

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

ROEMER ON EQUALITY OF OPPORTUNITY

Matthias Hild
California Institute of Technology &
Jet Propulsion Laboratory

Alex Voorhoeve
Departments of Philosophy and Economics
University College London



SOCIAL SCIENCE WORKING PAPER 1128

December 2001

Roemer on Equality of Opportunity

Matthias Hild

Alex Voorhoeve

Abstract

We critically discuss John Roemer's normative criterion of equal opportunity and advance opportunity dominance as an alternative criterion for the evaluation of policies. For Roemer, strict equality of opportunity obtains when people, irrespectively of circumstances beyond their control, have the same ability to achieve advantage through their free choices. To implement this idea, Roemer sorts people with similar circumstances into types and takes their free choices to be represented by their behaviour relative to other members of the same type or, as Roemer calls it, by their 'relative effort'. He then proposes that society should maximize the average advantage of all those whose circumstances cause them to be least well-off relative to others who have expended the same degree of relative effort.

We argue that typing and the relative effort metric conflate the factors for which we do and do not want to hold individuals responsible, whenever these factors are statistically correlated. Moreover, Roemer's rule for policy selection burdens the concept of equal opportunity with foreign distributive principles. Pointing to an inconsistency in Roemer's argument, we also note that his selection rule violates his own requirement of Pareto-optimality. In response to these difficulties, we advance the criterion of opportunity dominance which is Pareto optimal, maintains conceptual purity and does not conflate the factors for which individuals should and should not be held responsible. This criterion determines a set of candidate policies that are undominated in opportunity and from which a final policy must be selected by additional, conceptually distinct principles of distributive justice.¹

JEL classification numbers: D63

Key words: equality of opportunity, fairness

¹The order of the authors is alphabetical and does not in any way reflect a difference in contribution to this paper.

Roemer on Equality of Opportunity

Matthias Hild*

Alex Voorhoeve[†]

1 Introduction

Equality of opportunity figures prominently in debates on the nature of a just society and the achievement of some form of equal opportunity constitutes a professed goal of social policy in most Western democracies. Nonetheless, the precise meaning of this concept remains controversial. Conceptions of equality of opportunity range from the absence of discrimination on the basis of race or gender to the elimination of the influence of social background and natural differences on the attainment of advantage. In his pathbreaking work, John Roemer has in recent years attempted to overcome this impasse. He has advanced a precise definition of equality of opportunity and formulated an explicit rule for the selection of equal-opportunity policies (Roemer 1993, 1996, 1998). His proposal has been widely discussed and is being applied to the analysis of distributive issues in schooling, health care, taxation and development aid (Roemer 1999, Llavador and Roemer 2001). Despite of these achievements, we believe Roemer's proposal to be seriously flawed. This paper will present our main points of criticism, propose an alternative definition of equal opportunity and construct a selection mechanism that allows for a partial ranking of opportunity policies.

Before we proceed with our discussion, we propose the following abstract definition of equal opportunity. When we say that equality of opportunity with regard to a certain advantage should obtain for a group of people, we mean that certain factors should and other factors should not influence these people's differential attainment of advantage. Let us refer to the first set of factors as people's 'relevant characteristics' and to the second set of factors as their 'irrelevant characteristics'. Strict equality of opportunity thus obtains if differences in people's irrelevant characteristics do not result in differences in their attainment of advantage. To make this general concept concrete, we have to specify

*Address for correspondence: California Institute of Technology, Mailcode 228-77, Pasadena, CA 91125 (USA). Email: matthias@hild.org. Website: <http://www.hild.org>.

[†]Email: a.voorhoeve@ucl.ac.uk. Website: <http://www.ucl.ac.uk/~uctyaev>.

We thank Brian Barry, Alex Brown, Jerry Cohen, John Roemer, Peter Vallentyne, and Jo Wolff for helpful discussions. Predecessors of this paper were presented at the London University Graduate Conference, the Warwick Conference on Political Theory and the Analytical Philosophy National Postgraduate Philosophy Conference in Reading the participants of which we also thank for their comments.

not only the advantage in question, but also which characteristics we regard as morally relevant and irrelevant. It is in these specifications that disagreement among proponents of various forms of equal opportunity arise. Suppose, for example, that we are concerned with the advantage of life–time income and we wish to compensate people for those differences in advantage that result from their parents’ wealth. Suppose furthermore that we have a simple predictive model of life–time earnings in which individuals are described by only three characteristics, namely years of schooling, low or high IQ and poor or rich family background. Imagine we consider parental income as morally irrelevant, while we come to consider IQ and years of schooling as morally relevant in this context. Strict equality of opportunity then requires us to equalize the advantage of all individuals with identical relevant characteristics, regardless of their irrelevant features. In our example, this would mean the equalization of the attainment of advantage between poor and rich individuals with the same amount of schooling and the same intelligence. If we, however, regarded not only parental income but also IQ as morally irrelevant, we would reach a different conclusion. Strict equality of opportunity would then obtain when only the individuals’ amount of schooling, but not their intelligence or their parental income, determines their attainment of advantage. Different specifications of advantage and of relevant and irrelevant characteristics thus result in different normative conceptions of equality of opportunity and their specification involves an inescapable value judgement.

Recently, several authors have proposed a particular choice of the morally relevant and irrelevant characteristics (Arneson 1989, Cohen 1989). According to this conception, the characteristics of moral relevance are precisely those characteristics over which people have control. By contrast, people’s circumstances and other factors beyond their control (or ‘circumstances’ for short) should not be a source of differential levels of advantage. The motivating idea of this conception is that equality of opportunity demands people to be equally able to achieve the same outcomes through their free choices, but that inequalities of outcomes are permissible when caused only by differences in people’s free choices. Proponents of this conception of equal opportunity must address two crucial questions. Firstly, how do we determine when individuals, faced with different circumstances, have made similar choices and therefore have a claim to be made equally well–off? Secondly, which distributive mechanism can best compensate people for the effect of circumstances on their attainment of advantage? John Roemer’s work responds to both of these questions and, in a nutshell, makes the following proposal. We are to select some representation of what we regard as people’s circumstances as well as some measure of people’s actual choices which Roemer calls ‘effort’. We then sort people with similar circumstances into ‘types’. Since the expenditure of actual effort may be influenced by circumstances, Roemer takes individuals’ free choices to be represented not by their actual effort, but by their effort relative to other members of their type. For Roemer, strict equality of opportunity has been established when people in different types who expend the same degree of relative effort receive the same advantage. For practical purposes, he then proposes that society should choose a policy that maximizes the average advantage of those whose circumstances cause them to be the worst–off at a given level of relative effort (where the average is taken over all degrees of relative effort).

We shall argue that Roemer’s proposal is irrecoverably flawed. In Section 2, we take issue with Roemer’s representation of people’s relevant and irrelevant characteristics. We show that typing and the relative effort metric conflate people’s relevant and irrelevant characteristics in the presence of statistical correlations among these characteristics. By contrast, the simple formulation of the equal–opportunity criterion with which we started this paper suffers from no such flaw. Moreover, we note that Roemer’s measure of relative effort can be a product of an individual’s circumstances just as much as absolute effort. Roemer’s construction therefore falls short of its own ambition to create an interpersonally comparable measure of individuals’ free choices.

In Section 3, we take issue with Roemer’s rule for selecting an equal–opportunity policy. We argue his rule burdens the concept of equal opportunity with utilitarian and desert-based distributive principles that are foreign to the idea of equal opportunity. In our view, a criterion of equal opportunity should be purely concerned with compensating people for the effects of their irrelevant characteristics, and should steer clear of any commitment to principles governing the outcomes that people achieve due to their relevant characteristics. We also point to an internal inconsistency in Roemer’s policy rule that arises from the fact that his rule violates his own requirement of Pareto–optimality. In response to these problems, we advance a criterion of opportunity dominance which maintains this conceptual purity and is Pareto optimal. This criterion selects a set of candidate policies that are undominated in opportunity. If the set contains more than one policy, further choices have to be made on the basis of additional, conceptually distinct, principles of distributive justice. We suggest that future research should explore the combined effect of familiar distributive principles, such as utilitarian or desert–based principles, and the criterion of opportunity dominance.

2 Choice and Circumstance

2.1 The relevance of circumstance

On Roemer’s view, attainment of equal opportunity requires a distributional mechanism that compensates people for the influence of factors beyond their control, but allows differential outcomes insofar as they result from differences in people’s free choices. Equating morally relevant characteristics with the factors that are within a person’s control, this conception aims to incorporate a particular ideal of responsibility for outcomes into egalitarian theory. According to this ideal, people should be held responsible for the outcomes of exactly those choices that were free and unaffected by circumstances. Contemporary discussions of the notion of responsibility for outcomes cast doubt on the validity of this particular conception and emphasize that responsibility for outcomes does not have to be congruent with the notion of personal control (cf. Hart and Honoré 1959, Fleurbaey 1995, Scanlon 1998, chapter 6). In our opinion, the distinction between factors within and factors beyond a person’s control may at best act as a guide in the determination of moral relevance, but we must not expect the notion of moral relevance and of personal

control always to coincide. Nevertheless, in what follows, we shall discuss the merits of Roemer's proposal on its own terms and accept the premise that people should be held responsible exactly for the factors that are within their control.

Based on this premise, Roemer sets out to devise a method to determine what we can consider to be people's free choices and also to determine when individuals faced with different circumstances have made similar choices. Roemer (1998, 8–9) conceives of this distinction between choice and circumstance as involving two distinct questions. Firstly, what kinds of factors, such as family background and natural talents, are beyond a person's control? Secondly, given these factors and their influence on people's choices, what part of their behaviour can we attribute to their free choices? Roemer remains silent on the first question, leaving it to be decided by society through some political process. Once society has decided what to regard as factors beyond a person's control, Roemer constructs a 'typing' of the individuals concerned. For any list of circumstances, such as social class, parental income, race, natural talents etc., we are to divide society into types according to individuals' value of these factors, with people who have the same values of these factors falling into the same type. A type may thus consist of all individuals with sufficiently similar parental income or of all individuals with the same IQ. By construction, individuals within the same type are considered to be in the same circumstances and, thus, to have an equal ability to achieve advantage. Within each type, therefore, people's differential achievement is considered to be entirely due to the different choices they made.

The challenge is, then, to isolate the effects of circumstances on the attainment of advantage for people in different types. These effects can be either direct or indirect. If, for example, income is the advantage in question, the children of wealthy parents typically have an increased income directly through parental gifts or bequests. But circumstances can also affect the attainment of advantage indirectly, via their influence on the choices that people make. People's actual choices are dependent on their beliefs, preferences, the range of options open to them and on their deliberative process. All of these can be significantly influenced by factors beyond a person's control. Indirect influences of this kind are notoriously difficult to address since they imply that people's actual choices, due to their dependence on uncontrollable circumstances, are inappropriate grounds for assigning advantage. Roemer therefore proposes to 'go behind' people's actual choices, stripping them of the influence of circumstance to reveal their 'real choices' for which we can properly hold them responsible.

Roemer asks us to choose some observable measure of people's choices that is related to the advantage in question and which he calls 'effort'. He then ranks all individuals within a type according to their actual expenditure of effort. This ranking then allows Roemer to partition people within each type into quantiles of effort relative to their type. A person i belongs to the π -quantile of relative effort exactly if a fraction π of the entire population expends at most as much effort as i . For example, if a population is uniformly distributed between 0 and 2 hours of work, then the person who works for one hour will be at $\pi = 0.5$ in the distribution of relative work. Roemer argues that this measure of

type–relative effort reveals a person’s ‘real choices’, unaffected by circumstance. He thus uses the metric of relative effort to compare the choices of people belonging to different types and he considers people to have made the ‘same choice’ exactly if they occupy the same position in the relative effort distribution of their respective types. In effect, Roemer takes the individuals’ places in the relative effort distribution of their types to be their only relevant characteristic and concludes that strict equality of opportunity obtains when all people who expend the same degree of relative effort are equal in advantage (Roemer 1998, 15–16).

2.2 Typing misrepresents circumstance

In this section, we argue that Roemer’s method of typing leads to a conflation of the characteristics for which we do and do not want to hold individuals responsible when both kinds of characteristics are statistically correlated with each other. Suppose, for example, that we are concerned with health care insurance and that we wish to compensate people for their exposure to a persuasive and misleading advertizing campaign for cigarettes. Our example assumes that the cigarette manufactures targeted a particular subgroup of the population and that exposure E and non–exposure \bar{E} to this campaign was beyond the individuals’ control and is considered morally irrelevant. For simplicity’s sake, we also assume that up to the start of the campaign people’s decision to smoke was the result of an unimpeachable process of free choice. We therefore regard it as morally relevant whether a person was a smoker S or non–smoker \bar{S} before the campaign. Following Roemer’s methodology, we would now sort people into two types, the set t_E of people who were exposed to the campaign and the set $t_{\bar{E}}$ of people who were not. On Roemer’s method, any difference between the smoking behaviour of these two types would be attributed to the advertizing campaign and, consequently, would occasion a redistribution of social endowments. Now, imagine that those individuals singled out by the cigarette manufacturers for their advertizing campaign were exactly the individuals who did not smoke before the campaign. Hence, the set of individuals with characteristic E and the set of individuals with characteristic \bar{S} are identical. Imagine further that, although the advertizing campaign was effective and increased the proportion of smokers in the exposed type, people exposed to the campaign still tend to smoke less than members of the unexposed type of long–term smokers.

In this example, Roemer’s method would treat the exposed type t_E as if it were advantaged and the non–exposed type $t_{\bar{E}}$ as if it were disadvantaged by circumstance. To see this, suppose we had one person from each type, each smoking an equal amount of cigarettes. Since the average cigarette consumption in the exposed type is lower than in the non–exposed type, the person from the exposed type would be considered to have tried less hard to resist smoking than the person from the non–exposed type. Translated into a concrete policy, this might mean that the person from the exposed type would have to pay more for the same package of health care insurance. This is contrary to reason because, by assumption, both individuals had been unimpeachably free in their choices up to the start of the advertizing campaign and we set out explicitly to compensate for

differences arising from exposure to the advertizing campaign.

What this example shows is that Roemer's typing methodology cannot deal with correlation between factors for which we do and do not want to hold people responsible. The difficulty with his proposal occurs only because his formalism represents irrelevant characteristics, such as exposure to the advertizing campaign, by the set of individuals who exhibit these characteristics. This method goes astray when there exists a statistical correlation between relevant and irrelevant characteristics. In our example, we used an extreme form of correlation in which relevant and irrelevant characteristics were exhibited the very same set of individuals, but the same phenomenon also occurs of course for less extreme forms of correlation.

To avoid this problem, we must not represent irrelevant characteristics indirectly by the set of individuals who display these characteristics. Instead, once we are given the characteristics that are used to describe the individuals and to predict their expected outcome relative to our econometric model, we must sort these characteristics into a set of morally relevant characteristics and a set of morally irrelevant characteristics. Suppose that $u(\phi, y, z)$ is the (expected) outcome under policy ϕ of an individual with a combination y of relevant characteristics and a combination z of irrelevant characteristics. Strict equality of opportunity would then be achieved by a policy ϕ which, for any combinations z and z' of irrelevant characteristics, equalizes the functions $u(\phi, \cdot, z)$ and $u(\phi, \cdot, z')$:

$$u(\phi, y, z) = u(\phi, y, z') \tag{1}$$

for any combination y of relevant characteristics and any combinations z and z' of irrelevant characteristics. Our example used a simplistic econometric model in which individuals are described by only two characteristics, namely the irrelevant characteristic of exposure E or non-exposure \bar{E} and the relevant characteristic of voluntary smoking S or non-smoking \bar{S} in the past. Strict equality of opportunity requires a policy ϕ that, firstly, equalizes the outcome $u(\phi, S, E)$ and $u(\phi, S, \bar{E})$ of exposed and non-exposed past smokers and, secondly, equalizes the advantage $u(\phi, \bar{S}, E)$ and $u(\phi, \bar{S}, \bar{E})$ of exposed and non-exposed past non-smokers. (In our simple example of a perfectly targeted advertizing campaign, this is an empty constraint since there do not exist any individuals displaying both S and E or displaying both \bar{S} and \bar{E} .) However, equal opportunity does not require us to equalize the outcome $u(\phi, \bar{S}, E)$ and $u(\phi, S, \bar{E})$ of exposed past non-smokers and non-exposed past smokers. It is exactly this unwarranted compensation of the differences in the opportunity of exposed past non-smokers and non-exposed past smokers that is implied by Roemer's method of typing and that is easily avoided by our proposal.

2.3 Relative effort is also a product of circumstance

As we discussed above, Roemer’s reason for representing irrelevant characteristics by the set of individuals who display these characteristics is part of his method for comparing the choices of people in different circumstances and for determining when people in different circumstances are considered to have made the same choice. In our view, the problems of the typing methodology that we rehearsed above already discredit the measure of relative effort as the basis for assessing equality of opportunity. To return to the case of the advertizing campaign, it would be a mistake to conclude that an exposed and an unexposed individual who both smoke to the same degree relative to their own type are entitled to the same outcome. If we accepted this conclusion, we would have to transfer resources away from the very people whom we intended to compensate for their exposure to the campaign.

On a more fundamental level, we will now argue that Roemer’s measure of relative effort does not solve the problem it was designed to address. As we saw, Roemer suggests that absolute effort is an unsuitable basis for assessing equality of opportunity because “the distribution of effort is a characteristic of the type, and hence not something the person should be held accountable for” (Roemer 1998, p. 11). Roemer’s suggestion of moving to the metric of relative effort could only solve this problem if the distribution of relative effort was not also influenced by factors beyond individuals’ control. As a matter of fact, however, the distribution of relative effort is just as much a product of circumstances beyond people’s control as the distribution of absolute effort. Individuals base their choices to expend effort in part on the reward schedule for effort that is jointly determined by their circumstances and the policy they face. By implication, different circumstances and different policies therefore differentially affect an individual’s choice of effort. Under a policy that results in different reward schedules for individuals in different types, we therefore expect to see different distributions of absolute and relative effort for different types. Moreover, individuals typically differ in their response to the very same reward schedule due to different preferences for combinations of effort and reward. If these responses are distributed differently across different types, we also expect to see different distributions of absolute and relative effort for different types, irrespective of the policy chosen. The distribution of relative effort therefore typically differs across types just as much as that of absolute effort.

As an illustration, imagine a population from different ethnic backgrounds and a policy maker who is interested in establishing equal opportunity for life–time income. Since an important choice factor contributing to people’s income is their amount of schooling, we may decide to use years of post–mandatory education as a measure of effort and to sort people by types according to their ethnicity. Suppose that, some type *I* exhibits a large range of different effort levels with its members equally distributed over the levels of 1, 2, 3 and 4 years of post–mandatory education. We assume that types differ in their educational choice because of a culture that discourages higher education in type *II*. As a result, 75 percent of all members of type *II* pursue 1 year of additional education while only 25 percent of all members of type *II* pursue 2 years. Clearly, in this case, the

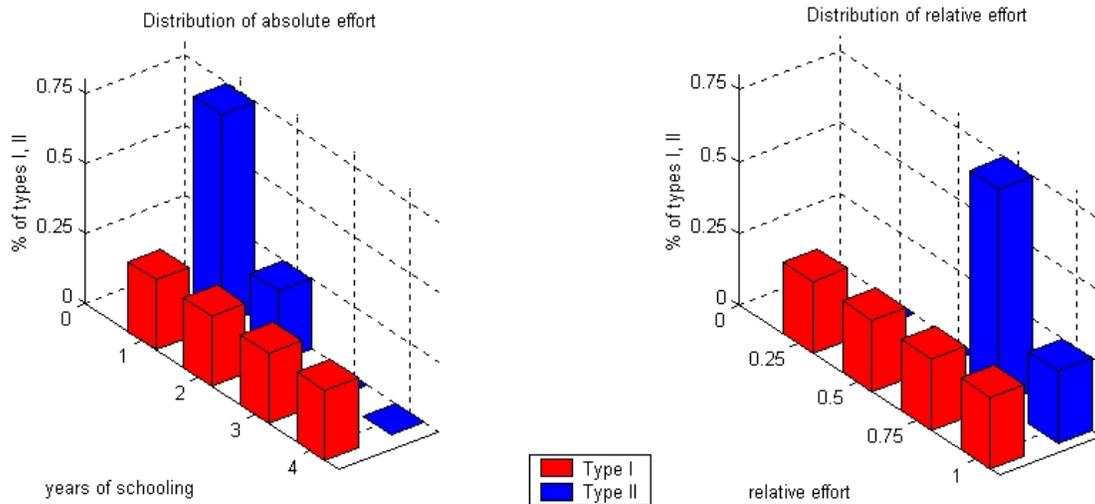


Figure 1: Distribution of absolute and relative effort

distribution of absolute effort is a characteristic of the type to which individuals belong. The same holds true, however, for relative effort. To see this, consider the distribution of relative effort. According to Roemer's definition, an individual i expends relative effort π if and only if π is the fraction of individuals of the type who expend a level of absolute effort of at most the amount expended by i . Figure 1 shows that both the distribution of absolute and the distribution of relative effort differ across type I and II .

Note that this simple example does not make any reference to the advantage received by the individuals. For all we have said, a policy that equalizes opportunities between the two types may already be in place, giving an equal life-time income to the individuals in different types who are at the same position in the relative effort distribution of their type. (In this example, this would mean that individuals in type II with 1 year of higher education obtain the same life-time income as individuals in type I with 3 years of higher education, and those individuals in type II who have 2 years of higher education obtain the same life-time income as those in type I who have 4 years of higher education.) There is hence no reason to suppose that the policy that equalizes opportunities between these two groups will produce the same distribution of relative effort in each type. A type-dependent variation of the distributions of both absolute and relative effort may occur regardless of individuals' advantage functions and regardless of whether or not these advantage functions equalize opportunity. In sum, relative effort can be just as much a product of circumstance as absolute effort and the appeal to relative effort therefore fails to solve the problem that Roemer set out to resolve. Roemer's construction therefore falls short of its own ambition to create an inter-personally comparable measure of individuals' free choices.

2.4 Typing misrepresents relevance

For each type t , Roemer assumes the existence of a function $u_t = u_t(\phi, e)$ that yields the (expected) advantage of an individual of this type depending on some relevant characteristic e (the effort variable) and on policy ϕ . Since the individuals in a given type share some, but not all of their personal characteristics, we cannot expect that they all enjoy the same advantage at a given level of effort and under a given policy. Suppose, for example, that we are interested in equal opportunity for life–time income among people from different social backgrounds. Suppose we choose years of post–mandatory schooling as our effort variable. We sort people into two types depending on whether they come from a ‘rich’ or ‘poor’ family background. Within each type, these people will differ in some additional characteristic, such as intensity of job–seeking activity, that influences their expected life–time income at a given level of schooling. Thus, the assumption that all individuals of a given type have the same function of advantage to effort will typically be false. Roemer addresses this problem by constructing the function u_t as the average of the advantage functions of the individuals in type t (Roemer 1996, p. 280, Llavador and Roemer 2001, p. 159):

$$u_t(\phi, e) = \frac{1}{I_t} \sum_{i \in t} u(\phi, e, x_i), \quad (2)$$

where I_t is the number of individuals in type t and x_i represents the characteristics of individual i other than the effort level e . In our example, this means that for any given choice of years of post–mandatory schooling, the advantage function is constructed by averaging the advantage attained by individuals with differing levels of job–seeking activity.

This construction leads to a misrepresentation of the influence of irrelevant characteristics when the effort variable is correlated with other relevant characteristics in different ways in different types. The present problem is structurally similar to the one discussed in section 2.2. The difference is that the present case does not involve correlations between relevant and irrelevant characteristics, but correlations within the set of relevant characteristics. The problem can thus arise even when relevant characteristics are not correlated with the characteristics used to define the type, as was the case in section 2.2.

This flaw can be illustrated by the following example. Suppose we have ten persons, A through J, and the three characteristics of years of schooling, job–seeking activity, and parental income. Suppose that these characteristics determine a person’s (expected) life–time income under the status quo by the following formula: life–time income equals 1 for short schooling + 2 for long schooling + 0 for low job–seeking activity + 8 for high job–seeking activity + 0 for low parental income + 1 for high parental income. Society decides that parental income is an irrelevant characteristic and that years of schooling is the pertinent effort variable. The example is illustrated numerically in the following table.

Individuals	schooling	job-seeking activity	parental income	advantage
<i>A</i>	short (1)	low (0)	poor (0)	1
<i>B</i>	short (1)	high (8)	poor (0)	9
<i>C</i>	long (2)	low (0)	poor (0)	2
<i>D</i>	long (2)	high (8)	poor (0)	10
<i>E</i>	short (1)	low (0)	rich (1)	2
<i>F</i>	short (1)	low (0)	rich (1)	2
<i>G</i>	short (1)	high (8)	rich (1)	10
<i>H</i>	long (2)	low (0)	rich (1)	3
<i>I</i>	long (2)	high (8)	rich (1)	11
<i>J</i>	long (2)	high (8)	rich (1)	11

Under Roemer’s approach, we would type according to parental income, with *A–D* falling into the ‘poor’ type, and *E–J* falling into the ‘rich’ type. To assess this status quo, we then construct an advantage function for each type by averaging the advantage received at every effort level. For type ‘poor’, this means that at effort level 1, advantage is 5, and at effort level 2 advantage is 6. For type ‘rich’, the numbers are 4.7 and 8.3 respectively. We then take the relative effort levels in each type and compare the advantage obtained in each type in order to judge which individuals are ‘disadvantaged’. The result is given in the following table.

Type	relative effort	advantage a la Roemer	Roemer’s verdict
poor	low (<i>A&B</i>)	5	‘advantaged’
	high (<i>C&D</i>)	6	‘disadvantaged’
rich	low (<i>E&F&G</i>)	4.7	‘disadvantaged’
	high (<i>H& I &J</i>)	8.3	‘advantaged’

We can see that because the two relevant characteristics of relative effort and job-seeking activity are correlated differently in the different types, Roemer’s method of determining the advantage function of a type yields the wrong verdict on which individuals are advantaged and which individuals are disadvantaged. For two individuals with the same level of job-seeking activity and the same amount of schooling, the advantage of the individual from a rich family is, by the assumptions of our example, one unit higher than the advantage of the individual from a poor family. Hence, individuals from rich families are advantaged in the status quo, regardless of their level of job-seeking activity or years of schooling. Instead, Roemer’s method concludes that ‘low-effort rich’ individuals are at a disadvantaged vis-à-vis ‘low-effort poor’ individuals. In addition, it overestimates the degree to which ‘high-effort rich’ individuals are advantaged vis-à-vis ‘high-effort poor’ individuals.

This example has been constructed in a such way that Roemer’s method does not cause the problems that we discussed in section 2.2. For the relevant characteristics are in the current example not correlated with the irrelevant characteristics used in typing. Moreover, the assumption that we represent people’s choices by relative effort is not

essential in this example. If, as a consequence of our earlier arguments, we replace relative by absolute effort, the example goes through unchanged.

A defender of Roemer's method could try to protest that we have not chosen the effort variable appropriately if we consider not only years of schooling but also job-seeking activity to be a morally relevant characteristic. This defender might argue that the effort variable must consist of all possible combinations of morally relevant characteristics. In our case, the effort variable would then consist of combinations of years of schooling and levels of job-seeking activity. This defense might not be welcome to Roemer since his method of calculation relative effort relies on some ranking of effort levels in the direction of increasing effort. It is not clear how such a ranking should be constructed for a multi-dimensional effort variable that is composed of several relevant characteristics. Does a combination three years of schooling and high job-seeking activity represent a greater effort than four years of schooling and low job-seeking activity?

Our earlier method for solving the problem of correlation between irrelevant and relevant characteristics also avoids the current difficulties. Since the notion of typing plays no role in our proposal, we do not require the construction of an aggregated advantage function for a type. Once characteristics have been sorted into relevant and irrelevant, strict equality of opportunity simply demands that we equalize the advantage of all individuals with the same relevant characteristics. This would correctly occasion a redistribution of advantage from individuals with a rich family background to those from a poor family background. The different ways in which this redistribution could take place is the subject of the following sections, where we will return to this example.

3 Selecting an Equal-Opportunity Policy

3.1 Utilitarianism for the disadvantaged?

We now turn to Roemer's method for selecting an equal-opportunity policy. It is important to note at the outset that strict equality of opportunity leaves open one important issue that must eventually be resolved if we wish to select a concrete policy. All that equality of opportunity demands is that individuals who share the same relevant characteristics should also receive the same outcome. This leaves open how much an equal-opportunity policy should spend on individuals with different relevant characteristics. Suppose in our previous example that we can carry out a redistribution without any loss of resources through lump-sum transfers among individuals with different parental income and/or effort levels. We could then achieve strict equality of opportunity by taxing each rich individual by 0.4 units of life-time income and giving 0.6 units to each poor individual. However, we could also achieve strict equality of opportunity by only taxing the 'short-schooling rich' by 1 unit. This would trivially equalize the advantage of poor and rich individuals with short schooling and we could then spend the as yet unused 3 units of advantage to equalize opportunity between poor and rich individuals with long

schooling. There may thus exist many different ways to establish strict equality of opportunity. Any method for choosing an equal–opportunity policy must therefore answer the question of how much of our budget we should spend on individuals with different relevant characteristics. In this section, we will discuss how this budget problem arises in the context of Roemer’s maximin approach and we will look at Roemer’s proposed solution to this problem. We will argue that criteria for solving the budget problem ought to be kept separate from a criterion of equal opportunity. We would otherwise blur the line between the concept of equal opportunity and other, conceptually distinct distributional criteria, such as utilitarianism or some notion of reward according to effort.

For each type t , Roemer assumes the existence of a function $u_t = u_t(\phi, e)$ that yields the (expected) advantage of an individual of this type depending on effort e and policy ϕ . Given the distribution of absolute effort levels within a type, we can now determine an individual’s relative level of effort π within this type. Roemer then calculates the function $v_t(\phi, \pi)$ that specifies the advantage of type t for the relative effort π under policy ϕ . For the sake of the argument, we will accept this construction in terms of relative effort in spite of our arguments in the previous section. Whatever we say in the present section applies unaltered if we replace relative effort by absolute effort or, indeed, by any other characteristic of individuals that we might deem to be relevant.

Suppose that we have sorted the population of individuals into centiles of effort within their type. Strict equality of opportunity would now require the choice of a policy ϕ within the budget such that

$$v_t(\phi, \pi) = v_{t'}(\phi, \pi) \quad \text{for all } \pi = .01, \dots, .99, 1. \quad (3)$$

Roemer approaches the problem of policy selection from the viewpoint of optimization theory. The idea of optimization is to define a relation of betterness and then to pick a policy that is best according to this relation. From this viewpoint, Roemer rehearses two problems with strict equalization. Firstly, there may not exist any feasible policy that achieves strict equalization. Secondly, strict equalization may lead to Pareto inferior allocations in which it is possible to make some individual better–off without making any individual worse–off. For these reasons (to which we will return below), Roemer replaces the requirement of equalization by the demand that we maximize the advantage of the type that is worst–off vis–à–vis other types at the same relative effort level, i.e. that we choose a policy ϕ that maximizes

$$\min_t v_t(\phi, \pi) \quad \text{for all } \pi = .01, \dots, .99, 1. \quad (4)$$

Roemer correctly points out that this is not a well–defined optimization problem unless we are given a budget constraint for each level of relative effort. Such a constraint must tell us how much of our budget we should spend on highly motivated individuals who expend a high amount of relative effort and how much we should spend on less

motivated individuals who expend a lower amount of relative effort. Roemer’s solution to the budget problem gives each of the 100 centiles equal weight in determining the overall policy. Thus, his equal–opportunity policy for the case of centiles is the policy ϕ that maximizes

$$\sum_{p=1}^{100} \min_t v_t \left(\phi, \frac{p}{100} \right) \quad (5)$$

Roemer’s criterion offers a mixture of maximin and utilitarian reasoning by requiring us to maximize the average, taken over all centiles of relative effort, of the advantage enjoyed by the least well–off type at that centile. The maximin element in Roemer’s criterion determines the advantage of the least well–off type at any level of relative effort. Note that different types may be least well–off at different centiles of relative effort. In the example from section 2.4, for instance, Roemer’s methodology considered the rich type to be at a disadvantage for the first fifty centiles of relative effort, while the poor type was considered to be at a disadvantage for the second fifty centiles. The utilitarian element in Roemer’s criterion then trades off the least advantage achieved at low levels of relative effort against the least advantage achieved at high levels of relative effort. It will therefore mandate a worsening of the situation for disadvantaged individuals with low levels of relative effort in exchange for an improvement in the situation of disadvantaged individuals at high levels of relative effort, if this increases the average least advantage across different centiles.

Needless to say, this solution to the budget problem has a plethora of competitors. Other criteria that appear equally defensible could recommend spending an equal portion of the overall budget on each combination of relevant characteristics (equal split between characteristics), spending a portion of the overall budget on each combination of relevant characteristics y that corresponds to the number of individuals with the combination y (a form of equal split weighted by the amount of individuals), or allocating the budget so as to maximize the advantage attained by those whose characteristics lead to the lowest absolute level of advantage (a form of leximinism). An additional constraint that Roemer (1993, 160; 1998, 104n26) himself occasionally advocates is based on a notion of desert and requires that individuals who expend a higher degree of relative effort should also attain a higher level of advantage. It should be clear that any of these criteria are conceptually distinct from the notion of equal opportunity. These criteria pertain to the advantage that individuals ought to receive on the basis of their morally relevant characteristics, while the notion of equal opportunity is concerned with the compensation of individuals for disadvantages that arise from their irrelevant characteristics. We would therefore like to see a formulation of a selection criterion that clearly sets such criteria apart from the concern for equal opportunity. We take up the task of formulating such a criterion in section 3.3, after clearing up a final inconsistency in Roemer’s approach in the following paragraphs.

3.2 Equality, Maximin and Leximin

Before we turn to a new formulation of the equal–opportunity criterion, we note an inconsistency in Roemer’s argument. Recall that Roemer abandoned strict equality (3) in favour of maximin (4) partly for reasons of Pareto–optimality. We start from the trivial observation that the argument for maximin based on Pareto optimality does not achieve its full purpose. Maximin recommends itself over equalization because it allows Pareto improvements of the egalitarian allocation. However, maximinism does not exploit all Pareto improvements that are possible. Consider a simple example in which, under a policy ϕ , black men and women enjoy an equal advantage that is lower than that of either white women or white men. Suppose that ϕ maximizes the advantage of the least well–off types, namely black men and black women. For simplicity, we may assume that this ordering holds for any level of relative effort. This policy may still be Pareto suboptimal if it is possible to find a policy ψ that increases the advantage of white women without decreasing the advantage of black men and women. This example illustrates the well–known fact that only the extension of maximin to leximin guarantees Pareto optimality (Sen 1970). Adapted to our context, the *leximin rule conditional on* level $\pi = .01, \dots, .99, 1$ of relative effort reads:

- (6) Maximize the advantage of the worst–off type at relative effort π . In case of a tie, maximize the advantage of the second worst–off type at relative effort π . Proceed analogously for all types.

Leximinism not only provides a choice rule for an optimal policy, but it also allows us to compare different policies relative to each other. According to leximin, a policy ϕ is *at least as good as* a policy ψ *conditional on* a level π of relative effort exactly when the advantage of the worst–off type under policy ϕ at relative effort π is at least as great as the advantage of the worst–off type under ψ at π and, in case of a tie, the advantage of the second worst–off type under ϕ at π is at least as great as the advantage of the second worst–off type under ψ at π etc. Note that this definition is relative to the parameter π and the budget question reviewed in the previous section is how we should treat this loose parameter in the suggested optimization problem.

At this point, an inconsistency arises in Roemer’s argument. The argument based on Pareto optimality leads to leximin rather than maximin, but Roemer’s treatment of the loose parameter π applies only to maximin and cannot possibly apply to leximin. When Roemer addresses the loose parameter π in the case of maximin, he suggests as a solution that we sum the values of the expression $\min_t v_t(\phi, \pi)$ over all possible levels $\pi = .01, \dots, .99, 1$. This treatment presupposes that the expressions over which the sum is taken yields a number, for otherwise the operation of summing is not defined. The minimum operator, of course, always yields a number; leximin, however, never yields a number. Leximin generates only an ordering of policies, but not a numerical measure of the amount by which one policy is better than another policy. Hence, the suggested treatment of the free parameter π does carry over from maximin to the case of leximin.

3.3 Opportunity dominance

Although there are disputes about the characteristics for which individuals deserve compensation, some form of equal opportunity is one of the few social goals in the current political discussion that find common support from opposing political ideologies. Political arguments therefore frequently appeal to equal opportunity as a principle that is thought to mitigate disagreements about how to compensate individuals for differences in outcomes. In this political context, it is important to keep our equal–opportunity criterion free from controversial assumptions about reward and compensation for outcomes. Only if this neutrality is preserved, can we explore by how much the satisfaction of some equal–opportunity constraint arbitrates the import of controversial compensatory principles, such as outcome–utilitarianism and outcome–leximinism. Generally speaking, our task is therefore to define a criterion that evaluates policies depending on how they compensate individuals for their irrelevant characteristics and yet remains neutral on the question of how we should distribute reward to individuals with different combinations of relevant characteristics.

With this purpose in mind, we recast the idea of the previous paragraphs in our preferred framework. Suppose we have a predictive model for the (expected) advantage that results from different combinations of individual characteristics. Suppose, furthermore, that society has sorted these characteristics into those that are morally relevant and those that are morally irrelevant. We could now apply the narrow criterion (1) of strict equality of opportunity that we formulated in section 2.2. For the purpose of policy selection, Roemer already raised two of the most common objections against this strict requirement. He pointed out that, on the one hand, there may not exist any policy that achieves strict equalization, while, on the other hand, equalizing policies may not be Pareto optimal. Although Pareto–optimality is itself not an intrinsic part of the concept of equality (Sugden 1984, Cohen 2000), it seems to be almost universally regarded as a desirable property of rules for policy selection. We will therefore use the leximin principle from the previous section to define when one policy compensates for irrelevant characteristics better than another policy, conditional on some combination of relevant characteristics. In a second step, we will then construct a partial ranking of possible policies on the basis of these conditional comparisons.

A policy ϕ is *at least as good as* a policy ψ *conditional on* relevant characteristics y exactly when ϕ leads to at least as great an advantage as ψ for individuals with the relevant characteristics y and the worst possible irrelevant characteristics z and, in case of a tie, ϕ leads to at least as great an advantage as ψ for individuals with relevant characteristics y and the second–worst irrelevant characteristics z' etc. A policy ϕ is *better than* a policy ψ *conditional on* relevant characteristics y exactly when ϕ is at least as good as ψ conditional on y and, moreover, for individuals with the relevant characteristics y and some irrelevant characteristics z'' , ϕ leads to a greater advantage than ψ . We noted above that it is not clear how these conditional orderings could be synthesized into an overall, unconditional ordering. The desire for such a synthesized ordering resulted from Roemer’s optimization approach that seeks to define a relation of

unconditional betterness and then to pick a policy that is best according to this relation. Abandoning this idea, we merely synthesize a partial ranking of the available policies. We will therefore not be able to use this ranking to pick some best policy but only to dis-select dominated policies.

To make such partial comparisons, we consider whether or not a policy is better than another conditional on any combination of relevant characteristics. We say that a policy ϕ *dominates in opportunity* a policy ψ if and only if, for every combination of relevant characteristics y , ϕ is at least as good as ψ conditional on y and, for some combination of relevant characteristics y' , ϕ is better than ψ conditional on y' . The allocations between which this criterion remains undecided are precisely those for which a budget decision has to be made. Suppose that ϕ and ψ are both undominated in opportunity. In particular, ϕ and ψ must then not dominate each other in opportunity. On the one hand, there must hence be some of characteristics y such that ϕ leads to a greater advantage than ψ for the worst irrelevant characteristics z (or, in case of a tie, the second-worst irrelevant characteristics etc.). On the other hand, there must also be some relevant characteristics y' such that ψ leads to a greater advantage than ϕ for the worst irrelevant characteristics z' (or, in case of a tie, the second-worst irrelevant characteristics etc.). We now need to decide whether to weigh the concerns of individuals with relevant characteristics y higher than those with relevant characteristic y' and to choose ϕ rather than ψ . This is exactly the budget problem from above.

To apply the notion of opportunity dominance, we start from the set of feasible policies. The feasibility of a policy will be determined by the overall available budget, the implementability of the policy and similar practical considerations. In a second step, we eliminate all policies from this set that are dominated in opportunity. The remaining policies are then considered as the final candidates for an opportunity policy, requiring a budgetary choice to be made. We note that the set of policies that are undominated in opportunity is never empty (for any finite set of feasible policies). Moreover, the set of candidate policies is Pareto optimal in the sense that for any of the candidate policies there is no other feasible policy that could make some combination of characteristics better off without making any other combination of characteristics worse off.

We apply this idea to the example of section 2.4 where life-time income under the status quo was determined by the following formula: life-time income equals 1 for short schooling + 2 for long schooling + 0 for low job-seeking activity + 8 for high job-seeking activity + 0 for low parental income + 1 for high parental income. Suppose society again regards parental income is an irrelevant characteristic and years of schooling and job-seeking activity as relevant characteristics. Suppose first that our feasible policy instruments allow us without any costs to redistribute life-time income depending on irrelevant characteristics. We may thus allocate a lump sum transfer to individuals from a poor family background. Any such lump sum that is given to the four poor individuals must be paid for by the six rich individuals. The opportunity-dominance criterion then eliminates all but one policy that is undominated in opportunity. An easy calculation shows that this policy gives a lump sum transfer of 0.6 units of income to individuals

with low parental income. Hence, each individual with high parental income is required to pay 0.4 units of income and thus keeps $1 - 0.4 = 0.6$ units of income.

In general, the set of policies that are undominated in opportunity may contain more than one element and we thus may again encounter the budget problem. Suppose, in the previous example, that the set of feasible policies allows us without any costs to redistribute life-time income depending not only on parental income, but also on years of schooling. Going to extremes, we may then decide to collect the one additional unit of advantage from all three rich individuals with short schooling. This would equalize the advantage of poor and rich individuals with short schooling. We could then draw on these additional three units when equalizing the opportunity between poor and rich individuals with long schooling. Our first undominated policy would then allocate a lump sum transfer of 1.2 units of advantage to poor individuals with long schooling. This leaves an additional 0.2 units and hence a total of 1.2 units of advantage for rich individuals with long schooling. We could also go to the opposite extreme and collect the one additional unit of advantage from all three rich individuals with long schooling. This would now equalize the advantage of poor and rich individuals with long schooling. Our second undominated policy would then give 1.2 units of advantage to poor individuals with short schooling. This again leaves an additional 0.2 units and hence a total of 1.2 units of advantage for rich individuals with short schooling.

Given this primary constraint of opportunity dominance, we may now vary the compensatory principle by which we select a particular undominated policy according to its outcomes. As we vary this compensatory principle from outcome-utilitarianism to outcome-leximinism and other distributive criteria, we will be able determine the magnitude of the remaining disagreement about the compensation for differential outcomes. Our hope would be that, at least in some important cases, the appeal to opportunity dominance will restrict the range of eligible policies sufficiently to marginalize the influence of these disagreements. Consider, for instance, the common objection against utilitarianism that it is insensitive to inequalities of outcomes. This objection may become less telling if the utilitarian outcome-criterion is subordinated to a selection from the set of policies that are undominated in opportunity. In the same vein, we should explore the effect of other constraints, such as the common requirement that greater effort should be rewarded by greater advantage. Roemer, for instance, suggests to constrain the selection to those policies that yield a positive reward to effort in each type (Roemer 1993, 160; 1998, 104n26). In our example, the decision that longer schooling should lead to a higher life-time income would exclude the second of the above policies. Nonetheless, we would still have to decide by how much life-time income should increase with longer schooling.

4 Conclusion

We conclude with a summary of our constructive suggestions. In sections 1 and 2, we argued that every concrete application of the idea of equal opportunity involves an

inescapable value judgement about which characteristics we consider relevant and irrelevant. We suggested that the distinction between factors within and factors beyond a person's control may act as a guide in the determination of moral relevance, but that we must not expect the notions of relevance and of personal control to coincide. Depending on the application and the advantage in question, relevant characteristics may involve people's natural or developed talents, people's actual choices, or some notion which might approximate Roemer's idea of relative effort. Any concrete application of the idea of equal opportunity therefore starts from a value judgement that classifies the characteristics by which we describe the individuals and predict their expected advantage. Individuals will then be held responsible for what are considered relevant characteristics, but not for characteristics that are considered morally irrelevant.

In section 3, we argued that a criterion of equal opportunity should steer clear of any controversial assumptions about how individuals should be rewarded for differences in outcomes due to their relevant characteristics. Any such commitment would further compromise the ability of a concrete principle of equal opportunity to muster support from a wide spectrum of political ideologies. We have therefore proposed a formulation that keeps the opportunity criterion conceptually distinct from principles governing the compensation for relevant characteristics. In this formulation, we are able to vary the criterion for the compensation of outcomes and gauge the magnitude of the disagreements that remain in the presence of an opportunity criterion. This strategy enables the political debate to be conducted on two distinct issues. Firstly, which factors should be considered irrelevant for certain social advantages? Secondly, on what basis do we choose from the set of feasible policies that are not dominated in opportunity? We hope that this improved focus can help to settle political disagreements surrounding the implementation of opportunity policies.

References

- [1] R.J. Arneson. Equality and equality of opportunity. *Philosophical Studies*, 56:77–93, 1989.
- [2] G.A. Cohen. On the currency of egalitarian justice. *Ethics*, 99:906–944, 1989.
- [3] G.A. Cohen. Why not socialism? unpublished, 2000.
- [4] M. Fleurbaey. Equal opportunity or equal social outcome? *Economics and Philosophy*, 11:25–55, 1995.
- [5] H.L.A. Hart and A.M. Honoré. *Causation in the Law*. Clarendon Press, Oxford, 1959.
- [6] H.G. Llavador and J.E. Roemer. An equal-opportunity approach to the allocation of international aid. *Journal of Development Economics*, 64(1):147–171, 2001.

- [7] J. Roemer. A pragmatic theory of responsibility for the egalitarian planner. *Philosophy and Public Affairs*, 22:146–166, 1993.
- [8] J. Roemer. *Theories of Distributive Justice*. Harvard University Press, Cambridge (Mass.), 1996.
- [9] J. Roemer. *Equality of Opportunity*. Harvard University Press, Cambridge (Mass.), 1998.
- [10] J. Roemer. Equity in health care delivery and finance. unpublished, 1999.
- [11] T.M. Scanlon. *What We Owe to Each Other*. Harvard University, Cambridge (Mass.), 1998.
- [12] A. Sen. *Collective Choice and Social Welfare*. Holden–Day, San Fransisco, 1970.
- [13] R. Sugden. Is fairness good? A critique of Varian’s theory of fairness. *Nous*, 18:505–511, 1984.