist. Furthermore, the fairness approach is an improvement on the rather general approach of welfare economics in that it actually specifies the characteristics of the solution allocation rather than allowing everything to depend on an *unspecified* social welfare function. After all, economic welfare theory is really giving us much more than we have asked for. An economic welfare function gives us a complete *ordering* of social states, when all we are really after is an answer to the question, what is the best state? It may be possible to answer the second question without answering the first. This is both a strength and a weakness of the fairness approach—a strength in that the fairness criterion provides a reasonable specification of a desirable social state and a weakness in that it does not provide any clues about how to compare nonfair allocations.⁹

FAIRNESS AND PRODUCTION

There is, however, an important criticism of the fairness model as it has thus far been presented: namely, that it ignores the question of production. For, as Nozick cogently points out, the question of *who*

9. But see H. Varian, "Two Problems in the Theory of Fairness," where I suggest a way to use the fairness idea to make second-best comparisons.

has contributed to the formation of the social product is very important when we consider the question of how to divide that product. Indeed the problem of fair division when there is no asymmetry in the agents' contribution to the social product is in fact quite simple. The real issue in the question of distributive justice is how to discriminate with respect to different contributions.

A first approach in extending the concept of equity to the production case might go something like this: let us imagine now that an agent's bundle consists not only of his goods but also of his labor contribution. We describe a consumer's position not only by how much he consumes but also by how much he works. Given a description of the technology, we can then determine the set of all feasible allocations of goods and labor and identify those which are pareto efficient. The equity concept makes sense just as before: we can ask each agent if he prefers the consumption-labor bundle of any other agent. If all agents answer no, we call the resulting allocation *fair*.

So far, so good. The concept of fairness previously introduced can be extended in a natural way to the production context. As before, we would like a proof of existence. Here, unfortunately, we run into problems. If we try to apply the old method of "divide and trade" we immediately run into a thorny conceptual issue: what do we mean by equal division when labor is present? Should we correct for ability? Should we give one agent some of the other agent's labor? Or what? Even if we resolve this dispute, the solution will be of no avail in determining the existence of a fair allocation in the production case, since it can be shown that in general a fair allocation will *not* exist in this case.¹⁰

The problem is that agents' abilities may not coincide with their tastes and, unfortunately, abilities cannot be transferred. However, it is possible to *partially* transfer ability: even if I cannot produce as much of a product as you can with the same amount of time or effort, I may be able to produce as much if I work longer or harder. There is some *degree* of substitution that can allow for a new extension of the concept of equity. Suppose we again ask each agent to compare his consumption-labor bundle with the consumption-labor bundles of the other

10. For explicit examples, see Varian, "Equity, Envy, and Efficiency"; and E. Pazner and D. Schmeidler, "A Difficulty in the Concept of Fairness," *Review of Economic Studies*, July 1974, pp. 441-443.

agents. Only this time each agent evaluates each of the other agents' bundles not on the basis of how much time each actually worked to produce his bundle but rather on the basis of how much time each agent would have to work to produce what each other agent produced. Equivalently, each agent compares his consumption-*output* bundle to the consumption-*output* bundle of each other agent. If, then, one agent prefers another agent's position, he is saying that he prefers the *complete* position of the other agent: he would rather consume what he consumes *and* produce what he produces. Hence the first agent has a legitimate complaint about the distribution of the social products. (If one agent cannot possibly produce what another agent produces, then no complaint against him can be made. Of course, problems arise with handicapped agents, etc. This seems to me to be a secondary issue that can be handled in a variety of other ways—by insurance, for example.)

If an allocation is such that each agent prefers his consumption-output bundle to that of every other agent, I will say that allocation is wealth-equitable; if the allocation also happens to be efficient, I will say that the allocation is wealth-fair. Happily, this definition does allow us to prove a general existence theorem; the basic idea of the proof is as before. One equally divides the total bundle of consumption goods, excluding labor, and allows each agent to start from this position of "equal wealth" and trade via the market mechanism to a market equilibrium. There is no correction for the different abilities of the agents; they simply buy and sell their labor at the market-wage rate. It is possible to show that the resulting market equilibrium is wealth-fair by an argument similar to the original argument.

This is the formalization of the concept of "people's capitalism" I discussed earlier. Property—in particular, productive property—is privately owned. All of the incentives of capitalism are present, as are all of the liberties, with one exception—the liberty to transfer wealth to others. Nozick would no doubt object to this scheme on the grounds discussed on page 64 of his article: it ignores the *right* of agents to give gifts. But it seems that this may be a small price to pay for such a desirable allocation. Agents give up the right of transfer, but receive instead the right of equal opportunity in the economic domain.

It is interesting to note that a wealth-fair allocation is immune to

arguments such as Nozick's Wilt Chamberlain example (p. 57). The liberty to trade does not upset the pattern of the wealth-fair allocation. When Wilt Chamberlain reaches maturity, he owns some share of society's wealth as well as his own labor power. He can sell that labor power to other agents or keep it for himself; either way is wealth-fair. The state does not interfere with such decisions; the state need only interfere with an agent's transfers when they take place outside of the market: for example, when an agent gives gifts or bequests that upset the initial endowment of goods.

Another approach to the question of fair allocation in this production context has been suggested by Pazner and Schmeidler. Again we start by considering the set of all possible efficient allocations. According to the second basic theorem of welfare economics, we can associate with each one of these a set of *market prices*—marginal products if you will—which will support these allocations as market equilibria. These prices represent the “market” evaluation of the value of each good, including of course the particular goods which are the labor contributions of each agent. Thus at each efficient allocation we can associate with each agent an implicit evaluation of his consumption bundle, which includes an evaluation of his leisure (i.e. nonlabor) time. In this way, we can associate with each agent a number, representing his *implicit income* at each efficient allocation. A possible criterion of fairness in this situation is to choose the allocation that gives *equal* implicit incomes to each agent. This allocation is called an *income-fair* allocation.

Will such an allocation always exist? Again the answer is yes. This can be shown by considering another way of dividing society's original bundle. Now we give each agent an equal division of all goods and an equal share of *each* other agent's labor. This could be done by giving everyone a ticket that would give him complete control over, say, one hour of every other person's time during some period. Given this initial endowment, the agents trade to a market equilibrium. Each agent's initial endowment has the same value since each agent holds an identical bundle—and therefore the final allocation must give everyone the same value bundle. It is this concept of equity that Nozick attacks on pages 68 and 69. Here the relationship between equal incomes and ownership of property rights in other agents is very explicit. Nozick

contends that such a redistribution of “natural assets” is grossly unjust. (See especially pp. 107–126.)

Regardless of the desirability of such a scheme, it clearly is also immune to Nozick’s Wilt Chamberlain example. For when Wilt Chamberlain reaches maturity, all of the agents can decide how to sell his labor time. If they sell it to fans of Wilt Chamberlain, fine; if they sell it to Wilt, fine. Either way, all agents’ *incomes* will be the same. The allocation will still be income-fair.

We now have three candidates for a just allocation when production is allowed: the original notion of the fair allocation, which has the defect that it might not always exist; the notion of the wealth-fair allocation; and the notion of the income-fair allocation. The last two concepts are especially interesting in that they make the central problem of justice in the production context very explicit: I refer to the problem of rewarding ability.

The wealth-fair allocation says that society will take the responsibility for an equal distribution of *goods* but that each person is entitled to complete control over his own time. Whether he is able or not, he is entitled to choose when, where, and how much he works, regardless of the social consequences. In a sense one can say that this concept favors the able at the expense of the nonable.

The income-fair allocation does exactly the reverse. It asks for a total correction for differences due to ability; it does this by ensuring that each agent has an equal share of *labor power*. If an agent is very able, his wage will be high, and consequently he will find it expensive to purchase his own leisure. In a sense, one might say that at this allocation, the able are exploited by the unable.

My own ranking of the desirability of the concepts is as follows: first, the fair allocation. Even though there may not exist a fair allocation in a productive economy, there *might* be one. And if there is, it seems ethically satisfying—no agent prefers any other agent’s position, and the economy is efficient. (It is possible to show that a fair allocation *does* exist when all agents have the same tastes, even though they may have different abilities. It is not hard to see that this is simply an equal division of goods and labor.)¹¹

11. See Varian, “Equity, Envy, and Efficiency,” p. 72, for a nonconstructive proof.

If, however, a fair allocation is impossible, which of the two other allocations, the wealth-fair or the income-fair, is more desirable? I believe I would choose the wealth-fair allocation—not, I hasten to add, for exclusively moral reasons but rather because it is so easy to organize. The income-fair allocation seems rather impractical by comparison.

Finally, I wish to point out an important fact about the fairness approach to distributive justice: it is quite compatible with the *form* of Nozick's entitlement theory. Let us specify each of the three principles of acquisition, transfer, and rectification for the concept of wealth-fairness.

(1) Agents acquire at birth (upon reaching maturity?) an initial endowment of an equal share of society's resources. Upon death, each agent's property reverts to the state to be distributed equally to new generations. Agents are entitled to their own endowments of natural assets, whatever they may be.

(2) Agents can transfer ownership of goods and services only through the market mechanism. If necessary, the state should serve as a watchdog to prevent monopolistic interference with the market. Other transfers, at least those of appreciable magnitude, are disallowed. Each good which an agent desires to sell is thus made available to all interested purchasers; there are no "private deals."

(3) We know that under such a competitive market arrangement, the resulting allocation should be wealth-fair—no agent will prefer any other agent's consumption-output bundle to his own. If some agent does prefer some other agent's consumption-output bundle to his own, he has a legitimate complaint about the allocation, and thus there are possible grounds for rectification.

Each of these principles needs to be explained in greater detail in order to provide a satisfactory theory of distributive justice. What are we going to do about acts of God, children, mistakes, small gifts, lies, malicious envy, and so on? If these questions can be answered in a satisfactory way, the idea of fairness may provide a very attractive theory of justice that combines the considerations of both procedural justice and distributive justice.