# Model-based Recursive Partitioning to Estimate Unfair Health Inequalities in the United Kingdom Household Longitudinal Study

Paolo Brunori [a,*], Apostolos Davillas [b], Andrew M. Jones [c], Giovanna Scarchilli [d]

[a] *London School of Economics, International Inequalities Institute & University of Florence - London School of Economics and Political Science, Houghton Street, London, WC2A 2AE, UK*
[b] *Department of Economics, University of Macedonia, Greece*
[c] *University of York, UK*
[d] *University of Trento & University of Modena and Reggio Emilia, Italy*

## ARTICLE INFO

## ABSTRACT

We measure unfair health inequality in the UK using a novel data-driven empirical approach. We explain health variability as the result of circumstances beyond individual control and health-related behaviours. We do this using model-based recursive partitioning, a supervised machine learning algorithm. Unlike usual tree-based algorithms, model-based recursive partitioning does identify social groups with different expected levels of health but also unveils the heterogeneity of the relationship linking behaviors and health outcomes across groups. The empirical application is conducted using the UK Household Longitudinal Study. We show that unfair inequality is a substantial fraction of the total explained health variability. This finding holds no matter which exact definition of fairness is adopted: using both the fairness gap and direct unfairness measures, each evaluated at different reference values for circumstances or effort.

## 1. Introduction

According to Fleurbaey and Schokkaert (2009), differences in health status can originate from either fair or unfair sources. They argue that unfair health inequalities are differences in health status determined by circumstances beyond individual control such as sex, ethnicity or socioeconomic background in childhood. Under this distinction a society that wishes to eliminate unfair health inequality should compensate individuals suffering a poorer health status due to unfavourable biological, social and economic circumstances in childhood. On the contrary, a society may not want to compensate individuals for differences in their health that arise from choices and behaviours they can control and are held responsible for. This conception is not new in egalitarian theory. The idea that fairness can be achieved by removing inequality due to circumstances while letting individuals facing the rewards and costs of their responsible choice is rooted in the moral philosophical

---

* Corresponding author.
*E-mail addresses:* p.brunori@lse.ac.uk (P. Brunori), a.davillas@uom.edu.gr (A. Davillas), andrew.jones@york.ac.uk (A.M. Jones), giovanna.scarchilli@unitn.it (G. Scarchilli).

literature and in the economic social justice theory: see among others Rawls (1958, 1971); Sen (1980); Dworkin (1981); Cohen (1989); Fleurbaey (1995); Roemer (1998); Fleurbaey (2008). The distinction between legitimate and illegitimate sources of inequality is well established in the health economics literature, in particular through the distinction between need-related and non-need-related variation in defining equity in the use of health care (Wagstaff and Van Doorslaer, 2000).

Merging the goals of equality and individual responsibility, Fleurbaey and Schokkaert (2009) drew on two distributive principles to be met in order to realize a fair distribution of health: *reward* and *compensation*. When both principles are satisfied, on the one hand, individuals characterized by identical circumstances face the benefits and the costs of their choices, on the other, individuals behaving in the same way all achieve the same health status independently from their circumstances.[1] In this perspective these two principles define a fair distribution of health, measuring unfair inequality in health means to measure violations of both principles: an ideal measure of unfair inequality should be sensitive to inequality within individuals who make the same choices (compensation) but should also be insensitive to any inequality observed between individuals characterized by the same circumstances who make different choices (reward). The first property captures horizontal equity, with respect to effort, and the second reflects judgements about vertical equity in the reward for effort.

A possible empirical approach to measuring unfair inequality consists of deriving a counterfactual distribution that fully reflects only these unfair inequalities and then applying a suitable inequality index to that distribution. However, Fleurbaey (2008) has discussed the impossibility of constructing a distribution which is consistent with both principles, unless the effects of choices and circumstances are independent from each other; that is, the process generating health is additively separable in circumstances and choices. In the general case, to solve this incompatibility problem, Fleurbaey and Schokkaert (2009) proposed two families of measures of health inequality. Each of these is fully consistent with only one principle, reward or compensation, and partially satisfies the other principle for individuals in some reference conditions. The two measures are the *direct unfairness*, fully consistent with the reward principle and only partly consistent with the compensation principle, and the *fairness gap* which fully satisfies the compensation principle but is partly inconsistent with the reward principle. In practice, these measures parallel the concepts of direct and indirect standardisation used in the measurement of equity in the use of health care (Wagstaff and Van Doorslaer, 2000).[2]

In this paper we implement the Fleurbaey and Schokkaert (2009) measurement approach using an innovative statistical tool, model-based recursive partitioning (MOB). MOB is a tree-based supervised learning algorithm developed by Zeileis et al. (2010) and its use to measure unfair inequalities contributes to the growing methodological literature that uses data-driven techniques in the study of inequality of health opportunity (Li Donni et al., 2015; Brunori et al., 2019; Carrieri et al., 2020). These data-driven techniques offer a compromise between the data-hungry nonparametric approach, which partitions the sample into all unique combinations of circumstances and, hence, often suffers from a curse of dimensionality, and the parametric approach which assumes that the relationship between observed circumstances and the outcome can be captured by a linear and additive (regression) model. Tree-based approaches allow the selection of relevant circumstances, and the way that they interact with each other, to be data-driven.

The model we adopt allows the relationship between health outcomes and health-related behaviours (effort) to be estimated, allowing it to vary according to circumstances that are beyond individual control. The MOB algorithm first estimates a parametric link between health status and lifestyle on the entire sample. Then recursively tests whether partitioning the population based on circumstances and re-estimating the model on population sub-samples can reject the null hypothesis of parameters' stability and obtain a better interpolation of the data. The output of the MOB algorithm is a partition of the sample into socioeconomic groups that are homogeneous in terms of their circumstances, what Roemer (1998) calls "types". Such groups are heterogeneous both in terms of expected health and in terms of the relationship between health-related behaviours and the health outcome. This machine learning approach to estimate health inequalities represents an innovative contribution to the literature and, provided that proxies for relevant responsibility variables are observed, could be straightforwardly extended to other welfare domains such as education or income.

We apply the MOB algorithm to estimate the level of unfair health inequality. We base our estimate on the nationally representative UK Household Longitudinal Study (UKHLS) to present estimates of the two unfair inequality measures introduced by Fleurbaey and Schokkaert (2009): direct unfairness and the fairness gap. We show that unfair inequality is a substantial fraction of the total explained health variability. This finding holds no matter which exact definition of fairness is adopted: using both the fairness gap and direct unfairness measures. These are evaluated at different reference values across the full distributions of types and of degrees of effort. Moreover, as shown in Appendix C, the substance of our conclusions is not affected when adopting reasonable alternative definitions of effort.

It is beyond the scope of this paper to explore and identify the detailed causal mechanisms through which early life circumstances, such as parental socio-economic status (SES), may shape peoples later life health outcomes. But it is helpful to put our approach in the context of this broad literature (for reviews see, Currie and Almond (2011); Almond et al. (2018); Conti et al. (2020)). It has been argued that early-life circumstances affect later life health and inequality in health (Conti et al., 2016; 2020). The early-life period is a critical period for children's development, reflecting biological processes through

---

[1] In what follows we consider the terms 'unfair health inequality' and 'inequality of opportunity in health' as if they were interchangeable. Roemer and Trannoy (2015) discuss the near perfect overlap of the two definitions.

[2] This literature recognises the importance of reference values, embodied in the notion that "on average the system gets it right", and the implied tension between measuring horizontal and vertical inequity with respect to need (Wagstaff and Van Doorslaer, 2000; Gravelle, 2003; Sutton, 2002).

which early-life SES has long-lasting and potentially persistent effects on biological systems (Ben-Shlomo and Kuh, 2002; Pudrovska and Anikputa, 2014). Malnutrition and lack of a nurturing environment are important for childrens development and may lead to changes in brain architecture (Conti et al., 2016; Taylor, 2010). Another important pathway in the association between early-life SES and health outcomes in later life is via risky health behaviours. The development of cognitive and socio-emotional skills goes on to influence investment in human capital including educational attainment and health behaviours. Higher SES in childhood is associated with reduced risks of risky behaviours (linked to smoking, drinking, obesity, and physical inactivity); the latter may persist over an individuals later life and affect (cumulatively) their health (Pudrovska and Anikputa, 2014). For example, evidence shows that childhood socioeconomic circumstances predict persistent smoking behaviours in women using British birth cohort data, highlighting the role of childhood socioeconomic position on persistent smoking in adulthood (Jefferis et al., 2004).

The paper is structured as follows, in Section 2 the metrics proposed by Fleurbaey and Schokkaert (2009) are introduced. Section 3 explains how the MOB algorithm can be used to estimate unfair inequalities. Section 4 presents the data and the empirical results. Section 5 concludes.

## 2. Fleurbaey-Schokkaert model and measures

Consider a population of $N$ individuals over which a distribution of the health outcome $H$ is defined. We assume that individual health is determined by three types of traits: a finite set of lifestyle related factors over which individuals have control ($\mathbf{E}$), which are called "effort" variables, a set of social factors for which individuals cannot be held responsible ($\mathbf{C}$), which are called "circumstances", and age ($\mathbf{A}$). We use an age-adjusted measure of health so we can abstract from $\mathbf{A}$. The individual health outcome is generated by a function of circumstances and effort variables:

$$H = g(\mathbf{C}, \mathbf{E}) \tag{1}$$

All the possible combinations of circumstance values, taken one at a time from $\mathbf{C}$, define a partition of the population into *types*. Individuals belonging to the same type are characterized by identical circumstances. Similarly, all the possible combinations of values taken one at time from $\mathbf{E}$ define a partition of the population into *tranches*. Individuals belonging to the same tranche exert exactly the same effort.

An important normative and empirical issue concerns the definition of the responsibility variables. While Fleurbaey and Schokkaert (2009) do not explain how responsible choices can be measured, considering it a normative choice that belongs to the political decision-maker, John Roemer goes a little further suggesting that the degree of effort exerted must always be orthogonal to circumstances. In Roemer's view, if individuals belonging to different types face different incentives and constraints in exerting effort, this is to be considered a characteristic of the type and should be included among circumstances beyond individual control.

For example, consider the frequency of eating fruit as a measure of effort. An individual with more educated parents may find it much easier to eat regularly fruit, while an individual who grew up in a less favourable environment may find it harder to eat fruit and avoid junk food. Roemer believes that the distribution of effort is, indeed, a characteristic of the type:

> "Thus, in comparing efforts of individuals in different types, we should somehow adjust for the fact that those efforts are drawn from distributions which are different, a difference for which individuals should not be held responsible."
>
> Roemer (2002) p. 458

Roemer therefore distinguishes between the 'level of effort' and the 'degree of effort' exerted by an individual. The latter is the morally relevant variable of effort and is identified with the quantile of the effort distribution for the type to which the individual belongs. In the example of effort exerted by an individual, the relevant measure is not the number of fruit portions eaten but rather the quantile of the type-specific distribution of fruit portions eaten.[3] Other authors have suggested that when measuring unfair health inequality individuals should be held fully responsible for their choices (see Roemer and Trannoy (2015) for a discussion). However, following the prevalent approach in this literature we will define the degree of effort exerted consistently with Roemer's proposal (the empirical difference between the two approaches is discussed by Jusot et al. (2013)).

In our model, health is determined solely by observable circumstances and effort. We are therefore ignoring health variability within cells, groups of individuals sharing the same observed efforts and circumstances. Empirically we easily observe individuals sharing the same circumstances and exerting the same effort, but obtaining a different health outcome. How then should we consider such unexplained variation? Is it more likely that this inequality arises from unobservable effort or unobservable circumstances? Is it simply the randomness inherent in many health outcomes? Or is it a reflection of measurement error which is convenient to ignore, that is replacing all outcomes in the cell with their mean? The answer depends on our beliefs about the observability of circumstances and effort; Lefranc et al. (2009) consider within-cell inequality to be due to randomness or "luck", a source of unfair inequality. On the contrary, the majority of the empirical

---

[3] An alternative way of addressing this issue, purging the influence of circumstances on effort, is to replace the observed level of effort with the residuals from a regression of effort on circumstances (e.g., Jusot et al. (2013); Carrieri et al. (2020)).

studies of income inequality consider variation within cell as due to effort. Checchi and Peragine (2010), for example, claim that this inequality is due to limited observability of effort and therefore should be attributed to effort.

In what follows we explicitly recognize that, to a large extent, health variability cannot be predicted by observable variables. We focus solely on the part of the limited health variability that can be predicted by observable circumstances and efforts and are agnostic about the unexplained variation. We will assign to each individual in type $k$ exerting effort $j$ the average outcome of cell $k, j$. To evaluate whether within-cell inequality is or is not to be considered unfair health inequality is beyond the scope of this approach.

Using this framework Fleurbaey and Schokkaert (2009) have proposed two types of measures of Unfair Inequality ($UI$)[4]. To quantify $UI$ the authors suggest a two-step method: first, starting from a distribution of health outcome ($H$), a counterfactual distribution ($\tilde{H}$) is derived, which reproduces only unfair inequality and does not reflect any inequality arising from choice and effort of individuals; second, inequality is measured for this counterfactual distribution.

In order to construct a measure of inequality in health that is sensitive to the problem of responsibility, Fleurbaey and Schokkaert (2009) present two conditions:

> Condition 1 (Reward, no influence of legitimate differences). A measure of unfair inequality should not reflect legitimate variation in outcomes, i.e. inequalities which are caused by differences in the responsibility variable.
> Condition 2 (Compensation). If a measure of unfair inequality is zero, there should be no illegitimate differences left, i.e. two individuals with the same value for the responsibility variable should have the same outcome.
>
> Fleurbaey and Schokkaert (2009) p. 75.

Putting together both of these requirements, we can state that a counterfactual distribution consistent with the compensation and the reward principles is a distribution that:

1) fully reflects the inequality in outcomes between individuals with the same effort (within-tranche inequality);
2) does not reflect any inequality in outcomes between individuals characterized by same circumstances (within-type inequality).

Any inequality measure applied to such distribution would be a measure of unfair inequality consistent with both the reward and the compensation principle. Fleurbaey and Schokkaert (2009) address the potential conflict between the principles of compensation and reward. They propose two $UI$ measures, each one fully consistent with one of the two principles and maintaining consistency with the other at a reference degree of effort or a reference type, respectively:

*Direct unfairness ($UI_{DU}$):* choose a reference value for the vector of responsibility variables $\tilde{E}$, with $\tilde{h}_i^{k,j} = g(C, \tilde{E})$. In the counterfactual distribution the health of an individual $i$ belonging to type $k$ is the health attained by an individual in type $k$ that exerts the reference degree of effort. Inequality in the counterfactual distribution, $\tilde{H}_{DU}$, is unfair inequality.

*Fairness gap ($UI_{FG}$):* choose a reference type $\tilde{C}$, with $\tilde{h}_i^{k,j} = g(\tilde{C}, E)$. Then $\tilde{H}_{FG}$ is obtained by taking the difference between the individual's health in the initial distribution and the health of individuals who exert the same effort but who have the reference circumstances. Unfair inequality is inequality in $\tilde{H}_{FG}$.[5]

$UI_{DU}$ measures inequality in a counterfactual distribution obtained by removing any inequality due to effort. All individuals belonging to the same type have the same value in $\tilde{H}_{DU}$. Hence $UI_{DU}$ is a measure of unfair inequality fully consistent with the principle of reward (no influence of legitimate differences). On the other hand, $UI_{DU}$ is consistent with the principle of compensation for the reference degree of effort: if all individuals with the reference level of effort obtain the same outcome inequality in $\tilde{H}_{DU}$ is zero. However, $UI_{DU}$ fails to satisfy the principle of compensation for all other effort tranches.

Symmetrically, $UI_{FG}$ measures inequality in a counterfactual distribution obtained by isolating inequality within tranches. It is a measure fully consistent with the principle of compensation: inequality in $\tilde{H}_{FG}$ is zero only if all individuals in the same tranche obtain the same outcome. Moreover, $UI_{FG}$ is consistent with the principle of reward for the reference circumstance; $UI_{FG}$ is insensitive to changes in inequality within individuals characterized by reference circumstances. However, $UI_{FG}$ fails to satisfy the principle of reward for individuals not belonging to the reference type.[6]

Summing up, we can estimate two sets of measures: compensation consistent measures ($UI_{FG}$), and reward consistent measures ($UI_{DU}$). These measures depend on either a reference effort or a reference combination of circumstances therefore we estimate a range of measures and we discuss their sensitivity to different reference values.

---

[4] Their proposal originates from a number of contributions on fair allocation and distributive justice (Fleurbaey, 2008; Fleurbaey and Maniquet, 2012). In these contributions the authors developed a theory of "responsibility-sensitive egalitarianism" whose ambition is to generalize the egalitarian ideal allowing individuals to be held responsible, to some degree, for their achievements.

[5] This index is equivalent to the measure of horizontal equity, based on indirect standardisation, that is typically used in the literature on equity in the delivery of health care (Wagstaff and Van Doorslaer (2000))

[6] Note that these measures differ from the ex-ante and ex-post inequality of opportunity measures inspired by Roemer (1998) and often adopted in empirical studies (Checchi and Peragine, 2010; Roemer and Trannoy, 2015). Ex-ante $UI$ is a reward-consistent measure of $UI$ obtained imposing: $\tilde{h}^{k,j} = \lambda_k = \mu_k$, where $\mu_k$ is the average outcome of individuals in type $k$ (see Property 1). Ex-post $UI$ is a compensation-consistent measure of $UI$ obtained imposing: $\gamma_j = \mu_j$, where $\mu_j$ is the average outcome of individuals in tranche $j$ (see Property 2). Ex-ante and ex-post $UI$ fail to satisfy both the principle of compensation and the principle of reward respectively, unless $g$ is additively separable in $E$ and $C$. However, because they are relatively easier to estimate and to decompose, they are very popular in the empirical literature about inequality of opportunity in income and consumption as well as applications to health inequality (Rosa Dias, 2009; Jusot et al., 2013; Davillas and Jones, 2021).

## 3. Empirical definition of $UI_{DU}$ and $UI_{FG}$ using model-based recursive partitioning

Estimation of $UI_{DU}$ and $UI_{FG}$ requires relevant circumstances beyond individual control to be observed and types to be defined. Ideally, a measure of unfair inequality should consider all the potential sources outside individual control. However, this would require considering a wide and complex set of circumstances, which brings with it the risk of noisy and upwardly biased estimates (Brunori et al., 2019). Traditionally, in empirical studies on unfair inequalities the relevant circumstances have been included in the model through normative decisions. In the nonparametric approach the population is partitioned into a parsimonious number of types and in the parametric approach the relationship between circumstances and the outcomes have been implicitly modelled as additive and fixed using linear regression. For these reasons, coupled with the fact that some circumstances may be unobserved, estimates have been interpreted as a *lower-bound* estimate of the real level of unfair inequality (Rosa Dias, 2009; Li Donni et al., 2015; Jusot et al., 2013; Carrieri and Jones, 2018).

A number of more recent empirical applications instead rely on data-driven semiparametric techniques to explore the information on social groups which is relevant to the formation of unfair inequalities. These are semiparametric in the sense that relationship between health outcome and effort is assumed to take a (linear) parametric form, while the definition of types is nonparametric. On one side, finite mixture models (FMM)[7] have been adopted to study the latent type membership of each individual given their observed circumstances (Li Donni et al., 2015; Carrieri et al., 2020; Brunori et al., 2021). The FMM approach relies on an *a priori* selection of the circumstance variables that influence the probability of belonging to each type. On the other side, tree-based methods have been adopted to perform a data-driven selection of the relevant circumstances and the interactions between them on the basis of model fit (Brunori et al., 2018; Brunori and Neidhöfer, 2020). The estimation approach proposed in this paper, model-based recursive partitioning (MOB), is an extension of the tree-based techniques applied with a specification of types that echoes the semiparametric mixture approach (Carrieri and Jones, 2018; Carrieri et al., 2020).

Consider again equation (1): individual health outcomes, $h_i$, are attributed to two sets of observable variables: a number of behaviours and a set of circumstances for which individuals are not held responsible, respectively **E** and the **C**. The isolation of the unfair health inequality requires the estimation of a model for health. For the sake of simplicity, and following Carrieri and Jones (2018), assume that behaviours can be summarized by a scalar index of lifestyle ($e$) and that its effect on health can be modelled using a linear regression:

$$h_i = \beta_0 + \beta_1 e_i + \epsilon_i \tag{2}$$

We can assume that this simple relationship is not independent from **C**. The relationship linking efforts and health can be affected by the circumstances though two channels: the intercept, $\beta_0$, and the slope, $\beta_1$.[8] A different intercept can be interpreted as the direct contribution of circumstances to health: independently from the choices made having favourable circumstances may improve individuals' health. Heterogeneity in the slope instead means that the contribution of lifestyle to health outcomes may be also affected by circumstances.

The final model can be represented as a weighted sum of sample splits performed to derive $k = 1, \ldots, K$ different models associated with each subgroup parameters $\beta_{(k)}$:

$$g(h_i | \mathbf{c}_i, e_i, \boldsymbol{\beta}_{(1)}, \ldots, \boldsymbol{\beta}_{(K)}) = \sum_{k=1}^{K} \pi_k(\mathbf{c}_i) \cdot g(h_i | e_i, \boldsymbol{\beta}_{(k)}) \tag{3}$$

Note that this representation of the individual health model as a function of efforts and circumstances can be either associated with both the FMM and the MOB approaches to estimation. Depending on which of the two methodologies is chosen, the weight $\pi_k(c_i)$ and the $K$ subgroups will be identified with a different estimator.

All of the specifications considered here begin with equation (2), that assumes a linear relationship between the outcome and effort. Effort may include a list of observed effort factors (Carrieri and Jones, 2018; Davillas and Jones, 2020) or these may be combined into a scalar latent variable as in this paper. The constant returns to effort implied by a linear relationship could be relaxed by using powers or transformations of effort in the regression. Circumstances are introduced by allowing the slope and intercepts of equation (2) to vary. The parametric approach makes these a linear function of observed circumstances. If only the intercept varies with circumstances then the regression model becomes a linear function of both circumstances and effort (e.g. Carrieri et al., 2020). If the slope coefficient also varies with circumstances then the regression model would include interaction terms between circumstances and effort. The non-parametric approach (e.g. Carrieri and Jones, 2018) takes equation (2) and estimates it separately for sub-samples for each type. These types are defined *a priori* by the analysts when they select the list of relevant circumstance variables and the categorical levels of these variables. This non-parametric approach suffers from a curse of dimensionality which limits the range of circumstance factors that can be accommodated. Semiparametric methods address this issue and the FMM and MOB approaches both serve this purpose. The

---

[7] Mixture models in statistics are a broad family of probabilistic models for observing latent subgroups in a population, including latent class analysis (LCA) as a specific case.

[8] In the empirical application we consider higher order polynomials for effort, with the chosen specification selected by cross validation. So, although this is the parametric part of the specification, the estimation does allow for a considerable degree of flexibility. Note also that the MOB specification allows for interactions with circumstances through the heterogeneity of parameters across types.

FMM model assumes a discrete set of latent types and that the probability of belonging to these types is captured by a multinomial logit specification (e.g. Carrieri et al., 2020). This is typically applied using a linear index of the circumstance variables within the logit functions. In contrast the MOB uses a decision tree to define the types and, as such, offers more scope for capturing interactions between the different circumstance variables that shape class membership.

A central aim of this paper is to implement the fairness gap and direct unfairness measures proposed by Fleaurbaey and Schokkaert (2009) and hence to construct counterfactual distributions at different reference levels of circumstances or effort. Reference levels of circumstances can be conveniently handled when the sample is split into types as in the nonparametric and semiparametric specifications (FMM and MOB). Reference levels of effort are conveniently handled when effort is condensed into a scalar latent variable, as done by the use of principal components analysis in this paper.

We opt for the use of the MOB to estimate the indirect relation between circumstances and behaviours, and to allow the health response to effort be estimated varying across meaningful social groups. Tree-based techniques are data-driven and rely on *decision trees* which, in statistics, can be used to visually represent the "decisions", or if-then rules, that are used to generate predictions of a single outcome variable or a model. Moreover, tree-based methods tend to be more parsimonious then FMM in terms of parameters resulting in less conservative (more fine grained) partitions in types. To get and intuition of how the two approaches, our tree-based approach and the approach based on latent classes, differ in practice, Appendix D summarizes the results one would obtain performing our empirical analysis using FMM models.

There are essentially two key components to build a decision tree: the *features* to split on the prediction sample, and the *rule* to stop splitting the sample. The MOB is a particular tree-based method which takes as input a set of partitioning variables and whose splitting rule relies on the estimated parameters of a model.

This model is initially estimated on the entire sample, afterwards, a statistical test is performed to verify whether there are any possible sample splits on the partitioning variables which achieve a better fit of the model. The outcome of this process is a set of models estimated on $K$ sub-samples of the original population (terminal nodes).

We briefly summarize here how a MOB is obtained from data (see Zeileis and Hornik (2007); Zeileis et al. (2008) and Zeileis et al. (2010) for details). The MOB uses the vector **C** to search for ways of splitting the sample into non-overlapping subgroups. If estimating the response of health to lifestyle into two sub-samples yields statistically different parameters and improve out-of-sample prediction, then the split is performed. The procedure is then repeated in the resulting sub-samples.

The parameter instability is detected by means of Generalised M-fluctuation tests. The test is based on a partial sum process of the estimation scores which captures instabilities (Zeileis and Hornik, 2007; Hothorn and Zeileis, 2015). It can be understood as a generalization of the type of test used to detect structural breaks in time series analysis. In the case of the MOB algorithm, the test is performed on the partial sum of residuals across the space defined by partitioning variables. The fluctuation test statistic is distributed as a $\chi^2$ and we can compute the p-value for testing its significance. If the fluctuation test statistic is higher than a certain threshold, the hypothesis of stability of the model parameters is rejected and algorithm splits the sample and re-estimates the model on the distinct subgroups.

Schematically, Zeileis et al. (2010) illustrate the steps of the MOB algorithm as follows:

1. Set a confidence level $(1 - \alpha)$ to be used as tuning parameter;
2. Fit the model - for example: $h_i = \beta_0 + \beta_1 e_i$ - on the entire sample;
3. Test whether there is any partitioning variable causing parameter estimates for the model to be unstable;
4. If the null hypothesis of parameters stability across possible sub-samples cannot be rejected, stop;
5. If the $p - value$ of the fluctuation test statistics is instead lower than the critical Bonferroni-adjusted $\alpha$, select the variable associated with the most statistically significant source of instability;
6. Compute the exact splitting point which optimises the objective function of the estimation according to the selected partitioning variable;
7. Split the node into child nodes and restart the procedure from (2) on the two subsamples.

The depth of the estimated tree depends on the tuning parameter $\alpha$ which determines the $p - value$ threshold for rejecting the null hypothesis in the instability test. The value of $\alpha$ can be set to a specific value or can be selected by a machine-learning technique ensuring that MOB stops splitting the sample when no further split would result in a better out-of-sample fit of the data.

The outcome of the algorithm is a partition of the population into types according to the composition of the terminal nodes. Individuals belonging to each type share the same circumstances and the same parameters for equation (2). The partition into types and the associated set of parameters allows the counterfactual distributions $\tilde{H}_{DU}$ and $\tilde{H}_{FG}$ to be computed. The counterfactual distribution $\tilde{H}_{DU}$ is obtained by choosing a reference degree of effort $\tilde{e}$ and then predicting $\tilde{h}_i^{k,j} = \hat{\beta}_0^k + \beta_1^k \tilde{e}$. The counterfactual distribution $\tilde{H}_{FG}$ is obtained by choosing a reference type $(R)$ and then predicting $\tilde{h}_i^{k,j} = (\hat{\beta}_0^k + \beta_1^k e_j) - (\hat{\beta}_0^R + \beta_1^R e_j)$. $UI_{DU}$ and $UI_{FG}$ are then obtained by computing a suitable inequality measure of the counterfactual distributions.

## 4. Data and estimates

The data comes from three waves of the UKHLS panel. The survey contains information about demographic characteristics, a rich set of information about individuals socioeconomic background in childhood, ethnicity, and place of birth among

**Fig. 1.** Timeline for the study design. *Note: Circumstances may be observed in multiple waves.*

other things. These provide our measures of circumstances that are used to construct types. Moreover, the survey contains questions about health-related behaviours, that are used to measure effort with a scalar index of lifestyle, and a number of measures of health outcomes.

Together with a panel study, like UKHLS, cohort studies can be used to measure unfair health inequality and have been used in the inequality of opportunity in health literature given the availability of long-running longitudinal data and information on participants socio-economic background in childhood (Rosa Dias, 2009; 2010). However, unlike cohort studies that sample elderly people (The English Longitudinal Study of Ageing) or young people (e.g., birth cohorts), the larger sample size and the wider age range of general household surveys, like the UKHLS, make them a valuable complement to cohort studies.

Given the aim of our study to partition the population into socioeconomic groups that are homogeneous in terms of their circumstances, the availability of a sample of all ages is of particular interest. Across generations and birth cohorts, a different set of opportunities may be open to the members of the society as a result of societal changes - allowing for a sample that incorporates people of all ages, rather than focusing on specific birth cohorts, is a useful addition to the literature. Moreover, our dataset allows for a set of lifestyle indicators (along with risk aversion variable that is used in our sensitivity analysis in Appendix C) to be collected before the health outcome of interest in our analysis.

We are aware that our analysis is limited by the role of any unobserved circumstance variables and the role of the public healthcare services and consumption. As the scope of our analysis is neither to identify the efficiency of public health services nor the role of genetic predisposition on peoples health later in life, our analysis is limited to this extent and irrespective of the dataset to be employed.

Fig. 1 shows the study design and indicates at what moments in time and to which waves the observations of the different variables used in the analysis correspond. Circumstances relate to fixed individual characteristics and to measures of parental background, health-related behaviours are measured at Waves 2 and 5, and the health outcomes are measured in the subsequent follow-up at Wave 6.

Our chosen health outcome (H) is measured at UKHLS Wave 6 (2014–2015). We use the Short Form 12 (SF-12), a well validated, self-administered health measure based on a set of 12 questions on respondent's health (Ware et al, 1995). For this study, we use the Physical Component Score (PCS-12), to capture respondents' physical health. The PCS-12 score has values between 0 and 100, and it has been standardized in order to have a mean of 50 and a standard deviation of 10; higher values indicate better physical health functioning. The PCS-12 is a reliable instrument developed to measure physical health in large surveys with higher values of sensitivity and specificity compared to other brief health scales (Ware et al., 2001; Ziebarth, 2010). It has been used in the literature as a robust self-reported measure of physical health (e.g., Eibich (2015); Guber (2019); Schmitz (2011); Ziebarth (2010)). The health measure has been adjusted for individual age (at the time of the interview) in order to control for the age-specific variability in health. The age-adjustment is performed by regressing individual health status on 5-year age classes between 14 and 100. To remove all the age-class fixed effects from total health variability we use the residuals as our measure of health status.

The full set of observed circumstances (**C**) beyond individual control that are considered as candidate variables in the MOB algorithm are: ethnic groups (the relevant categories have been summarised into the following levels: UK white; Irish white; other white; mixed: white with Asian/African/Arab; Asian: East and Middle East; Black: African, Caribbean, other; other ethnic groups), place of birth (a dichotomous variable indicating whether born in the UK or not), father and mother's skill levels in the main occupation (unemployed or four skill levels in occupation), mother and father's education (did not go to school, left school without qualifications, some qualification, post-school qualifications, university degree or higher), mother and father's activity status (working, unemployed, deceased, not living in the household). Note that all information about parents relate to when the respondent was 14 years old. We include sex as an additional source of unfair health inequality. The tree structure implicit in the MOB algorithm allows for a full set of interactions between the categories of these circumstance variables. However, as it is a data-driven technique, it guards against the curse of dimensionality and the risk of over-fitting that would be likely with a fully saturated nonparametric specification.

Table 1 shows the frequencies of each circumstance category in the sample. Fig. A.2 in the Appendix shows the most frequent patterns of missing values for circumstances and the health outcome. The most frequent missing information is parental education but note that for 4567 observations of the potential maximum sample to be used in our analysis, the

**Table 1**
Descriptive statistics: circumstances.

| Circumstance Category | Frequency (%) |
|---|---|
| Ethnic group | |
| UK white | 82.06 |
| Irish white | 2.11 |
| other white | 2.52 |
| Mixed (white with Asian/African/Arab) | 1.87 |
| Asian (East and Middle East) | 7.05 |
| Black (African, Caribbean, other) | 3.04 |
| other ethnic group | 0.26 |
| missing | 1.10 |
| Female | |
| yes | 55.95 |
| no | 44.04 |
| missing | 0.01 |
| Born in the UK | |
| yes | 86.27 |
| no | 11.32 |
| missing | 2.42 |
| Mother education | |
| did not go to school | 1.15 |
| left school without qualifications | 25.73 |
| some qualification | 13.59 |
| post-school qualification | 15.79 |
| university degree or higher | 6.86 |
| unknown | 3.94 |
| missing | 32.95 |
| Father education | |
| did not go to school | 1.88 |
| left school without qualifications | 29.29 |
| some qualification | 19.16 |
| post-school qualification | 10.94 |
| university degree or higher | 4.54 |
| unknown | 0.95 |
| missing | 33.24 |
| Mother's occupational skill level | |
| unemployed | 38.24 |
| high skill | 6.65 |
| up-mid skill | 5.84 |
| mid skill | 17.97 |
| low skill | 9.35 |
| unknown | 2.00 |
| missing | 19.94 |
| Father's occupational skill level | |
| unemployed | 5.24 |
| high skill | 10.94 |
| up-mid skill | 26.56 |
| mid skill | 15.70 |
| low skill | 5.93 |
| unknown | 6.93 |
| missing | 28.70 |
| Mother activity status | |
| working | 53.36 |
| unemployed | 38.24 |
| deceased | 1.33 |
| not living in hh | 0.67 |
| missing | 6.40 |
| Father activity status | |
| working | 80.47 |
| unemployed | 5.24 |
| deceased | 3.75 |
| not living in hh | 3.18 |
| missing | 7.37 |

*Source:* UKHLS Wave 6

only missing information is the SF-12 Physical Component Score. Appendix A shows the prevalence and covariance of item non-response for all variables used in the analysis.

To implement the specification in equation (2), a composite scalar index of lifestyle is created. Specifically, all our lifestyle indicators are summarised by a scalar index obtained by Principal Component Analysis (PCA) explained in some detail in

**Table 2**
Spearman correlation with effort.

| Behaviours | $\rho$ |
|---|---|
| Sport activity | 0.5275*** |
| Smoke intensity | −0.5593*** |
| Ex smoker | 0.1545*** |
| Fruit per week | 0.6456*** |
| Vegetables per week | 0.5410*** |
| Days walked at least 10 minutes | 0.6033*** |
| Drink $\geq$ 5 days per week | −0.0139*** |
| % of total variance explained | 43.7 |

*Source: UKHLS Waves 2 and 5.* Note: Signif. values: *** ($p < 0.001$).

Appendix B. For those lifestyle indicators that respondents are observed in both Waves 2 and 5 (and different responses are obtained) the more risky level of health behaviour is used in the PCA. The choice of using a summary measure of lifestyle is based on two main considerations. The first is to keep the MOB as parsimonious as possible and to avoid over-fitting the data. Second, we consider lifestyle as an intrinsically unobservable latent pattern of behaviour. On the one hand, each specific behaviour we observe is correlated with this lifestyle, on the other, specific behaviours may be a rather imperfect measures of the overall pattern. Note that in Appendix C we summarize an extensive sensitivity analysis performed across alternative ways of using information about health-related behaviours to measure effort.

The following indicators of health-related behaviours are included in our analysis to proxy efforts: current smoking status (non-smoker, up to 10 cigarettes per day, 10–19 cigarettes per day, 20+ cigarettes per day), a dummy variable for ex-smoker, number of days each week eating fruits (never, 1 - 3 days, 4 - 6 days, every day), number of days each week eating vegetables (never, 1 - 3 days, 4 - 6 days, every day), days per month walked at least 10 minutes (28 categories based on the frequency of walking habits during the days of a month), a dichotomous variable for drinking alcohol five or more days per week. We also account for a self-assessed measure of sports activity, which is an eleven categories scale from 0 to 10, with 0 being "doing no sport at all", and 10 being "very active through sport".

Table 2 shows the correlation of the lifestyle variable with the observed behavioural variables involved in the analysis. The sign of the correlation is positive for healthy habits such as non-sedentary lifestyle and healthy diet, whilst it is negative for heavy drinking and intensity of smoking.

As shown in Table 3, a non-negligible share of missing information concerns alcohol intake (about 23% in Wave 2, and 17% in Wave 5).

Fig. A.1 in Appendix A shows the most frequent combinations of missing data for effort variables. Interestingly about half of the missing information concerns only that aspect of lifestyle. Therefore, for respondents reporting complete information about all other effort dimensions but missing alcohol we impute drinking behaviour by multiple imputation using observed behaviours as imputers (Van Buuren and Groothuis-Oudshoorn, 2011). The final sample includes all respondents with complete information, obtained by merging the three UKHLS waves and, after imputation, this is made up of 18,016 adults. Although the final sample size is large relatively to similar empirical analysis, the item non response represents an issue and caution should be exercised in generalising the results to the entire UK population.

All of the circumstances and the scalar index of lifestyle are then used to estimate the model-based tree. The algorithm is tuned by 5-fold cross validation. We tested different critical values for the Bonferroni-adjusted p-value ($\alpha = 0, \alpha = 0.001, \alpha = 0.01, \alpha = 0.05, \alpha = 0.1$) and different health-effort polynomial link specifications (degree 1 to 4). Moreover, in order to guarantee sufficient degrees of freedom for each type, we impose a minimum number of 200 observations per terminal node. The output of the MOB specification with the smallest out-of-sample prediction error is shown in Fig. 2, it is obtained with $\alpha = 0.1$ and assuming a linear relationship between our measure of lifestyle and physical health rather than higher order polynomials.

The selected tree is made of 11 splits and 12 types. Circumstances used to partition the population are: ethnic group, sex, father's activity, mother's activity, mother's education, father's education, place of birth. Each terminal node contains a scatter plot in which lifestyle is on the horizontal axis and health outcome is on the vertical axis. All type-specific regression models have highly significant regression coefficients and a positive slope (the healthier the lifestyle the higher the expected health). The fitted model explains about 10% of the total health variance in the sample. In what follows we estimate how much of this explained variability is to be considered unfair.

Table 4 reports for each type: the average health status, the average effort exerted, the two parameters ($\beta_0$ and $\beta_1$) and the population share of each type.

In terms of average health, the worst-off type is type 1 made up of mixed race, other ethnic and Asian women whose mother did not work. This group represents about 4% of the sample and has an expected health outcome of −4.728 (not far from the 25th percentile of the entire PCS-12 distribution). The best-off type is type 12 made up of white or black men whose mother left school with at least some qualification and whose father has at least a post-school qualification (or for a few respondents is unknown). This type represents slightly more than 7% of the sample and their average health is 2.871 (clearly above the population mean 0.1964).

**Table 3**

Descriptive statistics: life-style behaviours.

| Life-style categories | Freq (%) |
|---|---|
| Sport activity | |
| no sport at all | 30.76 |
| 1 | 8.25 |
| 2 | 8.98 |
| 3 | 8.85 |
| 4 | 7.27 |
| 5 | 9.74 |
| 6 | 5.54 |
| 7 | 5.68 |
| 8 | 3.79 |
| 9 | 1.50 |
| very active through sport | 2.04 |
| missing | 7.58 |
| Current smoking status | |
| not smoking | 79.21 |
| up to 10 cigarettes per day | 6.56 |
| up to 20 cigarettes per day | 8.61 |
| more than 20 cigarettes per day | 5.62 |
| missing | 7.55 |
| Ex-smoker | |
| no | 62.02 |
| yes | 30.44 |
| missing | 7.55 |
| Fruit per week | |
| never | 8.61 |
| 1–3 days per week | 31.83 |
| 4–6 days per week | 17.69 |
| everyday | 34.36 |
| missing | 7.51 |
| Vegetables per week | |
| never | 2.59 |
| 1–3 days per week | 24.23 |
| 4–6 days per week | 27.22 |
| everyday | 38.47 |
| missing | 7.51 |
| Days per month walked at least 10 minutes | |
| 0 | 19.57 |
| [1-10) | 26.14 |
| [10-20) | 12.94 |
| [20-30) | 33.81 |
| missing | 7.55 |
| Alcoholic drink $\geq$ 5 days per week | |
| yes | 7.26 |
| no | 76.36 |
| missing | 16.37 |

*Source:* UKHLS Waves 2 and 5

**Table 4**

Types description.

| Type | Av. $h$ | Av. $eff$ | % Pop. | $\beta_0$ | SE | $\beta_1$ | SE |
|---|---|---|---|---|---|---|---|
| 1 | −4.728 | 3.153 | 3.96 | −9.991*** | (0.991) | 1.668*** | (0.290) |
| 2 | −2.606 | 3.093 | 2.02 | −6.310*** | (1.169) | 1.197*** | (0.346) |
| 3 | −2.400 | 3.042 | 6.97 | −8.306*** | (0.702) | 1.940*** | (0.204) |
| 4 | −0.755 | 3.695 | 1.76 | −6.082*** | (1.634) | 1.441*** | (0.418) |
| 5 | −0.608 | 3.542 | 1.12 | −8.405*** | (1.651) | 2.201*** | (0.434) |
| 6 | −0.063 | 3.587 | 3.84 | −3.702*** | (0.966) | 1.014*** | (0.249) |
| 7 | 0.082 | 3.172 | 17.19 | −7.077*** | (0.428) | 2.257*** | (0.120) |
| 8 | 0.380 | 3.494 | 15.20 | −8.067*** | (0.534) | 2.417*** | (0.140) |
| 9 | 0.487 | 3.480 | 25.48 | −5.737*** | (0.371) | 1.788*** | (0.097) |
| 10 | 1.172 | 3.351 | 1.59 | −3.302*** | (1.218) | 1.335*** | (0.334) |
| 11 | 1.494 | 3.424 | 13.57 | −5.095*** | (0.459) | 1.924*** | (0.122) |
| 12 | 2.871 | 3.584 | 7.26 | −1.725*** | (0.485) | 1.282*** | (0.123) |

*Source: UKHLS Waves 2, 5 and 6.* Note: In the first column types rank is determined by their average health (second column), the third column reports the average effort and the fourth the share of observations in each type. The other columns contain models' parameters. Signif. values: *** ($p < 0.001$)
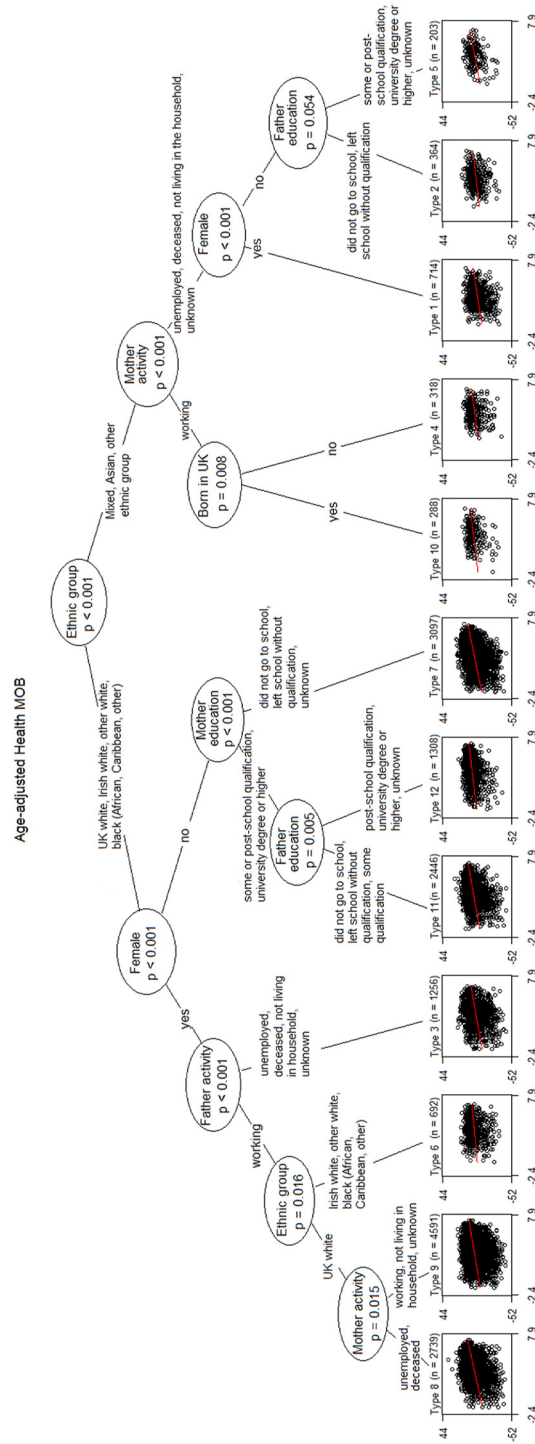
**Fig. 2.** MOB tree diagram.

**Fig. 3.** Opportunity sets by types: health - level of effort profiles. *Source: UKHLS Waves 2, 5 and 6.* Note: health-effort relationship is shown on the entire effort range of variation..

In general, the splitting rules selected by the MOB algorithm are consistent with what might be expected: ethnicity, place of birth, sex and parental background all play some role. A more advantaged socioeconomic background, mother's labour force participation, being born in the UK, and being white are predictive of a better health outcome. Less obviously, being either a white or black male is predictive of a better outcome. In terms of the parameters estimated type 1 and 12 are also the types with the lowest and highest intercepts. Type 6 has the lowest return to effort ($\beta_1 = 1.014$). This type is made of women that define themselves as non-UK white or black and whose father was working during their adolescence. Women that define themselves as UK white whose father was working, but whose mother was not (type 8), have the highest return ($\beta_1 = 2.417$), a gradient that is two-and-a-half times that of type 6. Note that slopes heterogeneity is a source of clash between compensation and reward discussed in Section 2 that justifies the need of considering two families of unfair inequality measures.

Fig. 3 shows the fitted regression lines for each type. These can be interpreted as the opportunity set (or health constraint) faced by individuals belonging to different types.

What emerges is that having favourable circumstances will produce a fixed advantage (higher intercept) but it will not necessarily imply a higher return to a healthy lifestyle (higher slope). That is, there is a correlation between the intercept and the types' rank in terms of expected health. But there is not a monotonic relationship between slopes and intercepts nor between slopes and expected outcome.

Having estimated the opportunity sets individuals face is not sufficient to obtain the two counterfactual distributions necessary to estimate *UI*. The counterfactual distributions will depend on these parameters and also on the type-specific distributions of effort that define the degree of effort. An initial intuition regarding the role of effort in determining the different type-specific health outcomes is provided by Figs. 4 and 5(a). Fig. 4 shows the distribution of effort in the 12 types, ranked according to their average health. The effort distribution in better-off types is more dispersed and higher than the overall average (dashed vertical line). The between-type variability of effort is limited ranging between 3.040 and 3.695 (the 39th and 55th percentile of the distribution in the population). There is also a moderate negative correlation between the average effort exerted and return to effort (-0.1478). So both individuals with more favourable circumstances and with lower return to effort tend to have healthier lifestyles.

However, focusing on the type-specific empirical cumulative distribution function (ECDF) of effort and health what is striking is the clear dominance in terms of expected health condition for better off types accompanied by absence of dominance in terms of effort.

Consider for example Figs. 5(a) and 5(b) where both ECDFs are shown for the two extreme types. Type 1 made of women with Asian or mixed origin, and an absent or non-working mother, and type 12 made of white men with both parents with at least post-school qualification.

**Fig. 4.** Distribution of effort across types *Source: UKHLS Waves 2, 5 and 6.*



(a) Effort                                                           (b) Health

**Fig. 5.** Empirical Cumulative Distribution Functions *Source: UKHLS Waves 2, 5 and 6.*

While the effort ECDFs cross, with individuals in the least favourable type behaving better at the bottom of the distribution (5 a), health ECDFs show a clear dominance of type 12 over type 1, with a particularly marked difference in expected health especially in the left tail of the distribution (5 b).

Finally, adopting John Roemer's view about what is the morally-relevant measure of effort, we remove the variation of effort systematically correlated with types by comparing individuals considering the degree of effort they exerted. $\hat{H}_{DU}$ and $\hat{H}_{FG}$ are therefore constructed ignoring the absolute level of effort (first component of the PCA) and comparing instead individuals belonging to the same quantile of the type-specific distribution of the same variable. This transforms Fig. 3 into Fig. 6. The distribution described by these segmented lines together with the types' population shares provides all the information needed to estimate $UI_{DU}$ and $UI_{FG}$.

The two measures of health unfair inequality are calculated for the 12 possible reference types and for 10 possible reference responsibility values (effort tranches) defined by the deciles of the scalar lifestyle index within each type. For both measures we calculate confidence intervals by bootstraping observations by types. This implies fixing the structure of the

**Fig. 6.** Distribution of health inequality: health - degree of effort profiles *Source: UKHLS Waves 2, 5 and 6.*

(a) $UI_{FG}$                    (b) $UI_{DU}$



**Fig. 7.** Unfair health inequality. *Source: UKHLS Waves 2, 5 and 6.* Note: In 7(a) reference types are sorted by increasing type-specific expected health. Confidence intervals are obtained from 200 stratified bootstrap samples.

tree and then resampling each type 200 times. This procedure is likely to underestimate the level of uncertainty about point estimates. A more robust approach would consist in estimating a different MOB for each sample. However, the need to set a reference type to calculate $UI_{FG}$ requires to fix the structure in types. Fig. 7a reports our estimates for $UI_{FG}$ based on the 12 reference types. Types are ordered according to their average health status (labelled below) but the expected outcome does not affect the value of $UI_{FG}$. Its value is entirely determined by the slope of the regression line estimated for the reference type. The flatter the regression line the more health variability is reproduced in the counterfactual distribution. In the extreme case in which the line has slope zero, health is independent from the degree of effort in the reference type and all health inequality is to be considered unfair. After all, if choices do not play a role, what sort of inequality can be justified? In our case, when type 6 is the reference ($\beta_1 = 1.014$ and the average slope of the resulting broken line in Fig. 6 is also the

flatter) close to 50% of the explained variability is to be considered unfair inequality. Moreover, no matter what reference type is selected $UI_{FG}$ is never lower than 30%.

Fig. 7(b) reports estimates for direct unfairness for ten reference effort tranches (deciles in ascending order). The ten unfairness measures are significantly smaller than the compensation-consistent measures and their value follows a U-shaped pattern. Unfair inequality is higher when the reference effort is at the two extremes of the lifestyle spectrum (close to 30% and 25% of the explained variance respectively). Fig. 6 shows that this pattern is driven by the outcomes for the worse-off types converging on those of the better off types as effort increases from the lower deciles to a more healthy pattern of behaviour in the middle deciles. This is due to the less dispersed distribution of effort in the worse-off types, who appear to catch-up with more advantaged types simply because the average effort exerted in the left tail of the distribution increases at a faster rate. This pattern is then reversed for individuals in the highest effort tranches. For individuals that adopt the healthiest lifestyle a clear social gradient is visible with two types lagging behind (1 and 2) in terms of health status. The comparison between the two extreme types is striking; no matter how healthily they behave, individuals in type 1 have a predicted health outcome below that of the worst-behaving individuals who have the most favourable circumstances (type 12). For type 1 there is no level of effort that could compensate for their adverse circumstances (no matter how badly an individual in type 12 behaves she has a higher predicted health).

We must underline here that our results are to a large extent dependent on the available data and on the specificity of our estimation approach. For this reason we have performed a number of sensitivity analysis presented in more details in Appendix C and D. The use of alternative measures of effort produces different MOB both in terms the partition and in terms of type-specific regression models. Nevertheless the emerging figure is rather consistent across specification in terms of variable used and in the magnitude of the level of unfair inequality. As robustness check we also estimated $UI$ estimating a FMM with and without including health as a concomitant variable. Although circumstances showing strongest association with latent classes' membership are similar to circumstances used as splitting variables in MOB, FMMs produce a remarkably less fine grained partition in types and, consequently, a more conservative estimate for both families of unfair health inequality.

## 5. Conclusions

This study aims to provide both a methodological innovation for the measurement of unfair health inequality, as well as new evidence on health inequalities measured in the UKHLS. The methodological innovation is the adoption of the MOB algorithm to estimate the health-to-lifestyle relationship while considering the different socioeconomic backgrounds in childhood. Moreover, a normatively defined responsibility-sensitive framework is adopted to measure direct unfairness and the fairness gap *à là* Fleurbaey and Schokkaert (2009). Among the main features of the use of MOB in the measurement of unfair health inequality is its ability to capture those socioeconomic characteristics which are fundamental to determine a change in the conditional distribution of the outcome in the health-to-lifestyle model.

The empirical application uses data from the UK Household Longitudinal Study (Waves 2, 5 and 6) considering all observations for which data on physical health status, relevant circumstances beyond individual control, and health-related behaviours are observed. We show that circumstances beyond individual control are a clear source of unfair health inequality. However, this is mostly driven by a fixed advantage for better-off types. Moreover, while on average individuals characterised by more favourable circumstances tend to have a healthier lifestyle, this seems not to be due to systematic heterogeneity in the return to effort across types.

The evidence we find for the dominant contribution of circumstances and for the lack of systematic heterogeneity in the return to effort echoes earlier findings in the literature. Carrieri and Jones (2018) decompose Gini and variance measures of explained inequality in various biomarkers from the Health Survey for England. Effort is proxied by measures of smoking, diet and alcohol consumption. They find that the indirect contribution of effort to these decompositions, which is attributable to differences in the slope coefficients for effort variables across the 36 types that are defined using a nonparametric approach, is tiny in comparison to the direct contribution of circumstances (attributable to heterogeneity in the intercepts). Overall, they find that circumstances are the leading determinant of inequality in cholesterol, glycated haemoglobin, and in an overall index of ill-health, while effort only plays a substantial role for fibrinogen. In the Carrieri et al. (2020) application of a FMM specification to data from UKHLS their decomposition analysis shows that about two-thirds of the total inequality in a measure of allostatic load can be attributed to the direct and indirect contribution of circumstances and that the direct contribution of effort is small. The decomposition analysis shows that about 50% of the total inequality in the composite health outcome is attributed to the direct contribution of demographic and parental circumstances (the intercepts). Though differences in the return to effort, circumstances exert an indirect contribution to the total inequalities of around 13%. The direct contribution of efforts is much less important, with a contribution of around 3%.

When adopting a reward-consistent approach, and measuring $UI_{DU}$, a clear pattern emerges; when the reference degree of effort is at the two extremes the level of unfairness detected is higher. This result is driven by the interactions of types' direct contribution to health (the intercept), the return to a healthier lifestyle (the slope) and the type-specific distribution of effort being more compressed for less advantaged types. The combined effect makes between-type inequality lower for individuals exerting an intermediate degree of effort.

Overall, our results show that the variation in physical health can only be partially explained by observed lifestyle and childhood socioeconomic background in the UKHLS. Indeed, there are many aspects which are not included in the model

even though they have an impact on health status. Some of these are likely to remain unobservable, such as genetic endowments, others, however, could fit in the Fleurbaey and Schokkaert (2009) framework and, given suitable data, could be taken into account, such as healthcare consumption and the role of public healthcare services.

A final important limitation for this type of analysis concern the marked sensitivity of the result to apparently negligible choices in the definition of behaviours. The robustness checks performed on alternative, and equally plausible, measures of effort show that estimated unfair inequality varies in a range between -28% and +23% with respect to our baseline. This must warn readers about the need for a careful evaluation of all aspects of measurement when adopting this approach.

## Declaration of Competing Interest

## Acknowledgements

## Appendix A. missing data

No data.



**Fig. A1.** Missing circumstances and outcome. *Source: UKHLS Wave 6* Note: The missing values (NA) are shown for the following variables: health (sf12pcs_dv), father activity status at respondent's age of 14 (f_actstat), mother and father skill in occupation (m_skill_occ, f_skill_occ), mother and father education (mother_ed, father_ed).

**Fig. A2.** Missing efforts. *Source: UKHLS Waves 2 and 5* Note: The missing values (NA) are shown for the following variables: fruit units eaten per week (wkfruit), vegetable units eaten per week (wkvege), days walked at least 10 minutes (daywlk), ex-smoker (smoke_ex), sport activity (sportact), drinking alcohol at least 5 days per week (drink_alot).

## Appendix B. Principal Component Analysis for health-related behaviours

Our estimation approach relies on a summary concept of lifestyle which is defined aggregating information on a set of observable health-related behaviours. Each behavioural variable is described in Table 3.

We propose to summarise the overall health-related behaviours of people through a Principal Component Analysis. With this technique we extract the information on the correlations among the set of variables in order to summarise them in a single variable which maintains the highest variability across individuals in the dataset. Because all measures of behaviours are categorical the PCA has been conducted after computing the polychoric transformation of the mixed data to obtain a meaningful covariance matrix[9]. The polychoric correlation in statistics is used to estimate the correlation among two theorised continuous and normally distributed latent variables, which are observed as ordinal variables (Drasgow, 1986). This analysis is generally used together with factor analysis when working with variables representing self-reported rating scales with a small number of response options or Likert scales (Uebersax, 2006).

Table B.1 shows the polychoric correlation pattern among the behaviours considered in the definition of the lifestyle.

The correlations all have the expected signs: smoking intensity is negatively correlated with the intensity of healthy behaviours such as non-sedentary life, vegetable and fruit consumption. Heavy drinking is also negatively correlated with fruit consumption and sport activity, but shows a low but positive correlation with walking and vegetable consumption.

The polychoric correlation matrix is formally the input of the PCA[10]. The final effort definition we derive is the first component obtained with the analysis.

Fig. B.1 summarizes the results of the PCA. The Figure shows the scatterplot of the factor loadings of the first and second component respectively on the horizontal and vertical axes. The intensity of their contribution to the total variability is represented on a scale which ranges between 8 to 25%. The sign of the correlation of behaviours with the first component of the PCA appears to be coherent.[11]

The resulting first component of the PCA (Fig. B.1, at the bottom of y-axis) accounts for the 45.3% of the total variability of all effort dimensions.

---

[9] We use a statistical package implemented in R for performing the polychoric correlation from Fox (2019).

[10] We use the R built-in function *prcomp()* for computing the PCA.

[11] Given the positive correlation of the first PCA component with the risky behaviours, the lifestyle variable has been multiplied by (-1) in order to obtain a measure associated with having a healthier lifestyle.

**Fig. B1.** PCA for lifestyle and observed behaviours *Source: UKHLS Waves 2 and 5.*

**Table B1**
Polychoric correlations among health-related behaviours.

|  | Sport activity | Smoke intensity | Ex-smoker | Fruit per week | Vegetables per week | Days walked at least 10 min |
|---|---|---|---|---|---|---|
| Smoke intensity | −0.2132 |  |  |  |  |  |
| Ex-smoker | −0.0493 | −0.3310 |  |  |  |  |
| Fruit per week | 0.1729 | −0.4025 | 0.0902 |  |  |  |
| Vegetables per week | 0.1317 | −0.2637 | 0.1051 | 0.5401 |  |  |
| Days walked at least 10 min | 0.3201 | −0.0621 | −0.0377 | 0.1370 | 0.1434 |  |
| Drinking ≥ 5 days per week | −0.0434 | 0.1611 | 0.2063 | −0.0065 | 0.0759 | 0.0248 |

*Source:* UKHLS Waves 2 and 5

## Appendix C. Alternative effort measures

Although the definition of health-related behaviours is constrained by data availability, the effort measure adopted in this study is not the only possible option.

In this appendix we summarize and extensive sensitivity analysis performed over a range of alternative ways of measuring effort.

More specifically we repeat the entire analysis using the following alternative definitions of effort:

1. When observing a different behaviours in wave 2 and 5, instead of selecting the worst behaviour we select the more healthy behaviour.
2. When observing a different behaviours in wave 2 and 5, instead of selecting the worst behaviour we select randomly one of the two.
3. Add to the list of behaviours a variable of risk aversion available for a sub-sample of the respondents. Specifically, at UKHLS Wave 1, respondents are asked whether they are fully prepared to take risks or try to avoid taking risks in general; responses are collected on a scale from zero to 10, with zero referring to avoiding taking risks and 10 to being fully prepared to take risks. Note that the inclusion of this variable, due to high prevalence of missing, implies a reduction of 43% in sample size.
4. Remove a single behaviour at a time from the lifestyle measure.

**Table C.1**

Correlation of effort definitions with health.

| Effort definition | $\rho$ |
|---|---|
| Using worst behaviour | 0.267*** |
| Using best behaviours | 0.179*** |
| Choosing randomly the wave | 0.249*** |
| Including risk attitude | 0.259*** |
| Without Sport activity | 0.206*** |
| Without Smoke intensity | 0.218*** |
| Without Ex-smoker | 0.225*** |
| Without Fruit per week | 0.226*** |
| Without Vegetable per week | 0.223*** |
| Without Day walked at least 10 minutes | 0.234*** |
| Without Drinking $\geq$ 5 days per week | 0.218*** |

*Significance levels:* *p<0.05; **p<0.01; ***p<0.001 *Source:* UKHLS Waves 2 and 5

**Table C.2**

Correlation between Effort and alternative effort definitions.

| Effort definition | $\rho$ |
|---|---|
| Using best behaviour | 0.719*** |
| Choosing randomly the wave | 0.836*** |
| Including risk attitude | 0.980*** |
| Without Sport activity | 0.930*** |
| Without Smoke intensity | 0.607*** |
| Without Ex-smoker | 0.935*** |
| Without Fruit per week | 0.852*** |
| Without Vegetable per week | 0.924*** |
| Without Days walked at least 10 minutes | 0.929*** |
| Without Drinking $\geq$ 5 days per week | 0.918*** |

*Significance levels:* *p<0.05; **p<0.01; ***p<0.001 *Source:* UKHLS Waves 2 and 5

**Table C.3**

Unfair health inequality using alternative definitions of effort.

| Effort definitions | N. Types | $\overline{DU}$ | $\overline{FG}$ | $\%\Delta(DU)$ | $\%\Delta(FG)$ |
|---|---|---|---|---|---|
| Using worst behaviours | 12 | 2.26 | 5.75 | – | – |
| Using best behaviours | 12 | 2.54 | 4.15 | 12.39 | −27.83 |
| Using random behaviours | 13 | 2.34 | 4.44 | 3.54 | −22.78 |
| Including risk attitude | 9 | 1.92 | 4.98 | −15.04 | −13.39 |
| Without Sport activity | 13 | 2.79 | 4.77 | 23.45 | −17.04 |
| Without Smoke intensity | 10 | 2.71 | 4.50 | 19.91 | −21.74 |
| Without Ex-smoker | 8 | 2.15 | 5.98 | −4.87 | 4.00 |
| Without Fruit per week | 12 | 1.98 | 5.89 | −12.39 | 2.43 |
| Without Vegetable per week | 11 | 2.42 | 6.80 | 7.08 | 18.26 |
| Without Days walked at least 10 minutes | 12 | 2.32 | 5.29 | 2.65 | −8.00 |
| Without Drinking $\geq$ 5 days per week | 12 | 2.46 | 4.90 | 8.85 | −14.78 |

*Source:* UKHLS Waves 2 and 5. Note: The last two columns $\%\Delta(DU)$ and $\%\Delta(FG)$ show the percentage variation between model estimated with a certain effort definition and the baseline model adopting the effort definition used in the paper.

Table C.1 shows how the different definitions of effort relate with the outcome variable of the analysis. The strongest correlation is found in the first line is our baseline effort choice. However, all definitions are consistently and significantly correlated with health.

Effort definitions are also strongly correlated one to each other. The correlation ranges between 0.84 and 0.98 with the notable exception of the effort definition obtained removing smoking intensity behaviour from the list of behaviours considered ($\rho = 0.607$). This result is indeed not surprising given the relevant weight assigned to the smoking behaviour in shaping the PCA outcome in our baseline estimation (Figure B.1).

Table C.2 shows the correlation pattern between the effort definition adopted in the paper and the various alternative efforts.

Indeed using different definition of effort would affect our conclusions. Table C.3 reports number of types, average $UI_{DU}$, average $UI_{DU}$, and the relative difference with respect to our preferred baseline specification. Note that depending on the effort definition considered the adjustment is both downward and upward with our baseline close to the mid-range for both measures.

As additional robustness check we provide the estimation output of the MOB in the case in which each health-related behavioural variable is included in the estimation as a separated predictor.

**Table C.4**

Estimation output on MOB with health-related behaviours.

| names | Type 1 | Type 2 | Type 3 | Type 4 | Type 5 | Type 6 | Type 7 | Type 8 | Type 9 | Type 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| (Intercept) | −7.6400*** | −6.6900*** | −5.9980*** | −3.5920*** | −4.9480*** | −6.0110*** | −6.9920*** | −3.3900*** | −5.5660*** | −3.7610*** |
| Sport activity=1 | 1.3600 | −0.8170 | 4.0450** | 2.9120** | 0.2480 | 2.3090*** | 3.2420 | 3.3390*** | 3.1850** | 2.4650*** |
| Sport activity=2 | 2.3150 | 2.0940 | 4.4270** | 2.9430** | 2.6050* | 3.6460*** | 0.2090 | 3.3970*** | 4.0520*** | 3.1790*** |
| Sport activity=3 | 3.6220* | 2.1710 | 2.0730 | 2.4690** | 3.4920*** | 4.1340*** | 3.2050 | 4.0300*** | 4.3390*** | 3.7700*** |
| Sport activity=4 | 2.7710 | 5.4610* | −0.2960 | 4.1460*** | 4.9440*** | 4.8040*** | 6.6810*** | 4.6360*** | 4.6900*** | 3.9570*** |
| Sport activity=5 | 4.0260* | 3.6750 | 4.1960** | 2.6230** | 5.3900*** | 5.0670*** | 2.2920 | 4.0610*** | 6.9850*** | 5.1950*** |
| Sport activity=6 | 6.3590* | −1.7980 | 4.5360* | 4.1880*** | 4.3690* | 6.7070*** | 4.3150 | 5.3630*** | 5.8570*** | 5.5200*** |
| Sport activity=7 | −1.1720 | 4.0580 | 4.7130** | 3.8060*** | 5.3560*** | 6.7910*** | 4.7790 | 4.8110*** | 6.6600*** | 5.9200*** |
| Sport activity=8 | 5.8660 | 10.2190 | 5.9310* | 3.5370* | 5.9350* | 7.2160*** | 8.0800* | 4.5210*** | 5.4770** | 5.3980*** |
| Sport activity=9 | 12.0500 | −1.2020 | 2.2790 | 4.7090 | 3.5190 | 8.5430*** | 1.3670 | 5.6840*** | 8.2940* | 6.7930** |
| Sport activity=10 | 0.7090 | −5.2730 | 10.2770** | 7.6580*** | 8.2670* | 6.3250*** | 10.1820 | 5.7980*** | 11.4870*** | 5.9620** |
| Walking habits per month | 0.0620 | 0.0760 | 0.1060** | −0.0140 | 0.0970** | 0.1820*** | 0.2390*** | 0.0870*** | 0.0790* | 0.0540** |
| Drink alcohol ≥ 5 times per week | 4.7310 | 2.8410 | 2.6190 | −0.3530 | 0.9060 | 2.0060*** | 0.9660 | 1.2710* | 0.3350 | −0.7150 |
| Ex-smoker (dummy) | 0.4300 | 0.0550 | 0.7670 | 1.1010 | 0.0690 | −0.7810* | −1.1940 | −0.0620 | −1.0040 | −0.7860* |
| (0,10] cigarettes per day | 0.9570 | −2.6620 | −3.8270* | −1.2780 | −2.8820* | −1.3050 | −2.4760 | −0.7670 | −0.9410 | −0.0000 |
| (10,20] cigarettes per day | 3.3190 | 4.4920 | −5.0540** | −3.6340*** | −2.9190** | −1.8330*** | −1.9180 | −1.9490*** | −3.5030** | −2.4380*** |
| > 20 cigarettes per day | −9.0640* | | −5.4920 | −0.7690 | −2.4020* | −3.5560*** | −7.4380*** | −3.8680*** | −7.9370*** | −5.4380*** |
| Units of fruit per week | 0.7440*** | 0.0600 | −0.3560 | 0.2040 | −0.1740 | 0.1390* | 0.0850 | 0.1740* | −0.0250 | 0.0560 |
| Units of vegetable per week | −0.1390 | 0.1210 | 0.6900*** | 0.2440 | 0.2880 | 0.3700*** | 0.5150 | 0.1900* | 0.5520** | 0.2580* |

*Significance levels:* *p<0.05; **p<0.01; ***p<0.001 *Source:* UKHLS Waves 2 and 5

With this model we treat 'sport activity', 'days walked at least 10 minutes' as cardinal variables (expressed in days per month). 'Drink alcohol at least 5 times per week' and 'ex smoker' are dummy variables, and 'smoking intensity', 'fruit per week' and 'vegetable per week' are categorical variables.

The tree obtained has a rather similar structure to our preferred MOB. Similar circumstances appear shaping health opportunities and the number of types is only marginally smaller (10 types). However, the key weakness of such approach emerges from Table C.4 containing the regression output of the MOB. Coefficients are not statistically significant in about 50% of the cases. Moreover, insignificant coefficients tend to be erratic in magnitude and sign. Relying on this type-specific returns to effort would produce not simply noisy but even meaningless results.

## Appendix D. Latent Class and Finite Mixture Model Estimation

In order to exemplify how our estimation approach is able to unveil aspects of health inequality that would not be visible using alternative well established methods, we repeat the entire empirical exercise using two Finite Mixture Models (FMM). The identification of latent types to analyse unfair inequality in health is not new in the literature representing an obvious competitor for our approach (Frick et al., 2014).

In their analysis of data from the Health Survey for England from 2003 to 2012, Carrieri and Jones (2018) adopt a non-parametric approach to define types. This illustrates how the curse of dimensionality becomes a constraint when compared to semiparametric approaches such as the FMM and MOB specifications. With only four circumstance variables considered birth cohorts (3 categories), education (3 categories), sex (2 categories) and deprivation (2 categories) 36 distinct types are generated, some with quite small cell sizes making reliable inference difficult. Using data from Understanding Society (UKHLS), Carrieri et al (2020) find that a latent class model with only three latent types best fits the data, with the corresponding types characterised in terms of differences in their observed circumstances. The specification of these latent types assumes that the probability of class membership is a (multinomial logit) function of age-sex groups, along with parental occupational status and parental education to proxy childhood socioeconomic status. The number of latent classes is selected by in-sample measures of goodness of fit (AIC and BIC). These earlier studies are not directly comparable to the current analysis as the health outcomes are different and are standardised for age in this paper. Also, ethnicity is not used as a circumstance variable in the earlier work. However the importance of parental activity status and education in defining types is comparable. The application of MOB generates 12 types which is substantially more than the 3 types selected by the FMM specification in Carrieri et al., (2020) but much lower than the 36 types in the nonparametric analysis of Carrieri and Jones (2018).g

We consider two variants of the finite mixture models. The first identifies types using latent classes of a latent class model as suggested by Li Donni et al. (2015). The second includes health as concomitant variable to identify types as suggested by Carrieri et al. (2020). In both case we take advantage of the algorithm implementation developed by Linzer and Lewis (2011).

To select the most appropriate number of types we consider an increasing number of latent types and, for both methods, we select the number of types producing the smallest BIC value. Although both models belong to the family of finite mixture models and both identify latent classes, for sake of clarity we call LCA the former and FMMCV the model that includes health status as a concomitant variable. The two methods produce rather parsimonious partitions. Six LCA latent types and four

**Fig. D1.** Distribution of ethnicity (left) and father's occupation (right) across LCA types *Source: UKHLS Waves 2, 5, 6..*

FMMCV latent types. Notably, and differently from MOB, sex does not appear to correlate with the probability to belong to any of the latent classes defined by LCA and FFMCV.

For LCA the worst-off type in terms of expected health is type 1, it accounts for 8.5% of the population and contains over 75% of non-UK born respondents, vast majority of them had a non-working mother, parents with mainly no/low education, and over 60% define themselves as Asian. This type is only marginally better-off in terms of expected health then type 2. This is a small type (3.8% of the sample) with an extremely particular characteristics: respondents report the absence of their father from the household (both because deceased or for other reasons) and in some cases also the absence of their mother.

The best-off type is type 6 (12.4% of the sample). Made of respondent mainly born in UK, reporting high skill occupation for both their parents. Respondents in type 6 were born in the most favourable circumstances: mother's education is high for the relative majority of them, and over 80% fathers are reported having post-school qualification. This type contains mainly UK white but it also contains a small fraction of Irish white, black and mixed (Fig. D.1).

The other type containing a majority of non-UK born is type 4. The largest ethnic group in this type is 'other white'. Parental education is prominently high (post school qualification/tertiary), moreover a substantial share of mothers and fathers are reported in high and mid skill occupations. The expected health in type 4 is far above the value in type 1 (the other type made of first generation immigrants) but far below the expectation for individuals with similar socioeconomic characteristics born in the UK (type 6).

Finally, type 5 and type 3 make together over 60% of the sample. They are made of mainly UK-born white respondents and represent a sort of middle and lower middle class respectively. Type 5 is characterized by relatively worse socioeconomic circumstances: father and mother with mostly no school-qualification (while in type 3 the largest share has post-school qualification). Moreover, type 5 has a large share of respondents reporting parents working in lower skill occupations and larger share of non-working mothers.

Fig. D.1 shows the ethnic composition of the six LCA types and the father's occupation composition.

Fig. D.2 clarifies how different methods partition the population into partly similar structures of Roemerian types.

The most clearly emerging pattern is the large overlap between the two partitions based on FMMCV and LCA. The four types defined by the FMMCV model almost perfectly overlap with four of the six types defined by the LCA model (1, 3, 5, 6). This makes the comparisons of the two partitions straightforward. Individuals in LCA type 4 are all found in FMMCV type 1 and 4. While respondents in LCA type 2 belong to the largest extent to FMMCV type 2 and 3. The largest part of individuals born outside the UK belong to type 1. A non-negligible share of non-UK born are found also in FMMCV type 4, coming from LCA type 4, they are born outside UK but define themselves as 'white'.

Moving backward in Fig. D.2 to the MOB partition a rather complex picture emerges. The small worst-off types represent a sub-partitions of LCA type 1 and type 2. Similarly, LCA type 3 is made to a large extent by individuals in MOB type 3, 7, 8, 9. Moreover, LCA type 5 is made of respondents in MOB type 8, 9, 11, 12. But some substantial re-ranking arises for other MOB types. LCA type 4 for example is made of individuals belonging to MOB types with very different expected health outcome (including the two extreme types 1 and 12).

Table D.1 and Table D.2 show a rather consistent picture in terms of the model estimated in each terminal node. Although the range of variation of both coefficients is smaller than what estimated with MOB the sign, magnitude, and significance appear consistent across all types of both models. The main difference between the two models is that FMMCV produces a more parsimonious partition than LCA. However, both partitions of latent types are far more conservative than the partition obtained with MOB. Not surprisingly this leads to a substantial underestimation of the share of unfair inequality: compared

**Fig. D1.** Types composition across estimation methods *Source: UKHLS Waves 2, 5, 6..*

**Table D.1**
LCA model summary.

| Class | Intercept | p-value | Coefficient | p-value | Avg. Health | Sample |
|-------|-----------|---------|-------------|---------|-------------|--------|
| Type 1 | −7.1453 | 0.0000 | 1.4888 | 0.0000 | −2.81 | 1,544 |
| Type 2 | −8.7572 | 0.0000 | 2.3063 | 0.0000 | −2.78 | 675 |
| Type 3 | −7.4491 | 0.0000 | 2.4182 | 0.0000 | −0.359 | 6,169 |
| Type 4 | −5.1358 | 0.0000 | 1.5636 | 0.0000 | 0.236 | 1,009 |
| Type 5 | −4.8443 | 0.0000 | 1.865 | 0.0000 | 1.07 | 6,376 |
| Type 6 | −2.9852 | 0.0004 | 1.4499 | 0.0000 | 2.19 | 2,243 |

**Table D.2**
FMMCV model summary.

| Class | Intercept | p-value | Coeffiecient | p-value | Avg. Health | Sample |
|-------|-----------|---------|--------------|---------|-------------|--------|
| Type 1 | −7.491 | 0.0000 | 1.6284 | 0.0000 | −2.56 | 2,172 |
| Type 2 | −7.6551 | 0.0000 | 2.4303 | 0.0000 | −0.58 | 6,430 |
| Type 3 | −5.0839 | 0.0009 | 1.9054 | 0.0000 | 0.917 | 7,133 |
| Type 4 | −1.4555 | 0.0000 | 1.1631 | 0.0000 | 2.76 | 2,281 |

with what obtained with our preferred MOB specification $UI_{FG}$ is 85% and 79% lower using FMMCV and LCA respectively, while $UI_{DU}$ is 19% and 35% lower.

Table D.1 and Table D.2 report intercept, coefficients and the expected outcome for the two sets of latent types.

## References

Almond, D., Currie, J., Duque, V., 2018. Childhood circumstances and adult outcomes: act ii. J Econ Lit 56 (4), 1360–1446.
Ben-Shlomo, Y., Kuh, D., 2002. A life course approach to chronic disease epidemiology: conceptual models, empirical challenges and interdisciplinary perspectives. Int J Epidemiol 31 (2), 285–293 PMID: 11980781.
Brunori, P., Hufe, P., Mahler, D.G., 2018. The roots of inequality: Estimating inequality of opportunity from regression trees. The World Bank.
Brunori, P., Neidhöfer, G., 2020. The evolution of inequality of opportunity in Germany: a machine learning approach. ZEW-Centre for European Economic Research Discussion Paper (20–13).
Brunori, P., Peragine, V., Serlenga, L., 2019. Upward and downward bias when measuring inequality of opportunity. Soc Choice Welfare 52 (4), 635–661.
Brunori, P., Trannoy, A., Guidi, C.F., 2021. Ranking populations in terms of inequality of health opportunity: a flexible latent type approach. Health Econ 30 (2), 358–383.
Carrieri, V., Davillas, A., Jones, A.M., 2020. A latent class approach to inequity in health using biomarker data. Health Econ 29 (7), 808–826.
Carrieri, V., Jones, A.M., 2018. Inequality of opportunity in health: adecomposition-based approach. Health Econ 27 (12), 1981–1995.
Checchi, D., Peragine, V., 2010. Inequality of opportunity in italy. The Journal of Economic Inequality 8 (4), 429–450.
Cohen, G.A., 1989. On the currency of egalitarian justice. Ethics 99 (4), 906–944.
Conti, G., Heckman, J.J., Pinto, R., 2016. The effects of two influential early childhood interventions on health and healthy behaviour. The Economic Journal 126 (596), F28–F65.

Conti, G., Mason, G., Poupakis, S., 2020. Developmental origins of health inequality. In: M. Jones, A. (Ed.), The Oxford Encyclopedia of Health Economics, Vol. 1. Oxford University Press, New York, US, pp. 314–359.

Currie, J., Almond, D., 2011. Human capital development before age five. In: Handbook of Labor Economics, Vol. 4. Elsevier, pp. 1315–1486.

Davillas, A., Jones, A.M., 2020. Ex ante inequality of opportunity in health, decomposition and distributional analysis of biomarkers. J Health Econ 69, 102251.

Davillas, A., Jones, A.M., 2021. The first wave of the covid-19 pandemic and its impact on socioeconomic inequality in psychological distress in the uk. Health Econ 30 (7), 1668–1683.

Drasgow, F., 1986. Polychoric and polyserial correlations. In: The Encyclopedia of Statistics, Vol. 7. Wiley, pp. Pp.68–74.

Dworkin, R., 1981. What is equality? Part 1: equality of welfare. Philosophy & Public Affairs 185–246.

Eibich, P., 2015. Understanding the effect of retirement on health: mechanisms and heterogeneity. J Health Econ 43, 1–12.

Fleurbaey, M., 1995. Three solutions for the compensation problem. J Econ Theory 65 (2), 505–521.

Fleurbaey, M., 2008. Fairness, Responsibility, and Welfare. Oxford University Press.

Fleurbaey, M., Maniquet, F., 2012. Equality of opportunity: the economics of responsability. World Scientific Pub Co Inc.

Fleurbaey, M., Schokkaert, E., 2009. Unfair inequalities in health and health care. J Health Econ 28 (1), 73–90.

Fox, J., 2019. Package polycor. R package version 3.3. 0, URL http://CRAN. R-project. org/package= polycor.

Frick, H., Strobl, C., Zeileis, A., 2014. To split or to mix? tree vs. mixture models for detecting subgroups. In: COMPSTAT 2014 21st International Conference on Computational Statistics. Citeseer, p. 379.

Gravelle, H., 2003. Measuring income related inequality in health: standardisation and the partial concentration index. Health Econ 12 (10), 803–819.

Guber, R., 2019. Making it right? social norms, handwriting and human capital. Labour Econ 56, 44–57.

Hothorn, T., Zeileis, A., 2015. Partykit: a modular toolkit for recursive partytioning in r. Journal of Machine Learning Research 16, 3905–3909.

Jefferis, B.J., Power, C., Graham, H., Manor, O., 2004. Effects of childhood socioeconomic circumstances on persistent smoking. Am J Public Health 94 (2), 279–285.

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 Econ 22 (12), 1470–1495.

Lefranc, A., Pistolesi, N., Trannoy, A., 2009. Equality of opportunity and luck: definitions and testable conditions, with an application to income in france. J Public Econ 93 (11–12), 1189–1207.

Li Donni, P., Rodríguez, J.G., Dias, P.R., 2015. Empirical definition of social types in the analysis of inequality of opportunity: a latent classes approach. Soc Choice Welfare 44 (3), 673–701.

Linzer, D.A., Lewis, J.B., 2011. poLCA: An R Package for Polytomous Variable Latent Class Analysis. J Stat Softw 42 (10), 1–29. https://www.jstatsoft.org/v42/i10/.

Pudrovska, T., Anikputa, B., 2014. Early-life socioeconomic status and mortality in later life: an integration of four life-course mechanisms. Journals of Gerontology Series B: Psychological Sciences and Social Sciences 69 (3), 451–460.

Rawls, J., 1958. Justice as fairness. Philos Rev 67 (2), 164–194.

Rawls, J., 1971. 1971: A Theory of Justice. Cambridge, MA: Harvard University Press.

Roemer, J.E., 1998. Theories of distributive justice. Harvard University Press.

Roemer, J.E., 2002. Equality of opportunity: a progress report. Soc Choice Welfare 19 (2), 455–471.

Roemer, J.E., Trannoy, A., 2015. Equality of opportunity. In: Handbook of Income Distribution, Vol. 2. Elsevier, pp. 217–300.

Rosa Dias, P., 2009. Inequality of opportunity in health: evidence from a uk cohort study. Health Econ 18 (9), 1057–1074.

Rosa Dias, P., 2010. Modelling opportunity in health under partial observability of circumstances. Health Econ 19 (3), 252–264.

Schmitz, H., 2011. Why are the unemployed in worse health? the causal effect of unemployment on health. Labour Econ 18 (1), 71–78.

Sen, A., 1980. Equality of what? The Tanner lecture on Human Values 1, 197–220.

Sutton, M., 2002. Vertical and horizontal aspects of socio-economic inequity in general practitioner contacts in Scotland. Health Econ 11 (6), 537–549.

Taylor, S.E., 2010. Mechanisms linking early life stress to adult health outcomes. Proceedings of the National Academy of Sciences 107 (19), 8507–8512.

Uebersax, J.S., 2006. The tetrachoric and polychoric correlation coefficients. Statistical methods for rater agreement web site.

Van Buuren, S., Groothuis-Oudshoorn, K., 2011. Mice: multivariate imputation by chained equations in R. J Stat Softw 45, 1–67.

Wagstaff, A., Van Doorslaer, E., 2000. Equity in health care finance and delivery. Handbook of Health Economics 1, 1803–1862.

Ware, J., Kosinski, M., Keller, S., 2001. SF-36 physical and mental health summary scales, 2nd edition. Boston: The Health Institute, New England Medical Center.

Zeileis, A., Hornik, K., 2007. Generalized M-fluctuation tests for parameter instability. Stat Neerl 61 (4), 488–508.

Zeileis, A., Hothorn, T., Hornik, K., 2008. Model-based recursive partitioning. Journal of Computational and Graphical Statistics 17 (2), 492–514.

Zeileis, A., Hothorn, T., Hornik, K., 2010. Party with the mob: model-based recursive partitioning in r. R package version 3.0.0.

Ziebarth, N., 2010. Measurement of health, health inequality, and reporting heterogeneity. Social Science & Medicine 71 (1), 116–124.