LSE International Inequalities Institute

WORKING PAPER 147

# Inherited Inequality, Meritocracy, and the Purpose of Economic Growth

**Francisco H. G. Ferreira**
International Inequalities Institute, LSE

**Paolo Brunori**
International Inequalities Institute, LSE

DECEMBER 2024

**Francisco H. G. Ferreira**
International Inequalities Institute, LSE

**Paolo Brunori**
International Inequalities Institute, LSE

Our working papers series is available to download for free on our website: www.lse.ac.uk/III

International Inequalities Institute
The London School of Economics and Political Science, Houghton Street, London WC2A 2AE

**E** Inequalities.institute@lse.ac.uk
**W** www.lse.ac.uk/III

International Inequalities Institute

**Inherited Inequality, Meritocracy, and the Purpose of Economic Growth**

**Francisco H. G. Ferreira and Paolo Brunori**[1]

Chapter for the *Oxford Handbook of Income Distribution and Economic Growth*

Edited by Gordon Anderson and Sanghamitra Bandyopadhyay

2 December 2024

**Abstract:** This is a chapter about inequality of opportunity and a closely related concept: inherited inequality. It does five things. First, it reviews the dominant economic model of inequality of opportunity, including its two main uses: the proposal of social objective functions and the measurement of inequality of opportunity. Second, it dispenses with two epistemically and normatively demanding assumptions that underlie the model and defines the closely related concept of inherited inequality. Although in practical terms the two are very similar, the latter rests on simpler, less demanding – and thus more solid – normative foundations. Third, it reviews recent advances in the measurement of inequality of opportunity and inherited inequality, focusing on data-driven solutions to model specification challenges. These methods are illustrated using UK data from 2009 to 2019. Fourth, the chapter proposes amending the standard static social objective functions proposed thirty years ago, towards a dynamic version that is better suited to addressing the implications – and conditioning the nature – of economic growth. Finally, the chapter discusses the differences between inequality of opportunity and meritocracy, and their possible roles in a fair society and growing economy.

**Keywords**: Inequality of opportunity; economic growth; inherited inequality; meritocracy

**JEL classification**: D63, O40

---

[1] Ferreira is at the London School of Economics and IZA. Brunori is at the University of Florence and the London School of Economics. We are grateful to Annaelena Valentini for excellent research assistance and Michael Vaughan for valuable feedback.

## 1. Introduction

In a well-known study of cognitive development among Ecuadorian children, Paxson and Schady (2007) document large and growing vocabulary differences between children with more and less educated mothers. Using the Spanish version of the Peabody Picture Vocabulary Test, they found that children whose mothers had 12 or more years of schooling scored around the international norm of 100 as they grew from three to five years of age. In contrast, those children in the same communities whose mothers had only 0-5 years of formal schooling saw a decline in their relative scores from approximately 85 to 65. [2]

Using survey data from 2004, Ferreira and Gignoux (2010) found that in Turkey, at age 15, school enrolment for girls with mothers without any formal schooling was just under 30%, whereas it was over 90% for boys with mothers with some secondary schooling or more. In a recent comparative study of educational achievements in Latin America, Fernández et al. (forthcoming) found large premia for private schools (relative to public schools) in PISA reading test scores across the region, reaching a high of 108% in Brazil. The same study documented that children from more educated parents (or from richer households, see de la O et al, forthcoming) are much more likely to attend those private schools. In the year 2000, whereas the indigenous population made up 31% of the Guatemala's prime-age population (aged 30-49), they represented 70% of the country's poorest income decile (Ferreira and Gignoux, 2008).

These differences are examples of inequality of opportunity: the situation that arises whenever predetermined circumstances that individuals cannot themselves control – such as family background, race or ethnicity, biological sex, place of birth, etc. – are predictive of valuable outcomes and advantages, such as income, educational attainment and the like. The concept is definitely not new. Indeed, it has long enjoyed political salience across many countries and cultures. In the United States, for example, President Franklin D. Roosevelt famously claimed that "We know that that equality of ability has never existed and never will, but we do insist that equality of opportunity still must be sought".[3]

Besides its broad political appeal, there are at least three important reasons why inequality of opportunity is an important concept, which matters separately from and in addition to inequalities in specific outcomes. First, the ideal of equality of opportunity has also garnered considerable philosophical support. In *A Theory of Justice,* one of the most influential works of modern political philosophy, John Rawls's (1971) proposed two principles of justice, the second of which he called the Equal Opportunity Principle. Dworkin (1981) and Arneson (1989) have also been influential in suggesting that "equality of resources" or "equality of opportunity for welfare" should be the normative objectives of just societies. The economist

[2] Scores are normalized, with 100 corresponding to the vocabulary recognition rates in a reference group taken to be a healthy and otherwise normal population.
[3] F.D. Roosevelt's address at Little Rock, Arkansas, 10 June 1936.

Amartya Sen (1980) reached similar conclusions in his 1979 Tanner Lectures at Stanford University, aptly entitled *"Equality of What?".* Although he may not have used the word "opportunity", Sen argued that societies ought not to seek to equalize specific final outcomes (or means), such as income or consumption. Instead, they ought to provide equal capabilities for individuals to lead lives (and enjoy 'functionings') of their own choosing. The concept of capabilities, much as that of opportunities, seems to refer to potential choice sets, often predetermined by local and historical circumstances, from which individuals can make their own choices. This notion became sufficiently dominant in political philosophy that Gerald Cohen (1989) refers to it as "the currency of egalitarian justice".

Second, there is also positive empirical evidence of a broader consensus on the need to implement equality of opportunity, certainly relative to any consensus on the need to seek equality of outcomes. A 2012 Gallup poll, for example, found that 70% of respondents in the US believe it is "extremely" or "very important" for the government to implement policies that enhance equality of opportunity for people to advance. In contrast, only 46% had the same opinion about reduction of income and wealth gap between the rich and poor (Pew Research Center, 2012).

These findings from opinion surveys are corroborated by various lab experiments, where participants often seem to distinguish between different sources of inequality, and are more averse to – that is, willing to compensate for – inequalities arising from circumstances agents have or had no control over, and less averse to those they interpret as arising from actions and choices that individuals themselves made. As an example, Cappelen et al. (2010) report on an experiment in which over 200 students were hired to perform a task that involved typing as many words as possible during either a 10 or 30-minute period. Participants chose either the short or the long task. They were given essentially identical documents to copy and were paid according to the number of words they got right, but at two different, randomly allocated 'wage rates': some were paid about USD 0.08, and others were paid USD 0.16 per correct word. Participants were then matched in pairs and told the working time, word output and payment for each. They are then given the opportunity to propose some (or no, or full) redistribution of earnings within the pair. In most cases, participants choose to redistribute some but not all earnings. In particular, a majority chooses to compensate participants who had been randomly allocated a low wage, but many fewer choose to compensate for the duration or quality of work, both of which were seen as being within the individual's control.[4]

Third, besides these widespread normative concerns, there are also efficiency reasons that might justify paying particular attention to inequality of opportunity. While the well-known efficiency-equity trade-off suggests that attempts to reduce inequality may hinder economic

[4] This paragraph draws on a similar summary in Ferreira (2022).

growth, empirical evidence is mixed, and more than one study finds a negative relation between inequality and economic prosperity (see Ferreira, 2023, for a summary). It has been suggested that one reason explaining the ambiguous relationship between inequality and growth is that inequality measures often mix, without distinguishing between, two types of inequality: inequality due to rewards of effort and choices, and inequality due to discrimination and lack of opportunities. Partially confirming this hypothesis, some studies have found a negative and statistically significant relationship between measures of inequality of opportunity and growth (see, e.g., Marrero and Rodriguez, 2013).

This chapter offers a brief survey of recent developments in the theory and empirics of inequality of opportunity, focusing on developments in the last ten years.[5] In addition, we offer three small novel contributions to the literature. First, we suggest a slight variation on the concept of inequality of opportunity, namely "inherited inequality". This is not intended as mere semantics but, instead, as an attempt to do away with one of the two normative principles underlying the theory of Equality of Opportunity (EOp), namely the Principle of Reward. We will argue that this principle has been controversial and has not aided wider acceptance of the EOp approach. A relatively simple reformulation eliminates the need for it, while also highlighting the common ground between IOp and the concept of intergenerational immobility.

Second, we consider the relationship between the normative ideals of equality of opportunity and meritocracy, which we feel have often been conflated and confused in popular discourse. We highlight key differences between them, while also suggesting that both may have a role to play in a fair society. Third, we revisit the important question of what the ultimate policy objective (or "allocation rule", in an older terminology) should be for a just, dynamic society. It is this final part that bears most closely on implications for economic growth, the distribution of incomes, and the relationship between the two. To illustrate the theoretical discussion, we draw on an empirical analysis of equality of opportunity in the UK, employing 11 waves of the Understanding Society survey (UKHLS) between 2009 and 2019.

The chapter proceeds as follows: Section 2 briefly reviews the dominant economic model of equality of opportunity and its implications both for the measurement of inequality of opportunity and for optimal allocation. Section 3 suggests dispensing with the Principle of Reward and introduces the closely related concept of inherited inequality. Section 4 reviews some selected novel developments in the practice of measuring IOp. Section 5 introduces our definition of meritocracy and contrasts it with EOp. Section 6 revisits the question of the ultimate policy (or development) objective, seeking to incorporate important considerations

[5] This time horizon is chosen in part to minimize overlap with Ferreira and Peragine (2016), who wrote a similar chapter for an earlier Oxford Handbook, namely the *Handbook of Well-Being and Public Policy*, edited by Matthew Adler and Marc Fleurbaey.

from other theories, such as the importance of equal liberties; of process fairness; and of poverty eradication. Brief concluding remarks are offered in Section 7.

## 2. The dominant economic model of inequality of opportunity[6]

Inspired by the debates in political philosophy on luck egalitarianism and the role of individual responsibility, economists – particularly those working on welfare economics and on social choice in the 1990s – soon offered formal theories of equality and inequality of opportunity (EOp and IOp, respectively); of how it should be conceptualized, modelled, and measured. One avenue, which Ferreira and Peragine (2016) call the *direct approach*, was to model a person's opportunities as a set of possible outcomes or 'goods', from which the individual may choose. One can then propose different rankings of opportunity sets; and indeed of distributions (or profiles) of opportunity sets. Pattanaik and Xu (1990), for example, characterise a simple cardinality ranking, in which one set is preferred to another if the former contains a greater number of elements than the latter. Kranich (1996), Ok (1997), Weymark (2003), and others built on these ideas to suggest criteria to compare and rank distributions of opportunity sets and thus to measure inequality of opportunity.

A different approach was taken by John Roemer (1993), Dirk van de Gaer (1993) and Marc Fleurbaey (1994, 1995), who proposed to infer information on the distribution of opportunities *indirectly* from relationships observed between outcomes, circumstances and, possibly but not necessarily, effort variables. To do so, this branch of the literature had to make a few basic assumptions, the most fundamental of which is what we will call the *Classifiability Assumption,* namely that a desirable individual outcome $y \in \mathbb{R}$ - where $y$ can be taken to represent an income or wealth level, educational attainment or achievement, etc. – is a function of two kinds of variables *only*: circumstance variables and effort variables.
In other words,

$$y = g(C, e) \tag{1}$$

where circumstances $C \in \Gamma$, $e \in \mathrm{E}$ and $g: \Gamma \times \mathrm{E} \rightarrow \mathbb{R}$. The distinction between the two sets is that circumstance variables are those which influence or determine $y$ (henceforth 'income') but lie entirely outside the sphere of responsibility (or control) of the individual. Conversely, effort variables are those which the individual can to some extent control and are thus considered to be within her sphere of responsibility.[7] To reiterate, the fundamental assumption is that if something contributes to determining $y$, it can be classified as either a circumstance or an effort.[8]

---

[6] This section draws (and elaborates a little) on Section 25.3 in Ferreira and Peragine (2016).

[7] Classifiability does not imply or require separability. Effort levels can be – and generally are – influenced by circumstances. The only requirement is that they should not be fully determined by them. There must be some individual locus of responsibility $\psi$, so that $e = h(C, \psi)$.

[8] The role of random events, or luck, has often been debated in this context. One view is that pure (or 'brute') luck might be considered a circumstance. On the other hand, 'option luck', where individuals actively opt to incur risk, may be viewed as within their sphere of responsibility. See chapter 6 in Fleurbaey (2008) for an in-

We then have a population of individuals $l \in N = \{1, \dots, l, \dots., L\}$, each of whom is fully characterized by the triple $(y_l, C_l, e_l)$. *C* denotes a vector of circumstance variables, each element of which is taken to be discrete. In principle, *e* can be an analogous vector of effort variables, although for simplicity we follow the convention of treating *e* as a discrete non-negative scalar that summarizes the relative effort or responsibility expended by the individual ($e \in \mathbb{N}$). The population can then be partitioned along two dimensions: into *types* $T_i$, which comprise all individuals displaying the specific realization of circumstances $C_i$, and into *tranches* $T^j$, which comprise all individuals expending effort $e_j$. If, without loss of generality, there are *n* types and *m* tranches, then society can be represented by two *n x m* matrices, $Y_{ij} = [y_{ij}]$, whose typical entry is:

$$y_{ij} = g(C_i, e_j)$$

and $P_{ij} = [p_{ij}]$, whose typical entry denotes the proportion of the population with circumstances $C_i$ and effort $e_j$.[9]

In addition to the epistemic *Classifiability Assumption*, the indirect approach to EOp posits two normative principles, which express its conception of fairness:[10]

- The *Compensation Principle* states that outcome ($y$) differences due to circumstances are unfair and should be either eliminated or compensated for.
- The *Reward Principle* states that outcome differences due to differences in effort levels are fair and should be (at least partly) preserved.

This simple framework, resting on the classifiability assumption and these two normative principles, has allowed researchers to analyse different aspects of distributive fairness, among which the two most dominant are also those most relevant to us. First, it has inspired normative allocation rules, or policy objectives and, second, it has inspired measures of inequality of opportunity, which seek to quantify the extent to which observed income distributions $F(y)$ deviate from these principles.

The literature contains a wide range of both proposed allocation rules and measures of inequality of opportunity, and this is due, at least in part, to the fact that the two foregoing normative principles, simple though they might seem, can be interpreted in a number of different ways, some of which are mutually inconsistent. There are at least two main forms of the Compensation Principle – namely ex-ante and ex-post – and a variety of Reward

depth discussion of the role of luck. Although there are exceptions, the dominant approach has been to apply *classifiability* to different kinds of luck or randomness, just as it applies to other variables.

[9] To clarify notation, we use $l$ to index individuals, *i* to index types, and j to denote tranches. Each individual $l$ perforce belongs to a type and to a tranche. If $l \in T_i \wedge l \in T^j$, then $y_l = y_{ij}$.

[10] See, for example, Fleurbaey (1995) and Fleurbaey (2008), building on Cohen (1989) and Arneson (1990).

Principles (Liberal, Utilitarian, Minimal, Inequality Averse, etc.). We do not propose to dwell on the details here, other than to briefly introduce the distinction between ex-ante and ex-post compensation.[11]

The ex-ante version of the compensation principle argues that inequalities due to different circumstances should be eliminated (conceptually) prior to the realization of effort, by equalizing the value of the opportunity sets faced by all. While this approach, originally associated with van de Gaer (1993), has the advantage that observing effort realizations is not necessary, it does require an evaluation of opportunity sets which, as noted earlier, are also difficult to observe. Most of the literature has followed van de Gaer in using the expectation of the outcome conditional on circumstances $C_i$ as a measure of the value of the opportunities open to people in type $T_i$. Under the assumption that the cross-section distribution $F(y|C_i)$ is a good approximation to the probability distribution $\mathrm{F}(y|C_i)$ faced by individuals in $T_i$, this leads to using the observed conditional type mean as a value of its opportunity set, and to equalization of type means as the objective of compensation.

By contrast, the ex-post version argues that inequalities arising from different circumstances should be eliminated (conceptually) after the realization of effort, by ensuring that all those exerting equal efforts should receive the same outcomes. This version has the opposite observability requirements: it does not need us to observe or evaluate opportunity sets, but it does require effort levels to be observed or, at least, identified. Most of the literature has followed Roemer (1993), in eschewing any attempt to empirically observe effort or responsibility and, instead, identifying a person's relative degree of effort as their rank (or quantile) in the distribution of outcomes conditional on circumstances $C_i$ – that is the type-specific conditional distribution $F(y|C_i)$.[12] So, whereas ex-ante compensation would associate perfect equality of opportunity with the condition

$$E(y|C_i) = E(y|C_k), \forall\, C_i, C_k \in \Gamma \qquad (2)$$

the ex-post version of the principle would associate it with the (clearly more demanding) condition:

$$F(y|C_i) = F(y|C_k), \forall\, C_i, C_k \in \Gamma \qquad (3)$$

[11] Fleurbaey and Peragine (2013) provide an excellent discussion of ex-ante and ex-post compensation, and of how the two principles relate both to each other and to different reward principles. See also Fleurbaey (2008).

[12] This is known as Roemer's "charity assumption". In addition to eliminating the need to observe effort levels explicitly, it also automatically ensures that the average effect of a specific set of circumstances on effort levels is treated as a circumstance. The quantiles of type conditional distributions are relative degrees of effort, and are invariant in type means. Although this is the most frequently adopted empirical approach, some contributions do assume observability of effort. Bourguignon et al. (2007) is probably the first example, but, most examples of selecting an observed variable to proxy for effort occur in studies of inequality of opportunity in health (e.g. Jusot et al. (2013), Brunori et al., 2022).

Each of these two versions of the compensation principle has given rise to alternative optimal allocation rules. In each case, a combination of monotonicity (or the Pareto principle) with extreme inequality aversion led to the adoption of (different) Rawlsian maximin solutions. In the ex-ante version (van de Gaer, 1993), societies are urged to organize policies and institutions so as to maximize the expected incomes of those who belong to the 'poorest' type:

$$W_V = \min_{T_i} E(y|C_i) \tag{4}$$

Although most of the empirical literature has not focused on such welfare functions, there has been some work on describing the worst-off types, or 'opportunity deprivation profiles' as they are called by Ferreira and Gignoux (2011).

In the ex-post version, the objective is to raise as much as possible the lowest incomes in every tranche which, given the charity assumption, corresponds to the lowest incomes at each quantile, across all types: $\min_{T_i} F^{-1}(p|C_i)$. However, since that is not a well-defined objective unless stringent separability conditions hold, Roemer (1993) suggest that society should seek to maximize the average across quantiles of the lower envelope of the quantile functions:

$$W_R = \int_0^1 \min_{T_i} F^{-1}(p|C_i)\, dp \tag{5}$$

Geometrically, (5) corresponds to the area under the lower envelop of the type-specific quantile functions (QFs), which are the inverse of the type's conditional distribution functions (CDFs). Figure 1 below illustrates using two types from the UK Understanding Society survey for 2009.[13] Panel A plots their CDFs and Panel B inverts those and plots the quantile functions: incomes as a function of the percentile of the distribution. The Roemerian maximand $W_R$ is the area shaded in light blue, below the lowest of the QFs at each quantile (notice the crossing at the very first quantiles). [14]

The two versions of the compensation principle have also led to different approaches to measuring inequality of opportunity. As Ferreira and Peragine (2016) note, the various different approaches to measuring IOp can all be described in terms of two steps. In the first step, all fair inequality is removed from the matrix $Y_{ij}$. This operation yields a counterfactual matrix $\tilde{Y}_{ij} = [\tilde{y}_{ij}]$, which contains only unfair inequality, or inequality of opportunity. In the

---

[13] The example is based on a simple partition that divides the entire population in two types by estimating a transformation tree of depth 1. The algorithm, described in more detail in section 4, splits the sample into two subsamples: one group includes respondents reporting a mother with at most primary education, the other contains respondents whose mothers had secondary or higher education.

[14] Note that if the poorest type is first-order stochastically dominated by all others, this objective becomes identical to the ex-ante maximand proposed by van de Gaer (1993), given in Eq. 4.

second step, an appropriate, axiomatically derived inequality index $I(.)$ is applied to the counterfactual distribution $F(p_{ij}\tilde{y}_{ij})$.

Figure 1: example of CDFs and QFs for two types, and Roemer's social evaluation function

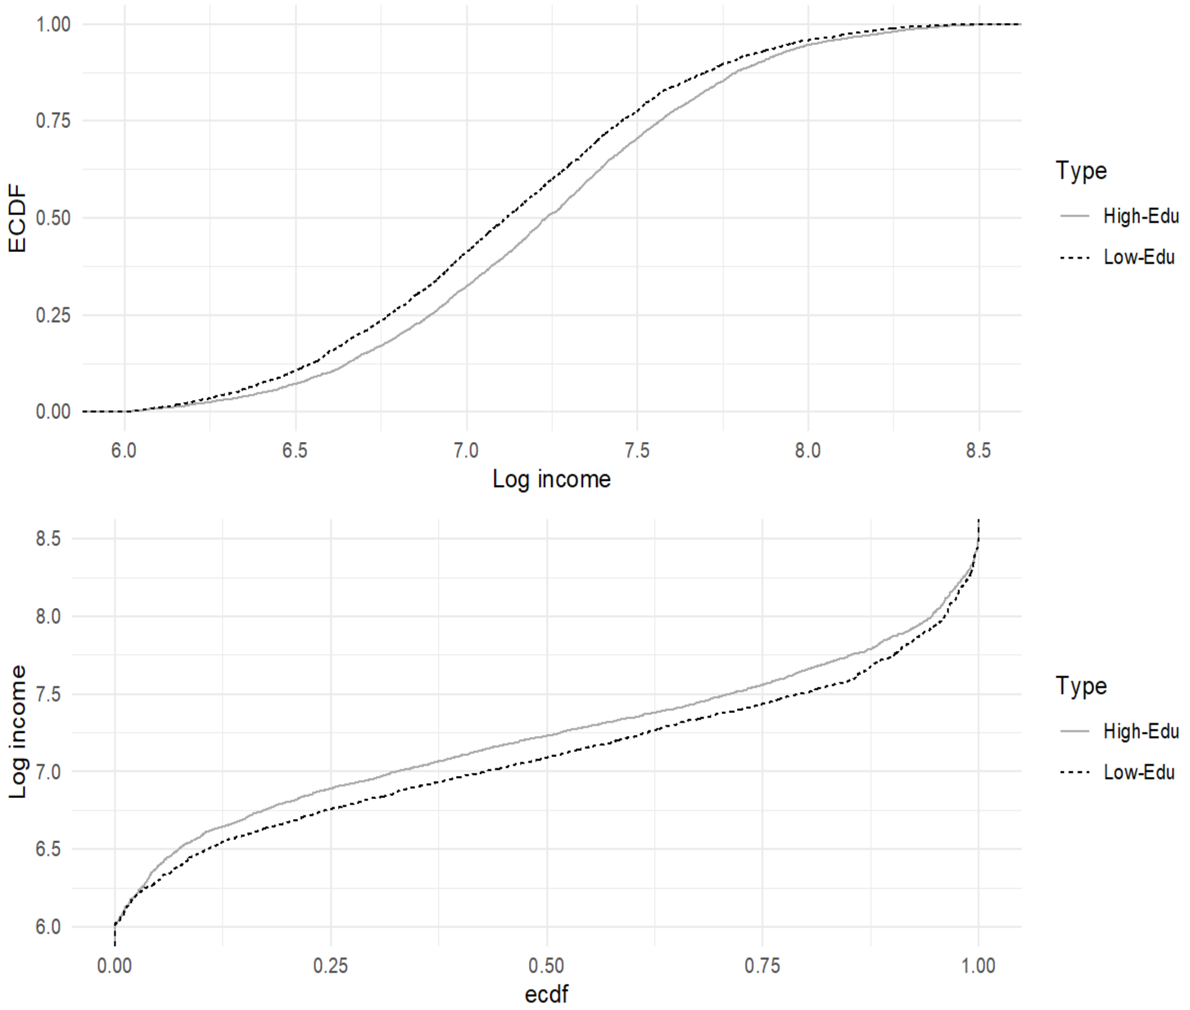


*Source: Own elaboration using data from UKHLS (2009).*

Many alternatives have been suggested in the literature for each of the two steps. Naturally, the first step is the one where the normative principles of EOp manifest. For example, one approach consistent with the ex-ante compensation principle (as implied by Eq. 2) is the between-types approach of, e.g., Ferreira and Gignoux (2011), where the typical entry of the counterfactual matrix $\tilde{Y}_{ij}$ is set as $\tilde{y}_{ij} = E(y|C_i)$. In words, the between-type approach assigns to each individual her type's mean income, as a measure of the value of her opportunity set. Inequality within types, which is considered fair by (various versions of) the Reward Principle, is thereby removed, and the inequality index applied over $\tilde{Y}_{ij}$ captures only the inequality between types, which is judged unfair because it reflects differences in circumstances.

Similarly, a measure consistent with the ex-post compensation principle (as implied by Eq. 3) is the within-tranches approach of, e.g., Checchi and Peragine (2010), where the typical entry of the counterfactual matrix is $\tilde{y}_{ij} = \frac{\mu}{E\left(y \middle| y = F^{-1}\left(q^j\right)\right)} F^{-1}\left(q^j \middle| C_i\right)$. In words, the within-tranches approach assigns to each individual in type $T_i$ and tranche $T^j$ her own income, $F^{-1}\left(q^j \middle| C_i\right)$, multiplied by the ratio of the overall population mean, $E(y) = \mu$, to the mean income in tranche $T^j$.[15]

Many other ways to eliminate fair inequality and to construct the counterfactual matrix $\tilde{Y}_{ij}$ have been suggested, each corresponding to a particular view of fairness and, often, to particular versions of the compensation and reward principles. These include, among others, the direct unfairness and the fairness gap methods, both proposed by Fleurbaey and Schokkaert's (2009).[16] That said, the between-types and within-tranches approaches remain the most frequently used in practice.

## 3. Inherited Inequality

This basic model of equality and inequality of opportunity, built around the *classifiability* assumption, the principle of compensation and the principle of reward, has been quite influential. At the time of writing, John Roemer (1998) has 4,336 citations in Google Scholar. Measures of IOp have been computed for at least 72 countries, accounting for 67% of the world population.[17] The World Bank dedicated one of its influential World Development Reports to this topic (World Bank, 2005), and many other international organizations have followed suit.

Nevertheless, the concept is not uncontroversial, and there have been a number of critiques, both of the theory and the associated empirics (e.g. Balcazar, 2015, and Kanbur and Wagstaff, 2016). Some of the scepticism is directed at the *classifiability* assumption, and particularly at the notion that there is such a thing as a locus of individual responsibility $\psi$ (see footnote 6). The discussion quickly becomes philosophical: is there such a thing as free will? Or is every decision a person takes driven by her previous history up until that point, and thus pre-determined by either nature or nurture, both of which are or reflect circumstances?

This view, often called "causal determinism" has a long history and many prominent adepts among philosophers, who question whether it is possible to disentangle what individuals

[15] This transformation is consistent with any scale-invariant inequality measure. For a translation-invariant inequality measure, the expected income in tranche j should be subtracted from, instead of dividing, the unconditional expectation.

[16] The interested reader is referred to the original paper, or to Ferreira and Peragine (2016) for a summary description.

[17] See www.geom.ecineq.org

obtain because of their free choices from the direct or indirect effect of genetics and environment. As suggested by Fishkin in his volume, *Bottleneck: a new theory of equality of opportunity*: "This project turned out to be like peeling away layers of an onion. [...] There is no way to separate a person from the accumulated effects of her interactions with her circumstances, including her opportunities, because the product of those accumulated interactions is the person" (Fishkin, 2014 p. 64.) .

Others object to the Principle of Reward: If a fellow human being is homeless and starving, is society devoid of any responsibility, even if the person's own actions and choices led him to this situation? At the other extreme, are the enormous rewards bestowed by the market upon the world's most successful CEOs unquestionably fair, and to be preserved? The theory can, to some extent, accommodate some of these critiques. In the case of the Reward Principle, in particular, different versions have been proposed to address the apparent excesses of liberal or natural reward. Inequality-averse reward, for example, permits some compensation of differences within a type. In addition, some authors have proposed modifications to exempt very low incomes or wellbeing from reward considerations. Examples include Bourguignon, Ferreira and Walton (2007) and Hufe et al. (2022), who suggest that rewards should be treated differently when addressing poverty.

In both cases – scepticism about *classifiability* and discomfort with the Reward Principle – it is the idea of effort and responsibility – the very role for individual responsibility in an egalitarian theory that appealed to Arneson and Cohen – that seems to make some scholars uncomfortable. It is difficult to fully observe or measure the exercise of responsibility throughout a person's life, and this poses difficulties. Yes, Roemer's charity assumption is both elegant and powerful, but it is based on a number of underlying assumptions, which are not necessarily easy to either confirm or believe.[18] These kinds of questions have prevented many social scientists from fully embracing the model in Section 2 as a tool for assessing unfair inequality.

In this section, drawing in part on Brunori et al. (2023), we propose an epistemically economical variant of the theory which dispenses with the classifiability assumption and the reward principle entirely. Indeed, it requires no concept of effort or responsibility at all, and it weakens the compensation principle. Nonetheless, *for the purpose of measuring unfair inequality*, it leads to precisely the same exercises that have been used for measuring IOp, with a much less demanding set of underlying assumptions. Seen from this new perspective, we argue that the empirics of IOp provide a meaningful assessment of unfair inequality with a considerably reduced burden of assumptions.

[18] The underlying assumptions include: monotonicity of outcome on effort, limited effect of luck and full observability of circumstances (see Lefranc et al. (2009) for a discussion).

This variant approach seeks to measure inherited inequality. It defines a set $\mathrm{H} \subset \Gamma$ of circumstances which are inherited at birth. $\mathrm{H}$ is a strict subset of $\Gamma$, because some factors that lie beyond a person's responsibility are not inherited at birth. Accidents, earthquakes, recessions, and many other events that take place during a person's life but which she cannot control help influence her outcomes. So, the vector of inherited characteristics $H \in \mathrm{H}$ that each individual has consists of a subset of the elements of C. The only normative principle we need to underpin this variant approach to the measurement of unfair inequalities is the:

- *Weak Principle of Compensation (WPC)*: In a fair of society, a person's outcome $y$ should be independent of inherited characteristics *H*. In other words, $y \perp H$.

This immediately implies that, in a fair society:

$$F(y|H_i) = F(y|H_k), \forall\, H_i, H_k \in \mathrm{H} \qquad (3')$$

Which, in turn, of course implies:

$$E(y|H_i) = E(y|H_k), \forall\, H_i, H_k \in \mathrm{H} \qquad (2')$$

If we then define inherited types $T'_i$, which comprise all individuals displaying the specific realization of inherited characteristics $H_i$, the measurement approaches previously defined now go through with respect to a counterfactual matrix $\widetilde{Y'}_{ij}$. Any such matrix that was previously defined with respect to types $T_i$ and circumstances $C_i$ can now be adapted to inherited types $T'_i$ and inherited characteristics $H_i$. But the interpretation of the ex-post measures, which quantify deviations from (2') no longer require taking a position with respect to efforts or responsibility. They simply measure deviations from the independence requirement in the Weak Principle of Compensation.

Whereas the vector *C* of circumstances in the EOP model is determined implicitly by the *classifiability* assumption, the set $\mathrm{H}$ may be selected explicitly, and in one of two ways, depending on the purpose at hand. First, if a society should choose to adopt the WPC as a normative principle to guide its government's actions, then the elements of $\mathrm{H}$ should arise as the result of a process of democratic deliberation, analogous to that envisaged by Sen (1985) for the choice of functionings that should comprise capability sets. Indeed, it is possible that different societies may take different positions on the composition of $\mathrm{H}$. It seems likely that race, ethnicity, sex, and certain elements of family or class background would be included by most. But there may be different opinions regarding the components of a person's genetic make-up, for example, which should be treated as inherited characteristics in this normative setup.[19]

Second, in a research context and in the absence of a socially or politically determined specification for the set $\mathrm{H}$ of inherited characteristics to which specific outcomes, such as

[19] See Harden (2021) for a comprehensive discussion of the emerging use of genetic data to understand inequalities.

income, should be orthogonal, its composition may be left to the discretion of the researcher. In this latter case, it is quite likely that it will be jointly determined by the researcher's own normative views and by the data that is available. In the UK-based illustrations in this chapter, we selected self-reported ethnicity, sex, country of birth, parental education, and parental occupation levels as inherited circumstances.

As noted above, the notion of inherited inequality is conceptually more restrictive, or conservative, than that of inequality of opportunity: $\mathrm{H}$ is a strict subset of $\Gamma$. Yet, in terms of measurement, three observations are in order. First, since the criteria implied by the WPC (equations 2' and 3') would be identical to those implied by the three assumptions underlying IOp (equations 2 and 3) if the two sets were the same, the measurement methods are also identical. They consist of following the same two steps, where a counterfactual distribution $\widetilde{Y'}_{ij}$ is constructed by removing all fair inequality, and then an inequality index is applied to it.[20]

To estimate the extent to which a population differs from the orthogonality norm, $y \perp H$, is effectively to estimate the extent to which inherited characteristics *H* predict incomes *y*. Indeed, under the orthogonality condition laid out in the WPC, no prediction function $f(H)$ is informative of y. Given any information on a person's inherited characteristics *H,* the best predictor of her income is simply $E(y)$.[21] It is then natural to measure inherited inequality by the extent to which inherited circumstances *H* can predict *y*. The concept of unfair inequality that corresponds to the Weak Principle of Compensation is the inequality that can be predicted by the set of inherited circumstances, H. A counterfactual matrix $\widetilde{Y'}_{ij}$ can therefore be constructed by letting its typical entry $\widetilde{y'}_{ij} = \hat{f}(H)$, where $f$ is a prediction function of the form:

$$y = f(H), f \in \mathcal{F}$$

By construction, this matrix removes all fair inequality. Then an (absolute) measure of inherited inequality is simply given by $I\left(\widetilde{y'}_{ij}\right)$, where $I(.)$ is some axiomatically derived inequality index, as before. $I\left(\widetilde{y'}_{ij}\right)$ is the unfair inequality corresponding to the WPC. Also, just as measures of IOp are often expressed as functions of ratios to total income inequality, relative measures of inherited inequality can similarly be written as:

[20] The only difference is, once again, of interpretation. The columns of the matrices $Y'_{ij}$ and $\widetilde{Y'}_{ij}$ correspond simply to discretized quantiles of the conditional distributions, with no association to effort or responsibility required.

[21] This is because $y \perp H \Leftrightarrow F(y|H_i) = F(y|H_k), \forall\, H_i, H_k \in \mathrm{H} \Longrightarrow E(y|H_i) = E(y|H_k) = E(y), \forall\, H_i, H_k \in \mathrm{H}$

$$I_n = \phi\left(\frac{I(\widetilde{y'}_{ij})}{I(y)}\right) \tag{6}$$

The second observation is that, in practice, very few empirical studies of IOp have used circumstance variables that are not inherited at birth. One example is Hufe et al. (2017), who include some variables affecting children while they are young. The vast majority of studies have relied on variables that are determined by the time a child is born, such as those used in the empirical analysis contained in this chapter. In these cases, $\mathrm{H} = \Gamma$ and the measurement of IOp and Inherited Inequality are the same in practice. The point is merely that the meaning of the measures thus obtained no longer requires taking normative positions on either the precise boundaries between responsibility and non-responsibility factors (the classifiability assumption), or on the extent to which personal responsibility is deserving of reward, in order to construct a meaningful measure of the extent to which inequality is inherited.

Third, the definition of inherited inequality is also closely related to that of intergenerational mobility or, more precisely, its inverse: intergenerational persistence. Suppose a researcher decides that the only inherited characteristic that really matters is parental income. Then *H* = $\{y_p\}$. If, in addition, the researchers chooses a prediction function of the form $y = e^{\alpha+\beta \log y_p + \varepsilon}$, then $\widetilde{y'}_{ij} = e^{\widehat{\alpha}+\widehat{\beta}\log y_p + \frac{\sigma^2}{2}}$. $\hat{\beta}$ is of course an estimate of the intergenerational elasticity of income, a common measure of mobility, estimated through the Galtonian regression $\log y = \alpha + \beta \log y_p + \varepsilon$. Another common measure of mobility, the correlation coefficient between $\log y$ and $\log y_p$, is given by $\hat{\beta}\sqrt{\frac{var \log y_p}{var \log y}}$, an example of $\phi\left(\frac{I(\widetilde{y'}_{ij})}{I(y)}\right)$, when the inequality index is the variance of logarithms.

Another way of seeing this is to note that this correlation coefficient between the child's income and the parent's income is the square root of the $R^2$ of the Galtonian regression which, of course, captures the share in the variance of the child's income that is predicted (or 'explained') by the variance in the parent's income. This is but an example of the share of total inequality that can be predicted by inherited circumstances, as in Equation (6). In that sense, certain measures of intergenerational mobility can be seen as specific examples of measures of inherited inequality, under the restriction that the only inherited circumstance that should be considered is parental income.[22]

[22] And, in this specific example, that the inequality index is the variance of logarithms. The variance of logarithms is not a great measure of inequality, in that it does not satisfy the Pigou-Dalton transfer axiom, but alternatives are available.

## 4. Recent advances in measurement

If one is interested in the full extent to which inherited circumstances predict incomes, and thus in the extent of inherited inequality, there is no *a priori* reason for restricting the set of inherited characteristics to parental income alone.[23] Most data sets that contain information on parental income also contain relevant additional information, on variables such as sex, ethnicity, or place of birth. For many countries, reliable information on parental income may not be available at all, but data on other aspects of family background, such as the education levels or occupations of parents, may exist. Once such information is analysed and a suitable set $\mathrm{H}$ of inherited characteristics is selected, the following step is to choose a prediction function $f(H), f \in \mathcal{F}$, to predict current incomes. Given that choice, counterfactual incomes $\widetilde{y'}_{ij} = \hat{f}(H)$ can be estimated, and (5) can be computed. The fundamental empirical problem for assessing the extent of – that is, for measuring – inherited inequality or inequality of opportunity is therefore one of model selection.

Since the first attempts to estimate inequality of opportunity, many different models $f(H)$ have been used. In early work by Bourguignon et al. (2007) and Ferreira and Gignoux (2011), for example, linear regressions of the form $f(H) = \alpha + H\beta + \varepsilon$ were often used. Alternatively, Checchi and Peragine (2010) and others used a non-parametric alternative, in which the sample was partitioned into empirical types $\hat{T}_i$ according to various elements of H, and $f(H)$ was computed from the means and distributions within $\hat{T}_i$. In a simple example, $f(H) = E(y_l | l \in \hat{T}_i)$.

But the very fact that we must distinguish between empirical types $\hat{T}_i$ and theoretical types $T_i$ illustrates the model selection problem that we are concerned with. The theoretical types $T_i$, introduced in Section 2, are the rows of the full population matrix $Y_{ij}$. In a sample, two problems may arise: first, certain types, particularly if rare, may not be observed at all, even if they exist in the population. Second, even with a limited number of inherited characteristics, each containing multiple categories, the fully interacted partition may be too fine for the moments of each cell to be estimated with any precision in a sample.

To see this, consider the sample of the 2009 wave of Understanding Society, the first wave of the longitudinal survey of the UK which we use to illustrate this section. We observe sex (2), country of birth (25), self-reported ethnicity (5), father's and mother's education (4 each), father's and mother's occupation (10 each). The number in parentheses following each circumstance is the number of categories it contains. The number of possible types in this partition is 400,000, the product of all these numbers. Considering that our sample size varies

[23] Of course, one might only be interested in the narrower question of the association between incomes across two or more generations. This is a perfectly legitimate question in itself, which the literature on intergenerational mobility seeks to answer. We are interested here in the broader question of how much inequality can be predicted by a full range of inherited circumstances.

between 4,700 and 10,987, one can immediately see that a fully saturated model would result in a large number of cells, some of which would simply be too small for moments to be estimated reliably. In a parametric setting, this would be equivalent to a fully saturated linear regression, in which all possible interactions were included. Such a model would clearly be overfitted.

This implies that empirical studies of IOp or inherited inequality must select a number of empirical types $\hat{T}_i$, denoted $n(\hat{T}_i)$, with $n(\hat{T}_i) \leq n(T_i)$. In practice, that inequality is typically strict, and the difference can be large. The choice of the empirical partition $\|\hat{T}_i\|$ is an important component of the model selection problem, and it involves a trade-off between two different kinds of bias that work in opposite directions. The first is an omitted variable bias: selecting a partition $\|\hat{T}_i\|$ with too few empirical types (a low $n(\hat{T}_i)$) leads to an underestimate of IOp or inherited inequality, relative to the true theoretical partition (Ferreira and Gignoux, 2011). On the other hand, overfitting the sample data and choosing too large a $n(\hat{T}_i)$ can lead to an upward bias in estimates of IOp (Brunori, Peragine and Serlenga, 2018)

Selecting an empirical partition is a key part of the model selection problem of choosing $f(H)$, but it is not the full problem. One could define ex-ante and ex-post analogues to the WPC, and choose different functional forms for the prediction function, for a given $\|\hat{T}_i\|$. The ex-ante version would focus on differences between type means (deviations from Eq. 2'), whereas the ex-post version would focus on differences between full distributions (deviations from the stronger requirement given by Eq. 3').

Linear regression models of the form $f(H) = \alpha + H\beta + \varepsilon$ have continued to be used, and the choice of interaction terms included in each specification determines $\|\hat{T}_i\|$. Panel data applications using fixed effects specifications have been proposed to estimate upper - as well as lower – bounds for IOp (Niehues and Peichl, 2014). A number of examples of this and other approaches were summarized by Ferreira and Peragine (2016) and Ramos and van de Gaer (2021). More recently, the prediction nature of the model selection problem has led to the use of machine learning approaches in the measurement of IOp or inherited inequality. These are often designed to solve the two components of the model selection problem – the empirical partition and the functional form of the prediction function – in an integrated fashion, and we illustrate them below.

To do so, we use the 2009-2019 waves of Understanding Society – The UK Household Longitudinal Study (UKHLS). We focus on these 11 waves because earlier surveys did not contain information about net household income.[24] To obtain a sufficient sample size, we

[24] The net household monthly income is contained in the variable w_fihhmnnet1_dv+. It is derived as the sum of net monthly incomes from all household members and it can be decomposed into the six subcomponents: net labour income (w_fihhmnlabnet_dv+), miscellaneous income (w_fihhmnmisc_dv), private benefit income

consider the survey's cross-sectional component, which necessitates accounting for the risk of life-cycle bias by restricting the analysis to individuals between the ages of 35 and 50. This age range is selected because it is a period in which earnings are assumed to be close to permanent income, minimizing the risk of biasing estimates (Nybom and Stuhler, 2016). The outcome of interest is equivalized net household income (excluding deductions). The equivalence scale adopted is the OECD square root scale, which adjusts total household income by dividing it by the square root of the household size (OECD, 2013). Monetary values are expressed in 2015 pounds.[25] The data is trimmed by dropping all observations reporting incomes below the value of the 0.1% percentile (74 negative incomes) and the top 99.9%, dropping individuals with an equivalized monthly net income above £35,2327 (N=127).

The circumstances considered are country of birth (24 countries plus a residual group of "other"), ethnicity (five self-reported categories: "Black", "White", "Asian", "Mixed", "Other"), father's and mother's occupation (9 categories plus "out of the labour force" using the Standard Occupational Classification 1990 (SOC90) of the father's/mother's job, converting from the 2-digit version to the 1-digit version)[26], and father's and mother's education (Primary, Lower secondary, Upper Secondary, Tertiary).

Two data-driven approaches recently introduced to the empirical measurement of IOp or inherited inequality are illustrated in Figures 2 and 3 for the first UKHLS wave (2009), the same figures for 2019 can be found in Appendix. Figure 2 shows a conditional inference tree (CIT) for the UK, which consists of a sequence of binary splits of the sample, each of which is selected by choosing the inherited characteristic that is most closely correlated with income, and then the specific split among categories that maximizes the statistical significance of the difference between the two branches (by minimizing the p-value of a suitable test of differences in means). Each split generates two nodes, and the process is repeated at each node, until the differences between the means of any possible downstream nodes are no longer significant (at some predetermined significance level). These conditional inference trees, initially proposed by Hothorn et al. (2006) and first applied to IOp estimation by Brunori, Hufe and Mahler (2023), provide a flexible way to choose an empirical partition $\left\|\widehat{T}_i\right\|$ which is driven by the patterns in the data and which avoids overfitting by means of cross-validation techniques designed to maximize prediction out of sample. These trees are known to be unbiased, but high-variance, estimators, and sets of trees – known as random forests – can improve their statistical robustness.

(w_fihhmnprben_dv), investment income (w_fihhmninv_dv), pension income (w_fihhmnpen_dv), and social benefit income (w_fihhmnsben_dv).

[25] Using the CPI from Consumer Price Inflation Team, Office for National Statistics. https://www.ons.gov.uk/economy/inflationandpriceindices/datasets/consumerpriceindices

[26] 1 - Managers And Administrators, 2- Professional Occupations, 3 - Associate Professional And Technical Occupations, 4 - Clerical And Secretarial Occupations, 5 - Craft And Related Occupations, 6 - Personal And Protective Service Occupations, 7 - Sales Occupations, 8 - Plant And Machine Operatives, 9 - Other Occupations.

The tree in Figure 2 splits the UK sample into 9 types. As in the first illustration in Figure 1, the first split is determined by mother's education, separating the tree in two sub-trees with very different expected income (all types in the left sub-tree have an expected income below the population mean). The second key circumstance is country of birth, selected by the algorithm to define half of the splits. The worst-off group consists of first-generation immigrants from Bangladesh, Canada, Pakistan, Poland, and Turkey, whose mothers had at most a primary education. Their expected monthly equivalized income is £1,225, which is 58% of the mean income in the population (£2.128 in 2009). Interestingly, the best-off group also includes immigrants, in some cases from the same countries (Poland and Canada). This exemplifies how circumstances interact in complex ways to shape opportunities. Type 14, with an expected income 57% above the average, are also first-generation immigrants, but their mothers have upper secondary or higher education. Between-type inequality is about 7 Gini points.

Figure 2: Conditional Inference Tree for UKHLS 2009

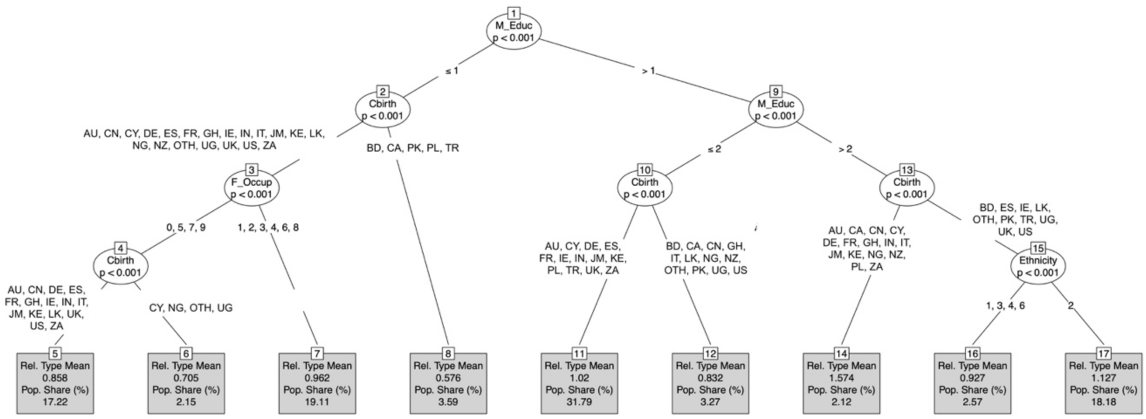


*Source: Own elaboration using data from UKHLS (2009). Note: CIT minimum number of observations per type is set to 200, maximum Bonferroni-adjusted p-value to perform a split is 0.001.*

As noted earlier, conditional inference trees are designed to select the partition $\|\widehat{T}_i\|$ by finding the most significant differences between type means. And the final summary measure of inherited inequality is computed over the (population-weighted) means of the 'leaves' or terminal nodes of the tree (and for some suitable average in the case of random forests). In other words, $I\left(\widetilde{y}'^{CIT}_i\right)$ is computed over $\widetilde{y}'^{CIT}_i = E\left(y_l \middle| l \in \ \widehat{T}^{CIT}_i\right)$. This makes CITs particularly well-suited to the weaker, "ex-ante" version of the WPC, which measures deviations from Eq. 2'.

A different machine learner, namely the transformation tree (TrT), is better suited to the stronger, "ex-post" version of the WPC, in that it measures deviations from Eq. 3'. Transformation trees were developed by Hothorn and Zeileis (2021) and first applied to the measurement of IOp and inherited inequality by Brunori, Ferreira and Salas-Rojo (2023). In essence, they are analogous to CITs, except that splits at each level of the tree seek to maximize statistically significant differences between the full cumulative distribution functions in each node, rather than just the mean.[27] Figure 3 displays the TrT for the UK in 2009. The same figure for 2019 is reported in Appendix.

Figure 3: Transformation Tree for UKHLS 2009

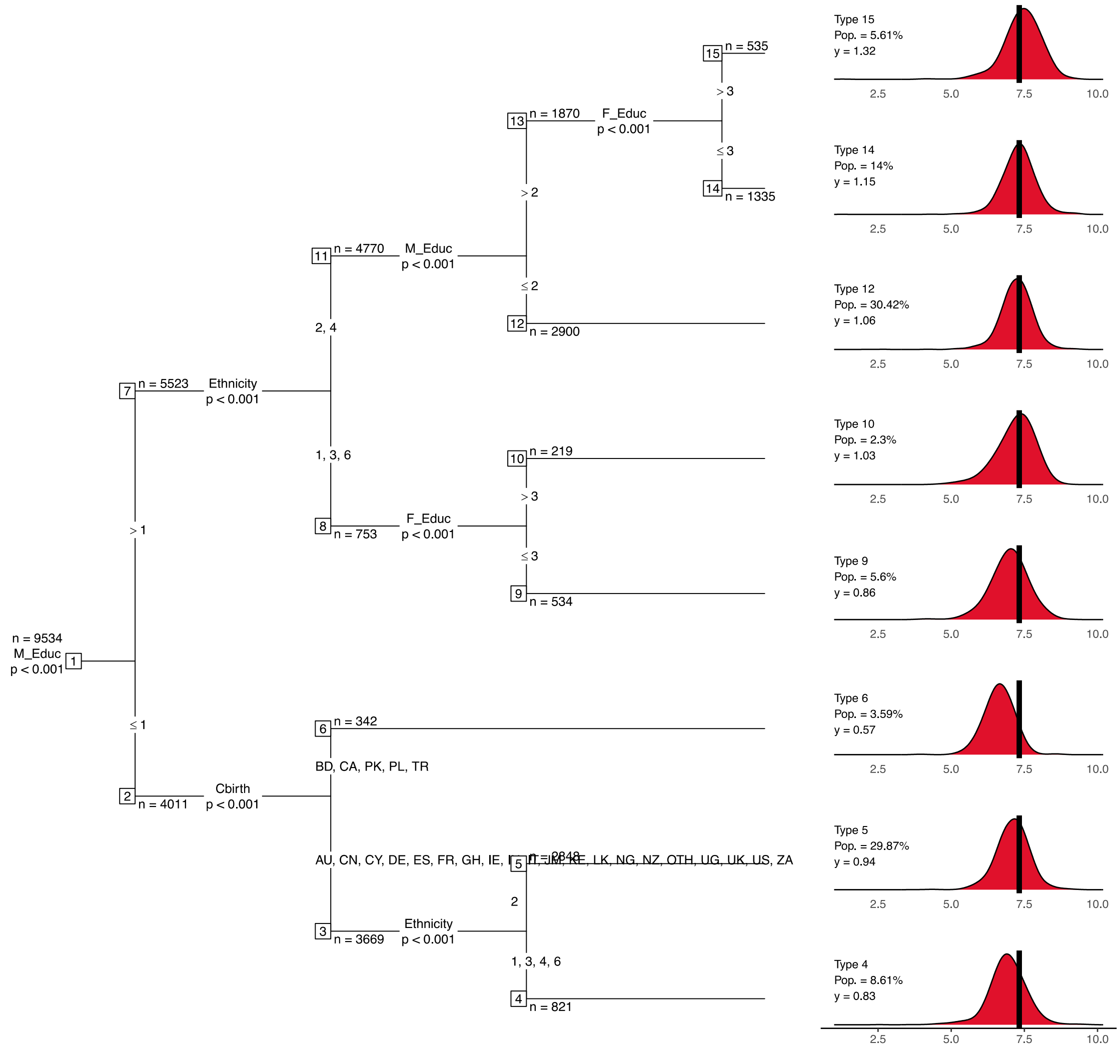


*Source: Own elaboration using data from UKHLS (2009). Note: TrT's Bernstein polynomial approximation order is 7, minimum number of observations per type is set to 200, maximum Bonferroni-adjusted p-value to perform a split is 0.001.*

Transformation trees have similarities, but also differ from conditional inference trees: while the first split in Figure 3 is exactly the same as in Figure 2 – mother's education primary or

[27] This is achieved by testing for parameter stability across Bernstein polynomials which are used to flexibly parameterize each CDF. See the references in the main text for details.

above – country of birth is used only once, to separate respondents with low mother's education into a better-off group (including UK natives) and a worse-off type (exactly the same observed in the CIT). Father's education is used twice, together with mother's education and ethnicity, to define the partition into empirical types. The best-off type is made up of individuals with highly educated mothers and fathers, reporting to be "white"[28] Figure 4 displays the expected CDFs (ECDFs) for the eight empirical types given by the terminal nodes of the TrT in Figure 3.

Figure 4: Ex-post inherited inequality visualization based on a TrT for the UKHLS (2009)

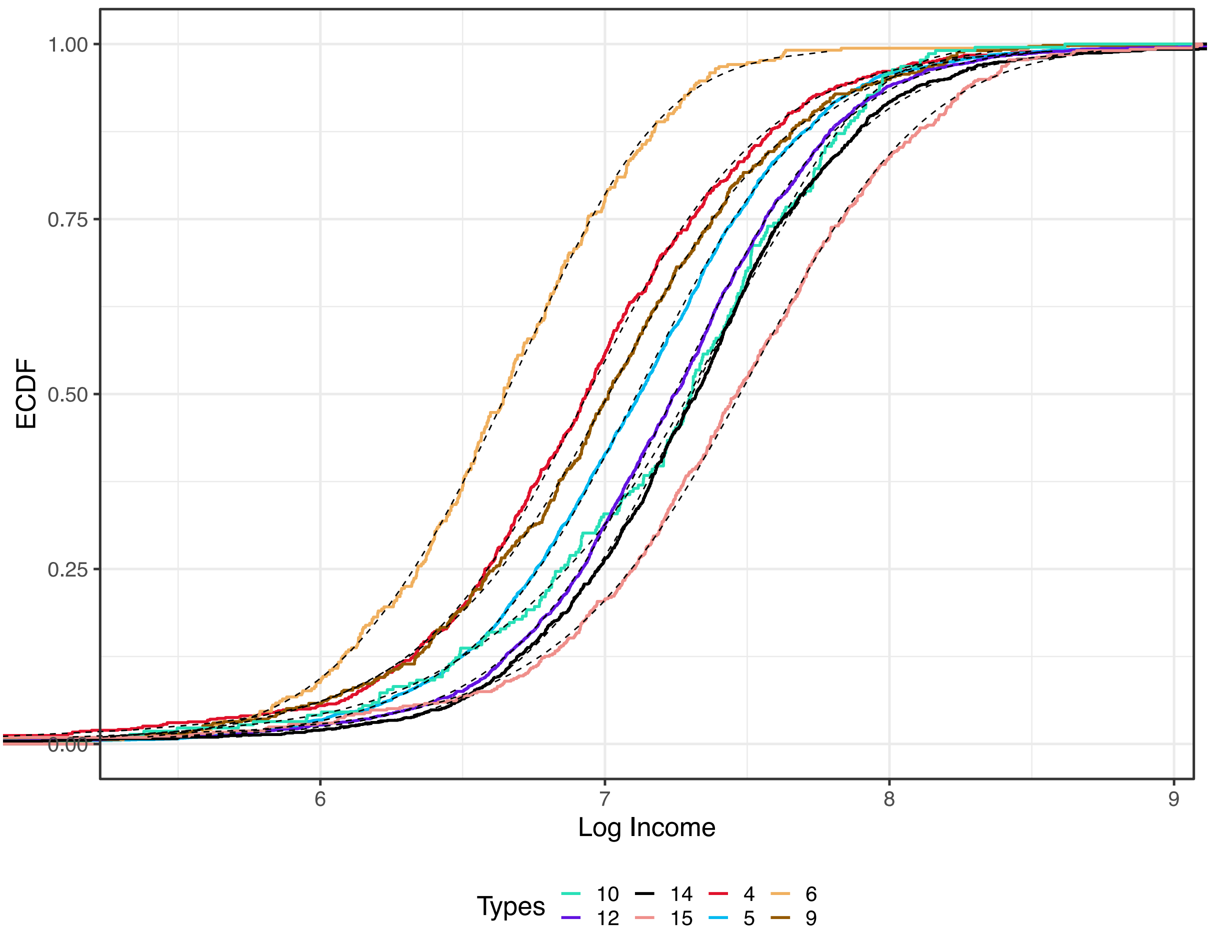


*Source: Own elaboration using data from UKHLS (2009). Note: TrT's Bernstein polynomial approximation order is 7, minimum number of observations per type is set to 200, maximum Bonferroni-adjusted p-value to perform a split is 0.001.*

The ECDFs reported in Figure 4 allow us to visualize ex-post inherited inequality (or IOp). Any horizontal gap in Figure 4 is a violation of the compensation principles – weak or strong – and is in fact aggregated into our second measure of ex-post IOp (about 8 Gini points in this example). Inspecting Figure 4 the reader can also appreciate aspects of the conditional

[28] This type also includes individuals reporting "other" as their ethnicity, but over 95% of the individuals in this type report being "white".

distributions that would be invisible from an ex-ante perspective. Consider type 10: it is among the worst-off for low quantiles, but its ECDF then crosses most of the others, and it becomes the second best-off around the 80th percentile. This suggests a conditional distribution with a first moment not far from the population average (the expected outcome is only 3% higher than average) but with a much higher variance (and left tail in particular). Interestingly, this is a type consisting of individuals reporting medium-high education for their parents but reporting to belong to "black, "Asian", and "other" ethnic groups.

Finally, the functions shown in Figure 4 can be inverted to show Roemer's original maximand: the area under the lower envelop of the quantile functions reported in Figure 5. The integral of the lower envelop is £ 1,256 per month, corresponding to 59% of the average in the population. Because some types can be quite small, one can also consider a 'robust' version of the lower envelope. Instead of considering only the worst-off type in each quantile, it considers enough types to have at least 10% of the quantile's population. The area under the robust lower envelope in Figure 5 is £ 1,493 , around 70% of the population average.

Figure 5: Roemer's lower envelope for the UKHLS (2009)

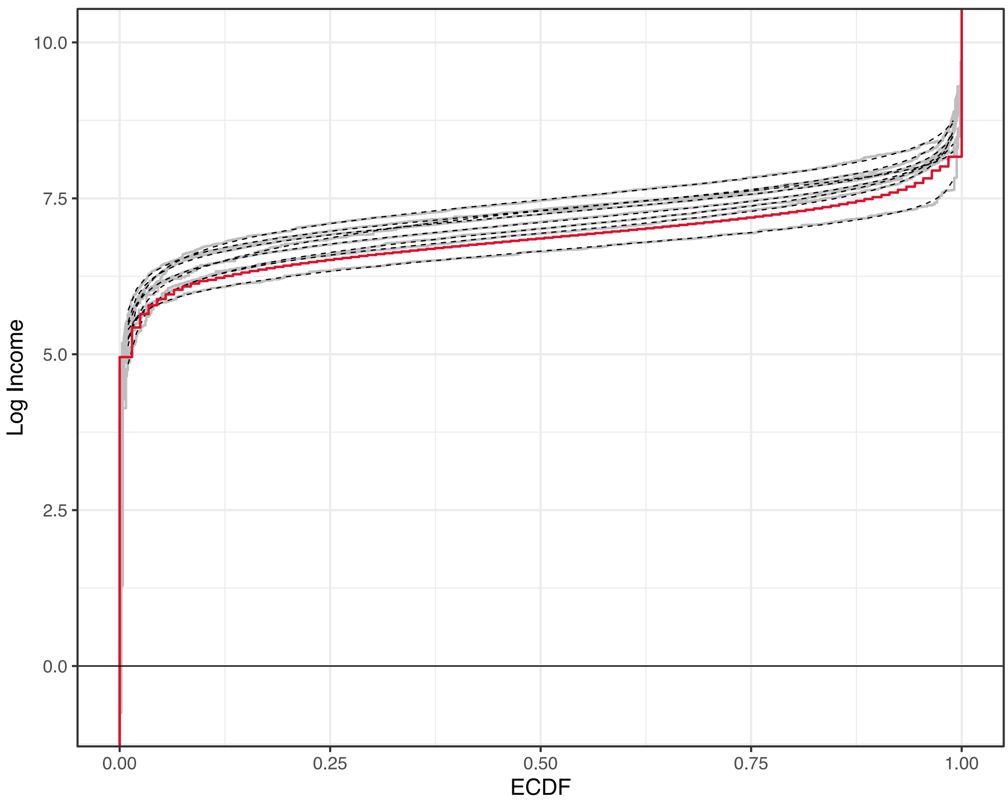


*Source: Own elaboration using data from UKHLS (2009).*

*Note: the red line identifies the robust version of the lower envelope.*

The results reported in Figures 2–5 are based on single trees – one conditional inference tree and one transformation tree. While these results are rather interesting and intuitive to interpret, they suffer from a significant limitation: high variance. Regression and classification trees are well-known to be low-bias, high-variance machine learning algorithms (James et al., 2013). The common solution to this limitation is to obtain predictions from a large collection of trees. Averaging results across many trees mitigates the large variance expected when predicting from a single tree. Conceptually, the use of bagging or random forests is highly attractive, as it acknowledges that a single type partition based on survey data is unlikely to accurately describe the full structure of opportunities in society. In contrast, using the average prediction from a large collection of varied structures may provide a better approximation of the data-generating process.

For this reason, in Figure 6 we compare the trajectories of ex-ante and ex-post estimates of inherited inequality in the UK from 2009 and 2019, by using random (CIT) forests and bagging of TrTs. All estimates are available in Table A2 in Appendix. Although the exact method of aggregating information from single-tree estimates differs between the ex-post and ex-ante approaches, the general approach is as follows: for all waves, we draw 200 samples of the minimum sample size (4,700 observed in 2019). For each sub-sample, we estimate ex-ante and ex-post IOp and calculate the average and standard deviation of these estimates across the 200 re-samplings. The bounds reported in Figure 6 are obtained by subtracting and adding 1.96 times the standard deviation to the mean estimates. These bounds are meant to provide a sense of the variability of the estimates but should not be strictly interpreted as confidence bounds (and have zero variability for the sample in which N=4,700).[29]

What emerges is a picture of relative stability, with perhaps some evidence of a slight decline in ex-ante inherited inequality. For both series, estimates range between 7 and 9 Gini points, or slightly less than one third of total inequality estimated in the same samples.

[29] A minimum common sample size is adopted to prevent differences in sample sizes to affect the depth of the trees and hence the level of IOp estimated. See Brunori, Ferreira, and Salas-Rojo (2023) for a discussion of the comparability of inherited inequality estimates. While the TrTs for the last wave are exactly the same for all iterations, random forests are based on random selection of variables and observations at each split and do show some variability even when estimated on the exact same sample.

Figure 6 Ex-ante and ex-post inherited inequality in the UK: 2009-2019.

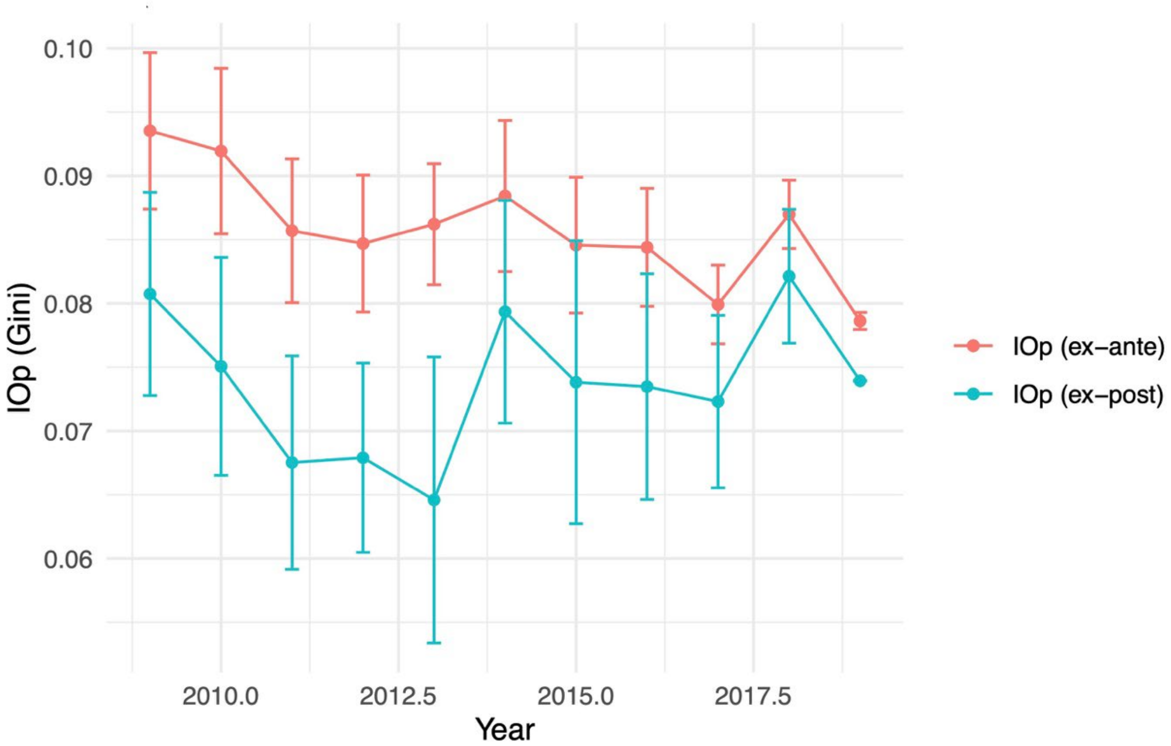


*Source: Own elaboration using data from UKHLS (2009). Note: ex-ante IOp is the average of 200 conditional inference random forests, average results produced with 200 CIT are reported in Table A2 in Appendix. Ex-post IOp is the average of 200 TrT (see Brunori, Ferreira, Salas-Rojo (2023) for details of the statistical procedure adopted).*

Finally, in Figure 7, we show the evolution of both the van de Gaer social objective function, $W_V$ (defined in Eq. 4) and the Roemerian social objective function, $W_R$ (defined in Eq. 5). Recall that the former is simply the expected income of the worst-off type, while the latter corresponds to the area below the lower envelope of the quantile functions. To ensure robust measures, we take two steps: first, we aggregate the worst-off types to cover at least 10% of the population; second, we average the results over 200 draws of the minimum common sample size without replacement. As a benchmark, we also plot the trend of the average income in the same samples.

Income is monthly equivalized disposable income expressed in 2015 CPI. Figure 7 shows a stagnant trend in the average income across waves. The trend in average incomes is consistent with what shown by Fisher (2023): we observe a limited increase in household incomes across wave. However, when considering the specific income definition and to some extent the sample selection a stable trend emerges (see Table A1 in Appendix). $W_R$ and overall mean income move in parallel in absolute terms, translating into an improvement of the objective function in relative terms, with the area below the lower envelope increasing from 70% in the first wave to 76% in 2019. $W_V$ also moves in parallel until 2018, but a drop in the last wave causes it to diverge overall, with the ratio between the two falling from 85% in the first wave to 78% in the last. $W_V$ and $W_R$ are almost identical in the last wave: this is always the case when one type, representing at least 10% of the population, is stochastically dominated by all others types, then the integral of the quantile function is exactly the mean of the type. In 2019 a single type (made of non-white respondents with a mother with no more than primary education) contains 11.4% of the population and its distribution is dominated by the distribution of all other types (see figure A.1 and A2 in Appendix).

Figure 7: Evolution in mean income and two opportunity-egalitarian social objective functions, UKHLS (2009-2019)

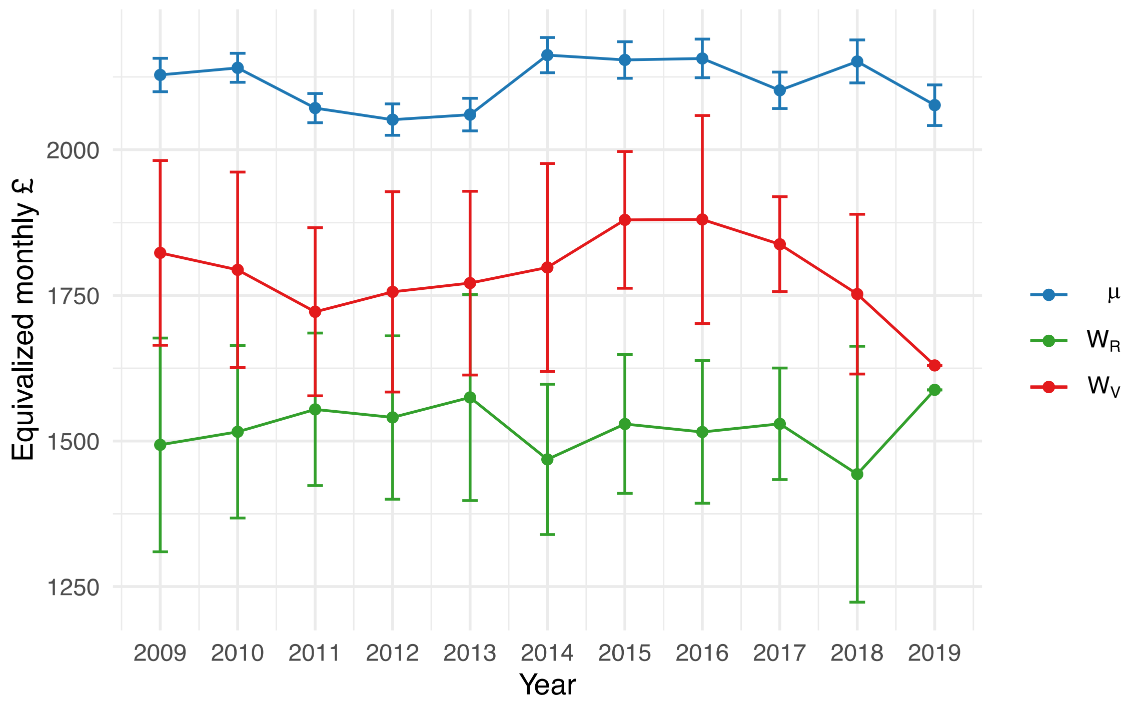


*Source: Own elaboration using data from UKHLS (2009). Note: $W_R$ is the area below the robust version of the lower envelop (Figure 5) calculated across TrTs estimated on 200 subsamples (N=4.700). $W_V$ is the average income of individuals belonging to the worst off types (where worst-off types are defined as types with lowest expected income populated by eat least 10% of the population). $\mu$ is the average income point estimates and 95% confidence interval is instead calculated on the entire sample.*

## 5. Aiming for opportunity-inclusive growth

As the foregoing discussion around Figure 7 illustrates, the social objective functions proposed by van de Gaer and Roemer, respectively $W_V$ and $W_R$ are useful indicators and can be tracked over time. But in a Handbook examining the links between distribution and growth, it is hard to be satisfied with such a static formulation of the problem, which avoids any intertemporal trade-offs a society might have to make if it cares about long-run outcomes. In addition, the social objectives described in Equations 4 and 5 suffer from (at least) the following shortcomings:

I. They are completely silent about non-income (or non-economic) aspects of justice and well-being and, therefore, about the kinds of policies that may be deployed in seeking to achieve the target allocation. In particular, they are silent on whether there are individual rights and freedoms which must be preserved or enhanced as a matter of priority, ranking more highly in the hierarchy of social objectives than the allocation of material resources.[30] It is possible, for example, that policies that might be quite effective in reducing inherited inequalities over the course of generations – such as prohibiting endo-type marriages, or even requiring high degrees of inverse assortative mating – may be unacceptable in a society that values certain personal freedoms. It is also possible that prohibiting certain forms of association (e.g. among industrialists or bankers) or certain forms of speech (e.g., against these objective functions for society), might assist in attaining the target allocation, and yet they too may not be tolerable in a free society.

II. A longer-term horizon is useful not only for bringing intertemporal substitutions – such as saving today to enjoy a higher standard of living tomorrow – into sharp relief. It is also important to ensure that future generations are taken into account when maximizing the wellbeing of the worst-off types, so that climate and environmental constraints are respected.

III. As noted in Section 3, there may be arguments against society tolerating certain levels of deprivation, such as extreme poverty, regardless of whether individuals find themselves in those situations through the exercise of (low) effort or (poor) responsibility decisions.

Building on what Bourguignon, Ferreira and Walton (2007) called the "equitable development policy", one may seek to incorporate these considerations into a more dynamic (and constrained) version of Eq. 5, such as the one given in Eq. 7:

$$\max_{\phi \in \Phi} \int_t^{\infty} e^{\delta(t-s)} \int_0^1 \min_{T_i} F_{is}^{-1}(p, \phi | C_i) dp ds \\ s.t.\ x_{ls} \geq z_s, \forall l, s \qquad (7)$$

In Eq. 7, notation is as before, with the following addition. *s* denotes (continuous) time and *t* is the initial moment, when the optimization takes place. $\delta, (0 \leq \delta \leq 1)$, is the discount rate: the higher it is, the lower the value placed on future periods (or generations) compared to the present. $z_s$ is a socially determined poverty or deprivation line at time s. $\phi$ are

[30] Such a hierarchy is present, for example, in Rawls's (1971) *A Theory of Justice*, which has helped shape, in one way or another, most contemporary visions of social justice. Indeed, Rawls's Principle of Equal Liberty takes precedence over his Equality Principle, which is more narrowly focused on the allocation of 'primary goods'.

government policy constellations (meaning a specific value of the full vector of policy variables available to the government). The set $\Phi$ is the set of all feasible and socially permissible policies.

Let us briefly consider how such a social objective function might address the shortcomings of $W_V$ and $W_R$, listed above. First, the introduction of time, with a discount rate which may be chosen to represent the time or generational preferences of the society at hand, clearly brings dynamic considerations to the fore. The economy is not static. If agents save, invest, and innovate, economic growth may take place, raising standards of living. Of course, there may be many alternative kinds of growth, with different drivers, different sectoral compositions, and different distributional consequences (or different growth incidence curves).[31] The maximand in (7) mandates policymakers to prioritize those forms of growth-cum-redistribution which would reach the highest possible present discounted value (first integral) for the stream of average incomes in the lower envelope of type-conditional quantile functions (second integral), as this average evolves over time.

The introduction of time, and the infinity upper-bound of the first integral also mandate a concern not only with our future selves, but also with our children and the generations that will follow, which addressed shortcoming II above. This concern will obviously weaken as the discount rate rises. In the limit, for policies with very long-term implications, it may well be defensible to set $\delta \to 0$.

Concerns I and III are addressed by each of the constraints to the problem. Beginning with the simplest, concern III with the abhorrence of extreme deprivation regardless of responsibility, is addressed by the constraint requiring that no incomes ever be allowed to fall below some socially-determined poverty line $z_s$. The subscript *s* indicates that the poverty line may evolve over time, potentially growing as an economy becomes more affluent. The choice of $z_s$ may also be seen as a kind of relative weight placed on social aversion to poverty – regardless of cause – and aversion to inherited inequality or inequality of opportunity. For a given state of technology and other feasibility constraints, raising $z_s$ leaves fewer resources and more limited policy space for pushing up the lower envelope of quantile functions (above $\min_{T_i} F_{is}^{-1}(p, \phi|C_i) = z_s$).

Concern I, which is perhaps the most important concern a political philosopher would have with respect to a such a simple, "economicist" maximization programme, is synthetically addressed in the restriction that permissible policy vectors should belong to a set $\Phi$, inclusion in which reflects not only technological feasibility and the usual participation, incentive compatibility, and administrative capacity constraints (which, together, determine a budget

[31] See Ravallion and Chen (2003), who introduce the very useful concept of growth incidence curve (GIC), and Ferreira (2012) for a discussion of how the GIC mediates the relationship between growth, changes in poverty, and changes in inequality.

constraint) that apply to taxation and redistribution, but also a set of individual rights and freedoms that the state cannot violate in pursuing its economic objectives.

These rights and freedoms may impose both requirements and restrictions on state actions. If the right to the freedom of association is included, then eliminating unions may not be possible. If there are restrictions on how many hours workers may be required to work per week, then planning for labour supply is correspondingly restricted. If a universal right to public education is prescribed, then public spending to provide it is mandated, whether or not it would be optimal in an unconstrained optimization. And so on.

Simply writing $\phi \in \Phi$ is obviously an empty shell in the absence of decisions about the list of rights and freedoms that $\Phi$ should contain, and that list can make an enormous difference to the state's ability to pursue its optimization problem. If, for example, the right to all fruits of one's labour is enshrined in absolute terms, then labour income taxation may be ruled out. If there is a prescribed right to fully bequeath wealth to one's descendants, then the taxation of inheritance may be excluded, and so on. The extent of and limits to personal rights and freedoms have long been debated, e.g. between John Rawls and Robert Nozick, and we do not propose to revisit that issue here. We suggest that the composition of $\Phi$ should be debated democratically (and periodically), and then respected. But we believe it is nonetheless important to be explicit about the rightful existence of such limits on the state's decision-making powers, and on the links between the rights one accords to citizens and the expectations one places on public action. At its most elementary, if society establishes rights that curtail most of the government's revenue raising ability, then it cannot simultaneously establish universal rights to very costly services. Politics and philosophy have the dominant role in setting up the "constitution" within $\Phi$. But they are not immune to arithmetic.

## 6. Meritocracy and IOp

Among the rights individuals may be granted by such a set of rules are certain rights to fair process. Some of these may govern elections of public officials, others the procedures to be followed by the justice system, and others yet the powers and limits to the regulation of markets. Some might concern the fair process for allocating scarce goods which, for one reason or another, society decides not to entrust (or entrust entirely) to markets. One such fairness criterion is equality of opportunity itself, with an emphasis on eliminating the role of inherited circumstances in the allocation of the scarce goods.[32]

But there are other criteria with normative merit, which have been and continue to be used. One of them is need, which is the dominant criterion used, for example, in the allocation of

[32] Recall that, since $H \subset \Gamma$, equality of opportunity (Equation 3) implies satisfaction of the Weak Principle of Compensation and Equation 3'. So, we revert to referring to the stronger, more demanding criterion in this Section.

organs for transplant. In many health and hospital systems, hearts and livers that become available are allocated to those most in need, without regard to whether, say, the children of more and of less educated parents are represented equiproportionately among recipients. This may well be a social preference, even as society would like to pursue the programme in Equation 7 to determine its income or wellbeing growth path. In other words, even as equality of opportunity may play a guiding normative role for the allocation of economic resources such as incomes, wealth, or consumption expenditure over time, it is certainly possible for society to choose other criteria – such as need – in specific realms of policy and allocation, by writing such procedures into the set $\Phi$.

Another criterion commonly mentioned is meritocracy and, confusingly, this one is often conflated with equality of opportunity – the idea being that one should provide equal conditions "at the start of the race" (equal opportunities), and then let the winners enjoy their rewards (merit). In the remainder of this section, we suggest that this view is erroneous and that what most people think of as meritocracy is not consistent – and indeed will generally clash with – equality of opportunity as defined above and in the literature.

To do so, we insert a minor modification into Equation (1). We define productivity $\pi = \pi(C, e)$ as a function of circumstances and efforts, and outcome y as a function of productivity: $y = h(\pi), h' > 0$.[33] This is, of course, perfectly consistent with Eq. (1), which can be rewritten as:

$$y = h\big(\pi(C, e)\big) = g(C, e)$$

But now that circumstances and efforts affect incomes through productivity, it is easy to see how a new allocation criterion might arise, giving rise to the two following definitions:

a. Equality of opportunity: $F(y|C_i) = F(y|C_k), \forall\, C_i, C_k \in \Gamma$
b. Meritocracy: $y_l = y_m \Leftrightarrow \pi_l = \pi_m, \forall l, m \in N$

Definition (a) for equality of opportunity is the same as before. It was introduced in Eq. (3), and rewritten for the subset of inherited circumstances in Eq. (3'). Under the full assumptions of the EOp theory, it states that individuals should attain the **same reward** when they exert the **same (relative) degree of effort**. (Under the weaker assumptions of inherited inequality, it states that we should not observe differences in the distributions of rewards across groups defined by inherited characteristics.) Definition (b) for meritocracy states that individuals should attain the **same reward** when they have the **same (expected) productivity**.

To illustrate the sharp difference between the two, consider the example of a society consisting just of two types, the Advantaged (A) and the Disadvantaged (D). Recalling that $L$ denotes the total population size, and with obvious subscripts, note that $L_A + L_D = L$.

[33] Under uncertainty, let $\pi$ denote expected productivity.

Further, let the conditional distributions of both productivity ($G(\pi)$) and income ($F(y)$) for the Advantaged first-order stochastically dominate the corresponding distributions for the Disadvantaged. Note that the monotonicity of $h(\pi)$ implies that:

$$p = G_i(\pi) = F_i(y), i = A, D$$

Now consider the problem of allocating $L_j(< L)$ scarce goods, which are vacancies or positions in coveted public schemes. These could be, for example, scarce places in kindergartens, universities, or public sector jobs.

If the scarce $L_j$ positions are to be allocated according to the Equality of Opportunity criterion, then definition (a) (plus the Pareto principle) require that those exerting the greatest degree of effort in each type should be given the positions until they run out, starting from the 'hardest working' and moving down, so that the last people to be given jobs in the two types exert the same degree of effort:

$$L_J^{EOp} = (L_A + L_D)(1 - p^*)$$

Graphically, the EOp selection process can be represented by Figure 8, which used the same CDFs as in Figure 1, for illustration only.

Figure 8: Allocation of scarce goods under the EOp criterion

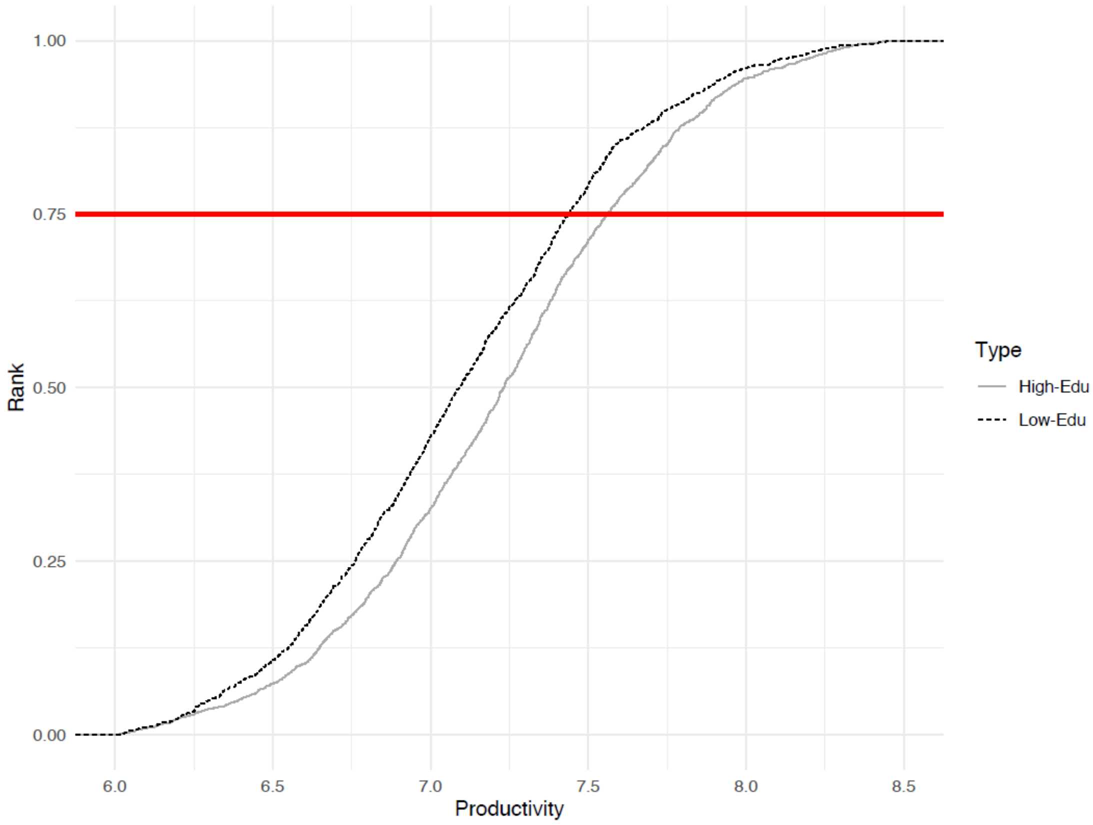

By contrast, if the scarce $L_j$ positions were to be allocated according to the Meritocracy criterion, then definition (b) (plus the Pareto principle) require that those with the greatest expected productivity in each type should be given the positions until they run out, starting from the most productive and moving down so that the last people to be given jobs in the two types have the same productivity:

$$L_J^M = L_A\big(1 - G_A(\pi^*)\big) + L_D\big(1 - G_D(\pi^*)\big)$$

Graphically, this meritocratic selection process can be represented by Figure 9:

Figure 9: Allocation of scarce goods under the Meritocratic criterion

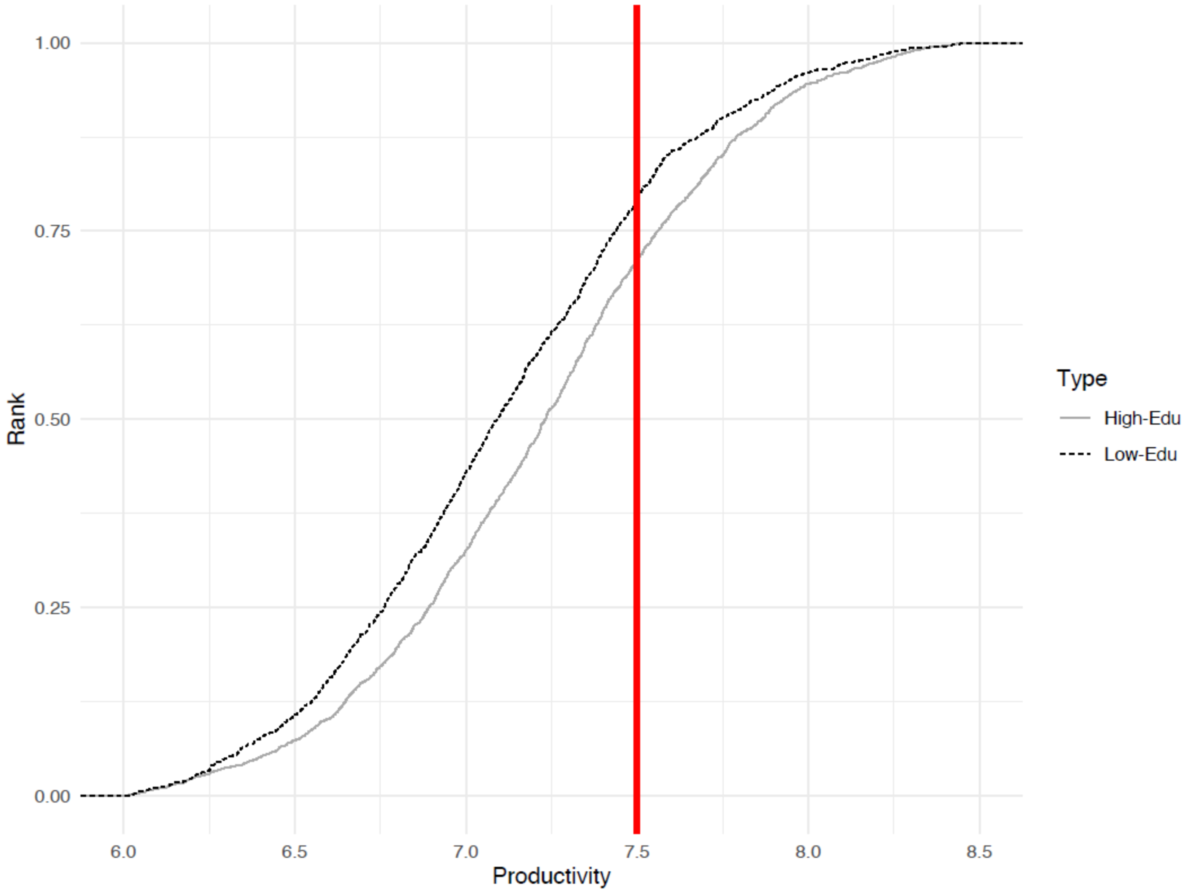


Clearly according to these standard definitions, equality of opportunity and meritocracy yield very different solutions to this allocation problem. The EOp solution awards vacancies to the same proportions of the two types, that is: it divides the available vacancies into the two types in accordance with the groups' population shares. If the disadvantaged group is less productive, as in this example, then the marginal person from this group to be awarded a vacancy will be less productive than the marginal person from the advantaged group, as can be seen in Figure 9. Conversely, the meritocratic solution equalizes productivity on the margin

– since that is its criterion – but will therefore allocate a greater proportion of vacancies to the advantaged group.

Of course, meritocracy is not the worst regime under which the Disadvantaged could live. In a Discriminatory world, where institutions unfairly favour the relatives, children and other members of the advantaged group, one could have an allocation such as that in Figure 10 below. In that situation, the marginal person from the advantaged group to be given a slot is less productive than the marginal disadvantaged candidate, and the proportion of scarce positions going to the disadvantaged is even lower than under meritocracy. Given the prevalence of racial, gender, and other forms of discrimination around the world, this might explain why the disadvantaged often asked only for meritocracy: let the "best" (as in, "expected to be most productive at the job") person win, regardless of her gender of the colour of her skin.

Figure 10 Allocation of scarce goods under Discrimination

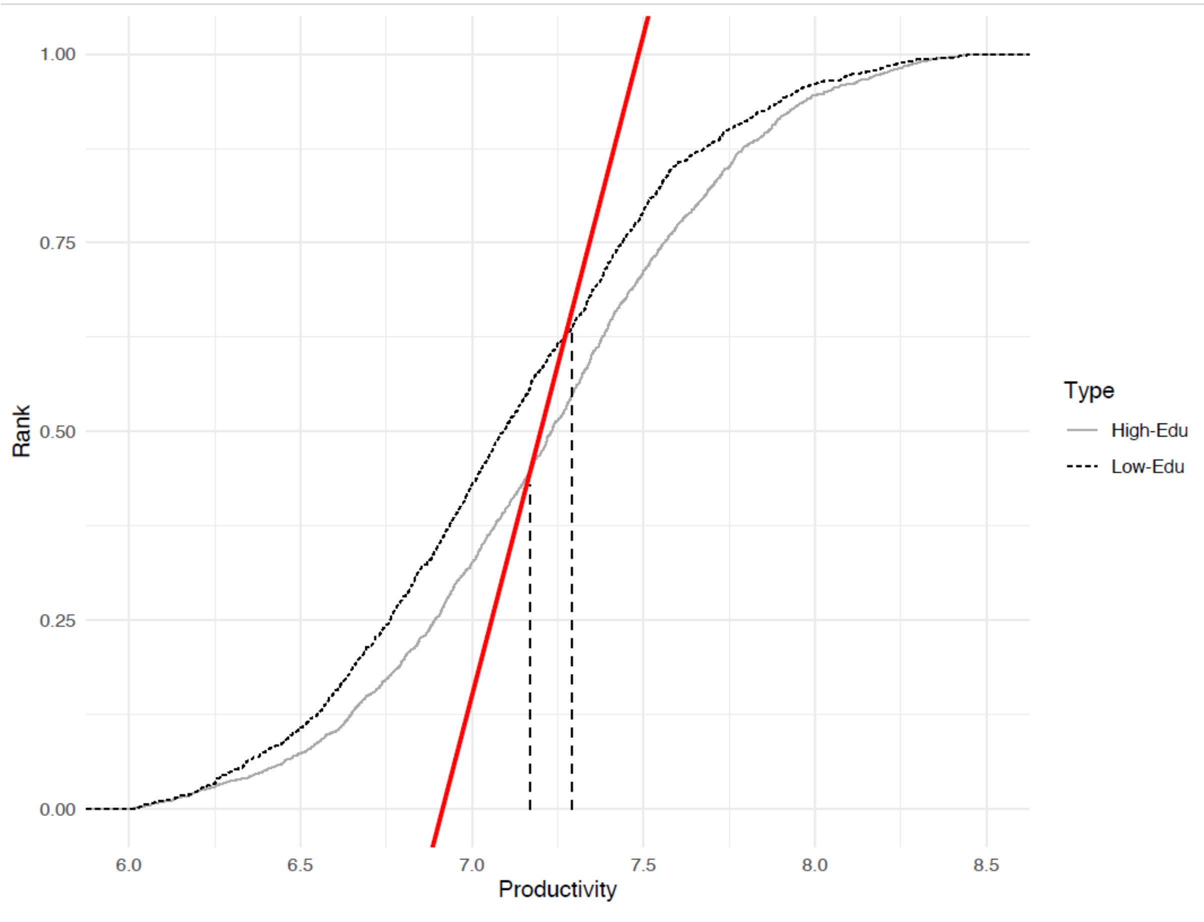


In Section 5 we proposed a modified dynamic version of the well-known Roemerian allocation rule arising from an EOp perspective (while upholding the Pareto principle and thus avoiding the levelling down objection). We find it compelling that one can arrive at essentially the same idea from very weak normative requirements, such as the Weak Principle of Compensation and Pareto. As a general principle to guide economic policy in pursuit of the 'right kind' of economic growth, we favour the EOp perspective and suggest that it will generally dominate Meritocracy in terms of fairness.

But this does not necessarily mean that there is no role for Meritocracy in a fair society. Indeed, certain roles can be justified by two separate arguments. First, it is possible that society decides that there are certain allocation problems where meritocracy is preferred on the basis of process fairness. Consider, for example, a competition to allocate a limited number of brain surgeon positions at a public hospital, or the appointment of a few high-ranking civil servants in the Central Bank. Or yet the selection of athletes to represent the country in an international competition, such as the Olympic Games. For these purposes, society may (or may not) decide that meritocracy is the fair allocation process, without detriment of the overall programme in Eq. 6, and may include stipulations to this effect in the set $\Phi$.[34]

There are other contexts, however, where even for the allocation of limited positions, an EOp criterion is chosen over meritocracy. A good example would be the allocation of a limited number of pre-school vacancies in a high-quality programme. There are many arguments for providing those positions disproportionately in favour of disadvantaged children, who can count on fewer resources at home to aid their development. Consider the Ecuadorian example in the first paragraph of this chapter.

Second, and separately from considerations of process fairness imposed as restrictions on Eq. 7, there may be "instrumental" reasons, internal to the optimization process itself, for meritocracy to prevail in certain domains. Those occur where the gain in efficiency arising from allocating opportunities based on expected productivity are large enough to outweigh the costs of not targeting the disadvantaged directly. The choice of a (well-regulated) market economy as a system to allocate most goods and services is itself such a choice. Profit-making firms have incentives to hire meritocratically (which is not to say that they always do). In a centrally planned system, one might allocate good jobs directly to the disadvantaged, hoping for rapid compensation of inherited injustices. Yet, recall that the state operates subject to individual participation and incentive compatibility constraints as well and, in the long run, it may be that the children of the disadvantaged will be better off if markets are allowed to operate in a (large?) number of domains.

Nor is it necessarily the case that individual public allocation problems must be solved exclusively according to either EOp or meritocracy. One could imagine a hybrid system that combines – or weighs – elements of the solutions in Figures 8 and 9, to arrive at something like Figure 11 below:

[34] Similarly, there may be other domains where other fairness criteria may be chosen, such as in the case of the need criterion for the problem of allocating organs for transplant.

Figure 11: Allocation of scarce goods under a combination of the EOp and Meritocratic criteria

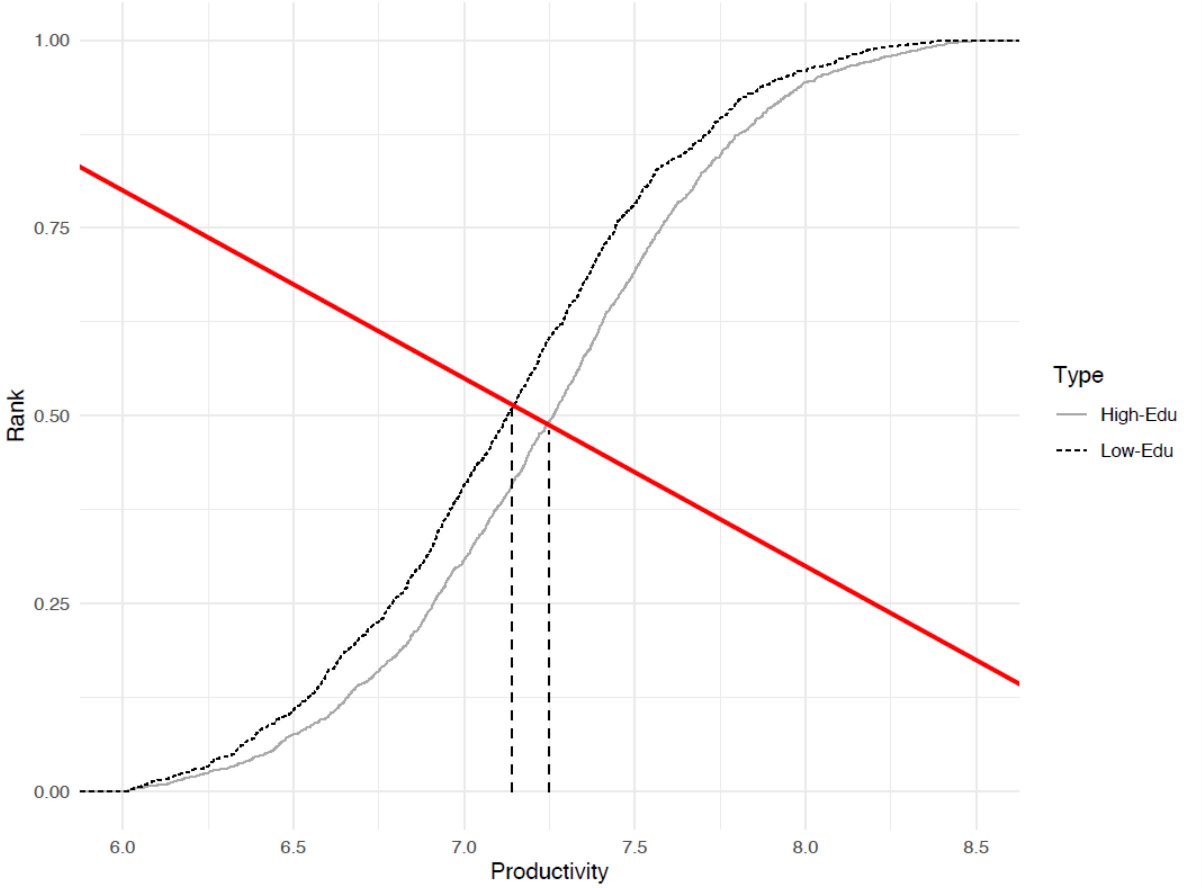


In such a hybrid solution, the marginal productivity of the selected candidates from the advantaged group are still higher than that of the disadvantaged, but not by as much as in pure EOp. And the proportion of vacancies going to the disadvantaged is still lower than that going to the advantaged, but not by as much as in pure meritocracy. Situations in which static efficiency (allocate good to the person likely to use it best) must be traded-off against dynamic efficiency (e.g., provide role models today so that highly productive members of the disadvantaged groups will apply tomorrow) may be particularly well-suited to these forms of compromise. Possibly, these solutions might be appropriate for selection processes that take place somewhere in between the kindergarten and the Central Bank entrance exam, such as undergraduate college admissions. Indeed, affirmative action policies used by universities in many parts of the world, where entrance requirements are lower for members of disadvantaged groups, are specific examples of this hybrid approach.

## 7. Conclusions

After briefly discussing some of the main reasons why the concept warrants serious consideration by economists and other social scientists, this chapter reviewed the dominant economic model of inequality of opportunity. This model relies on a classifiability assumption (that all determinants of advantage can be classified as either 'circumstances' or 'efforts'), and on two key normative principles: the Principle of Compensation and the Principle of Reward. Historically, the model has been used for two main purposes: proposing optimal policy or allocation rules that meet certain societal objectives; or measuring the extent of inequality of opportunity in a given society. While we did not provide a comprehensive compilation of the multiple proposals arising from different versions of the Compensation and Reward principles, we did briefly review the original allocation proposals by van de Gaer (1993) and Roemer (1993), as well as two common approaches to measurement, namely between-types and within-tranches inequality.

Because the classifiability assumption and the Principle of Reward have attracted some resistance from different quarters, for both epistemic and normative reasons, we then suggested a closely related concept that yields very similar allocation rules and measures of unfair inequality, while resting on a much more economical – and possibly less controversial – normative foundation. Dispensing with both the classifiability assumption and the Principle of Reward, the concept of inherited inequality relies only on a weaker version of the Principle of Compensation (WPC), namely that in a fair society, incomes (or whichever advantage variable one chooses to focus on) be distributed independently of a set of inherited characteristics, which are given at birth. This much weaker requirement yields allocation rules and measures of inequality which are closely analogous to those of IOp, differing only in that the set of inherited characteristics is, in general, a subset of the set of circumstances in the theory of IOp. As it happens, in many empirical applications the two sets actually coincide, in which case so do the measures of inherited inequality and IOp.

The chapter then reviewed recent advances in the measurement of IOp and inherited inequality, focusing on novel solutions to the model selection problem that has plagued the literature. The WPC requires that, in a fair situation, inherited characteristics should not be predictive of incomes. Therefore, the extent to which those characteristics are, in fact, predictive of income is a good measure of the extent of inherited inequality. The model selection challenge arises from the fact that the optimal specification of the prediction model must trade off a (downward) omitted-variable bias against an (upward) overfitting bias. We reviewed recent data-driven, machine learning approaches to solving this problem, including conditional inference trees, random forests, and transformation trees. We compared these different methods in the context of an application to the UK, between 2009 and 2019.

Although that empirical application allowed us to track the actual levels of the social objective functions proposed by van de Gaer (the 'min of means') and Roemer (the 'mean of mins'), further consideration of these objectives revealed some conceptual shortcomings, including the fact that both formulations are purely static, ignore non-economic objectives and constraints and, in their pure form, do not incorporate unconditional objections to absolute deprivation. We suggested an alternative formulation, which explicitly incorporates economic growth and is a close variant of the proposal by Bourguignon, Ferreira and Walton (2007). We argue that this 'equitable development policy' provides a clearer statement of what economic growth should be for, and what constraints should be taken into consideration when promoting it. It is also, therefore, a richer basis with which to assess the 'normative quality' of actual episodes of economic growth.

Finally, we turned to a comparison of the normative criteria for the allocation of scarce goods (such as educational or professional vacancies) that are implied by the equality of opportunity perspective, and by meritocracy. For very intuitive definitions of both approaches, we showed that the implications differ considerably between them. In a setting where more advantaged groups are better able to equip themselves to become productive, meritocracy will typically lead to an under-representation of the disadvantaged in coveted positions (but higher overall productivity), whereas EOp would lead to a better representation of the disadvantaged (but lower aggregate productive efficiency).

Contrary to popular perception, equality of opportunity is not the same as meritocracy. It is more radically egalitarian in its prescriptions, and we endorse it as the main anchor of our proposed social objective for growth and development. That said, we suggest that there may be specific public policy problems within that broader program where meritocratic criteria may play a useful role – either for efficiency reasons or for fairness-in-process reasons.

We hope that this review of the recent literature, including some novel suggestions, convinces the reader that any discussion of the relationship between income distribution and economic growth should consider the part of inequality which is inherited at birth – transmitted from generation to generation – both because it is positively important and because it is central to determining what kind of growth and development our societies want to pursue.

## References


Arneson, Richard (1989): “Equality of Opportunity for Welfare”, *Philosophical Studies*, 56: 77-93.

Balcázar, C. F. (2015). Lower bounds on inequality of opportunity and measurement error. *Economics Letters*, 137, 102-105.

Bourguignon, François, Francisco H.G. Ferreira, and Marta Menendez (2007): “Inequality of Opportunity in Brazil”, *Review of Income and Wealth*, 53 (4): 585-618.

Bourguignon, François, Francisco H.G. Ferreira, and Michael Walton (2007): “Equity, efficiency and inequality traps: A research agenda”, *Journal of Economic Inequality*, 5 (2), 235-256.

Brunori, P., Davillas, A., Jones, A. M., & Scarchilli, G. (2022). Model-based recursive partitioning to estimate unfair health inequalities in the United Kingdom Household Longitudinal Study. *Journal of Economic Behaviour & Organization*, 204, 543-565.

Brunori, P., F. H. G. Ferreira and P. Salas-Rojo (2023). “Inherited inequality: a general framework and an application to South Africa”. International Inequalities Institute Working Paper 107, London School of Economics.

Brunori, Paolo, Paul Hufe and Daniel Mahler (2023): “The Roots of Inequality: Estimating Inequality of Opportunity from Regression Trees”, *Scandinavian Journal of Economics*, 125 (4), 900-932.

Brunori, Paolo, Vito Peragine, and Laura Serlenga. 2018. “Upward and Downward Bias When Measuring Inequality of Opportunity.” *Social Choice and Welfare*, 52, 635-661.

Cappelen, Alexander, Erik Sorensen and Bertil Tungodden (2010) “Responsibility for what? Fairness and individual responsibility”. *European Economic Review* **54** (3): 429-441.

Checchi, Daniele, and Vito Peragine (2010): “Inequality of Opportunity in Italy”, *Journal of Economic Inequality*, 8 (4), 429-450.

Cohen, Gerry A. (1989): “On the Currency of Egalitarian Justice”, *Ethics*, 99, 906-944.

De la O, Ana, Cecilia Rossel and Pilar Manzi (forthcoming): “The Political Economy of Opting Out from Public Services in Latin America”. *Oxford Open Economics*.

Dworkin, R. (1981). What is equality? Part 2: Equality of resources. *Philosophy and Public Affairs*, 10(4), 283-345.

Fernández, R. C. Pagés, M. Székely and I. Acevedo (forthcoming): “Education inequalities in Latin America”, *Oxford Open Economics*.

Ferreira, F. H. G. (2012) “Distributions in Motion: Economic Growth, Inequality, and Poverty Dynamics”. Chapter 13 in Philip Jefferson (ed.) *Oxford Handbook of the Economics of Poverty*. Oxford: Oxford University Press.

Ferreira, F. H. G. (2022): “Comment: Not all inequalities are alike”, *Nature,* 606: 646-649.

Ferreira, F. H. G. (2023). “Is there a ‘new consensus’ on inequality?” III Working Paper (101). International Inequalities Institute, London School of Economics, London, UK.

Ferreira, F., & Gignoux, J. (2008). “The measurement of inequality of opportunity: Theory and an application to Latin America.” World Bank Policy Research Working Paper 4659.

Ferreira, F., & Gignoux, J (2010). “Inequality of Opportunity for Education: Turkey”. Chapter 6 in R. Kanbur and M. Spence (eds.) *Equity in a Globalizing World, Commission on Growth and Development*, Washington, (pp.131–56).

Ferreira, F., & Gignoux, J. (2011). The measurement of inequality of opportunity: Theory and an application to Latin America. *Review of Income and Wealth*, 57, 622-657.

Fisher, P., & Hussein, O. (2023), Understanding Society: the income data. Fiscal Studies, 44, 377–397

Fishkin, Joseph (2014*): Bottlenecks: A New Theory of Equal Opportunity.* New York, NY: Oxford University Press.

Fleurbaey, Marc (1994): “On fair compensation”, *Theory and Decision*, 36, 277-307.

Fleurbaey, Marc (1995): “Three Solutions for the Compensation Problem”, *Journal of Economic Theory*, 65(2), 505-521.

Fleurbaey, Marc (2008): *Fairness, Responsibility and Welfare*, 1st Edition. Oxford: Oxford University Press.

Fleurbaey, Marc, and Vito Peragine (2013): “Ex ante versus ex post equality of opportunity,” *Economica* 80, 118-130.

Fleurabey, Marc and Schokkaert, Erik (2009): Unfair inequalities in health and health care. *Journal of Health Economics*, 28(1):73-90.

Harden, Kathryn Paige (2021): *The genetic lottery: why DNA matters for social equality.*” Princeton University Press. New Jersey.

Hothorn, T., K. Hornik and A. Zeileis (2006): “Unbiased recursive partitioning: A conditional inference framework”, *Journal of Computational and Graphical Statistics*, **15** (3): 651-674.

Hothorn, Torsten, and Achim Zeileis (2021). “Predictive Distribution Modeling Using Transformation Forests.” *Journal of Computational and Graphical Statistics*, 30: 1181-1196.

Hufe, P., Kanbur, R., & Peichl, A. (2022). Measuring unfair inequality: Reconciling equality of opportunity and freedom from poverty. *Review of Economic Studies*, 89(6): 3345-3380.

Hufe, P., Peichl, A., Roemer, J. Ungerer, M. (2017). Inequality of income acquisition: the role of childhood circumstances. *Social Choice and Welfare* 49, 499–544.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). An introduction to statistical learning: With applications in R. Springer.

Jusot, F., Tubeuf, S., & Trannoy, A. (2013). Circumstances and efforts: How important is their correlation for the measurement of inequality of opportunity in health? *Health Economics*, 22(12), 1470-1495.

Kanbur, R., & Wagstaff, A. (2016). How useful is inequality of opportunity as a policy construct? In K. Basu & J. E. Stiglitz (Eds.), *Inequality and growth: Patterns and policy* (International Economic Association Series). Palgrave Macmillan.

Kranich, Laurence (1996): “Equitable opportunities: an axiomatic approach”, *Journal of Economic Theory*, 71, 131-147.

Lefranc, Arnaud, Nicolas Pistolesi, and Alain Trannoy (2009): “Equality of opportunity and luck: definitions and testable conditions, with an application to income in France”, *Journal of Public Economics*, 93, 1189-1207.

Marrero, G., & Rodriguez, J. (2013). Inequality of opportunity and growth. *Journal of Development Economics*, 104, 107-122.

Niehues, J., & Peichl, A. (2014). Upper bounds of inequality of opportunity: Theory and evidence for Germany and the US. *Social Choice and Welfare*, 43(1), 73-99.

Nybom, M., & Stuhler, J. (2016). Heterogeneous Income Profiles and Lifecycle Bias in Intergenerational Mobility Estimation. The Journal of Human Resources, 51(1), 239–268.

OECD (2013): OECD Framework for Statistics on the Distribution of Household Income, Consumption and Wealth, available at: https://www.oecd-ilibrary.org/

Ok, Efe A. (1997): “On opportunity inequality measurement”, *Journal of Economic Theory*, 77, 300-329.

Pattanaik, Prasanta K., and Yongsheng Xu (1990): “On ranking opportunity sets in terms of freedom of choice", *Recherches Economiques de Louvain*, 56, 383-390.

Paxson, Christina and Norbert Schady (2007) "Cognitive Development among Young Children in Ecuador: The Roles of Wealth, Health, and Parenting," *Journal of Human Resources*, **42**(1).

Pew Research Center. (2012). For the public, it's not about class warfare, but fairness. Available at: https://www.pewresearch.org/politics/2012/03/02/for-the-public-its-not-about-class-warfare-but-fairness/

Ramos, X., & Van de Gaer, D. (2021). Is inequality of opportunity robust to the measurement approach? *Review of Income and Wealth*, 67(1), 18-36.

Ravallion, Martin, and Shaohua Chen (2003): "Measuring pro-poor growth", *Economics Letters*, Elsevier, 78(1), pages 93-99.

Rawls, John (1971): *A theory of justice*, Cambridge, MA: Harvard University Press.

Roemer, John (1993): "A Pragmatic Theory of Responsibility for the Egalitarian Planner", *Philosophy & Public Affairs*, 10: 146-166.

Roemer, John (1998): *Equality of Opportunity*, Cambridge, MA: Harvard University Press.

Sen, Amartya (1980): "Equality of what?" in S. McMurrin (ed.) *The Tanner Lectures on Human Values*, Salt Lake City: University of Utah Press.

Sen, Amartya (1985), *Commodities and Capabilities*, North-Holland, Amsterdam.

Van de Gaer, Dirk (1993): "Equality of opportunity and investment in human capital", Ph.D. Dissertation, Katholieke Universiteit Leuven.

Weymark, John A. (2003): "Generalized Gini indices of equality of opportunity", *Journal of Economic Inequality*, 1, 5-24.

World Bank (2005): *World Development Report 2006: Equity and Development*. Washington, DC: World Bank.

# Appendix

Table A1: descriptive statistics of the UKHLS samples (wave 2009-2019)

| Year | N | | Mean income | | Age | | Inequality (Gini) | |
|---|---|---|---|---|---|---|---|---|
| | Original | Sample | Original | Sample | Original | Sample | Original | Sample |
| 2009 | 14,389 | 9,531 | 2,052.2 | 2,128.4 | 42.52 | 42.53 | 0.3129 | 0.3091 |
| 2010 | 15,043 | 1,0982 | 2,058.6 | 2,140.7 | 42.73 | 42.74 | 0.2974 | 0.2961 |
| 2011 | 13,500 | 9,231 | 2,001.3 | 2,071.6 | 42.82 | 42.86 | 0.2870 | 0.2867 |
| 2012 | 12,491 | 8,367 | 1,981.4 | 2,051.8 | 42.88 | 42.94 | 0.2873 | 0.2884 |
| 2013 | 11,593 | 7,722 | 2,002.5 | 2,060.3 | 42.92 | 43.03 | 0.2893 | 0.2897 |
| 2014 | 11,693 | 7,771 | 2,081.8 | 2,162.4 | 42.93 | 43.06 | 0.2912 | 0.2921 |
| 2015 | 11,093 | 7,152 | 2,077.9 | 2,154.2 | 42.87 | 43.06 | 0.2886 | 0.2920 |
| 2016 | 10,328 | 6,543 | 2,087.6 | 2,156.8 | 42.83 | 43.03 | 0.2939 | 0.2949 |
| 2017 | 9,158 | 5,723 | 2,037.2 | 2,102.1 | 42.82 | 43.05 | 0.2838 | 0.2794 |
| 2018 | 8,569 | 5,270 | 2,060.6 | 2,151.6 | 42.92 | 43.16 | 0.2919 | 0.2942 |
| 2019 | 7,807 | 4,700 | 2,004.6 | 2,076.6 | 42.89 | 43.18 | 0.2806 | 0.2833 |

*Source: Own elaboration using data from UKHLS (2009). Note: the first columns refer to the complete sample while the right columns refer to the analytical sample with complete information about circumstances used in the analysis.*

Table A2: Equality and inequality of opportunity in the UKHLS (wave 2009-2019)

| | Ex-ante | | | | Ex-post | | |
|---|---|---|---|---|---|---|---|
| Year | # types | IOp CIT | IOp forest | $W_V$ | # types | IOp TrT | $W_R$ |
| 2009 | 7.15 | 0.0662 | 0.0927 | 1823.1 | 6.47 | 0.0808 | 1493.4 |
| 2010 | 8.01 | 0.0710 | 0.0919 | 1793.9 | 6.78 | 0.0749 | 1515.9 |
| 2011 | 7.70 | 0.0672 | 0.0857 | 1721.9 | 6.68 | 0.0673 | 1554.4 |
| 2012 | 6.42 | 0.0559 | 0.0847 | 1756.2 | 6.27 | 0.0679 | 1540.4 |
| 2013 | 7.82 | 0.0646 | 0.0862 | 1771.1 | 6.57 | 0.0652 | 1574.8 |
| 2014 | 7.69 | 0.0592 | 0.0884 | 1797.9 | 7.76 | 0.0794 | 1468.4 |
| 2015 | 7.81 | 0.0563 | 0.0846 | 1879.6 | 7.45 | 0.0738 | 1529.2 |
| 2016 | 7.12 | 0.0569 | 0.0844 | 1880.2 | 7.60 | 0.0728 | 1515.6 |
| 2017 | 5.96 | 0.0512 | 0.0799 | 1837.9 | 7.21 | 0.0724 | 1529.5 |
| 2018 | 8.08 | 0.0581 | 0.0870 | 1752.2 | 8.49 | 0.0821 | 1443.0 |
| 2019 | 8.00 | 0.0495 | 0.0786 | 1629.8 | 8.00 | 0.0741 | 1587.9 |

*Source: Own elaboration using data from UKHLS (2009).*

Figure A1: CIT for 2019

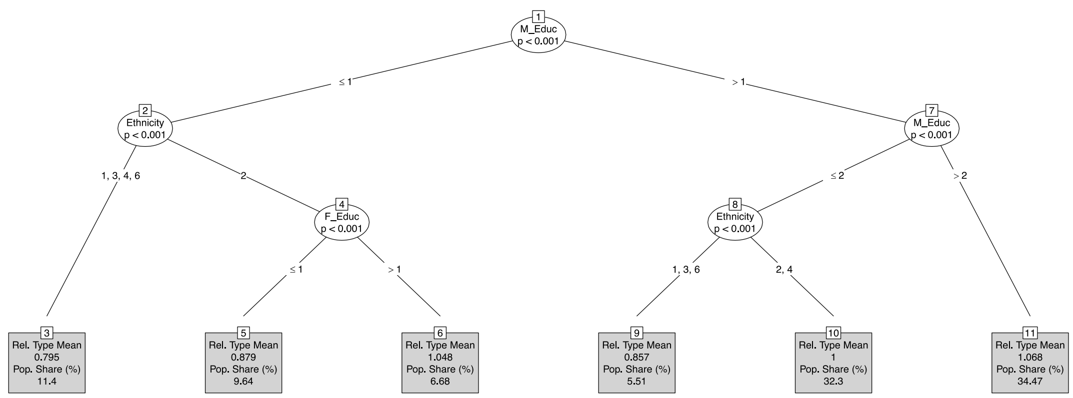


*Source: Own elaboration using data from UKHLS (2009).*

Figure A2: TrT for 2019

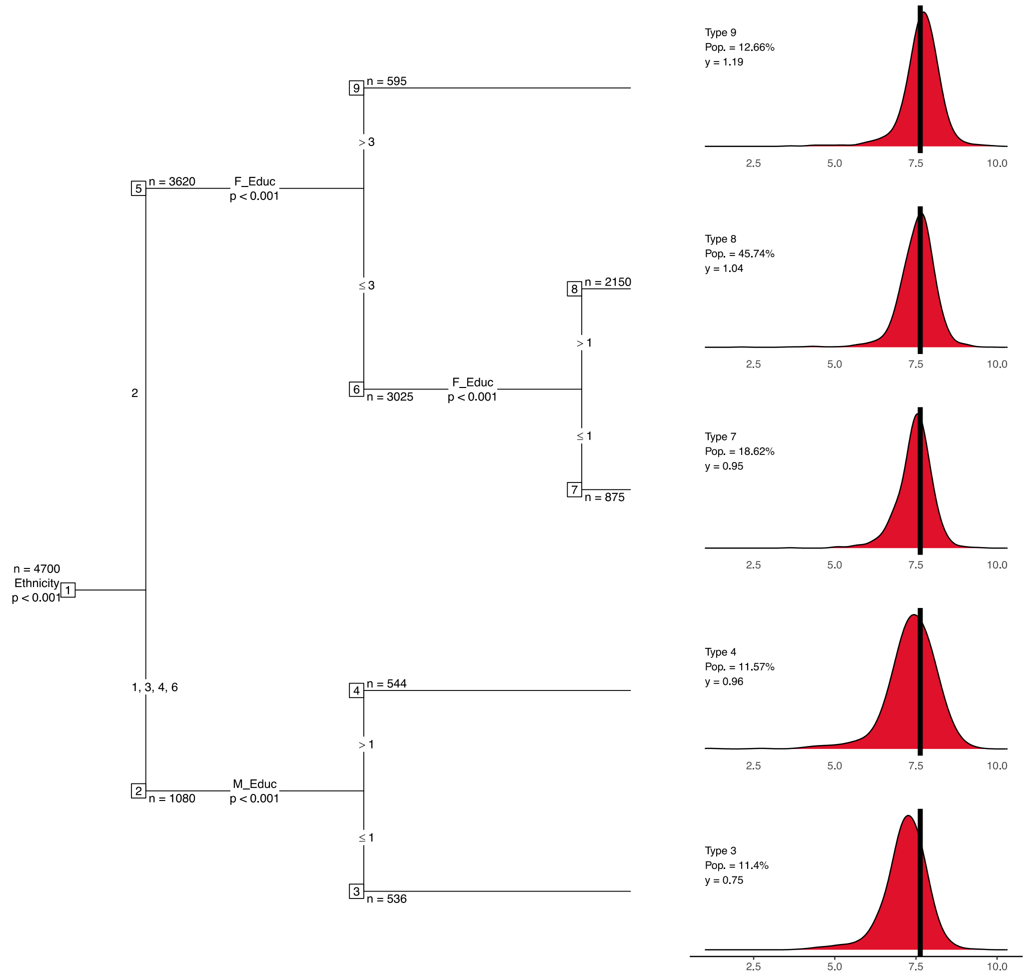


*Source: Own elaboration using data from UKHLS (2009).*

Figure A3: Quantile functions for TrT 2019

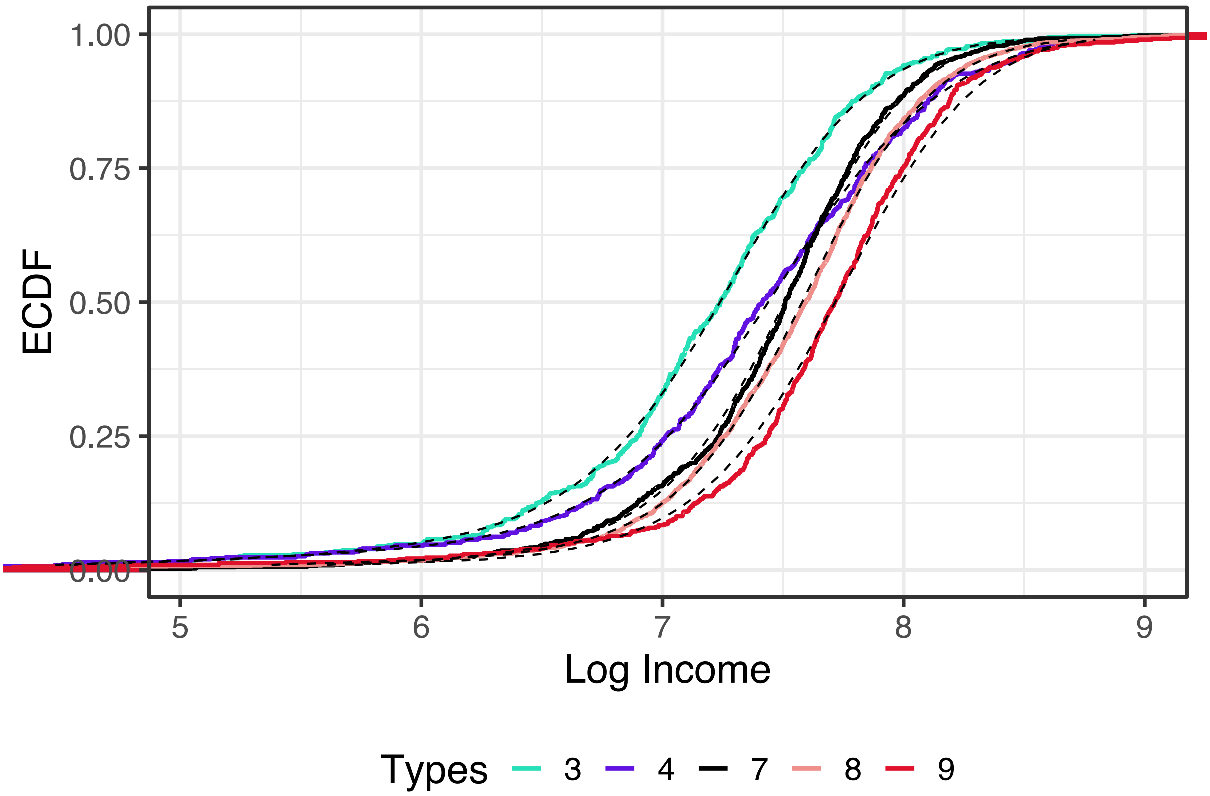


*Source: Own elaboration using data from UKHLS (2009).*