

Epistemological Issues in the Theory of Equality of Opportunity

John E. Roemer
Yale University

Abstract. EOp (Equality of opportunity) is an algorithm for measuring the degree of opportunity equality in a society. The researcher must identify the kind of advantage of concern (e.g., income, life expectancy, wage earning capacity), and what circumstances are important. She then partitions society into types, where each type consists of persons with approximately equal circumstances. The set of feasible policies must be delineated. If income is the kind of advantage under study, policies consist of different income tax functions. The algorithm holds people responsible for their effort, but not responsible for their circumstances. Compatibilism is discussed, as is the concept of an age of moral consent. Children, who have not reached the age of moral consent, are responsible for neither their nature nor their nurture. The scope of EOp requires us to recognize that EOp concerns itself with treating those who are competing for social positions fairly, but does not address the issue of justice among those who consume what those competitors produce. Both are important.

Keywords: equal opportunity, meritocracy, circumstances, effort, type, responsibility, compatibilism.

I partition my remarks into seven sections.

I. THE EOP APPROACH

Consider a population whose members are competing for some valuable kind of *advantage* (G.A. Cohen's nomenclature - 1989), such as income, life expectancy, wage-earning capacity, or a state of good health. Call this kind of advantage the *objective*. To some degree, the achievement of the objective depends on *circumstances*, natural (such as genetic) characteristics of the individual and aspects of her environment over which she has no control (e.g., the socio-economic status of her family, her race and sex). Apart from these, her own *effort*, the degree to which she consciously decides to succeed in pursuing the objective, will matter. Finally, her success depends upon state policy, which may compensate persons with disadvantageous circumstances in order to improve their chances of acquiring the objective.

We partition the population into *types*. A type is the set of persons who have, as far as we can tell, very similar circumstances. Suppose, for example, we choose as circumstances the income of the family in which the individual was raised, her sex, her race, a rural versus urban background, and her IQ. These comprise five components. Suppose her race, sex, and rural v. urban background are binary variables, and we measure IQ and parental income as each having four possible levels. Then the number of *types* is $2^3 4^2 = 128$. If we have a population of 2^{20} souls,

about a million, this means on average a type will possess $2^{13} \gg 8000$ members. We will thus, with some confidence, be able to speak of a well-defined distribution of the objective within each type. The fact that the average size of a type is 8000 implies we can be fairly confident that the empirical distribution of income – if income is the objective – within a type, is something not due to chance, but is intimately associated with the circumstances of that type.

The *data* of the problem consist of the objective, the distribution of circumstances, a typology, the distribution of the objective in each type, and a set of social policies.

I have not said that it is these circumstances that *cause* the income distribution to be what it is, but I think we must also be able to assert causation. That is, we must have a theory of how variation in these circumstances causes variation in income. This theory will come from social science.

Let us continue with the example. What causes income to vary *within* types? One can be quite confident that, even if the number of types is fairly large – like 128 – there will be substantial variation of the objective within most types. The EO theory asserts that this variation is due to *differential effort* within types. Effort, as I said, comprises the conscious actions individuals have taken to succeed in attaining the objective. There may as well be *episodic luck* – that is, the luck of being in the right place at the right time, or perhaps the effect of circumstances that we have not included in our list. I will say more about this below.

Returning to effort. The theory assumes that the more effort a person expends, the higher the level of the objective he achieves. I am thus discounting useless effort, which may be misguided because a person has an incorrect theory of causation. This means that, except for luck or useless effort, we may assume that the higher a person is in the distribution of income of his type, the more effort he has expended. In other words, if a person sits at the 60th centile of the income distribution of his type, then he sits at the 60th centile of the effort distribution of his type. This is the case *if* we can assume that all individuals in the type have faced the same *policy* (say, state financing of education), and because they all, by hypothesis, have the same circumstances. If all persons in a type have not faced the same policy, then the differences in the application of policy must itself be circumstances.

Thus, we can define the *degree of effort* of an individual as the centile of the income distribution of his type at which he sits. But how can we compare the efforts of two persons in different types? We must understand that the *distribution of effort* within a type is a *circumstance* of that type: that is, it is a fact of nature, not of any individual. That distribution is not associated with any individual, but with the circumstances and policy that define the type. So it is a circumstance for that type, and as such, persons should not be responsible for being in a type with a ‘low’ distribution of income – a distribution with, let us say, a low mean.

Because of this close association of the distribution of effort with the circumstances of the type, we cannot compare efforts easily across types. What we need is a measure of effort that sterilizes effort of its type-specific aspects. I say a measure that does this is the *centile* of the effort distribution of her type at which an individual sits. This is, as I've said, also the centile of the objective distribution of the type at which he or she sits.

So if a man, Jack, a member of a type of white men with advantaged social backgrounds, and a woman, Jill, a member of a type of black women with disadvantaged backgrounds, both sit at the 60th centiles of the respective income distribution of their types, I say they have expended *equal degrees* of effort, even though Jack may have spent many more hours studying than Jill did. This is the consequence of sterilizing the effort distribution of its type-specific characteristics. Of course, this view is only sensible if we have a *causal theory* explaining why Jill's circumstances made it harder to study than Jack's did. The fact that they sit at the same centile of their respective income distributions means *relative to others with the same circumstances*, they have performed equally well.

What is the optimal equal-opportunity policy? It is a policy that will compensate disadvantaged types with resources that improve their income distributions. To be specific, we'd like to find a policy that renders the income distributions of the types to be as close as possible to each other. In real life, we can never find a policy that will equalize all the income distributions *at the highest possible level*; given the space of feasible policies, there will be an optimal one. We need not here discuss the statistical details of how one measures the closeness of income distributions to each other: that's a technical issue, on which there is ample literature.

It is also desirable that we include as many circumstances as we can measure, for if the partition of the population into types is too coarse, we will implicitly assign to effort some achievement that is, in reality, due to an unmeasured circumstance. We now have some data sets with quite fine partitions of populations into types. Sweden, for example, has a statistical registry of all residents that permits us to have over a thousand types, with meaningful income distributions for each type. Sweden, however, is virtually unique in this regard. The United States and the UK also have very good data sets, but that's about it. There are no developing countries for which data sets currently exist that allow one to make a fine partition of the population into types. This means that when we measure the fraction of income inequality that is due to circumstances, we are in fact only computing a lower bound to the true number, because the paucity of data on circumstances means we will falsely assign to 'effort' effects that are in reality due to unmeasured circumstances.

Returning to the issue of episodic luck: I think the only way to deal with this is to increase the set of circumstances that we measure. There are some for dealing with luck, but it's beyond my scope to discuss these today.

II. COMPATIBILISM AND RESPONSIBILITY

The materialist thesis says that any action a person takes has a correlated brain state. Thought and action, in theory, can be read from brain states. *Compatibilists* say that the materialist thesis and the thesis that responsibility is meaningful are consistent. One can believe the materialist thesis and still assert that there are actions for which one can hold a person responsible. Incompatibilists say the two are inconsistent: if the materialist thesis is true, responsibility makes no sense.

This is, at least, my understanding of these terms.

I believe it's probably the case that most philosophers are compatibilists. What thinking supports the compatibilist view? I think the justification of the view is as follows. If one has contemplated the action, one was in a calm and sober state, one's powers of thought were whole, then one is responsible for the action. One way of countering this view would be to say that evolution has endowed us with a theory of responsibility that is very useful for survival, but is false. That is, having this theory of responsibility will help us avoid dangerous people, because we won't wish to associate with people who 'behave badly.' It's easier to justify ostracizing or disassociating with a person who responsibly takes bad actions, than one who is 'out of control.' One might desire to help a person who is out of control, even if it is dangerous to do so.

I think evolution may very well have endowed us with a theory of responsibility, but I don't say it's false: I say it is part of who we are. Of course, society may, over time, hold people responsible for a smaller set of actions than it did earlier in history, as we learn how circumstances cause behavior – that is, as social science matures.

I am a compatibilist. I generally hold people responsible for actions that they appear to have arrived at by calm, conscious thought, even though I recognize that there eventually may be a way of explaining those actions as due to circumstances for which the individual was not responsible.

A few years ago, I was writing a paper measuring equality of opportunity in several countries, and one of my collaborators proposed using brain scans of people that were available in our data set as circumstances. I strongly opposed doing so. Why? As a compatibilist, I believe that *every* action a person takes, responsible or not, has an associated brain state. Thus, showing that particular brain states were associated with the action tells us nothing about whether we should hold the person responsible for the action. I am not saying that brain-scan data are irrelevant, but at present, I think that using such data to excuse an individual from responsibility is only permissible if we have a causal theory of action.

III. THE PRIVACY ISSUE

I know of two philosophers (there may be more) who have objected to the EOp approach on the grounds that acquiring information about circumstances may violate

privacy. These are Elizabeth Anderson and Norm Daniels. I quote from Moriarty (2005), who writes that Daniels writes that ‘responsibility tests’ are ‘intrusive and demeaning and violate concerns about liberty and privacy.’¹

This concern reflects a poor understanding of how the application of the EOp approach proceeds. As I said, there are four decisions a social scientist must make: what is the objective, what are the circumstances, what is the typology, and what is the *policy space*.

I will focus now on the policy space – here is an example (see Roemer 2013 for the details). There is a population with two types: Rich and Poor. For the purposes of the example, we assume this is the single circumstance – being rich or poor. We will be concerned with the objective of life expectancy, and as policy makers in the Ministry of Health, we are not to be concerned with how people became Rich or Poor. Rather, from the Ministry’s vantage point, income level is a circumstance that it believes should not, from the ethical viewpoint, affect a person’s life expectancy. We wish to design a policy that will equalize life expectancies in the two types.

There are two diseases that afflict this population. The Rich suffer from cancer; the Poor suffer from cancer and tuberculosis. Within each type, the probability of developing cancer or TB is a function of *life-style quality*: how much one exercises, whether or not one smokes, diet, and so on. Life-style is the effort variable. There is, unsurprisingly, a different distribution of life-style quality in the two types. The Rich have a better distribution of life-style quality than the Poor: their mean life-style quality is high, and therefore they suffer a lower rate of cancer than the Poor. Only the Poor contract TB, and this probability is decreasing in their life-style quality.

Suppose the Ministry has a fixed budget for medical care. There are various treatments for each disease; the more a hospital spends on the treatment of a disease, the higher will be the life expectancy of the population receiving that treatment. Question: How much should the Ministry spend on each disease treatment? The data one needs to address fully the allocation problem include what the life expectancy will be within each type as a function of how much is spent on each occurrence of the disease. Knowing the life-style of a person who presents with cancer or TB may tell us something about the effectiveness of treatment.

Horizontal equity is a property of a policy under which anyone who presents with a certain disease (with a particular level of severity, etc.) receives the same treatment, regardless of his circumstances *or effort*. One might well want to consider only policies that are horizontally equitable – policies that will use the same treatment for any person who comes to the clinic with a given disease and a given level of severity. There are two reasons: first, privacy may be violated if the medical practitioner has to get details of

[1] Anderson’s citation is similar (Anderson 1999, 289): “equality of fortune, in attempting to ensure that people take responsibility for their choices, makes demeaning and intrusive judgments of people’s capacities to exercise responsibility and effectively dictates to them the appropriate uses of their freedom.”

life-style quality and circumstances, and second, because this may put the practitioner in the position of triaging patients, which might reduce the trust between practitioner and patient.

In other words, one might not use certain policies that could improve the life-expectancy of the Poor, making it closer to that of the Rich, because there is a trade-off between doing so and other values one has – such as patient privacy and the quality of the relationship between practitioner and patient.

That is my first response to Daniels and Anderson. My second response is that, yes, we may want to use information about patients that violates privacy – indeed, we may decide this is essential – but that need only be done with a *small random sample* of patients. For this random sample, one acquires the private information, and then one experiments with different treatments. The experiment enables us to compute what the optimal treatments for each disease would be both by discriminating in the treatment of patients based on their personal information, and by ignoring this information. One can then decide how much particular kinds of information are worth knowing. Of course, it would be even better from the purely clinical viewpoint to get the exact life-style quality of the patient by asking probing questions, but it may be prudent not to compromise the patient's trust in the doctor and treatment by doing so.

Thus, the critique from Daniels and Anderson can be rebutted by understanding that the policy space of the Ministry may be chosen so that clinicians need not compromise patient privacy.

Another example might be whether upper middle-class parents who read a lot to their children should be taxed for doing so. Aside from the difficulties of enforcing such a policy, this is not a policy that any society today would support. As I said, the EOp theory is meant to implement a policy that conforms to the society's views about responsibility. Very few societies would support a policy of 'levelling down' like taxing parents who read to their children. The right policy is to provide more teaching personnel so that all young children can be read to by loving adults.

IV. THE AGE OF CONSENT

Children are humans in the state of formation, which includes moral formation. We should think of there being an age of moral consent below which a person is not responsible for his behavior. More subtly, one might wish to hold children increasingly responsible as they get older, although not being fully responsible until the age of 18 or so.

For simplicity, let us ignore the continuous approach, and say that the age of consent is a certain age during adolescence, and a person is not responsible for her actions below that age. Of course, we teach children what responsibility means by commenting on the morality of their actions. Not holding the child responsible does not mean we fail to remark upon the moral quality of the child's behavior.

It follows from this view that effort – responsible choice – strictly speaking, does not exist for children below the age of consent. Every action of a child, in this binary view, is the consequence of circumstances. Both nature and nurture are clearly circumstances. This implies that if we are performing an EOp analysis of income acquisition among young adults, and if we had a complete biography of the individual at the age of consent, that biography should be treated as a circumstance! If the age of consent were 16, then everything that happened to the child, and that the child did before the age of 16, should be treated as a circumstance or as being caused fully by circumstances.

Now one might say that this view will result in holding people responsible for very little about their achievement of the objective – because what the child has accomplished by the age of 16 may severely narrow the degree of acquisition of the objective that the child can attain by the age of 30. I say it would not be correct, on this account, to conclude that we should lower the age of consent: rather, our societies should undertake effective(s compensation policies beginning early in childhood so that by the age of 16, virtually all children have acquired good skills and a good moral character. This should begin with state-financed early-childhood education, and other policies (such as paid parental leave) that will produce more successful children. Europe is surely much further along in having such policies than the United States.

But if, despite our attempts, there is a great deal of variation in accomplishment by the age of 16, which engenders a great deal of variation in income acquisition by the age of 30, then the EOp view says we must attempt to compensate those who are disadvantaged, through policy. Adopting the view that before an age of consent sometime in adolescence the biography of the child is a circumstance will have quite radical implications for redistributive policy among adults.

V. THE STATISTICAL VERSUS THE IDEAL APPROACH

Philosophers, when considering a theory such as the EOp theory, delight in constructing hypothetical examples showing that application of the theory will result in errors. A genre of such examples are ones in which, with a given list of circumstances, one will always mis-categorize a certain outcome as being due to the lack of effort, when it is in fact determined by a missing circumstance. I say this is holding social science to an ideal-theory standard.

Social scientists and policy makers are concerned with statistical outcomes. We want to get things approximately right. At some point, in delineating the set of circumstances and the policy space, we will cease to rely on science and instead use intuition. But this is not a defect of the EOp theory – it is true of all social science, and perhaps of all science. Theories of causation require models of the world, and a model is a simplification of an impossibly complex reality. To be able to construct a theory, one must prove theorems about what happens in one's model. A modeler is always open to

criticism that she has ignored an important aspect of reality in constructing the model. Since *a priori* she has only a vague theory of what is happening in the aspect of the world she is studying, she must make the decision about what features of the world to include in the model and what features to ignore, or, as we say what features to 'abstract away from.' Having no complete theory at the point of model-construction, that decision must be made with intuition, not science.

The same is true in analytic philosophy. To construct a theory (say, of distributive justice) one constructs a model that abstracts away from features of reality that one thinks *intuitively* are not essential for the theory. Then, if the theorems one derives from the model seem bizarre, one changes the model and starts again. (Rawls described this process as one of reflective equilibrium, but the idea is much older than Rawls: surely Euclid constructed his axioms for geometry by trial and error.) I view social science and philosophy as in part arts, because at crucial points they rely on the intuition of the practitioner. Perhaps philosophy is more of an art than social science, for the social scientist can often observe whether the conclusions of the model appear to hold in reality, while philosophers deal with claims that are not observable.

VI. THE SCOPE OF EOP

A major application of the EOP theory is to the fair treatment of individuals who are competing for desirable positions in society: admission to good high schools or universities or medical schools, or acquiring government contracts. There is less (or no) emphasis on what people who acquire positions can do for the rest of society. In the US, we talk a lot about access to Ivy League universities for disadvantaged students; we talk less (or very little) about what the graduating classes will do for society as a whole. Of course, everyone understands that one of the university's roles is to train a skilled cohort of workers. But this is not seen as the moral issue that equal opportunity is.

Roughly speaking, *meritocracy* is the view that those who will be most skilled in carrying out the requirements of social positions should be recruited to those positions, and EOP focuses instead on fairness to the cohort of applicants for those social positions. My own view is that both considerations are important, and neither side should ignore the importance of the other.

Why do we laugh when I propose that we equalize opportunities for short players to be recruited to professional basketball teams? After all, being short is almost surely something beyond control of the individual. The answer is because there are only a handful of basketball players, but there are millions of fans. Here the welfare of the fans trumps the fair treatment of short players. I am here supposing that basketball games would be less interesting if every team had some short players. I do think, however, that we should practice equal opportunity for short players to join high school basketball teams: for these players may thereby become successful coaches, if not professionals.

More generally, I think we must consider how distant the social positions we are concerned with are from the final provision of consumption goods and services to citizens. I think expenditures on early-childhood education should be *solely* concerned with equalizing the achievements of children. I think admission to high schools should be *largely* concerned with equalizing opportunities for children, but to some degree be concerned with preparing a highly skilled cohort of adults to produce complex goods and services for the consumption of the citizenry. I think admission to medical schools should be concerned with both merit and opportunity. But I think the standards for passing the surgery board exams should be entirely based on merit: because the outcome is producing doctors who immediately will provide surgeries consumed by citizens.

Now a militant meritocrat will respond that if we admit a substantial number of disadvantaged students to the Bronx High School of Science in New York City, and other such excellent institutions, that will eventually reduce the pool of highly accomplished medical-school applicants, which will then force us to lower the standards on the surgery boards, to produce a large enough cohort of surgeons. So the average quality of surgery will degrade.

I think the proper way to evaluate this claim is empirical: it is not, in my view, an ethical question. If it turns out that admitting a substantial number of disadvantaged students to medical schools does degrade the quality of surgery the society achieves, then we must conclude it is too late at the point of medical school admissions to rectify the injustice of the disadvantage – the rectification must begin at an earlier age². This is a similar conclusion to the one I reached in the example concerning the age of consent.

VII. RIGHT-WING VERSUS LEFT-WING USES OF EOP

The EOp theory that I have outlined attempts to equalize in a population the acquisition of an objective so far as it is influenced by circumstances beyond the control of individuals, but not to equalize the degree of objective acquisition to the degree that it is differentiated due to effort. The left-wing approach emphasizes the equalization-via-compensation-for-disadvantage aspect, and the right-wing approach emphasizes the differential-returns-to-effort aspect. Emphasizing the ‘desert’ aspect of EOp seems to me to be part of the right-wing approach. That is, the left-wing approach emphasizes the claim that surely people do not deserve the poor outcomes that are associated with

2] Only 10 black students passed the entrance exam for admission to the prestigious Stuyvesant High School in New York City in 2020. It is unclear, however, whether lowering the standards for admission for disadvantaged students, or giving crash courses teaching to the exam, are good policies, even if they succeed in raising the matriculation of black students. Much better, I conjecture, would be sharply to improve the public elementary schools that most disadvantaged students attend. Even that may be insufficient: eliminating homelessness and poverty of families may be necessary. Obviously these policies are much costlier than achieving a higher black matriculation by teaching to the test. For the effects to show will take time. Rectifying the costs of racism does not come cheap, but ‘no short-cuts’ may be the right slogan.

disadvantageous circumstances, and the right-wing approach says that surely people do deserve rewards for effort.

When I first proposed the EOp theory (Roemer 1993), the paper was entitled “A pragmatic approach for the egalitarian planner.” By pragmatic I meant that the social planner did not have a moral theory of responsibility, but rather would take the theory of moral behavior of the society in which she was working to determine the set of circumstances. I suppose it’s what Rawls would have called a political, not metaphysical, theory.

I continue to hold this view: that the theory is a tool with which a society can construct social policies that are consonant with its own view of the causation of behavior. Having the language of the theory – circumstances, effort, type, objective, and policy – sharpens discussions. For instance, in the affirmative action debate that has taken place over the last forty years in the US, there have been two kinds of objection to affirmative action policies. The first claims that scarce social positions (e.g., in medical schools) should be assigned on the basis of predicted merit, while the second claims that race is not the right circumstance upon which to focus – rather, the relevant one is socio-economic disadvantage. The language of the theory tells us that the first objection is against the theory of equal opportunity as such, while the second is only quibbling with what the set of circumstances should be. It is interesting to note that policies have evolved, in the main, to respond to the second objection rather than the first. That is, the circumstance of race has often been replaced in university admissions policy by one of socio-economic disadvantage. But recruiting for social positions on the basis of merit alone is now quite rare.

john.roemer@yale.edu

REFERENCES

- Anderson, E. 1999. What is the Point of Equality?. *Ethics* 109(2): 287-337.
- Brunori, P., Hufe, P. and Mahler, D. In press. The roots of inequality: Estimating inequality of opportunity from regression trees and forests. *Scandinavian Journal of Economics*.
- Cohen, G. A. 1989. On the Currency of Egalitarian Justice. *Ethics* 99(4): 906-44.
- Moriarty, J. 2005. The epistemological argument against desert. *Utilitas* 17: 205-21.
- Roemer, J. 1993. A pragmatic theory of responsibility for the egalitarian planner. *Phil. & Public Affairs* 22: 146-66.
- . 2013. Equity in health care delivery: Some thoughts and an example. *Research on economic inequality* 21: 471-89.