

Equality of opportunity: A progress report

John E. Roemer

Yale University, Department of Political Science, 124 Prospect Street, New Haven, CT 06520-8301, USA (e-mail: john.roemer@yale.edu)

Received: 5 January 2001/Accepted: 4 December 2001

1 Introduction

Probably the most universally supported conception of justice in advanced societies is that of equal opportunity. What does this mean, and how can it be formalized? The first treatment of this problem in social choice theory involved interpreting equality of opportunity as equality of opportunity *sets*, that is, rendering the sets of choices available to different individuals the same. In my view, this literature was too abstract – it failed to take sufficient cues from popular and long-standing conceptions of equality of opportunity. A consequence of this over-abstractness is that no empirical work, on what equal-opportunity policies should be in real-world situations, has followed from this work, as far as I am aware.

Those popular conceptions have, in contrast, influenced greatly the recent literature in egalitarian political philosophy, particularly the idea that in the equal-opportunity ethic, there is, as well as a desire to equalize something, an insistence that individuals be held responsible for what happens to them. This is popularly formulated in the ‘level-the-playing-field’ metaphor: equal-opportunity policy must create a level playing field, after which each individual is on his own – what outcomes finally occur will reflect individual effort, and outcome differentials are ethically acceptable, if the playing field was initially level, and if they are due to differential effort.

John Rawls (1971), in his path-breaking work, certainly understood, and supported, the idea that the social mandate to equalize outcomes across indi-

I am grateful to my collaborators, listed in the references, for permitting me to report some of the results of our joint work in this paper, which was read as the Arrow Lecture at the bi-annual meeting of the Society for Social Choice and Welfare, in Alicante, July 2000.

viduals was rightfully constrained by how responsibly individuals had behaved, but his model of souls contracting behind the veil of ignorance failed to capture the intuitions correctly (see Roemer 1996, Chapt. 5). Ronald Dworkin, in two articles published in 1981, came much closer to getting it right. Dworkin (1981a,b) distinguished between a person's preferences and his resources, where the latter comprised endowments which the individual had no role in choosing. He argued that the ethically attractive notion of equality entailed equalizing resources across individuals, but allowing differences to emerge in final conditions due to the exercise of choice following from differential preferences: thus, Dworkin said people should be held responsible for their preferences but not their resources. A moment's thought will convince you that it is not so easy to make the cut between preferences and resources, at least if one wants preferences to be attributes for which the individual is rightly held responsible. (For example, resources clearly influence preference formation, so if one is not responsible for resources like family background, should one be responsible for exercising preferences induced by childhood nurture?) Nevertheless, Dworkin's contribution was path-breaking, and other philosophers took Dworkin's work further. Dworkin even proposed an economic model of his idea, what economists would call a market for contingent claims behind a veil of ignorance conceived differently from Rawls's, in which souls know the preferences of the persons with whom they are associated, but not the resource bundles those persons will receive in the birth lottery.

The general structure of Dworkin's theory is that a person's attributes (endowments, preferences, actions) can be partitioned into two sets – those for which we think it is morally correct to hold him accountable or responsible, and those for which we think it is not so. Call the first set responsible factors and the second, arbitrary factors. An egalitarian should seek to equalize the final conditions (perhaps, welfares) of individuals to the extent that those conditions are due to differences in the sets of arbitrary factors, but allow differences in condition to the extent that such are due to differences in responsible factors.

It is this conception of 'responsible' and 'arbitrary' factors which, I think, is key to the notion of equal opportunity. I will now adopt a somewhat less general approach, and say that the outcomes individuals sustain are the consequence of circumstances, effort, and policy, where circumstances are aspects of an individual's environment and actions which are either beyond his control, or for which we (society) wish not to hold him responsible, effort comprises the actions he takes for which we do wish to hold him responsible, and policy is the instrument by which society (or the planner) influences outcomes – perhaps some allocation of a publicly owned resource. Equal-opportunity policy aims to level the playing field. What are the troughs in that metaphorical playing field? They are the disadvantages that some persons face, with regard to achieving high levels of welfare – or, more generally, the condition about which the planner is concerned – due to unfortunate circumstances. Thus, leveling the playing field means choosing that policy that will make it the case that an individual's final condition will be, as far as possible, only a

function of the effort he makes. In particular, equality of opportunity finds no moral bad in inequality of final condition across individuals ascribable to differential effort. In this way, it differs from the equality-of-outcome ethic.

2 A formalization

We can formalize this view as follows. Suppose the condition for which we desire to equalize opportunities – the opportunity equalisandum – is measured by a function $u(C, e, \varphi)$, where C denotes the individual's circumstances, e his effort, and φ the policy of the intervening agency or planner. u need not measure welfare in the economist's sense: it could be income, or life expectancy, or wage-earning capacity. I further specialize by assuming that an individual's *type* is the value of his C , that the population has a finite number of types, and that there is a continuum of individuals in each type. From now on, we denote types $1, 2, \dots, \tau$, with the generic element t . We denote the set of types by T , and the set of feasible policies by Φ .

We assume, as well, that, for any policy φ implemented by the planner, there will be a distribution of effort in each type, denoted by F_{φ}^t , which is a probability measure on the set of non-negative real numbers – the domain of effort. (Thus, we here make another assumption, that effort can be captured by a one dimensional variable.) Letting \mathbf{P} be the set of probability distributions on \mathbf{R}_+ , we take the mapping $F : T \times \Phi \rightarrow \mathbf{P}$ as a primitive of the model. An economist would say that effort responses are the consequence of utility maximization facing policies, but that does not here concern us, and there is no reason to further complexify by modeling that process.

Let me give two examples. Let u be wage earning capacity, a person's type be the socio-economic status (SES) of his parents, effort be the number of years of school he attends, and policy be the per-pupil educational expenditures of the state, which, we assume, may be varied according to the socio-economic characteristics of the neighborhood in which the school is located. In this conception, we do not hold a person responsible for his parents' SES, but we do hold him responsible for the number of years of school he attends. Equal-opportunity policy entails distributing an educational budget among schools so that, so far as possible, the wages individuals eventually earn are, on average, equal across SES types of child who attended school similar numbers of years. Intuitively, that policy would spend more per pupil in low SES neighborhoods than in high SES neighborhoods. The policy would be used to compensate individuals with disadvantaged circumstances so that, finally, at each level of school attendance, the wages of individuals, in the various types, would be equal.

Here is a second example. Let u be life expectancy, type be, again, the economic class of the parents of the individual, effort be a measure of 'life-style quality,' in the sense of exercising, eating healthily, not smoking, and so on, and policy be some allocation of medical care services to the population. Here, the equal-opportunity goal is to choose the policy so that the life expectancies of

individuals of different types are equal for cohorts of those types whose members have lived equally healthy life-styles.

Of course, the *distribution* of effort within a type, in these two examples, may well respond to the policy, which is why we predicated that distribution, above, on both the type and the policy.

There is, however, an inconsistency between the model, as I have thus far described it, and the intent of the modeler. Everyone knows that the distribution of effort is, indeed, a characteristic of the type. If schools spend \$5000 per student on upper middle class children, we will observe a different distribution of years of school completed, our proposed effort measure, than if they spend the same amount on poor, ghetto children. Our intent is to hold individuals not responsible for their circumstances, their type. Thus, in comparing efforts of individuals in different types, we should somehow adjust for the fact that those efforts are drawn from distributions which are different, a difference for which individuals should not be held responsible. We require an inter-type comparable measure of effort which factors out the goodness or badness of the distribution, because the distribution is a *characteristic of the type*, not of any individual. I propose, as a simple measure of the morally relevant degree of effort, the quantile of the effort distribution for his type at which an individual sits. Thus, we shall declare two individuals as having exercised the same *degree* of effort if they sit at the same quantile or rank of their type distributions of effort. Thus, in deciding how hard a person has tried, we compare him only to others with his circumstances.

To amplify on this choice: the distribution of effort that we observe within a type is the outcome of beliefs and preferences of the individuals within it. To a large extent, those beliefs and preferences are determined by circumstances. It is this part of beliefs and preferences that individuals should not be held responsible for. We therefore measure a person's effort by his rank in the effort distribution of his type, rather than by the absolute level of effort he expends.

This is an important part of the theory. Suppose the distributions of effort of the advantaged type are uniformly distributed on the interval $[1, 2]$, under some policy, while the distributions of effort of the disadvantaged type are uniformly distributed on the interval $[0.25, 1.25]$. It makes sense to say that someone in the latter type who exerted effort 1.25 tried very hard, while someone in the former type who exerted that effort did not.

Since we know the mapping F , we can compute the indirect outcome function $v^t(\pi, \varphi)$, which is the level of u for individuals of type t at the π^{th} quantile of the effort distribution for type t when the policy is φ , where π is any quantile in the interval $[0, 1]$. We have now amended the original proposal to say that equalizing opportunities means choosing policy to equalize outcomes, v^t , across types, at fixed levels of π . Thus, from a moral viewpoint, we declare that two individuals in different types have tried equally hard if they lie at the same rank of the effort distributions of their types, and our policy aims to render such persons equal in final condition, independent of their circumstances. How shall we formalize this goal?

Suppose we fix π , and choose that policy that equalizes outcomes across

types at that π – well, we don't really mean equalize, we mean maximize the minimum value of such outcomes. Thus we define

$$\varphi^\pi = \text{Arg Max}_\varphi \text{Min}_t v^t(\pi, \varphi). \tag{2.1}$$

If we were concerned only with the π -slice of individuals, then φ^π would be the equal-opportunity policy. Unfortunately, the set $\{\varphi^\pi \mid \pi \in [0, 1]\}$ will generally consist in a continuum of different policies. This is simply a reflection of the fact that we can generally not equalize something for an infinite number of populations at the same time. We require some compromise approach. My usual choice has been to form a social objective function in which the objective function of each effort slice of the population, $\text{Min}_t v^t(\pi, \varphi)$, is weighted by its size. Since each effort quantile of the population has the same size, the social objective becomes

$$\int_0^1 \text{Min}_t v^t(\pi, \varphi) d\pi,$$

and the equal-opportunity policy¹ is thus defined by:

$$\varphi^{\text{EOP}} = \text{Arg Max}_\varphi \int_0^1 \text{Min}_t v^t(\pi, \varphi) d\pi. \tag{2.2}$$

A graphical illustration of this optimization problem is presented in Fig. 1. Here, there are three types, and the particular policy entails distributing resources in amount x^t to individuals of type t . The three graphs are of the functions v^t as a function of π , for a fixed policy. Program (2.2) directs us to vary the policy to maximize the area under the lower envelope of the v^t functions, the heavy line in the figure.

It is useful to display some other well-known conceptions of justice in this environment. Let the population fraction of type t be p^t ; then the utilitarian policy is given by:

$$\varphi^U = \text{Arg Max}_\varphi \sum_{t=1}^{\tau} p^t \int_0^1 v^t(\pi, \varphi) d\pi,$$

¹ Two other possible compromises are as follows: to define the EOp policy by

$$\varphi^{\text{EOP}} = \int_0^1 \text{Arg Max}_\varphi \text{Min}_t v^t(\pi, \varphi) d\pi, \tag{2.2a}$$

or as:

$$\varphi^{\text{EOP}} = \text{Arg Max}_\varphi \text{Min}_t \int_0^1 v^t(\pi, \varphi) d\pi. \tag{2.2b}$$

(2.2a) declares the EOp policy to be the average of the policies $\{\varphi^\pi\}$, while (2.2b) declares it to be the policy that maximizes the average value of the objective of the worst off type. I have no strong preference for any of these alternatives over the others. Sometimes computational simplicity recommends one over the others.




Fig. 1. An illustration of the equal-opportunity objective

and the Rawlsian policy (the policy that maximizes the condition of the worst-off individual) is given by:

$$\varphi^R = \text{Arg Max}_{\varphi} \text{Min}_{t, \pi} v^i(\pi, \varphi).$$

You will note that the utilitarian program contains no ‘Min’ operator, and the Rawlsian program contains no integral operator, while the equal-opportunity program contains both operators. This reflects the fact that the equal-opportunity ideal, as I have formulated it, is egalitarian with respect to differences in outcomes due to differential type or circumstance, but utilitarian with respect to differences of outcome due to differential effort. It will come as no surprise that the equal-opportunity policy generally takes a midling stance between these two, being more egalitarian than utilitarianism and less egalitarian than Rawls. To the extent that differences in outcome are due to circumstance, it approaches Rawlsianism, but to the extent they are due to effort, it approaches utilitarianism.

There is another distinction worth noting between equal-opportunity, on the one hand, and utilitarianism and Rawlsianism on the other, a distinction which is not immediately observable from the formulae. Utilitarianism and Rawlsianism, as formulated here, are both welfarist², in the sense that the

² Since the outcome need not be welfare in the traditional sense, it would be more accurate to say these two philosophies are both consequentialist.

optimal policies can be computed knowing only the effect of policy on the *set* of outcomes, that is, the set of numbers $\{v^t(\pi, \phi) \mid t \in T, \pi \in [0, 1], \phi \in \Phi\}$. The utilitarian objective, for a given ϕ , is the average of this set of numbers, while the Rawlsian objective, for a given ϕ , is just the minimum of this set of numbers. In contrast, one cannot compute the equal-opportunity policy knowing only the set of outcomes – one must know, as well, the *functions* $\{v^t\}$. That is, much more information is required to compute the EOp policy: one must know *distributions* of outcomes *by type*, for each policy. In the social-choice terminology, equal-opportunity is non-welfarist, or non-consequentialist. This reflects, of course, the informal, popular view – how much society should help people depends on *how hard they tried*, not just on how badly off they are. The equal-opportunity ethic is not a welfarist ethic, and it is therefore missed by classical social-choice assumptions that preclude non-welfarist rules.

3 Solving for the equal-opportunity policy from the data

Thus far, I have assumed that effort is unidimensional – a poor assumption. In reality, both circumstances and effort are highly complex. I shall argue in this section that we can finesse the complexity of effort – in fact, we can compute the equal-opportunity policy without knowing the distributions of effort within types – as long as we are willing to assume that, whatever constitutes an increase in effort increases, *ceteris paribus*, the level of the outcome. (What is held fixed in the *ceteris paribus* clause is type, policy, and luck.)

Formally, I will suppose that if effort is in reality a vector $\mathbf{e} = (e^1, \dots, e^R)$, then there is for each type an index $f^t(\mathbf{e})$ such that we can write the outcome as a function of $f^t(\mathbf{e})$, circumstances, and policy. The econometrician, of course, standardly estimates the coefficients α^i that give the best fit of the data under the assumption that $f^t(e^1, \dots, e^R) = \sum \alpha^i e^i$. Our assumption is that all the α^i are positive and that outcomes are strictly monotone in f^t .

This assumption – that effort is positively related to the outcome – may strike the journeyman economist as bizarre. Consider the classical set-up where utility is a decreasing function of labor expended, and suppose we have an interpersonally comparable measure of utility. Suppose we want to equalize opportunities for utility. Is it not reasonable to interpret ‘labor expended’ as effort, as we standardly say, in which case the opportunity equalisandum is a *decreasing* function of effort? The answer is that labor is not the proper construal of effort in this application of the equal-opportunity model. Rather, we should think of individuals as differing with respect to a preference parameter which, *inter alia*, determines their labor response to prices and policies. Then ‘high effort’ people are ones whose parameter induces them to expend much labor; *ceteris paribus*, those people will have higher utility levels than others. In this case, the value of the parameter in question reflects the individual’s industriousness, as opposed to laziness, and an equal-opportunity policy would consider it morally all right that industrious people end up better off than lazy people. With this interpretation, the outcome is, indeed, positively related to ‘effort.’

I now proceed with the promised finesse of the effort data. I shall assume that all policies treat the members of any given type identically. It follows that any variation in outcomes within the type is ascribed to differential effort, and those with higher outcomes have expended greater effort³. Let G'_φ be the distribution function of outcomes in type t at policy φ ; it then follows that *those at the π^{th} quantile of the effort distribution are exactly those at the π^{th} quantile of the outcome distribution*, by the monotonicity of outcomes in effort. We can therefore write:

$$G'_\varphi(v^t(\pi, \varphi)) = \pi. \quad (3.1)$$

Now, assuming that the distribution function is strictly increasing, it has an inverse, and we can write:

$$v^t(\pi, \varphi) = G'^{-1}_\varphi(\pi).$$

It now follows that the equal-opportunity program (see (2.2)) can be written as:

$$\varphi^{\text{EOP}} = \text{Arg Max}_\varphi \int_0^1 G'^{-1}_\varphi(\pi) d\pi. \quad (3.2)$$

Hence, we can compute the equal-opportunity policy knowing just the distribution of the outcomes, for each type, as a function of the policy. Geometrically, (3.2) says, in the coordinate plane in which one graphs the distribution functions, the problem is to choose that policy that maximizes the area to the left of the *left-hand envelope* of the distribution functions for the various types, and bounded by the vertical axis, the horizontal axis, and the line $y = 1$. See Fig. 2.

To review, three assumptions have led to this conclusion, which considerably simplifies the computation of equal-opportunity policies from data. These are:

- (1) We identify the degree of a person's effort with his quantile, or rank, on the effort distribution of his type;
- (2) Effort is the residual determinant of outcomes once type and policy are fixed;
- (3) Greater effort increases the level of the outcome, *ceteris paribus*.

I have argued for (1) by observing that the *distribution* of effort is a characteristic of the type, not of any individual, and hence, if leveling the playing field means compensating individuals for being in a disadvantaged type, it means compensating them for being in a type with a bad effort distribution. I have argued for (3) as a matter of the definition of effort. Assumption (2) is conservative, both in the sense of conserving information, and in the political sense. For it attributes to effort, in all likelihood, more than we should. We never specify all the actual circumstances when we look at a problem – for

³ We prescind here from the occurrence of luck; our method is wanting in not distinguishing between luck and effort.




Fig. 2. The value of the EOp objective with two types

instance, referring to the earlier examples, a person's wage earning capacity has, as a relevant circumstance not only the SES of his parents, but his own natural talent, his sex, race, and so on. All these things should arguably be treated as circumstances. But if we define circumstances solely by SES status as above, and we adopt (2), then we will be attributing to effort part of the outcome (wage-earning capacity) that should be attributed to natural talent. I call (2) a conservative assumption because the politically conservative move is to attribute much to effort and little to circumstance, and this is exactly what (2) does.

In the applications below, do remember that we have made a politically conservative move in adopting (2), for that means that the equal-opportunity policies we shall compute are less compensatory, to disadvantaged types, than they should be – than they would be, if we delineated *all* the relevant circumstances. This is important, because, as you will see, the equal-opportunity policies that we compute are quite compensatory.

Even though we have proposed a set of assumptions that enables us to compute the EOp policy without observing effort, it remains the case that the EOp objective is non-welfarist. The remarks on this issue made above continue to hold.

4 Equalizing opportunities for wage-earning capacity

In this section and the next one, I shall present the results of two empirical studies. The first one treats the distribution of educational finance as the policy, and computes the policy that equalizes opportunities, among American

males, for the acquisition of wage-earning capacity. I report on my collaboration with the labor economist Julian Betts (1998).

Our aim is to use educational finance to equalize opportunities for wage-earning capacity among young men in the United States. We shall compute the equal-opportunity policy under several different typologies of the population. We begin by defining an individual's type by the level of education of his more educated parent, and we partition our population into four types, where, in the most disadvantaged type, the parent had less than eight years of schooling, and in the most advantaged type, she had at least some tertiary education. We need a data set that will allow us to compute the elasticity of wage-earning capacity of young men with respect to what was spent on their educations, and we need as well to be able to identify their types. The National Longitudinal Study of Young Men (NLSYM) is a panel data set containing all this information; in particular, it reports the per capita expenditures in the school district the young man lived in at age 16, and it reports his wage at age 30. Young men in our sample were aged 16 in the late 1960s. I should remark that we can compute these elasticities in the US, because there is great variation on per capita school expenditures by municipality, unlike in most European countries.

We take the outcome to be the logarithm of the wage at age 30, and we estimate econometrically the functions $v^t(\pi, x)$, where x is the per-pupil expenditure on schooling. The details are provided in Betts and Roemer (2001).

In the late 1960s, when our population sample was in secondary school, per capita per annum school spending in the US was about \$2500, in 1989 dollars. We take as our policy space the set $\{(x^1, x^2, x^3, x^4) \mid \sum p^t x^t = 2500\}$; thus, a policy expends different amounts of educational dollars on students of different types. The first row of Table 1 gives the equal-opportunity policy.

To have equalized opportunities given the budget of the late 1960s, the US would have had to spend about five times as much per capita on the education of the most disadvantaged children as on the most advantaged children. The last column computes what the average wage would have been had the equal-opportunity (EOp) policy been implemented as a fraction of the average wage under the 'equal resource' (ER) policy, the policy where $x^t = 2500$ for all t . We see, perhaps surprisingly, that total earnings would have *increased* under an equal-opportunity policy by 2.6%. (This is, of course, a partial equilibrium calculation.)

Table 1. Equal-opportunity educational policy with SES types

r	x^1	x^2	x^3	x^4	$w^{\text{EOp}}/w^{\text{ER}}$
\$2,500	5,360	3,620	1,880	1,110	1.026
\$4,330	7,310	4,750	3,610	2,510	1.023

Type 1: Parental education ≤ 8 years.

Type 2: $8 < \text{PE} < 12$.

Type 3: PE = 12.

Type 4: PE > 12 .

It is clearly unrealistic to suppose that the allocation recommended in the first row could have been implemented. We accordingly ask: How much would the educational budget have to increase so that, in the equal-opportunity allocation, no type receive less than \$2500 per capita? The answer is: to about \$4330 per capita, an increase of about 72%. Now real educational expenditures per capita have more than doubled in the US since that time, so it would have been economically feasible (if not politically so) to implement the equal-opportunity policy gradually, with a higher budget, and not reduce spending on any student.

The basic reason that the equal-opportunity policy requires such large variation in expenditures across types is that the elasticity of wage earnings with respect to prior educational expenditures in the US is very small – this is a statistical fact which is not completely understood by labor economists. It is certainly possible that more sophisticated analysis will eventually produce larger elasticities, in which case the equal-opportunity policy would exhibit smaller variance across types.

In the SES typology, we do not predicate type on race. We can ask: if the equal-opportunity policy of the first row of Table 1 were to have been implemented, what effect would it have had on the distribution of wages of black men? In the actual data, blacks comprised 38.1% of the lowest wage quintile, approximately three times what they would, were race uncorrelated with wages. Under the above equal-opportunity policy, we calculate that they would still comprise 35.3% of the lowest quintile! In other words, if we predicate type only on the socio-economic status of the family, and ignore race, then the equal-opportunity policy will do very little to reduce the racial inequality of wage earnings.

We therefore carried out another calculation, where this time we predicated type on both the education of the parents and race. Our four types are now:

- LB: black and $PE \leq 10$
- HB: black and $PE > 10$
- LW: white and $PE \leq 12$
- HW: white and $PE > 12$.

The equal-opportunity policies are presented in Table 2, for two values of the educational budget.

We see that the variance in allocations across these four types is even greater than in Table 1. Moreover, we note that the HB type receives much greater compensation than the LW type. One might leap to the conclusion that being

Table 2. Equal-opportunity policy with SES/racial typology

r	x^{LB}	x^{HB}	x^{LW}	x^{HW}	w^{EOP}/w^{ER}
\$2,500	8,840	16,260	2,610	679	0.980
\$4,480	11,100	23,860	3,920	2,500	0.977

a low SES white is a more advantaged situation than being a high SES black, but that conclusion is too hasty, because the equal-opportunity allocation takes account of the frequencies of the different types, as well, and the HB type is about one-fourth the size of the LW type, so its members are cheaper to subsidize. We note that there is an aggregate cost to implementing the equal-opportunity policy with this conception of circumstances – the average wage would fall by about 2%. (Note that less is invested in the ‘low black’ than in the ‘high black’ type at the EOp policy. This is due to extremely low estimated productivities of investment for the former type.)

Betts and I believe there is an important conclusion from these calculations, for the American debate about affirmative action. In recent years, American institutions have been moving away from policies which predicate affirmative action on race, to ones which predicate it on socio-economic status of the individual. The view motivating this move is that equal-opportunity policies should be color-blind. What our work shows is that predicating equal-opportunity educational policy only on the SES status of individuals does almost nothing to change the relative economic outcomes of blacks in US society. To affect that racial inequality, it appears that we must predicate equal-opportunity policy specifically on race. Of course, our results are, to be precise, only suggestive of this conclusion. Institutions might be able to choose a vector of characteristics that comes close to capturing the social consequences of race without including race as such.

5 Equalizing opportunities for income acquisition with the fiscal system

In the second application I will describe, we ask: To what extent do different countries equalize opportunities for the acquisition of income? Let us define a type by the educational level of the individual’s parents. In our study (Roemer et al; in press), my collaborators and I partition young men into three types, based on whether their parents had only elementary education, completed secondary schooling, or had at least some tertiary education. We observe the pre-fisc distribution of income in each type. In all eleven countries of our study – the US, Britain, France, Netherlands, Belgium, Germany, Spain, Italy, Denmark, Norway, and Sweden – these three distribution functions do not cross, and are ordered as you would surmise. We can view the system of income taxes and transfers in a country as an instrument for equalizing opportunities for income in this sense: that if that system performed its job perfectly, then the post-fisc distributions of income of the three types would be identical.

To paraphrase: the pre-fisc distributions of income differ in our three types, in the expected way, that is, the worst distribution is associated with the type whose parents are least educated. If equal-opportunity holds, then the *distribution* of income in a type should be independent of the type. This is just like saying that, if equal opportunity holds, the rows in the intergenerational mobility matrix should be the same. We ask: to what extent does the tax-and-transfer regime of country make it the case that these three distributions are




Fig. 3a–c. Empirical distribution functions of pre-fisc income, three educational types, for US (a), Spain (b) and Denmark (c)

the same? Equality of opportunity places no importance on decreasing the variance of each distribution, for that variance is attributable to differential effort, but it does try to make the three distributions identical.

Figure 3 presents the empirical distribution functions of pre-tax income for Spain, the US, and Denmark. We see that there is hardly any pre-fisc inequality in Denmark.

We observe incomes, not wages, in the panel data sets of these ten countries. We then attribute to every individual a utility function

$$u(y, L) = y - \alpha L^{1+1/\eta},$$




Fig. 3a–c. (Continued)

where y is post-fisc income and L is labor. η is the elasticity of labor supply w.r.t. the wage. We then estimate the mapping of pre-fisc into post-fisc income

$$y = (1 - a)x + c,$$

where x is pre-fisc income. In all countries this affine relationship holds almost perfectly: we also estimated a quadratic relationship, but the quadratic term adds almost no explanatory power to the regression, except for Belgium (and there are reasons to be suspicious, there). Thus, there is a system (a, c) characterizing each country, where a is the marginal tax rate and c is the transfer; there is, as well, an amount of government revenue per capita raised, call it g . We now define the universe of feasible policies Φ as the set of affine tax systems which are revenue neutral in the sense of raising g as government revenue, after transfer payments.

Our optimal taxation exercise is to compute the affine tax system in Φ which maximizes the value of the equal-opportunity objective. In this case, this turns out to mean ‘Maximize the average post-fisc income of the worst-off type,’ a simplification of the general form, since the distribution functions of income of the three types do not cross. We then compare the observed fiscal system to the optimal one, and we define a measure of the extent to which the observed system achieves equal-opportunity, denoted v .

Table 3 presents some results. We see that Denmark, Norway, Sweden, the Netherlands, and Germany all ‘overtax’ with respect to our equal-opportunity objective: they tax more than they should to equalize opportunities with respect to this conception of circumstances. Belgium essentially achieves the optimum. The worst performers are Italy, the US and Britain.

It is noteworthy that the optimal equal-opportunity tax rate for Norway, Denmark, Sweden, and Germany is zero: this means that, in the equal-opportunity policy, there is no redistribution of income; rather, each citizen

Table 3. Equal-opportunity tax policy in ten OECD countries

Country	a^{obs}	c^{obs}	a^{EOp}	c^{EOp}	ν
Belgium	0.531	148.9	0.535	158.	0.9996
Germany	0.364	5540.	0	-17477	Over tax
Denmark	0.440	41021	0	-53989	Over tax
Italy	0.232	2.69	0.819	21.3	0.160
Netherlands	0.533	10410	0.474	18736	Over tax
Norway	0.393	45526	0	-63170	Over tax
Spain	0.376	172.8	0.605	663.9	0.748
Sweden	0.524	46886	0	-30207	Over tax
Britain	0.364	38.43	0.709	171.5	0.615
US	0.243	2036.	0.647	13578	0.200

Key: a^{obs} = the observed marginal tax rate; c^{obs} = the observed average transfer; a^{EOp} = marginal tax rate in the EOp policy; c^{EOp} = average transfer in the EOp policy; ν = extent to which observed policy achieves equality of opportunity

would pay a constant lump sum to fund government goods and services – hence, the negative values of c^{EOp} . This result occurs because there is so little pre-fisc inequality of the distribution functions of the three types that any taxation is not worthwhile, because of the deleterious effect on labor supply.

There are at least three possible reactions to these results. The first is that the Nordic countries, the two Germanies and the Netherlands have moved towards an equal *outcome* ethic, more radical than an equal opportunity ethic. The second interpretation is that, in fact, these countries are indeed equalizing opportunities through the fisc, but they have a more comprehensive set of circumstances than, simply, parental education. The third is to note that we have restricted policy to a small set (a unidimensional policy space), in requiring that all types face the *same, affine* income tax policy. Were we to optimize over a higher dimensional policy space, then it would probably be the case that no country would have achieved full equality of opportunity.

After family background, a preeminent circumstance one might wish to include is natural talent. Our next experiment was to augment the set of circumstances to include a measure of the natural talent of the individuals. For four countries – the US, Denmark, Sweden, and the Netherlands—we were able to find a secondary data set with IQ information, from an IQ test taken in the early teenage years, and we are able to simulate, using that data set, IQ values for individuals in our original panel data set. We then partitioned our large sample into six types, crossing the three parental education types with above and below average IQ. IQ is sharply correlated with type and income: for instance, in the Netherlands, only 33% of the lowest income quintile of the type with least educated parents have above-average IQ, but 69% of the lowest income quintile of the most advantaged SES type have above-average IQ. We then recomputed the optimal equal-opportunity affine tax system, with respect to this more comprehensive set of circumstances.

With IQ as a circumstance, Netherlands is no longer overtaxing, but Swe-

Table 4. Contrast of EOp policies, with and without IQ as a circumstance, $\eta = 0.06$

Country	a^{EOp}	ρ	ν	ε
US 3 types ED.ST	0.647	0.69	0.200	0.955
US 6 types (with IQ)	0.723	0.61	0.165	0.941
DK 3 types ED.ST	0	0.828	OT	1.035
DK 6 types (with IQ)	0	0.710	OT	1.035
SW 3 types ED.ST	0	0.88	OT	1.046
SW 6 types (with IQ)	0.257	0.78	OT	1.027
ND 3 types ED.ST	0.470	0.830	OT	1.007
ND 6 types (with IQ)	0.700	0.69	0.767	0.959

den and Denmark still are. Thus, Denmark and Sweden, through the fisc, more than compensate individuals for the disadvantageous circumstances associated with being in a family with poorly educated parents, *and* with respect to being of below average intelligence (whether this is due to nature or nurture).

We do not study how Denmark and Sweden achieve this remarkable egalitarianism. There are, it would seem, three possible explanations: first, their relatively homogeneous populations; second, their good educational systems; but perhaps most importantly, the solidaristic wage policy, which has compressed wage differentials a great deal. (The historical explanation of the solidaristic wage policy may well be related to population homogeneity.)

Of course, were we to create four categories of individual by IQ, rather than two, then it might well be the case that Denmark would no longer be overtaxing. And similarly, if we decomposed parental education into six rather than three intervals, we would see less equalization of opportunity.

6 Conclusion

To recap: one, if not the, major accomplishment of egalitarian theory since Rawls's reinvention of the field thirty years ago, is the inclusion of considerations of responsibility. In this way, as the philosopher G. A. Cohen has said, egalitarian theory has incorporated the most appealing idea in the arsenal of the anti-egalitarian Right, that people should be held responsible for their accomplishments. I have argued that welfare economists should take on board the lessons of this recent philosophical amendment, and that in doing so, the recommendations we make, about social policy, will be considerably more redistributive than utilitarianism proposes, yet considerably less radical than Rawlsianism proposes. Moreover, I believe that the equal-opportunity approach is the *right* one – it corresponds to what most people intuitively believe, that persons should be compensated for certain kinds of bad luck, but should be held responsible for much of what they do. When we work out the details, it turns out that the conception of equal opportunity that I have presented often involves quite different prescriptions from common conceptions of equal opportunity. The common conception holds, for instance, that equal

amounts of public educational resource should be provided to all students, yet our recommendation is considerably more compensatory than this.

People who initially hear me explain the equal-opportunity theory often say, “Well, this sounds very nice, but how can we decide what *are* the circumstances and what *is* effort?” These people are skeptical that the theory can actually be applied to policy. I have tried to challenge that skepticism by showing, with two specific policy applications, that even taking a relatively conservative approach, of delineating only a small number of circumstances, and attributing all remaining variation in outcome to effort, produces policy recommendations that are quite strongly compensatory. Surely, it is hard to argue that a person should be held responsible for the consequences of growing up in a home with poorly educated parents. Thus, in my view, the difficult philosophical and neurophysiological problem of delineating exactly the cut between circumstances and effort need not, for most all practical purposes, be solved. Social science has much to say about resource distributions that equalize opportunities even before philosophers and students of the mind further refine our conception of personal responsibility.

References

- Dworkin R (1981) What is equality? Part 1: Equality of welfare. *Philos Public Aff* 10: 185–246
- Dworkin R (1981) What is equality? Part 2: Equality of resources. *Philos Public Aff* 10: 283–345
- Betts J, Roemer JE (2001) Equalizing opportunities through educational finance reform. Discussion Paper, Institute for Social and Policy Studies, Yale University
- Rawls J (1971) *A theory of justice*. Harvard University Press, Cambridge, MA
- Roemer JE (1996) *Theories of distributive justice*. Harvard University Press, Cambridge MA
- Roemer JE (1998) *Equality of opportunity*. Harvard University Press, Cambridge, MA
- Roemer JE, Aaberge R, Colombino U, Fritzell J, Jenkins SP, Lefranc A, Marx I, Page M, Ruiz-Castillo J, Pommer E, Sansegundo M, Traenes T, Trannoy A, Wagner G, Zubiri I. To what extent do fiscal systems equalize opportunities for income among citizens? *J Public Econ* (in press)