On Pareto efficiency and equitable allocations of resources

Florin Popa

This paper discusses some problems associated with the concept of Pareto efficiency: first of all, the difficulties of defining what is a Pareto improvement and a Pareto optimum in social sciences (and economics in particular), and secondly, the relation between Pareto efficiency and equity of distributions of resources.
In particular, we want to point out some ambiguities and contradictions of the concept of Pareto efficiency, which makes it hardly usable in real-world social contexts. When qualitative variables are involved, there are situations in which there is no straightforward way of deciding a preferable optimum within the Pareto set, other than taking each individual's preference and choice into account when establishing what is and what is not an improvement (which is in many cases impossible) or having a ruling authority to establish that (which is undesirable).
With this concept, we seem to be caught between the subjectivity of personal preference and the arbitrariness of dictatorship.
JEL code: A 10

A system of resource allocation is Pareto-inefficient as long as there are possible adjustments to it that make at least one individual better off while not making anybody else worse off. Usually, there are many possible Pareto improvements, which raises the question of what to choose (and on what grounds). When there are no more improvements to be made (meaning than any change in the distribution of resources would make at least one person worse off), the system is said to be Pareto optimal.

A trivial example of a Pareto inefficient allocation is one that does not use all the resources available. When there is more to be offered to the participants (either all or just some of them), improvements can be made until the 'stock' is emptied. But as the majority of allocations involve scarce resources, the real problem is distributional: it involves deciding how to distribute a fixed amount of resources (how to divide the pie). This, in turn, raises the question of what criteria we use to determine a given distribution of resources.

Pareto efficiency requires only that, compared with the initial allocation of resources, the modifications make at least one individual better off without harming anybody.

Determining the Pareto set is simple when there are only quantitative parameters to be taken into account: the problem is solved by solving a set of equations. Each equation illustrates a relation of interdependence of two or more parameters. In engineering,

for instance, drawing the Pareto frontier is part of product design. By considering only the potentially optimal solutions, the designer restricts the range of possible designs to a limited set, instead of taking into account the full range of all parameters.

Economics, on the other hand, operates with quantitative as well as qualitative parameters. Individual preferences cannot be considered purely quantitative variables, and collective preferences even less so. Moreover, getting from individual to collective preferences is problematic in many cases. Individual and collective utilities fall in the same category: any naïve utilitarian attempt to treat them as purely quantitative variables will break down in the face of their fundamentally subjective character and irreducible qualitative features.

Condorcet's 'paradox of voting' illustrates this situation: suppose we have a group which has to take collective decisions using majority rule. For the sake of simplicity, it has only 3 members. The ordering of preferences is the following:

X: a preferred to b, b preferred to c,
Y: c preferred to a, a preferred to b,
Z: b preferred to c, c preferred to a.

This means that:

X and Y prefer a to b,
X and Z prefer b to c,
Z and Y prefer c to a.

Therefore, there is a majority which prefers a to b and b to c, but this does not imply a preference of a over c (non-transitivity). When trying to aggregate preferences by majority rule, we get from individual preferences that are consistent and transitive to collective preferences that are no longer transitive.

Moreover, if we take into account ordinal utilities, we can compare different utilities of the same individual and obtain a hierarchy. But we cannot aggregate utilities of different individuals, since we do not know exactly what 'I value A over B' and 'you value A over B' means. There is no absolute point of comparison or common scale of evaluation. The order of preference is, by itself, not sufficient in order to aggregate utilities consistently.

On the other hand, cardinal utilities can, in principle, be used to compare utilities trans-individually. However, in order to treat economic variables as cardinal, most of the time we adopt simplifying assumptions. Even when we are talking about monetary value - a seemingly undisputed cardinal variable - there are still ordinal features that cannot be quantified. The marginal utility of different types of benefits (including money) decreases as the available quantity increases. This relation of indirect proportionality can only be described in ordinal terms: we can compare and order preferences, without being able to allocate clearly specified values to each.

*

At individual level, the stability of preferences seems to be an essential assumption of a broad explanatory model of economic behavior. For Gary Becker, it is one of the three characteristic features of the ‘economic approach’ to decision and action. The other two are ‘maximizing behavior’ (consistently choosing the options which maximize utility) and ‘market equilibrium’.

Apart from stability of preferences, James G. March lists other commonly used assumptions:

1 Preferences are absolute. The individual takes his own preferences for granted, as given, without questioning their relevance or 'rationality'.

2 Preferences are relevant (decisions are consistent with preferences).

3 Preferences are consistent and transitive. Inconsistency of individual preferences can be explained as imperfect information (the person has not analysed his/her own preferences well enough).

4 Preferences are well-defined. The individual can specify a straightforward and quantifiable relation between outcomes of his actions and preferences.

5 Preferences are exogenous. They determine decisions, but are not affected in turn by these decisions.

Most of these assumptions are systematically contradicted by real-life behavior of economic agents: they do not treat their preferences as stable, well-defined and exogenous. Individuals 'manage' their preferences imperfectly: they have a role in creating or modifying them, but at the same time they confound them, might try to avoid or suppress them or simply ignore them. The inconsistency or intransitivity of individual preferences cannot always be explained as lack of sufficient information. Lastly, decisions are not always consistent with preferences – which questions the assumption that we 'reveal' our preferences through our decisions and actions. The modification of preferences can be explained on the basis of psychological or sociological variables, technological development, or as accumulation of 'human capital' (for instance, taking part in a wine tasting course can modify the hierarchy of preferences, because the individual had become more efficient in generating utility out of this activity).

What is the standpoint for determining Pareto efficiency at collective level? The most sensible answer here is: the aggregation of individual evaluations regarding the advantages or disadvantages of each individual affected by a certain reallocation. In other words, the aggregation of individual orderings of preferences. But once we adopt this view, we are forced to recognize that there is no straightforward way of determining how Pareto efficiency may be brought about. There is no objective standpoint for judging what is an individual gain or loss, or what means that we (as a group) are better or worse off. This stems from the basic fact that preferences and utilities of different individuals cannot be compared in any straightforward way, since they value differently the same benefit or gain. There is no common 'objective' denominator of preferences and utilities; therefore, there is no trans-individual authority that can determine (or rather calculate) orderings of preferences.

On the other hand, there is no guarantee that such an aggregation of individual preferences will yield a consistent result, that can be used to determine the set of Pareto optimal point. Kenneth Arrow's general possibility theorem (sometimes called Arrow's paradox) aims to demonstrate that, under certain (plausible) conditions of aggregation, there is no social choice function (SCF) that can aggregate the individual hierarchies of preferences into a social hierarchy.

The conditions imposed on SCF are the following:

1. *non-dictatorship*: SCF should not take into account the ordering of preferences of only one individual, ignoring the others,
2. *unrestricted domain*: SCF must take into account all possible sets of individual orderings of preferences and must be able to rank all possible choices to one another,
3. *independence of irrelevant alternatives*: the social ordering of alternatives x and y for two individuals A and B is solely a function of individual orderings of preferences of A and B, and does not take into account orderings of irrelevant alternatives (for instance x and z),
4. *non-imposition*: any possible SCF is obtained on the basis of one set of individual preference orderings,
5. *monotonicity*: if an individual prefers x to y and all the other individuals either prefer x to y or are indifferent between the two, then the SCF should also express a preference of x over y.

Arrow's theorem states that, for situations with at least two participants and three options, it is impossible to define a SCF which respects all five conditions. In particular, it can be shown that a SCF complying with conditions 1, 2, 3 and 5 will have at least one preference ordering determined by a dictator.

Of course, the relevance of this result depends on accepting the five requirements imposed on the social aggregation of preferences. A lot of work has been devoted to adjusting or restating these conditions, in order to eliminate the paradox. However, the intuitive strength of the initial five conditions cannot be denied: they seem to embody common assumptions regarding the way in which we aggregate our preferences rationally.

Although economics uses frequently qualitative variables that can hardly be subjected to a fully quantitative treatment, the concept of Pareto efficiency is traditionally linked to markets, competition and equilibrium. According to the first theorem of welfare economics, a free competitive market will produce, under certain conditions, Pareto-efficient distributions. In other words, in conditions of Walrasian equilibrium, the invisible hand will produce Pareto-efficient distributions most of the time. If true, this would certainly be a strong reason in favor of market non-intervention. Kenneth Arrow and Gerard Debreu offered a mathematical demonstration of the theorem. However, it is based on some simplifying assumptions:

1 markets are perfectly competitive (no monopolies, oligopolies or other forms of coordination which distort the competition),
2 all possible goods are represented on the market,
3 transaction costs are negligible.

The use of simplifying assumptions is notoriously tricky: in the best case, they embody fundamental features of the analyzed phenomena, and leave out non-essential

features, whose presence or absence does not change the nature of the phenomenon itself. But there is no guarantee that certain assumptions will not leave out essential parameters. And if they do, then the application of results obtained on the basis of those assumptions to facts is problematic. It might well be the case here. Can we assume that there are real-world economic systems that can be described accurately using the simplifying model of a perfectly competitive market?

Perfect competition implies perfect and complete information, homogeneity (goods and services can be substituted for one another), independence (all market participants act independently), equal access to resources and technologies. It also requires assumptions regarding the rationality of agents: they analyze objectively, on the basis of perfect and complete information, all possible decisions and choose the one that maximizes utility. Social systems with a large number of agents do not come even near to respecting any of these criteria.

According to the second theorem of welfare economics, the best way to make an efficient distribution of resources sustainable is to subject it to market mechanisms. No central planning is better than the action of the invisible hand in preserving the Pareto efficiency of a certain distribution. There are two problems associated with this approach.

One of them is that it relies on simplifying assumptions about perfect competition even more than the first theorem. There is no guarantee that the efficiency will be preserved as long as there are even minor distortions. The second is that it completely ignores the way the distribution of resources was made. It just takes it for granted. If the government is the authority which makes all initial distributions, does it have the necessary (perfect and complete) knowledge about consumers' preferences and production processes?

If the government simply tries to distribute resources as the market would do, trying to balance supply and demand, why not let the market take over from the very beginning? What is the point of having a ruling authority at all? This would only appear justified in order to correct significant imbalances or distortions on the market – that is, implicitly recognizing that the simplifying assumption about the market equilibrium is false. On the other hand, if the government makes the distribution according to specific policy goals – in other words, it redistributes resources outside the market mechanism – the problem is where to draw the borderline of government intervention.

However, even if we assumed that – in a significant number of cases – equilibrium leads to Pareto-efficiency, or that market equilibrium preserves and reinforces the Pareto efficiency, this actually tells us very little.

First of all, it does not allow us to specify commonly accepted criteria for choosing between many possible Pareto-efficient distributions. If I decide to share a cake with two other persons and each gets a third, this is Pareto-efficient. But if they decide to share the cake between themselves and leave me out of it, this is also a Pareto-efficient distribution. Is there a way of deciding which allocation is better? As long as I am strictly interested in Pareto efficiency and leave equity out of the equation, there is no common criterion for choosing a point of optimum among the Pareto frontier. From a purely technical perspective, all allocations are efficient and there is no criterion for choosing

an optimum. Of course I can still hope that, through negotiation or persuasion, the participants will decide to go for the more equitable distribution (a third for each). But this solution is not taken for granted.

Suppose John and Peter work for the same company but have radically different levels of income – the first makes $100000 a year while the second makes only $30000 a year. Peter considers that, given his level of responsibility and complexity of tasks, he is underpaid. In other words, he questions the equity of salary distribution. But it doesn't mean that the distribution is Pareto inefficient: any change in the current distribution would surely make somebody worse off. There is no link between the Pareto efficiency and equity of one given distribution of resources. Actually, Pareto efficiency protects the status quo, as any modification to an existing distribution, in order to include newcomers, would imply a reduction in the current share of resources.

Who decides who is better or worse off? What is the criterion or ruling authority to determine this?

At first glance, it seems unproblematic: the more you get, the better off you are. The more they take away from you, the more they harm you. It should be obvious for anybody, given that we only have to compare numbers and quantities. But many individual decisions imply qualitative variables. Usually we try to convert qualitative variables into quantitative variables, on the basis of simplifying assumptions. This is what we do when assuming that, faced with the same set of alternatives, equally informed and equally capable individuals will generally take the same decision. Game theory relies heavily on this sort of 'quantification'. But this conversion does not always work and, even when it does, it is just on the basis of assumptions (regarding the homogeneity of preferences and utilities in different individuals). Assumptions are hypotheses and, as such, they are fallible.

Let's go back to John and Peter. They might have different levels of income, but they share a hobby: skiing. They both like to take at least a few days off to go skiing each winter. Now let's suppose each of them is offered a salary raise of $1000. Clearly there is a difference in the way they value the additional money: for John it might not change anything, since he already has a level of income that allows him to pursue his hobby in style. However, for Peter the salary raise might mean the possibility of taking the much desired ski trip, a possibility which was previously unavailable for financial reasons. In other words, the marginal value of the $1000 raise is significantly different for John and Peter.

Let's take one more step. Suppose somebody raises the question of redistributing the additional money, in order to offer an incentive to employees with lower levels of income. What would be a Pareto efficient redistribution? If John finally got only $500 and Peter got $1500, would the first feel worse off because he lost half of the raise? Or would it be too small a sum to bother about losing it? Or would he feel that the 'good deed' is well worth the $500 he gives away, that the moral satisfaction he gets from contributing to a more equitable distribution of salary surpasses his personal financial loss?

Obviously, these are questions only John is in a position to answer. He is the only authority to decide what would be an improvement or not in his case.

Secondly, Pareto efficiency does not necessarily lead to socially desirable allocations. There is no guarantee that the Pareto distributions will be equitable. An oligopoly can be Pareto efficient, just as a competitive market can be. Amartya Sen developed this criticism, stating (plausible) conditions under which systems of social choice will generate, most of the time, Pareto efficient but inequitable allocations.

Lastly, Pareto efficiency does not take into account the origin of resources or certain rights or entitlements associated with it, but only the present distribution of these resources. This preference for the status quo does not mean that the initial distribution of resources is not important. It only means that, for purposes of efficiency, it is irrelevant and should be considered as it is. This has two undesirable consequences. First, any inequity in the initial distribution will be preserved and perpetuated by an efficient market. Second, Pareto efficiency is unable to help us distribute resources that nobody has produced, such as land or newly discovered (natural) resources.

The concepts of *Pareto efficiency* and *Pareto optimum* are marked by an ambiguity which weakens their applicability to real-life social contexts. The first difficulty concerns the specification of the Pareto set, in dominantly qualitative contexts. Attempts to turn qualitative into quantitative variables are based on simplifying assumptions whose legitimacy is questionable. The second difficulty regards the possibility of specifying a criterion to choose a 'preferable' point of optimum among the Pareto set. The third one concerns the social consequences and ethical implications of this ambiguity: socially undesirable allocations and ignorance of rights/entitlements over resources or their origin.

Consequently, their relevance and use of these concepts is restricted to clearly specified domains based on quantitative (or quantifiable) parameters. The attempt to apply them to dominantly qualitative socio-economic variables can lead to counter-intuitive or inconsistent results.

Florin Popa
Licenţiat în Filosofie (Universitatea Bucureşti, 1999) şi Ştiinţe Sociale (Nottingham Trent University, Marea Britanie, 1997). Doctorand al Universităţii Bucureşti cu o teză despre criteriile/normele de raţionalitate ale deciziilor politice. Domenii de interes: teoria deciziei, raţionalitate, teoria jocurilor. Articole: "Competing Modes of Explanation in Social Sciences", Social Science Review (Nottingham, Marea Britanie, nr. 2/1997),
"Modele de decizie raţională: criterii vs. norme", Revista de Filosofie (în curs de apariţie).