

The Special Significance of Equality of Opportunity

Martin O'Neill, University of York

29 March 2013

martin.oneill@york.ac.uk

Comments always very welcome.

(This very much a draft, so the usual caveats obviously apply. Not for quotation, circulation, etc. Apologies for the lack of a full set of footnotes and references.)

I.

Introduction

My aim in this article is to examine the special significance of equality of opportunity. I want first to answer the question of whether equality of opportunity really is special, when compared to any other kind of distributive equality. My answer will be an affirmative one. Equality in the distribution of opportunity has an importance that is not derivative from a broader concern with egalitarian distribution and, therefore, a plausible view of distributive justice has to give a special place to a concern with equal opportunity. My hope is that, in answering this question of the significance of equality of opportunity, one can also thereby show some interesting features, and intriguing shortcomings, of views that do not grant it a special place. I will then try to defend the “special significance” view by relating it to broader issues regarding the self-understanding of the agents who are also the subjects to whom an account of justice is addressed. My account looks to provide a reconstructive justification for a view that is very close to Rawls’s view of the place of equality of opportunity, but is revisionary in a number of respects, not least of which is the rejection of Rawls’s commitment to the *lexical priority* of a principle of equality of opportunity over other distributive concerns.

II.

What Do We Mean When We Talk About Equality of Opportunity? Some Distinctions

Let us distinguish between two senses of equality of opportunity. Some views about the content of the demands of distributive justice take the form of general principles of equality of opportunity. For example, Richard Arneson has defended a distributive principle of “equality of opportunity for welfare”, while G. A. Cohen defended a distributive principle of “equality of access to advantage”, where the relevant idea of *access* is closely related to, but not coextensive with, an idea of opportunity. To abstract away from particular instances, consider, then a *standard luck egalitarian* view, which makes the following claim:

- (1) The central demand of distributive justice is a requirement to bring about equality of opportunity for X (where X is some general distributive currency or, in Cohen’s terms, ‘distribuendum’)

The *standard luck egalitarian view* holds that demands of justice are satisfied when and only when every relevant individual – every ‘subject of justice’ – has an equal opportunity to receive X, for some value of X to be determined. Clearly, in some

important sense, such a view is a view about the special significance of equality of opportunity, for on such a view, the demands of a principle of equality of opportunity *just are* the demands of (egalitarian) justice.

Compare this to what Rawls has in mind when he talks about equality of opportunity. When Rawls talks about fair equality of opportunity (FEO) he has in mind a specific distributive principle regulating, in particular, the distribution of “powers and prerogatives of offices and positions of power and responsibility”. Unlike the case above, Rawls’s principle is not a comprehensive distributive principle describing the demands of distributive justice in general, but is a more circumscribed, domain-specific principle, specifically regulating the distribution of one particular kind of social primary good. The name that Rawls gives to this particular kind of social primary good – “powers and prerogatives of offices and positions of power and responsibility” – is somewhat cumbersome for our purposes, so let’s shorten it to something punchier: *jobs*. (This is a simplification, but not one that will prove to be too excessive.)

Comparing it to the *standard luck egalitarian* claim (1), Rawls’s claim has something close to this form:

- (2) Among the demands of justice are that the distribution of *jobs* should be determined by a particular principle of fair equality of opportunity.

So, that’s a first difference between these two ideas of equality of opportunity: one is comprehensive and the second is ‘job specific’. The first provides a comprehensive account of the demands of (egalitarian) justice, while the second provides an account of how *jobs*, in particular, should be distributed.

This distinction in turns suggests a more general distinction between two kinds of theories of distributive justice. One kind of view about distributive justice seeks for a single distributive principle, the satisfaction of which would produce a just outcome. Call such a view a *single principle view* of distributive justice. Another variety of view about distributive justice suggests that the truth about distributive justice can only be captured by the conjunction of more than one principle. Call this a *multiple-principle view*. The set of views that I have called *standard luck egalitarian* views are *single-principle* views, but they are not, of course, the only such views. By way of stark contrast, the view that Michael Walzer defends in his book *Spheres of Justice* is very much a multiple-principle view, as Walzer thinks that there are sui generis principles of justice that are each appropriate to a different domain of social goods but, of course, Walzer’s is not the only such view.

On a single principle view, such as a standard luck egalitarian view, equality of opportunity in the distribution of jobs may be an important goal, but its value would have to be explained by the instrumental role that this variety of equality of opportunity played in terms of its broader consequences. For example, one might defend a principle of equality of opportunity with regard to the distribution of jobs because this was (part of) the best way of achieving overall equality of opportunity with respect to welfare or resources (or some other X). But its significance would not be special. Instead, it would be one justifiable policy commitment among others,

justifiable on instrumental grounds in terms of a broader and more comprehensive egalitarian aim.

Now, although this may strike some readers as being unfairly stipulative, I want for the sake of simplicity to make a terminological ruling. When, in this article, I talk about “equality of opportunity”, what I mean is equality of opportunity with regard to the distribution of *jobs* in particular. What I mean, therefore, is something at least closely akin (although, as we shall see, not identical) to what Rawls means by “equality of opportunity” when he talks about “fair equality of opportunity”. (Incidentally, I would suggest that, when people discuss ideas of equality of opportunity in quotidian political discussions, it is something like this idea of equality of job opportunities that they tend to have in mind. By contrast, the standard luck egalitarian view’s comprehensive principle of “equality of opportunity for X” feels, by contrast, to be more abstract, and more distant from day-to-day political discussions.)

My central question, then, is this: is equality of opportunity special? Does it have a particular, distinctive kind of significance? If not, then there are two possibilities; either there is no principle of equality of opportunity that we should seek to enact in order to create a just society (as we shall see, this is Arneson’s view), or otherwise there *is* such a principle, but its standing depends derivatively on certain instrumental effects that it may have. On the other hand, if equality of opportunity does have a special significance, as I shall claim that it does, then we have decisive reason to reject single-principle accounts of distributive justice which, by their very nature, fail to take that special significance seriously.

III.

Rawls’s Complex Double Principle of Fair Equality of Opportunity (and His Five Principles of Justice)

I began by contrasting the *standard luck egalitarian view* with Rawls’s view of equality of opportunity. But that naturally raises the question of what, precisely, Rawls’s view actually is with regard to equality of opportunity. This is worth investigating in some detail, not least because the structure of Rawls’s view is in some ways more complicated than has been commonly recognized.

John Rawls famously defends two principles of justice. Less famously, he actually endorses five lexically ordered principles of justice. Or, at any rate, I shall suggest that this is the best way to interpret Rawls’s view. Ranging over the third and fourth of those five principles is embedded a more complicated view about equality of opportunity than the one that is standardly attributed to Rawls.

It is well known that Rawls’s second principle of justice, which deals with distributive justice, rather than with the protection of basic liberties (which is done by the first principle) itself splits into two. So, we tell our students, rather often, that, although Rawls says that he has two principles of justice, he really has three. Of course, when we say that Rawls has three principles of justice, we’re over-simplifying. Talk of three Rawlsian principles is oversimplification because Rawls also endorses a zeroth principle of justice, which is to have lexical priority over the other three. The zeroth

principle of justice states that a society is first to look after the basic human needs of its citizens, before worrying about the basic liberties. That seems entirely sensible, given that the basic liberties are of no value to citizens whose lives are seriously imperilled by abject material want. A basic structure of legal and political institutions would be imposing an unreasonable set of policies on people, and would be doing so in a way that citizens would have good reason to reject, if it were to prioritize the protection of certain civil and political liberties even in circumstances where the predictable cost of doing so would be to reduce the population below the most basic level of material sustenance. (Obviously, there will be difficult cases at the margins, and difficulties associated with specifying the appropriate background baseline of basic needs. But I set aside those difficulties for now, as lying outside my present concerns).

The addition of the zeroth principle brings us up to from three lexically-ordered principle to four. But, in fact, there are five principles, because Rawls's view also gives a rather distinctive place to a principle guaranteeing the Fair Value of the Political Liberties (FVPL). In responding to the critiques mounted by "radical democrats and socialists (and by Marx)"¹ and in particular to the early lines of objection pushed on Rawls by Norman Daniels in "Equal Liberty and Unequal Worth of Liberty", Rawls makes clear that the basic constitutional protections associated with the first principle of justice are to be supplemented by whatever measures are necessary in order to safeguard the *fair value* of the political liberties (but not the other liberties). What Rawls means by the guarantee of their *fair value* (and note that he does not say *equal* value, but specifically only *fair* value) is that:

"(i) This guarantee means that the worth of the political liberties to all citizens, whatever their economic or social position, must be sufficiently equal in the sense that all have a fair opportunity to hold public office and to affect the outcome of elections, and the like. *This idea of fair opportunity parallels that of fair equality of opportunity in the second principle.*"²

One small point to make about this is that, despite Rawls's confident and seemingly unproblematic distinction between the political liberties and the other liberties, it is not clear that there is really any such distinction to be made. For among the political liberties, as well as easily demarcatable political liberties such as the freedom to vote are also, presumably, general liberties of (political) speech and (political) association that are indistinguishable in their structure from general freedoms of speech and association. The general liberties that allow individuals to pursue their own private conceptions of the good with likeminded fellow citizens just are the very same liberties of which citizens also make use when pursuing, with their fellow citizens, the project of engaging in political argument and debate, and determining how their polity is to be governed. It is therefore implausible that "the political liberties" really marks out a distinct category. Nevertheless, the set of basic liberties does, uncontroversially, have value to individual citizens in both the private and the political domains, or, to put things in what may be a clearer way, these liberties have value with reference both to individuals' private interests and their political interests. Thus, it would be more

¹ Rawls, JF, 148.

² JF, 149 (*my italics*)

felicitous to talk about “the fair political value of the basic liberties” rather than the “fair value of the political liberties”.

But what of my claim that we should understand Rawls’s commitment to this kind of fair political value of the basic liberties as constituting a distinctive fifth principle within his general architectonic of principles of justice? That claim of mine may seem implausible if one is impressed by one particular kind of reading of Rawls’s claims on this matter, which present FVPL as an *internal* feature of the equal basic liberties, properly understood, with the requirement to ensure the fair value of the political liberties being part of the proper understanding of the meaning of the basic liberties themselves. As Rawls puts it:

“(ii) When the principles of justice are adopted in the original position, it is understood that the first principle *includes* the proviso [of FVPL] and that the parties take this into account in their reasoning. The requirement of the fair value of the fair value of the political liberties, as well as the use of primary goods, *is part of the meaning* of the two principles of justice.”³

That may suggest that what Rawls is giving us is a particular interpretation of the idea of the basic liberties, rather than, as I’m suggesting, an additional principle to be understood alongside the traditional first Rawlsian principle of justice. If that picture were the most compelling one, then talk of a second, distinct, opportunity principle in Rawls might seem implausible, notwithstanding Rawls’s remark that the idea of the fair value of the political liberties “parallels that of fair equality of opportunity”, with the parallel principles operating, respectively, in the political and economic domains.

However, the idea of FVPL as simply an *interpretation* of the specific meaning of the equal basic liberties is less plausible if one starts to think about concrete cases where the fair value of the political liberties might come into question, and if one thinks about the institutional implementation of a regime of basic liberties and of mechanisms for ensuring their fair value. Rawls talks of the first principle of justice as applying at the stage of a constitutional convention, the second of his four-stage sequence for the application of principle of justice.⁴ This role fits with the generality and comprehensiveness of the basic liberties as defined in justice as fairness. But the idea of the fair value of the political liberties appears not as something that could be specified in a general way at the constitutional stage, but rather as a complicated matter of political judgement, demanding a range of different kinds of policies to work in harness. Among the policies that would be needed in order to deliver the fair value of the political liberties, Rawls mentions the public funding of elections, the regulation of campaign finance, guarantees of more equal access to public media (as in the regulation of party political broadcasts on television in the UK), and “certain regulations of freedom of speech and of the press (but not restrictions affecting the content of speech).”⁵ But there is no reason to think that this list need be exhaustive, and there is evidence that Rawls himself did not think that it was.

Just as one might need to limit socioeconomic inequalities in order to guarantee fair equality of opportunity across generations, so too one may have reason to limit

³ JF, 149 (my italics)

⁴ Rawls, JF, 48.

⁵ Rawls, JF, 149.

socioeconomic inequalities because of their effects on the fair value of the political liberties. Indeed, Rawls took the view that among the most significant justifications for the rejection of welfare-state capitalism, in favour of a socioeconomic regime that took the form of either liberal socialism or of a property-owning democracy, is that the distribution of control over property in those latter kinds of regimes would be consistent with the protection of the fair value of the political liberties in a way that capitalist welfare states could not be. It therefore seems as if FVPL is a wide-ranging principle, carrying implications for distributive justice that could only be worked through at the legislative stage and at the stage of application and implementation. This gives the idea of FVPL much more in common with the role in Rawls's theory of the principle of FEO than with the constitutional principle of equal basic liberties, and suggests that it is more natural to think of it as a supplement to that constitutional principle rather than an interpretation of it.

This "fifth principle" reading of FVPL, as opposed to the "internal" reading, is further buttressed by the very sharp distinction that Rawls goes on to draw between "the political process as a public facility", understood as "specified by the constitutional rules and procedures", as against the further issue of the need to "secure for each citizen a fair and roughly equal access" to the use of that public facility.⁶ The FVPL principle is needed to ensure access for all to that public facility under fair terms, as without that additional guarantee, the wealthy would colonize the limited space available to access that public facility, to the detriment of the less economically advantaged and, as Rawls puts it, "the difference principle is presumably not sufficient to prevent this."⁷ Note, furthermore, that the question of what would be needed in order to guarantee fair access to the political institutions of a just society (i.e. the "public facility" generated by the constitutional rules of the polity, as determined by the first principle of justice) could not be specified before the prior specification of the nature and scope of those institutions themselves. Rather than treating FVPL unproblematically a part of the meaning of the equal basic liberties, it makes more sense, conceptually and procedurally, and in terms of thinking through the "rules of regulation" and public policies that Rawls's principles would mandate, to treat FVPL as a distinct principle of justice, lexically prior to the principle of fair equality of opportunity in the economic domain, and to the difference principle, but lexically subsequent to the first principle of justice.

There is good reason, then, to take with seriousness Rawls's claim that FVPL "parallels... fair equality of opportunity in the second principle". Note that both principles relate to the distribution of opportunities to occupy particular kinds of "offices and positions of authority and responsibility,"⁸ or, to put things more simply, to the distribution of opportunities to occupy particular *jobs*. (Although, as we've seen, fair equality of opportunity for political jobs is a necessary but insufficient part of FVPL, the other aspect of the principle relating to fair opportunity for political influence, more broadly considered.) The central difference between the two principles is that FVPL ranges over jobs in the political domain, whereas FEO ranges over jobs in the economic domain, but both govern the distribution of positions of power within the shared institutions within which social cooperation takes place in conformity with the demands of justice. Both principles make demands on the

⁶ Rawls, JF, 150.

⁷ Rawls, JF, 150.

⁸ Rawls, JF, 58.

regulation of procedures for selecting individuals for positions of power – whether that involves formal requirements of non-discrimination in hiring to ensure the realization of “careers open to talents” or, more substantively, in requiring regulation of economic and social life so as to remove unjustifiable social barriers that may stand in the way of the careers of some groups of citizens.

In a sense, then, we can simplify and say that FVPL and FEO together form a hybrid, dual-domain principle of fair opportunity for *jobs*. This principle occupies the space in the Rawlsian architectonic in between the principle of equal basis liberties and the difference principle. I take it that, if pushed, the most detailed account of their relationship would be that FVPL should have lexical priority (or, at any rate, strong presumptive priority) over FEO, given that in some respects fair equality of opportunity for political jobs can be seen as a precondition for the right kind of distribution of political power across different groups within society, of a kind that could be a precondition for the legislative and practical realization of the principle of FEO itself. But one does not do much violence to the proper understanding of the structure of Rawls’s view if one says that he has a single job opportunity principle. One might therefore be reasonably unfussy about whether one takes the view that Rawls endorses four principles of justice, or five. The structure of the view should now be clear enough, under either notational variant.

IV.

Further Distinctions: Perfect vs. Fair Equality of Opportunity (vs. More Radical Alternatives)

I have outlined the distinction between comprehensive and domain-specific principles of equality of opportunity, and between single-principle and multiple-principle distributive views. I now want to make a distinction between three different *strengths* of principles of equality of opportunity. Here are three grades of equal opportunity principle, in ascending order of stringency.

Fair Equality of Opportunity

When Rawls talks about FVPL, he talks about citizens being “*sufficiently equal* in the sense of having a *fair opportunity* (to hold public office and influence the outcome of elections)”. Note that he does not say an “absolutely equal opportunity”, or anything that would communicate the same idea. What Rawls seems to have in mind here, just as also in the case of FEO, is that it would be impractical, and overly demanding, to institute a perfectly meritocratic scheme, not least because instituting such a scheme would require the abolition of the family. So it is certainly not true to say that either FEO or FVPL involve the neutralization of all morally arbitrary differences in social background, for such complete neutralization of differences in social background would require a level of intervention in the social order that Rawls finds unconscionable. Instead, the goal specified by these principles is what one might describe as a *sufficient* degree of neutralization of social contingencies, whereby concerted action is taken to limit the effects of social background on job opportunities (in either the economic or the political domains), but this kind of action is not pursued to the degree that it would involve restructuring social life, or the abolition of central social institutions, such as the family itself.

This *strength* of equality of job opportunity can be compared to a stronger version, which I will label *perfect equality of opportunity* (PEO) or *perfect meritocracy*.

Perfect Equality of Opportunity

Under PEO or perfect meritocracy, the effects of morally arbitrary social contingencies are completely neutralized with respect to the distribution of jobs in the economic and political spheres. In a perfect meritocracy, the determination of who gets attractive occupational positions would be determined entirely by facts about individuals' (natural) talents and their (natural) willingness to use them, such that there would be no meaningful correlations between an individual's social class (insofar as such a thing could be said to exist in a meritocratic world) and that individual's occupational destinations. As we shall see, some things that Rawls says seem to indicate that he is committed to something along the lines of this perfectly meritocratic view with respect to the distribution of jobs, but any such impression is in the end a misleading one, as Rawls's considered view draws back from endorsing this more radical flavour of meritocratic principle.

Radical or Cohenite Equality of Opportunity

Much more radical than either *fair* or *perfect* equality of opportunity is a third kind of view, with the same kind of strength as the more radical vision of equality of opportunity embodied in the *standard luck egalitarian view*. As this kind of view perhaps finds its clearest expression in G. A. Cohen's "internal critique" of Ronald Dworkin's theory of equality, I shall call it Cohenite equality of opportunity (CEO), or alternatively anti-meritocratic equality of opportunity (Cohen himself calls this kind of view "socialist equality of opportunity" in book *Why Not Socialism?*). On such a view, when applied specifically to the domain of jobs and their distribution, the allocation of occupational positions in the economic and political realms should be done in such a way that it is possible to neutralize the effects of both natural *and* social contingencies on outcomes. Such a view rejects meritocratic conceptions of equality of opportunity as providing no more justifiable a basis for the distribution of jobs than would a view that did no more than ensuring what Rawls called "careers open to talents", but which still allowed social contingencies to play out through the occupational structure of the economy and of the political domain.

On the anti-meritocratic CEO view, when applied specifically to occupational opportunities, the only factors which should influence the distribution of jobs are factors for which individuals can properly be held responsible, presumably therefore being limited to features of individuals' psychological and motivational make-up which are themselves under the active control of the will in the appropriate way (where "in the appropriate way" is, as on a view like Cohen's, to be specified by some underlying account of the nature of responsible agency). Now, although I describe this view as "Cohenite", I am not (or at least, not yet) looking to attribute such a domain-specific CEO view to Cohen himself. Cohen's inspiration comes in our understanding of the nature and strength of the anti-meritocratic conception of equality of opportunity, not in the specific elaboration of such a view with regard to the domain of jobs. Now, it is difficult, it seems to me, to arrive at a clear view of what a domain-specific anti-meritocratic account of equality of occupational opportunity would look like at the level of policy or of "rules of regulation" (in Cohen's sense). This is a matter to which we will return. In the meantime, though, I take it that at least that the

structure and aims of such a radical view of equality of opportunity for jobs is now perspicuous.

It should not be forgotten, of course, that it was Rawls himself who made the famous claim that both natural and social contingencies are “equally arbitrary from a moral point of view”, and, given this, one might naturally think that Rawls ought to endorse the strong and radical anti-meritocratic conception of equality of opportunity with regard to the distribution of jobs. As we’ve seen, though, he does not do so. Indeed, not only does Rawls not endorse the CEO/anti-meritocracy view, but he is not even committed to what might seem like the moderate PEO/perfect meritocracy view, rejecting both in favour of what might seem like the rather weak fare of (merely) fair equality of opportunity. A number of luck egalitarian theorists have taken the view that this is in some way proof that Rawls failed to have the courage of his own convictions. Alternatively, one might think it a simple argumentative misstep. As we shall see, though, it is neither. But, when assessing the plausibility of a view that stops at (mere) FEO, rejecting both PEO and CEO, it is important to bear in mind the rather attractive and straightforward line of argument that would claim that anybody who would begin down the road of moving from “careers open to talents” to FEO would then fail in following their line of thought further, towards these stronger and apparently more radical options, only by virtue of being either insufficiently brave or insufficiently consistent, or both.

V.

Three Questions about Equality of (Job) Opportunity:

Special or Derivative? Radical, Perfect or Fair? (Or Should It Just Be Rejected?)

With this range of distinctions now in view, it would be useful, before moving onto the detailed substance of assessing which view of equality of (occupational) opportunity is the right one, to take an overview of the available conceptual space. Here is a diagram representing a number of the available views. Our task is then to determine which of these views is the correct one.

Nature of Justification → Strength of (Occupational) Opportunity Principle ↓	Special Significance [Column A]	Derivative or Instrumental Significance [Column B]
Fair (FEO)	[1] a. Rawls’s view b. “The <i>Non-Intrinsic Egalitarian</i> Opportunity Principle”	[4] Luck Egalitarian Fair Meritocracy
Perfect (PEO)	[2] Anti-family Rawlsianism	[5] Luck Egalitarian Perfect Meritocracy
Radical (CEO)	[3] Seriatim Luck Egalitarianism	[6] “Double Luck Egalitarianism”

In each case, the labels in each box do not so much name the view associated with *any* theory occupying that position, as illustrate one such representative view. The view that I shall be defending is a version of view [1]. Note that what the diagram leaves out is one very important and common view, associated with the work of (among others) Richard Arneson. On view [7], which does not fit within this grid, there is simply held to be no general reason to privilege a principle of equal occupational opportunity.

Navigating in this conceptual space involves settling three questions. Firstly, do we have general reasons to endorse some form of equality of opportunity, or should we instead adopt view [7]? Secondly, if we have reason to adopt some commitment to equality of opportunity, is our support for that principle an instrumental one, dependent for example on the empirical consequences of a policy of equality of opportunity, as they relate to some other set of distributive aims, or is our support for that principle based on an independent justification of its special significance (i.e. should we be in Column [A] or Column [B])? Thirdly, insofar as we endorse a goal of equality of occupational opportunity, what flavour or strength should that commitment take?

VI.

The Negative Argument – Luck Egalitarianism, Burke and Natural Aristocracy: Considering the Kingdom of Burkeana

Imagine that you hold the *standard luck egalitarian* view. This, we shall recall, is a single-principle distributive view, holding that distributive justice is achieved when requirement (1) is met: i.e. where everyone has equal opportunity for X, where X is some general distributive currency. Given your commitment to the general claim that egalitarianism involves the view that we need to equalize opportunity for X, where the specification of X is a central theoretical problem yet to be faced, you believe in what we might call an *egalitarian sphere of compensation*. Equality of opportunity for X can be realised, even when individuals are unequal among the various sources of opportunity for X, as long as their aggregate levels of opportunity for X are equal. So, if you agree with Arneson (1989) that distributive justice requires equal opportunity for welfare, you would be untroubled by a situation where, for example, person A has less opportunity to derive welfare from her job than person B, but more opportunity to derive welfare from her hobbies, so long as their aggregate levels of opportunity for welfare were as high and equal as possible.

I take it that many who endorse a single principle opportunity of equality view, with an egalitarian sphere of compensation, would tend to see themselves as also endorsing some kind of equality of occupational opportunity, either of a ‘radical’ variety, or perhaps of a weaker sort. But insofar as their overarching distributive commitment is to a single principle view, with an egalitarian sphere of compensation, their tendency to endorse some kind of occupational equal opportunity principle is clearly one that can only be derivative of that main distributive commitment, and must depend on a number of empirical assumptions about the likely effects of such a principle or policy. Of course others, such as Arneson, who endorse a single-principle non-domain-specific equality of opportunity view, would simply deny that equality of occupational

opportunity had even instrumental or derivative significance. Both kinds of views, though, whether from ‘Column [B]’ or Arneson’s view [7] are going to face a certain degree of discomfort under certain kinds of empirical conditions. This can be made clearer with the aid of an example.

The Kingdom of Burkeana

Imagine a Burkean aristocracy, where a hereditary ruling caste with a well-developed sense of *noblesse oblige* controlled the basic features of both economic and political life. This caste are self-perpetuating and admirably efficient. Far though this may put us from more local possible worlds, let us stipulate further that, in the Kingdom of Burkeana, this ruling caste (a) do a much better job of governing wisely and of growing the economy than would an alternative, more classically meritocratic system, and (b) are also committed *standard luck egalitarians*, whose policies are directed towards guaranteeing, for the subjects of Burkeana, a high and equal level of opportunity for welfare.

Let us further specify that the Burkeanans aristocratic mode of organization is perpetuated by social convention, generated by informal kinds of class-based social power, and by the efficiency of the ruling caste in perpetuating themselves by developing their next generation. Burkeana’s aristocratic features are not sustained by legislation or formal discrimination and, indeed, the Kingdom has the full set of formal legal protections associated with Rawls’s principle of Equal Basic Liberties, as well as those associated with the rejection of formal discrimination through a commitment to ensuring the maintenance of “Careers Open to Talents”.

Our ruling caste derive tremendous personal welfare from their exercise of the duties of their station and so, in accordance with luck egalitarian principles, make sure to dampen down their opportunities for welfare in other domains. Instead, they live relatively ascetic lifestyles, devoting themselves to the arts of economic and political governance, and the cultivation of the next generation of aristocrats. The lower caste members of society, by contrast, although they have extremely restricted real opportunities to come to exercise meaningful “powers and prerogatives of office and positions of responsibility”, in either the economic or the political spheres, live in other respects extremely well. Burkeana is prosperous, and has made sure that citizens have plentiful leisure time, and plentiful resources to deploy in their leisure pursuits. The booming economy funds generous and efficient public services, including far-sighted and effective programmes devoted to the public health of the Burkeanan subjects. The Burkeanans live long, fulfilling and happy lives.

Is there anything wrong with this picture? (I’m not particularly interested in its empirical plausibility, but in what it tells us about the structure of our values and political commitments.) It seems to me that we should have a very strong sense that the subjects of Burkeana are treated unjustly. They are treated unjustly, I want to suggest, even if they are very happy (and have high and equal *opportunity for welfare*), and even if, were they were asked whether they would trade a reduction in their aggregate levels of opportunity for welfare, in exchange for greater job opportunities, they would refuse the offer. In what follows, I want to make good on that claim of injustice.

It is noteworthy that, when Rawls rejected a distributive view that gave fuller rein to the difference principle, unencumbered by the lexical priority of equality of fair equality of opportunity, the view he took himself to be arguing against was one he called “Natural Aristocracy”. The name is a curious misnomer, of course, as there is nothing *natural* about a “Natural Aristocracy”, in which a self-perpetuating social elite call the shots, and in which the abilities of their every next generation are a product of concerted social cultivation rather than natural good fortune. (Indeed, “Natural Aristocracy” sounds like it should be the name given to the kind of unrestricted meritocracy that Michael Young satirized in his 1958 book *The Rise of the Meritocracy*, rather than the name applied to a society such as Burkeana.) But what Rawls was worried about, in rejecting a Natural Aristocracy, and thereby in rejecting a (more or less) single principle distributive view (constricted only by the formal equal opportunities associated with having “careers open to talents”), which gave a centrality to the operation of the difference principle, was precisely the importance of avoiding a theory that would have no good reason not to reject the choice between Burkeana and a poorer but more open society.

In contrast, single-principle egalitarians have no good reason to reject Burkeana. Some, like Arneson, may very happily bite bullet here, “out smart”⁹ their opponents, and wonder what the force of this kind of objection is really held to be. Others may be more troubled, and instead want to point to the improbability of my empirical assumptions. Yet I think this is a case where an outlandish thought experiment can nevertheless illustrate an important normative truth. That truth is that, as citizens of a just society, our interest in economic and political opportunities is not exhausted by their potential contribution to our well-being. Rather, it is a *sui generis* interest in the exercise and development of particular kinds of significant capacities, as they play out in the shared institutions of our common life, and as such it is irreducible to any more general individual concern. Hence the need for a separate principle, subject to a form of special justification. And so, if we are moved by this thought, then we would see the views of column [B] on our diagram are just as implausible as the more sceptical view [7].

VII.

Seriatim Luck Egalitarianism – A Response to the Burkeana Objection?

One may take the view that, while Burkeana poses a troublesome objection to the plausibility of a single-principle egalitarian view with an undifferentiated sphere of distributive compensation, and thereby counts against the plausibility of *standard luck egalitarianism*, it may nevertheless not be a good objection to luck egalitarianism in its most plausible form. This could be a defensible view if the best way of capturing the central luck egalitarian intuition in a view of distributive justice was not to endorse (single-principle) *standard luck egalitarianism*, but instead a form of multiple-principle luck egalitarianism. Now, one might here say that, if so, this is by the by, as it would do nothing to undermine the claim that we have (at least) strong *prima facie* reason to reject views [4],[5],[6] and [7], and it would, indeed, be right to say that this sort of admission would not challenge the rejection of these views. Nevertheless, even within the views of Column [A], we face the task of deciding

⁹ In the J J C Smart sense.

between the various available strengths of a job opportunity principle – fair, “perfect” or radical/Cohenite. A multiple-principle version of luck egalitarianism could provide an interesting position by combining a sensitivity to the special significance of job opportunities with a commitment to a more radical form of equality of opportunity than the Rawlsian view would countenance, thereby occupying place [3] in our grid.

Let us therefore define a view that I shall call *Seriatim Luck Egalitarianism*.¹⁰ This view denies that there is a general sphere of egalitarian compensation, in favour of adopting a view that distinguishes between a number of distinct, ‘non-fungible’ distributive goods, each conceived as having a distinctive, non-derivate significance. According to seriatim luck egalitarianism:

- (3) The central demand of distributive justice is to bring about a luck-neutralizing but choice-sensitive egalitarian distribution with regard, serially, to a range of particular, specific distributive goods (a, b, c, etc., ...) (where a, b, c, etc., are a set of specific distributive goods, the size and constitution of which set is to be determined)

Now, of course not all versions of seriatim luck egalitarianism would give a special place to equality of opportunity in the standard sense, that is, to equality of job opportunities. But we can imagine a version of the view that responded to the troubling case of Burkeana by saying that, properly understood, job opportunities should be among the specific distributive goods (a, b, c, etc.) the distribution of which should be governed by a scope-limited luck egalitarian principle. So, on a version of seriatim luck egalitarianism which gave this sort of special place to the distribution of job opportunities, the operative claim would be that it is a specific demand of distributive justice that jobs be allocated in a way that is insensitive *both* to social contingencies of class background and family environment, but also to the natural contingencies embodied in unchosen differences in talent and ability.

Such a view presents an arresting and distinctive possible understanding of what justice demands. But consideration of the implications of such a view quickly reveal its peculiarities, and may plausibly have the effect of reducing the plausibility of position [3] in our grid, and of arresting and deflecting the confident march that one may be tempted to undertake from a starting point in meritocracy towards a more radical destination. This will become clear if we consider what a society would look like that realised a luck egalitarian form of equality of opportunity (i.e. CEO) with specific regard to the distribution of job opportunities. Consider, then, *Ceoniana*:

¹⁰ This view was suggested to me by reading a hitherto unpublished paper of Kasper Lippert-Rasmussen’s, on luck egalitarianism and social egalitarianism. Among the interesting features of Lippert-Rasmussen’s discussion, he shows that there are aspects of G. A. Cohen’s view which incline him more towards the camp of (a form of) seriatim luck egalitarianism, rather than seeing him, as is more usual, as a proponent of the standard luck egalitarian view. It is unfortunate in the extreme that Cohen’s positive view of distributive justice never really reached a settled, canonical form, despite the fascinating and suggestive work that was done as part of that view’s development. (See Cohen, “On the Currency of Egalitarian Justice”, *Ethics*, (1989), and “Expensive Taste Rides Again” in Justine Burley, ed., *Dworkin and His Critics*.)

Ceoniana

Ceoniana is a society in which the principles of public political morality accord with the demands of a form of *seriatim* luck egalitarianism, in which job opportunities (in both the economic and political spheres) are distributed in a luck egalitarian fashion. The Ceonians hold the view that neither natural nor social contingencies should have a (significant) effect on the distribution of “powers and prerogatives of offices and positions of power and responsibility”, which should, instead, be distributed in a way that is fully egalitarian, other than when a departure from equality reflects an instance of responsible choice. It is not easy to say quite what this commitment of principle would or could amount to at the level of policy.

The most obvious thing that could be said is that, under any plausible empirical assumptions, the organization of work and politics in such a society would be strikingly inefficient. Any systematic correlation between the level of power or responsibility that a particular citizen attained and her talent or ability to make skilful use of that power or responsibility would offend against the principles to which that society is committed. The one exception to this would be for cases where an individual’s level of talent and ability was the outcome of responsible choice, rather than constitutive luck. But it is hard to see how ability could be a pure matter of choice, unconnected to the morally unchosen distribution of native talent. One is reminded here of Rawls’s well-known claim that “The effort a person is willing to make is influenced by his natural abilities and skills and the alternatives open to him. The better endowed are more likely, other things equal, to strive conscientiously, and there seems to be no way to discount for their greater good fortune.”¹¹ And even if choice or “conscientious effort” could somehow be extricated from their unchosen connections and preconditions, one should bear in mind that, in Ceoniana, the distribution of job opportunities could be correlated *only* with the pristinely chosen element within any more complicated volitional state, and not with the adulterated amalgam of the chosen and unchosen. In any case, there would be no reason to expect, in Ceoniana, that those who would get particular jobs are those who would be good at doing those jobs. What might this look like, in policy terms? Perhaps the most plausible instantiation of such a principle would be in a job lottery, in which a demonstrably fair procedure would be used to allocate jobs in a way unsullied by connection to unchosen natural or social endowments. To reflect the residual place for *choice* in such a system, citizens might perhaps be asked to indicate which part of the economy they would like to work within, with specific positions within each sector then being allocated by a fair, random procedure.

Given that it is difficult to make sense of what the enactment of a luck egalitarian principle in the space of job opportunities would really involve, I want to make only three rather general claims. Firstly, if one adopts position [3], one is endorsing a view that is likely to be strikingly, heroically, inefficient; and so a society that looked to enact such a principle in a direct way would, for that reason, be a notably unattractive prospect. Secondly, and relatedly, it seems difficult to make sense of what would be normatively attractive about pursuing such a goal. Thirdly, even if we can, in a rather abstract sense, give some shape to our intuition that economic and political roles should not be occupied by those who are likely to do them well, given the unfairness

¹¹ Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), p. 312 (Rev. Edn., p. 274)

of the distribution of such abilities, one has to realise that the normative costs of such a policy are sky high. It is not only that such a society would be inefficient, but also that it would be very likely systematically to frustrate chances for individuals to do a good job with skill, thereby massively reducing opportunities for the competent exercise of valuable social functions, and (in a way that I will explore later) stunting the agency of many of its citizens.

Now, an advocate of *seriatim* luck egalitarianism may say that this line of attack is unfair, insofar as I am able to paint this rather bleak picture of Ceoniana only by eliding the distinction between a normative principle of justice and a concretely realizable policy or “rule of regulation”. According to this counter-objection, it is consistent with endorsing view [3] that one would still want to sign up for a meritocratic social system, precisely because other values – such as efficiency – are also valuable, and because, therefore, one needs to formulate rules of regulation by taking into account a plurality of relevant values, rather than seeing such rules as more-or-less direct enactments of specific principles. Oddly, this sort of move seems easier to motivate when one considers the *standard luck egalitarian* view, which, after all, has the general form of providing a highest-order principle of distributive justice, which might then naturally be tempered by contact with other desirable values that might be realised within a socioeconomic system. It seems a stranger move with regard to *seriatim* luck egalitarianism, which is a view that suggests a range of more specific recommendations, which seem in form much closer to “rules of regulation” than to highly abstract normative principles. But setting this point aside, it is worth returning to the background context here, as regards the question of whether we might really have good reason to give a principle of equality of opportunity some special standing in our account of what a just basic structure should actually do. What is interesting about the version of *seriatim* luck egalitarianism under consideration is that it looked like a view that really would defend a distinctive and radical version of an equality of opportunity principle. That initial impression is likely to be illusory if the view supports no more than a rather weak *prima facie* claim about the value of eliminating the role of morally arbitrary contingencies in the distribution of job opportunities, which is then to be comprehensively submerged to the point of *de facto* eradication in an appeal to general normative concerns of efficiency and (comprehensive, not job opportunity-specific) equality.

My suggestion is that this puts the advocate of *seriatim* luck egalitarianism in an unwelcome dilemma. Either the view is distinctive in giving a special place to a particularly radical principle for distributing job opportunities, in which case it is strikingly unattractive, as I hope the example of Ceoniana illustrates. Or, otherwise, despite the surface feature of distinctiveness, it collapses back into more familiar forms of luck egalitarianism, in which case it can say something distinctive about the distribution of opportunities only in a rather contingent and derivative way, and is hostage to fortune with regard to the kind of empirical possibilities that obtain in Burkeana. Put simply, if *seriatim* is distinctive enough to fit into position [3], then we have reason to reject it; if it collapses back into a more familiar kind of view, that advocates a position from column [B] or box [7], then we have already rejected it. My provisional claim, then, is that we have only two plausible views left standing: those associated with position [1] and position [2].

VIII.

Against Perfect Meritocratic Equality of Opportunity: Or, Why Rawls Does Not Believe in the *Lexical* Priority of Equality of Opportunity (and Neither Should Anyone Else)

Let us now consider position [2], the view that I have called “anti-family Rawlsianism”. On this view justice is understood to demand the institution of a *perfect* meritocracy with regard to the distribution of job opportunities. In other words, this view demands PEO rather than mere FEO. We have already considered why we might have good reason to halt the march from imperfect meritocracy all the way to luck egalitarian equality of opportunity, but it is also worthwhile to consider the question of whether we might also have reason (as Rawls thinks we do) to halt that march right at its first stage, in resisting the move even to a more perfect version of meritocracy.

I am going to suggest that, in fact, the forward march should be halted at its origin: Rawls is right that we do not have good reason to move beyond the imperfect meritocracy of FEO with regard to the distribution of job opportunities. But in understanding why this is the case, we see that we also have good reason to reject the familiar Rawlsian idea that a principle governing equality of opportunity should have *lexical* priority over a broader distributive principle. The best explanation for why we need not go beyond FEO leads us directly to an argument against lexical priority. What this means is that, if one wants to defend the view that best lays out our reasons for giving a special, non-derivate significance to a principle of equality of opportunity, we have to understand that a commitment to the *special* significance of equality of opportunity should not be a commitment to its absolute, lexical priority. This will be important when we come to a careful excavation of the positive case for view [1], i.e. for the special significance of *imperfect* meritocracy within a plausible view of what we owe one another as regards the organization of economic and political life. There is a considerable and important conceptual gap between taking equality of opportunity seriously as a distinct, and even presumptively dominant, concern within an account of what justice demands and the more extreme view that gives that concern an absolute lexical priority. One of the more problematic aspects of Rawls’s view is that his ready reliance on the slogan of *lexical priority* has come to obscure the more plausible views lurking in the near neighbourhood of his official position.

In order to make good on these claims, consider the passage where Rawls infamously discusses the possible prospect of advocating the abolition of the (traditional) family. In a perfect meritocracy – that is, a society enacting a principle of PEO with regard to job opportunities – citizens’ prospects with regard to ascending to positions of political or economic power and responsibility would have to be *perfectly* independent of features of their social and familial background. If we thought that we had decisive reason to guarantee PEO, while stopping short of the sort of ‘job lottery’ that might be associated with CEO, then seemingly the only plausible way of doing so would be to ensure that all individuals faced more-or-less indistinguishable social conditions in which their native talents and abilities came to be moulded and shaped into fully realised features of their adult selves, and in which the differential influences of the small material and social benefits and burdens associated with membership in differing families were unable to work their influence.

Abolishing broad variability in the circumstances of upbringing would seem to require the abolition of traditional family structures, given that differences between different families have such an overwhelming role, by a diversity of social, economic and psychological mechanisms, in the development of the individuals who grow up within those families. Hence, Rawls realises that he has to give serious consideration to the vertiginous possibility that his account of justice may demand the obliteration of traditional forms of family structure.¹² Consider, then, how – despite the fact that he does see the need to take this possible consequence of his view seriously – Rawls comes to reject this possibility:

“The consistent application of the principle of fair opportunity requires us to view persons independently from the influences of their social position. But how far should this tendency be carried? It seems that even when fair equality of opportunity (as it has been defined) is satisfied, the family will lead to unequal chances between individuals. Is the family to be abolished then? Taken by itself and given a certain primacy, the idea of equal opportunity inclines in this direction. [i.e. inclines towards PEO rather than FEO] But within the context of the theory of justice as a whole, there is much less urgency to take this course. The acknowledgement of the difference principle redefines the grounds for social inequalities as conceived in the system of liberal equality; and when the principles of fraternity and redress are allowed their appropriate weight, the natural distribution of assets and the contingencies of social circumstances can more easily be accepted. We are more ready to dwell upon our good fortune now that these differences are made to work to our advantage, rather than to be downcast by how much better off we might have been had we had an equal chance along with others if only all social barriers had been removed. The conception of justice, should it be truly effective and publicly recognized as such, seems more likely than its rivals to transform our perspective on the social world and to reconcile us to dispositions of the natural order and the conditions of human life.”¹³

There are, no doubt, some peculiarities in this very rich passage. For one, it is rather strange to talk about the difference between PEO and FEO in terms of the “removal of social barriers”, for what a state would have to do in order to move from a basic structure that satisfied FEO to one that satisfied PEO is not so much to *remove* any barriers that individuals might face, as to raise up a number of significant social barriers to the transmission of relative advantage in the labour market within social groups, and especially within family structures. Moreover, viewed from the perspective of the citizens rather than the perspective of the state, so to speak, it is not that unequal opportunities for the less advantaged are generated by actions that constitute the erection of social barriers (as when some rival group are actively prevented from doing something), as by a readiness to act in such ways that directly benefit those near and dear to us (as when more-advantaged parents give their

¹² What might take the place of the family? For an insightful discussion of possible alternatives, including mandatory state-run orphanages, see Véronique Munoz-Dardé, “Is the Family to be Abolished Then?”, *Proceedings of the Aristotelian Society*.

¹³ Rawls, *A Theory of Justice*, revised ed., 448.

children as engaging and stimulating a childhood as they can manage to generate¹⁴), to the indirect positional disadvantage of others.

Setting these peculiarities aside, what is most interesting here is the structure of Rawls's argument for the adequacy of FEO. The argument is that, as long as the difference principle is enacted *alongside* the principle of FEO, the demands of justice as fairness are satisfied, because those individuals – presumably those from unpropitious family and social background, who remain disadvantaged by those backgrounds even under the conditions of muscular but imperfect meritocracy that obtain under FEO – will be satisfied to look past their residual complaint about the failure to enact PEO if, and only if, the enactment of the difference principle means that these recalcitrant social inequalities end up being to the benefit of everyone in society, rather than solely to the benefit of those who were already socially advantaged. Note the consequences of this view for our assessment of a society in which FEO was enacted without the simultaneous realization of the difference principle. Under such a basic structure, those socially disadvantaged with regard to their job opportunities by the imperfections of FEO as a meritocratic principle would have a strong and residual complaint that they had not been treated fairly and as equals. Under the social system that Rawls calls “Liberal Equality”, and which we might call an imperfect meritocracy, the failure of FEO to do enough to correct for unfair social inequalities as they affected job opportunities would generate an unaddressed complaint, leading to a social system that was clearly and bluntly unjustifiable to those who lost out by virtue of the imperfections of FEO.¹⁵ Put simply, only when the difference principle is satisfied, and when socially disadvantaged individuals benefit from its operation, does FEO begin to look like a justifiable principle for the distribution of job opportunities. Without the real work done by the difference principle in forming part of the composite ‘offer’ that can be made to those who lose out from the effects of social stratification, FEO is an excessively weak and implausible principle of justice.

The problem for lexical priority then comes when we begin to take seriously the fact that, at the level of enactment in public policy, FEO and the difference principle are strongly rivalrous. In large part, this is because we would go seriously wrong in thinking about FEO if we saw it as an easily satisfiable requirement. Rather than identifying a modest aim that a just basic structure could easily meet, FEO is actually an extremely demanding standard, notwithstanding the fact that it represents merely an ‘imperfect’ form of meritocracy. Holding fixed the assumption that the state is not going to get into the business of disrupting family life, or fundamentally reconfiguring the social contexts in which children are to be raised, there is still an almost endless amount that can be done to move a society ever closer to realising the goal of enacting fair equality of opportunity for jobs. Looking specifically at FEO and child-rearing, the state might throw more and more resources into SureStart and HeadStart programs, parenting classes, early-years interventions, and compensatory social

¹⁴ See, for example, Annette Lareau, *Unequal Childhoods*.

¹⁵ This is not to say that a more stringent meritocracy or form of Liberal Equality, which substituted PEO for FEO, but unaccompanied by a principle of economic reciprocity such as the difference principle, would be adequate. But the conclusive complaints against such a view would be of a different kind, with their source in concerns of fraternity and reciprocity, even though one could at the same time say that such a society would at least do a good job of being true to what we might call the internal logic of the idea of meritocracy. Liberal Equality with FEO fails even in that latter regard.

programmes of all kinds. If FEO were really to be given lexical priority over any other distributive principle by a legislature looking to create a just society, then an ever greater share of social resources would be directed into schemes such as these. Given the fact that the labour-intensive programmes that might help us to approach the goal of full FEO are going to carry substantial costs, it is inevitable that, the closer one gets to FEO, the less social resources there will be for other purposes and, therefore, the worse-off the most disadvantaged are likely to be in terms of their overall index of social primary goods. A society that gives *lexical* priority to FEO can only do so at significant cost to its ability to provide non-FEO-related public services, and at a significant cost to the amount of the overall social product that will be left over for private consumption.

This problem becomes all the more troubling when we recall that FEO is a justifiable principle only as a constituent of a broader bargain, in which the least-advantaged can be offered the gains associated with what we might describe as a difference principle that does *real work*. To put things another way, FEO is a worthwhile goal only if, elsewhere in the social system, the less advantaged are compensated for FEO's structural shortcomings by the real and substantive benefits guaranteed to them by the difference principle. But, if FEO is given lexical priority, then it is highly likely that we would in so doing remove the capacity of the state to deliver those real and substantive benefits (whether in terms of the provision of public services or fungible resources). As we approach the asymptote of full FEO, the difference principle does less and less work, and the benefits that it can deliver to the less advantaged become ever less impressive. In short, an understanding of the justification for FEO within a broader, multiple-principle account of distributive justice precludes from the outset the possibility of FEO having the sort of lexical priority that Rawls canonically affords it.

Rawls himself showed a cautious and non-committal awareness of (something like) this problem, in the final statement of his view in *Justice as Fairness: A Restatement*. The view around which Rawls carefully edged there involved the jettisoning of the idea that giving a special priority to equality of opportunity need involve giving it absolute lexical priority:

“Some think that the lexical priority of fair equality of opportunity over the difference principle is too strong, and that either a weaker priority or a weaker form of the principle would be better, and indeed more in accord with fundamental ideas of justice as fairness itself. At present I do not know what is best here and simply register my uncertainty. How to specify and weight the opportunity principle is a matter of great difficulty and some such alternative may well be better.”¹⁶

One can only speculate here as to why Rawls was not simply prepared to bite the (small, easily digestible) bullet of admitting that there was no good reason to maintain the commitment to the lexical priority of FEO over the difference principle. Perhaps the cost, in terms of losing the clear, hierarchical architectonic embodied in the canonical statement of the principles of justice seemed too high. But, while a more moderate position may seem, in a sense, more boringly complicated, when held up

¹⁶ Rawls, *JF*, 163.

against the clear lines of a neater view, consistency with the underlying commitments of Rawls's view unavoidably urges the adoption of the messier view.

If we have to reject the lexical priority of FEO, where does this leave us? Well, it does not mean that we have to give up on the belief that FEO retains a special significance as a non-derivative, domain-specific, non-instrumental principle of justice. The correct view, as I shall argue in the next section, is that, although lexical priority simply has to go, FEO retains a kind of general "presumptive priority" over the difference principle (or any other broader principle of distributive justice).

IX.

The Positive Argument for the Last View Standing: A Non-Intrinsic Egalitarian Account of the Special Significance of Equality of Opportunity

In the foregoing discussion, I hope to have marshalled a range of considerations that count against the adoption of any of the views [2] through to [7], as presented in the diagram in Section V. But none of this is yet decisive. What we now need is a positive argument for giving general presumptive priority to FEO (or some other closely variant view that might fit in box [1]). Rawls himself does not make the positive case as clearly or explicitly as he might have done. What follows is offered up as a sympathetic reconstruction of a Rawlsian position; hopefully it will serve both as a plausible reconstruction of what Rawls might have said were he to have gone into more detail in his defence of FEO, and also as an independently plausible account of what we need to find space for an independent and non-derivate principle of imperfect meritocracy within our most plausible view of what justice demands.

The positive argument for giving general presumptive priority to FEO turns on three closely-related ideas. The first is that, in thinking about the distribution of job opportunities, we need to be attuned to the highest order interests of citizens; that is, their interests in exercising their two "moral powers" (in Rawls's terminology): their sense of justice and their ability to frame, revise and pursue a conception of the good. The second idea is that the principles of justice together should give expression to an ideal of how citizens might live together as equals, where the relevant idea of equality should be understood as encompassing constitutive values of non-domination, self-respect and the protection of the status and standing on equal terms of every individual.¹⁷ Thirdly, our thinking about what we owe each other in the joint imposition of political and economic institutions upon one another must have at its centre an idea of reciprocity, such that these institutions are understood as being to our mutual and collective advantage.

Taken together, FVPL and FEO provide a satisfactory context for the "adequate development and full and informed exercise of the two moral powers" (JF, 112). This meritocratic approach to equality of opportunity responds to citizens' interests in the exercise and development of their capacities, as they are deployed in the shared political and economic institutions of our common life. Related to the idea of citizens

¹⁷ See Rawls, JF, section 39; and my "What Should Egalitarians Believe?", *Philosophy & Public Affairs* (2008).

seen as equal, FEO and FVPL help to ensure the secure standing and self-respect of citizens, through allowing them to develop and deploy their talents as they are able, and within the bounds of reciprocity. These principles prevent the domination of one part of society by another over time in terms of setting the terms of productive activities or the political direction of their shared society (through ensuring new entrants to positions of economic and political power); and they prevent disparities of status between those equally endowed with talent.

Treating equality of opportunity as of merely derivative significance fails to treat citizens as equals in the right way, as it fails to protect the ground of their status and self-respect as individuals who will have an interest in developing and deploying their sense of justice and conception of the good through engagement in political activity and productive economic activity, and to protect them from ongoing domination by others in those spheres (contrast this with the way in which citizens are systematically dominated in Burkeana, no matter how well-off they may be in such a society). By contrast, a view which marches past meritocracy towards CEO (as in the political morality of Ceoniana) fails to connect in the right way with citizens' interests in deploying their actually existing moral powers; it sacrifices their standing as equals who each have a real and present capacity for agency, through an over-eagerness to push an excessively rigoristic idea of fairness beyond the bounds of reasonable reciprocity.

This defence of a central role for imperfect meritocracy (i.e. FEO or something like it) among the principles of distributive justice expresses, in its essence, a social egalitarian (or, in my terms, a *non-intrinsic egalitarian*¹⁸) ideal. If we think that among the valuable features of justifiable social relationships are that these relationships preserve each individual's standing and status as an equal citizen, protect each from domination by others, and maintain the conditions needed for fraternal relations of mutual reciprocity, then it is not enough only to consider each individual citizen's interests as consumer of goods, or as a site of greater or lesser welfare. We need instead to take seriously the challenge of discovering what we can guarantee for everyone in terms of each citizen's opportunities for genuine agency, in both political and economic life. In order to connect in the right way with each individual's sense of his or her own economic and political agency, these opportunities have to be indexed to what those individuals really can do well, rather than seeming just like another kind of opportunity for promoting their own well-being. What we are interested in is opportunities for the exercise of capacities for developed forms of agency, as such, not merely occasions for enjoyment. Furthermore, the provision of job opportunities to those who could not make particularly good use of them, as in Ceoniana with its job lotteries, would make the provision of those opportunities feel like a rather perplexing kind of gift, cast free from the constraints of reciprocity, rather than making them seem like what they should be, which is opportunities to develop and exercise one's capacities within the larger shared project of creating an economic and political system that is both justifiable to everyone within it, and is also to the benefit of everyone within it.

¹⁸ Martin O'Neill, "What Should Egalitarians Believe?"

X.

Subjects of Justification: Identity, Alienation and Moral Arbitrariness

In some ways, it may seem surprising that a commitment to such a seemingly cautious and compromised position as imperfect meritocracy in the distribution of jobs is something that we can justify to each other as an essential aspect of social justice. But this surprise may evaporate if we delve a bit deeper into the question of how such a view might be justified and to whom the justification must be offered. The endorsement of an important role for imperfect meritocracy (i.e. for FEO and FVPL) fits with a rather complex, if not to say unavoidably strained, view on the (merely partial) moral arbitrariness of talents, and of the relationship between individual citizens *qua* embodied, empirical agents and individual citizens *qua* moral subjects. This tense and ambivalent view about the subjects of justification can be seen as implicitly present in Rawls's discussion of these issues, and provides a powerful way of thinking about what we can justify to one another. Put simply, the view is that, in one sense, all talents and abilities are fully morally arbitrary; but, simultaneously, when viewed in another way or from a different standpoint, our talents and abilities are partially constitutive of who we are as real, empirical, embodied agents, and thereby their exercise and development is importantly connected with our sense of individual agency, our status, our self-understanding and our self-respect. We are properly pulled in opposite directions by importantly and opposing moral commitments that can and should simultaneously have a grip on us: in one way, contingent facts about our empirical identities are irrelevant to justice, while in another way they matter a great deal.

The central claim here, then, is this: you and me, conceived of as citizens with two moral powers, are *subjects of justification* who can and should have a morally complex, if not to say deeply ambivalent relationship to our own talents and abilities. Individual citizens who think that the socioeconomic system should be organized so as to deliver all available rewards to those best able to exercise their native talents – as in Michael Young's skeweringly satirical treatment of the citizens of his imagined meritocracy – commit a kind of unattractively short-sighted (if not to say almost monstrously crass) moral error by attaching much too much significance to what are, after all, a morally arbitrary set of characteristics, and of mistakenly reifying the surface features of their relationship to the world's opportunities into the whole truth about their nature as subjects and objects of moral justification in collective codeliberation. After all, not only are their particular talents and abilities in an important sense a matter of "constitutive luck"¹⁹ but, even if they were not, the fact that they as individuals occupy a particular place in the *distribution* of talents,²⁰ or that their talents or abilities happen to command a certain market price given the constellation of others' talents, abilities, wants and desires, can only seem to carry with them a radical contingency.

On the other hand, the individual whose self-conception is of someone wholly alienated from her talents and abilities seems to commit an opposite (if not quite equal) kind of error. After all, when we come to live our lives under the basic

¹⁹ Nagel, "Moral Luck".

²⁰ Rawls ref. to follow. Note that Rawls talks about the moral arbitrariness of our place in the distribution of talents, not of talents themselves. (Lending support to the sense that this position about 'ambivalent subjects' fits well with Rawls's own view.)

structure of a particular socioeconomic system, it is not the case that our commitment to social justice in any way removes the fact that the life we must live is our own. In Rawls's terms, we have two moral powers, not just one; having a sense of justice sits alongside having a highest-level interest in framing, revising and pursuing *our own* conception of the good, and in living a life in which that conception can find expression. Moreover, the life we live is not the life of an abstract individual, but is the life of a real, empirical, embodied and embedded individual, the contingency of whose talents and abilities does not make them any less our own. One need not deny the fact that one's talents and abilities are undeserved, nor that their origins may be in part lie in the irresolvable failures of social fairness that come with merely imperfect meritocracy (as in FEO and FVPL). Nevertheless, there one stands, as one contingently happens to have become, and with a deep-seated interest in living within a social, political and economic system in which one can live a life that involves the development and deployment of the talents and abilities that one happens to have, whether or not one happens to have them because of various kinds of natural or social luck, and whether or not the role attributable to social luck is something that might have been eradicated in a more stringently meritocratic society.

My suggestion is that we go wrong in our self-understanding, as individuals who want to live under conditions justifiable to others who are similarly situated, if we do not feel both kinds of pull. In this domain, lack of ambivalence can only look like a basic moral mistake. The hope, then, is to find a set of justifiable principles that reach a coherent balance between these competing self-conceptions, neither of them dismissible but neither of them the full truth. My suggestion is that a pure meritocratic society (whether perfect or imperfect, operating under PEO or FEO), with no place for the difference principle or any similar principle of reciprocity, gives too much justificatory significance to contingent facts about agents' self-constitution, and loses the balance in one direction. Conversely, a view with no room for a meritocratic element in job competition, whether of type [3] with its anti-meritocratic conception of equality of opportunity (i.e. CEO), or whether any of the other views [4]-[7] that either reject equality of opportunity outright or fail to give it a non-derivative significance, make the opposite kind of mistake. Such views fail to take sufficiently seriously the fact that providing propitious conditions for contingently constituted, empirical individuals to pursue their highest-order interests means allowing a special significance to the deployment of individuals' talents and abilities, contingent though they may be. My suggestion is that the combination of a meritocratic job opportunity principle with a broader commitment to a principle of reciprocity, such as the difference principle, governing broader matters of distributive justice, strikes the right balance, and gives due weight to these two deep, recalcitrant and opposed elements of our self-conception as subjects of justification. Neither too in thrall to arbitrary contingencies, nor too far abstracted from embodied agents' real interests, the domain-specific commitment to fair equality of opportunity gets things just right. If we are conducting a project of finding mutually justifiable terms of social cooperation between individuals who stand in a complicated and ambivalent relationship to their own talents and abilities, then this is exactly where we should end up.