

The Roots of Inequality: Estimating Inequality of Opportunity from Regression Trees and Forests*

Paolo Brunori, Paul Hufe, and Daniel Mahler

Abstract

We propose the use of machine learning methods to estimate inequality of opportunity and illustrate that regression trees and forests represent a substantial improvement over existing approaches: they reduce the risk of ad-hoc model selection and trade off upward and downward bias in inequality of opportunity estimates. The advantages of regression trees and forests are illustrated by an empirical application for a cross-section of 31 European countries. We show that arbitrary model selection may lead to significant biases in inequality of opportunity estimates relative to our preferred method. These biases are reflected in both point estimates and country rankings.

JEL-Codes: D31; D63; C38

Keywords: Equality of Opportunity; Machine Learning; Random Forests

***Brunori** (corresponding author): International Inequalities Institute, London School of Economics. Postal Address: London School of Economics, Houghton Street, London, WC2A 2AE, UK (paolo.brunori@unifi.it); **Hufe**: University of Bristol, IZA, CESifo (paul.hufe@bristol.ac.uk); **Mahler**: World Bank (dmahler@worldbank.org). We thank Chiara Binelli, Marc Fleurbaey, Torsten Hothorn, Niels Johannesen, Andreas Peichl, Giuseppe Pignataro, Dominik Sachs, Jan Stuhler, Dirk Van de gaer, and Achim Zeileis for useful comments and suggestions. Furthermore, we are grateful for the comments received from seminar audiences at Princeton University, the University of Perugia, the University of Essex, the World Bank, ifo Munich, the University of Copenhagen, Canazei Winter School 2018, the European Commission JRC at Ispra, the EBE Meeting 2018, IIPF 2018, and the Equal Chances Conference in Bari. Any errors remain our own.

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process which may lead to differences between this version and the [Version of Record](#). Please cite this article as doi: [10.1111/sjoe.12530](https://doi.org/10.1111/sjoe.12530)

1 INTRODUCTION

Equality of opportunity is an important ideal of distributive justice. It has widespread support in the general public and its realization has been identified as an important goal of public policy intervention (Cappelen et al., 2007; Alesina et al., 2018; Corak, 2013; Chetty et al., 2016). In spite of its popularity, providing empirical estimates of equality of opportunity is notoriously difficult. Next to normative dissent about the precise factors that should be viewed as contributing to unequal opportunities, current estimation approaches are encumbered by ad-hoc model selection that lead researchers to over- or underestimate inequality of opportunity.

In this paper we propose the use of machine learning methods to overcome the issue of ad-hoc model selection. Machine learning methods allow for flexible models of how unequal opportunities come about while imposing statistical discipline through criteria of out-of-sample replicability. These features serve to establish inequality of opportunity estimates that are less prone to upward or downward bias.

The empirical literature on the measurement of unequal opportunities has been flourishing since John Roemer's (1998) seminal contribution, *Equality of Opportunity*. At the heart of Roemer's formulation is the idea that individual outcomes are determined by two sorts of factors: those factors over which individuals have control, which he calls *effort*, and those factors for which individuals cannot be held responsible, which he calls *circumstances*. While outcome differences due to effort exertion are morally permissible, differences due to circumstances are inequitable and call for compensation.¹ Grounded on this distinction, inequality of opportunity measures quantify the extent to which individual outcomes are predicted by circumstance characteristics. They are usually calculated in a two-step procedure. First, researchers predict an outcome of interest from observable circumstances. Second, they calculate inequality in the distribution of predicted outcomes: the more predicted outcomes diverge, the more are circumstances associated with outcomes, and the more inequality of opportunity there is.

¹The distinction between circumstances and efforts underpins many prominent literature branches in economics such as the ones on intergenerational mobility (Chetty et al., 2014a,b), the gender pay gap (Blau and Kahn, 2017) and racial differences (Kreisman and Rangel, 2015). For different notions of equality of opportunity, see Arneson (2018).

Accepted Article

Current approaches to estimate inequality of opportunity suffer from biases that are the consequence of critical choices in model selection. First, researchers have to decide which circumstance variables to consider for estimation.² The challenge of this task grows with increasing availability of high-quality datasets that provide very detailed information with respect to individual circumstances (Björklund et al., 2012; Hufe et al., 2017). On the one hand, discarding relevant circumstances from the estimation model limits the explanatory scope of circumstances and leads to downward biased estimates of inequality of opportunity (Ferreira and Gignoux, 2011). On the other hand, including too many circumstances overfits the data and leads to upward biased estimates of inequality of opportunity (Brunori et al., 2019). Second, researchers must choose a functional form according to which circumstances co-produce the outcome of interest. For example, it is a well-established finding that the influence of socio-economic disadvantages during childhood on life outcomes varies by biological sex (Chetty et al., 2016; Dahl and Lochner, 2012). In contrast to such evidence, many empirical applications presume that the effect of circumstances on individual outcomes is log-linear and additive while abstracting from possible interaction effects (Bourguignon et al., 2007; Ferreira and Gignoux, 2011). On the one hand, restrictive functional form assumptions limit the ability of circumstances to explain variation in the outcome of interest and thus force a downward bias on inequality of opportunity estimates. On the other hand, limitations in the available degrees of freedom may prove a statistically meaningful estimation of complex models with many parameters infeasible.

This discussion highlights the non-trivial challenge of selecting the appropriate model for estimating inequality of opportunity. Researchers must balance different sources of bias while avoiding ad-hoc solutions. While this task is daunting for the individual researcher, it is a standard application for machine learning algorithms that are designed to make out-of-sample predictions of a dependent variable based on a number of observable predictors. In this paper, we use conditional inference regression trees and forests to estimate inequality of opportunity (Hothorn et al., 2006). Introduced by Morgan and Sonquist (1963) and later popularized by Breiman et al. (1984) and Breiman (2001), they belong to a set of machine learning methods that is increasingly integrated into the statistical toolkit of economists (Varian, 2014; Mullainathan

²Roemer does not provide a fixed list of circumstance variables. Instead he suggests that the set of circumstances should evolve from a political process (Roemer and Trannoy, 2015). In empirical implementations typical circumstances include biological sex, socioeconomic background, and race.

and Spiess, 2017; Athey, 2018). Trees and forests obtain predictions by drawing on a clear-cut algorithm that imposes only minimal assumptions about which and how circumstances interact in shaping individual opportunities. Thereby, they restrict judgment calls of the researcher and inform model specification by data analysis. As a consequence, they cushion downward bias by flexibly accommodating different ways of how circumstance characteristics shape the distribution of outcomes. Moreover, the conditional inference algorithm branches trees (and constructs forests) by a sequence of hypothesis tests that prevents the inclusion of noisy circumstance parameters. This feature reduces the potential for upward biased estimates of inequality of opportunity through model overfitting. Hence, regression trees and forests address the detrimental consequences of ad-hoc model selection in a way that is sensitive to both upward and downward bias in inequality of opportunity estimates.

To showcase the advantages of regression trees and forests we compare them to existing estimation approaches in a cross-sectional dataset of 31 European countries. We demonstrate that current estimation approaches overfit (underfit) the data which in turn leads to upward (downward) biased estimates of inequality of opportunity. These biases are sizable. For example, some standard methods overestimate inequality of opportunity in the Nordic countries while they underestimate the extent of inequality of opportunity in Germany and France. As a consequence, these countries appear close in terms of their opportunity characteristics. Hence, standard estimation approaches may yield misleading information about the level of inequality of opportunity in different societies to policymakers and the general public alike.

While we demonstrate the advantages of regression trees and forests for estimations of inequality of opportunity, they are not a panacea to empirical challenges in this literature. First, regression trees and forests cannot address one of the most relevant sources of downward bias in inequality of opportunity estimates: missing data on relevant circumstances. Second, although regression trees and forests are less prone to upward and downward bias, the remaining bias may nevertheless be substantial when samples are small. Therefore, we encourage applied researchers to exercise caution when estimating inequality of opportunity on data with small number of observations.

The remainder of this paper is organized as follows: section 2 gives a brief introduction to

current empirical approaches in the literature on inequality of opportunity. Section 3 introduces conditional inference regression trees and forests, and illustrates how to use them in the context of inequality of opportunity estimations. An empirical illustration based on simulated data and the EU Survey of Income and Living Conditions is contained in section 4. In this section we also highlight the particular advantages of tree- and forest-based estimation methods by comparing them to the prevalent estimation approaches in the literature. Section 5 concludes the paper.³

2 EMPIRICAL APPROACHES TO EQUALITY OF OPPORTUNITY

Theoretical Set-up and Notation. Consider a population $\mathcal{N} = \{1, \dots, N\}$ and an associated vector of non-negative incomes $y = (y_1, \dots, y_N)$. Income is determined by two sets of factors: *circumstances* beyond individual control and individual *efforts*. We define the $(P \times 1)$ -vector $\omega_i \in \Omega$ as a comprehensive description of the circumstances of $i \in \mathcal{N}$. Analogously we define the $(Q \times 1)$ -vector $\theta_i \in \Theta$ as a comprehensive description of the efforts that are exerted by $i \in \mathcal{N}$. The income generating function can be defined as follows:

$$y = d(\omega, \theta). \quad (1)$$

Based on the realizations of individual circumstances, the population can be partitioned into *types*. We define the type partition $\mathcal{T} = \{t_1, \dots, t_M\}$, such that individuals are member of one type if they share the same circumstances: $i, j \in t_m \iff \omega_i = \omega_j$.

Measurement. Opportunity egalitarians are averse to inequalities that are rooted in circumstances, however, they are indifferent to inequalities that originate from individual effort exertion. In spite of the intuitive appeal of this idea, the literature has suggested a variety of formulations that differ in their precise normative content—see Ramos and Van de gaer (2016) for an overview. In this work we exclusively focus on *ex-ante utilitarian* measures of inequal-

³In a parallel paper, Blundell and Risa (2019) apply machine learning methods to the estimation of intergenerational mobility. In particular, they assess the completeness of rank-rank estimates of intergenerational mobility as measures of equal opportunities. In contrast to their work, we directly estimate inequality of opportunity statistics. Therefore, our focus is not on downward bias that follows from focusing on one circumstance only (i.e. parental income) but on balancing both downward and upward bias if the set of available circumstances is large in relation to a given sample size.

ity of opportunity (Van de gaer, 1993; Checchi and Peragine, 2010). They are the most widely applied formulations in the empirical literature.⁴

According to the ex-ante utilitarian view, the value of a type's opportunity set is pinned down by the expected value of its outcomes, $\mathbb{E}[y|\omega]$. Thus, the distribution of opportunities in a population can be expressed by the following counterfactual distribution y^C :

$$y^C = (y_1^C, \dots, y_i^C, \dots, y_N^C) = (\mathbb{E}[y_1|\omega_1], \dots, \mathbb{E}[y_i|\omega_i], \dots, \mathbb{E}[y_N|\omega_N]). \quad (2)$$

From this distribution one can construct ex-ante utilitarian measures of inequality of opportunity by choosing any functional $I()$ that satisfies the following two properties:

1. $I(y^C)$ decreases (increases) through transfers from i to j if i is from a circumstance type with a higher (lower) expected value of outcomes than the recipient j .
2. $I(y^C)$ remains unaffected by transfers from i to j if they are members of the same type.

In most empirical applications $I()$ represents an inequality index satisfying the standard properties of anonymity, the principle of transfers, population replication, and scale invariance (Cowell, 2016).⁵ Examples of the latter are the Gini index or any member of the generalized entropy class. Note that the choice of $I()$ is normative in itself as it specifies the extent of inequality aversion at different points of the counterfactual distribution y^C . For example, the mean logarithmic deviation (MLD) values compensating transfers to the most disadvantaged types more than the Gini index. In this work we are agnostic about the normatively correct choice of $I()$. While we present our main results in terms of the Gini index, we provide robustness checks based on other inequality indexes in Supplementary Material S.6.

⁴The use of machine learning methods is not restricted to ex-ante utilitarian formulations and can be easily extended to alternative measures of inequality of opportunity.

⁵The β coefficient from intergenerational mobility regressions can also be interpreted as an ex-ante utilitarian measure of inequality of opportunity. In the intergenerational mobility framework, $\beta = \frac{E(y_{ic}|y_{ip})}{y_{ip}}$, where y_{ip} represents parental income as the sole circumstance. Hence, the functional applied to the distribution of conditional expectations can be written as $I() = \frac{1}{y_{ip}}$. Note that β decreases (increases) through transfers from children from advantaged (disadvantaged) backgrounds to children from less (more) advantaged backgrounds. However, β remains unaffected by transfers between children from parental households with equal y_{ip} .

Estimation. Given the measurement decisions described above, we require an estimate of the conditional distribution y^C . The data generating process (DGP) described in equation (1) can be rewritten as follows:

$$y = d(\omega, \theta) = f(\omega) + \epsilon = \mathbb{E}(y|\omega) + \epsilon. \quad (3)$$

$\mathbb{E}(y|\omega)$ captures unfair variation due to observed circumstances. The iid error term ϵ captures both fair (individual effort) and unfair (unobserved circumstances) determinants of individual outcomes; hence resulting measures of inequality of opportunity have a lower bound interpretation.

Estimating y^C is a prediction task in which the researcher tries to answer the following question: What outcome y_i do we expect for an individual that faces circumstances ω_i ? The precise form of $f(\cdot)$ is a priori unknown. In the vast majority of empirical applications, researchers address this lack of knowledge by invoking strong functional form assumptions. For example, they perform a log-linear regression of the outcome of interest on the set of observed circumstances and construct an estimate for y^C from the predicted values:

$$\ln(y_i) = \beta_0 + \sum_{p=1}^P \beta_p \omega_i^p + \epsilon_i, \quad (4)$$

$$\hat{y}_i^C = \exp \left[\beta_0 + \sum_{p=1}^P \hat{\beta}_p \omega_i^p \right]. \quad (5)$$

The literature refers to this estimation procedure as the *parametric approach* (Bourguignon et al., 2007; Ferreira and Gignoux, 2011).⁶

According to another procedure, the researcher partitions the sample into mutually exclusive types based on the realizations of all circumstances under consideration. An estimate for y^C is then constructed from average incomes within types:

$$\hat{y}_i^C = \mu_{m(i)} = \frac{1}{N_m} \sum_{j=1}^{N_m} y_j, \quad \forall j \in t_m, \quad \forall t_m \in \mathcal{T}. \quad (6)$$

⁶The logarithmic transformation is not innocuous as the marginal impact of circumstances on incomes may differ from their impact on log-incomes. Therefore, the predicted outcome should be obtained by applying the correction suggested in Blackburn (2007). This correction, however, is rarely implemented in empirical applications.

The literature refers to this estimation procedure as the *non-parametric approach* (Checchi and Peragine, 2010).

Both approaches face empirical challenges that are typically resolved by discretionary decisions of the researcher. For example, the parametric approach assumes a log-linear impact of all circumstances and therefore neglects the existence of interdependencies between circumstances and other non-linearities. To alleviate this shortcoming the researcher may integrate interaction terms and higher order polynomials into equation (4). However, such extensions remain at her discretion. Reversely, the non-parametric approach does not restrict the interdependent impact of circumstances. However, if the data is rich enough in information on circumstances, the researcher may be forced to reduce the observed circumstance vector to obtain statistically meaningful estimates of the relevant parameters.⁷ The necessary process of restricting the circumstance vector again remains at the researcher's discretion.

The previous discussion illustrates that common approaches leave the researcher to her own devices when selecting the best model for estimating the distribution y^C . In this paper, we provide an automated solution to this problem. Similarly, Li Donni et al. (2015) propose the use of latent class modeling to obtain type partitions that allow for estimates of y^C according to the non-parametric procedure outlined in equation (6). In their approach, observable circumstances are considered indicators of membership in an unobservable latent type. For each possible number of latent types individuals are assigned to types so as to minimize the within-type correlation of observable circumstances. Then the optimal number of types, M^* , is selected by minimizing an appropriate model selection criterion such as Schwarz's Bayesian Information Criterion (BIC). The latent class approach therefore partly solves the issue of arbitrary model selection. However, it has important drawbacks. First, it cannot solve the problem of model selection once the potential number of types exceeds the available degrees of freedom. In such cases, the latent class approach replicates the limitations of parametric and non-parametric approaches: the researcher must pre-select circumstances and their subpartition. Second, latent classes are obtained by minimizing the within-type correlation of circumstances while ignoring the correlation of circumstance variables with the outcome variable. As a consequence, they are

⁷Assume the researcher observes ten circumstance variables with three expressions each—a quantity easily observed in many data sets. The non-parametric approach would require the estimation of $3^{10} = 59,049$ group means.

likely to underfit the data leading to downward biased estimates of inequality of opportunity (Lanza et al., 2013).

In the following section, we will discuss how regression trees and forests address the outlined shortcomings of existing estimation approaches.

3 ESTIMATING INEQUALITY OF OPPORTUNITY FROM REGRESSION TREES AND FORESTS

Regression trees and forests belong to the class of supervised learning methods that were developed to make out-of-sample predictions of a dependent variable based on a number of observable predictors. As we will outline in the following, they can be straightforwardly applied to inequality of opportunity estimations and solve the issue of model selection.

First, we will introduce conditional inference regression trees. By providing predictions based on identifiable groups, they closely connect to Roemer's theoretical formulation of inequality of opportunity.⁸ Second, we will introduce conditional inference forests, which are—loosely speaking—a collection of many conditional inference trees. While forests do not have the intuitive appeal of regression trees, they perform better in terms of out-of-sample prediction accuracy and hence provide better estimates of the counterfactual distribution y^C .

3.1 Conditional Inference Trees

Trees obtain predictions for outcome y as a function of input variables $x = (x^1, \dots, x^k)$. They use the sample $\mathcal{S} = \{(y_i, x_i)\}_{i=1}^S$ to divide the population into non-overlapping groups, $\mathcal{G} = \{g_1, \dots, g_m, \dots, g_M\}$, where each group g_m is homogeneous in the expression of some input variables. These groups are called *terminal nodes* or *leafs*. The conditional expectation for observation i is estimated from the mean outcome $\hat{\mu}_m$ of the group g_m to which i is assigned. Hence, in addition to the observed outcome vector $y = (y_1, \dots, y_i, \dots, y_N)$ one obtains a vector of predicted

⁸Furthermore, their simple graphical illustration may be an instructive tool for comparisons of opportunity structures in different societies.

values $\hat{y} = (\hat{f}(x_1), \dots, \hat{f}(x_i), \dots, \hat{f}(x_N))$, where

$$\hat{f}(x_i) = \hat{\mu}_{m(i)} = \frac{1}{N_m} \sum_{j \in g_m} y_j. \quad (7)$$


The mapping from regression trees to equality of opportunity estimation is straightforward. If the input variables $x = (x^1, \dots, x^k)$ are circumstances only, each resulting group $g_m \in \mathcal{G}$ can be interpreted as a circumstance type $t_m \in \mathcal{T}$. Furthermore, \hat{y} is analogous to an estimate of the counterfactual distribution y^C that underpins the construction of ex-ante utilitarian measures of inequality of opportunity.

Tree Construction. Regression trees partition the sample into M types by *recursive binary splitting*. Recursive binary splitting starts by dividing the full sample into two distinct groups according to the value they take in one input variable $\omega^p \in \Omega$. If ω^p is a continuous or ordered variable, then $i \in t_l$ if $\omega_i^p < \tilde{\omega}^p$ and $i \in t_m$ if $\omega_i^p \geq \tilde{\omega}^p$, where $\tilde{\omega}^p$ is a splitting value chosen by the algorithm. If ω^p is a categorical variable then the categories can be split into any two arbitrary groups. The process is continued such that one of the two groups is divided into further subgroups (potentially based on another $\omega^q \in \Omega$), and so on. Graphically, this division into groups can be presented like an upside-down tree (Figure 1).

The exact manner in which the split is conducted depends on the type of regression tree that is used. In this paper, we follow the conditional inference methodology proposed by Hothorn et al. (2006). Conditional inference trees are grown by a series of permutation tests according to the following 4-step procedure:

0. Choose a significance level α^* .
1. Test the null hypothesis of density function independence: $H_0^{\omega^p} : D(y|\omega^p) = D(y)$, for all $\omega^p \in \Omega$, and obtain a p -value associated with each test, p^{ω^p} .
 \Rightarrow Adjust the p -values for multiple hypothesis testing, such that $p_{adj}^{\omega^p} = 1 - (1 - p^{\omega^p})^P$ (Bonferroni Correction).

FIGURE 1 – Exemplary Tree Representation



Note: Artificial example of a regression tree. Gray boxes indicate splitting points; white boxes indicate terminal nodes. The values inside terminal nodes show estimates for the conditional expectation y^C .

2. Select the variable ω^* with the lowest p -value, i.e.

$$\omega^* = \underset{\omega^p}{\operatorname{argmin}} \{ p_{adj.}^{\omega^p} : \omega^p \in \Omega, p = 1, \dots, P \}.$$

\Rightarrow If $p_{adj.}^{\omega^*} > \alpha^*$: Exit the algorithm.

\Rightarrow If $p_{adj.}^{\omega^*} \leq \alpha^*$: Continue, and select ω^* as the splitting variable.

3. Test the null hypothesis of density function independence between the subsamples for each possible binary partition splitting point s based on ω^* , and obtain a p -value associated with each test, $p^{\omega_s^*}$.

\Rightarrow Split the sample based on ω^* , by choosing the splitting point s that yields the lowest p -value, i.e. $\tilde{\omega}^* = \underset{\omega_s^*}{\operatorname{argmin}} \{ p^{\omega_s^*} : \omega_s^* \in \Omega \}$.

4. Repeat steps 1.–3. for each of the resulting subsamples.

In words, conditional inference trees start by a series of univariate hypothesis tests. The circumstance that is most related to the outcome is chosen as the potential splitting variable. If

the dependence between the outcome and the splitting variable is sufficiently strong, then a split is made. If not, no split is made. Whenever a circumstance can be split in several ways, the sample is split into two subsamples such that the dependence with the outcome variable is maximized. This procedure is repeated in each of the two subsamples until no circumstance in any subsample is sufficiently related to the outcome variable. Note that the depth of the resulting opportunity tree hinges on the level of α^* . The less stringent the α^* -requirement, the more we allow for false positives, i.e. the more splits will be detected as significant and the deeper the tree will be grown. In our empirical application we fix $\alpha^* = 0.01$, which is in line with the disciplinary convention for hypothesis tests. To illustrate the robustness of this choice we show comparisons to setting $\alpha^* = 0.05$ and choosing α^* through cross-validation in Supplementary Figure S.1.

A particular advantage of trees is that they avoid list-wise deletion of observations by implementing surrogate splits. In case of missing data, the algorithm searches for an alternative splitting point that mimicks the sample partition based on $\tilde{\omega}^*$ to the greatest extent. All observations that lack information on $\tilde{\omega}^*$ are then allocated to subbranches based on this surrogate splitting point.

3.2 *Conditional Inference Forests*

Regression trees provide a simple and standardized way of dividing the population into types. Therefore, they solve the model selection problem outlined in section 2. However, trees suffer from three shortcomings: first, the structure of trees—and therefore the estimate of y^C —is fairly sensitive to alternations in data samples. This issue is particularly pronounced if there are various circumstances that are close competitors for defining the first splits (Friedman et al., 2009). Second, trees assume a non-linear data generating process that imposes interactions while ruling out the linear influence of circumstances. Third, trees make inefficient use of data since some of the circumstances $\omega^p \in \Omega$ are not used for the construction of the tree. However, circumstances may possess informational content that can increase predictive power even if they are not significantly associated with y at level α^* . This becomes an issue if two or more important circumstances are highly correlated. Once a split is made using either of the two,

it is unlikely that the other contains enough information to cause another split. Conditional inference forests address all of these shortcomings (Breiman, 2001; Biau and Scornet, 2016).

Forest Construction. Random forests create many trees and average over all of these when making predictions. Trees are constructed according to the same 4-step procedure outlined in the previous subsection. However, two tweaks are made. First, given the sample $\mathcal{S} = \{(y_i, \omega_i)\}_{i=1}^S$ each tree is estimated on a random subsample $\mathcal{S}' \subset \mathcal{S}$. In our application, we randomly select approximately 60% of the observations for each tree, and estimate B^* such trees in total. Second, only a random subset of circumstances of cardinality \bar{P}^* is allowed to be used at each splitting point. Together these two tweaks remedy the shortcomings of single conditional inference trees. First, averaging over B^* predictions cushions variance in the estimates of y^C and smooths the non-linear impact of circumstance characteristics. Second, using subsets of all circumstance variables increases the likelihood that all observed circumstances with informational content will be identified as splitting variable ω^* at some point.

Predictions are formed as follows:

$$\hat{f}(\omega; \alpha^*, \bar{P}^*, B^*) = \frac{1}{B^*} \sum_{b=1}^{B^*} \hat{f}^b(\omega; \alpha^*, \bar{P}^*). \quad (8)$$

Equation (8) illustrates that individual predictions are a function of α^* —the significance level governing the implementation of splits, \bar{P}^* —the number of circumstances to be considered at each splitting point, and B^* —the number of subsamples drawn from the data. In our empirical illustration we fix $B^* = 200$ and determine α^* and \bar{P}^* by minimizing the *out-of-bag* error (MSE^{OBB}). Details on these choices and empirical procedures are disclosed in Supplementary Material S.1.

4 EMPIRICAL APPLICATION

In this section we illustrate the machine learning approach using harmonized survey data from 31 European countries. We compare the results from trees and forests with results from the prevalent estimation approaches in the extant literature; namely parametric, non-parametric

and latent class models. Comparisons are made along two dimensions.

First, we evaluate the different estimation approaches by comparing their out-of-sample mean squared error (MSE^{Test}). MSE^{Test} is a standard statistic to evaluate the prediction quality of estimation models.⁹ To calculate MSE^{Test} , we follow the machine learning practice of splitting our sample into a *training set* with $i^{-H} \in \{1, \dots, N^{-H}\}$ and a *test set* with $i^H \in \{1, \dots, N^H\}$. For each sample, we choose $N^{-H} = \frac{2}{3}N$ and $N^H = \frac{1}{3}N$.¹⁰ We fit our models on the training set and compare their performance on the test set according to the following procedure:

1. Run the model on the training data (for the specific estimation procedures, see section 3.1 for trees and forests, and section 4.2 for our benchmark methods).
2. Store the prediction function $\hat{f}^{-H}()$.
3. Calculate the mean squared error in the test set:

$$\text{MSE}^{\text{Test}} = \frac{1}{N^H} \sum_{i \in H} [y_i - \hat{f}^{-H}(\omega_i)]^2.$$

Second, we evaluate the different approaches by comparing inequality of opportunity estimates. To this end, we run the models on all data for a country, and apply the resulting prediction functions $\hat{f}()$ to obtain \hat{y}^C . Estimates of inequality of opportunity are derived by summarizing \hat{y}^C with the Gini index. Estimates for alternative inequality indexes are presented in Supplementary Material S.6.

4.1 Data

We base our empirical illustration on the 2011 wave of the European Union Statistics on Income and Living Conditions (EU-SILC). EU-SILC provides harmonized survey data with respect to

⁹Minimizing MSE^{Test} is equivalent to trading-off upward and downward biases of inequality of opportunity estimates in a given data environment: The more parsimonious the model, the higher the prediction bias (underfitting) and the stronger the downward bias in inequality of opportunity estimates. The more complex the model, the higher the prediction variance (overfitting) and the stronger the upward bias of inequality of opportunity estimates. We provide a thorough illustration of this mapping in Supplementary Material S.2.

¹⁰Note that the size of the training set for each country is constant regardless of the estimation method. Hence, any cross-method differences in prediction accuracy are not driven by differences in sample size.

Accepted Article

income, poverty, and living conditions. It is the official reference source for comparative statistics on income distribution and social inclusion in the European Union. In its 2011 wave EU-SILC covers a cross-section of 31 European countries. For each country, it contains a random sample of all resident private households. Data is collected by national statistical agencies following common variable definitions and data collection procedures. We use the 2011 wave since it contains an ad-hoc module about the intergenerational transmission of (dis)advantages. This module allows us to construct finely-grained circumstance type partitions. Observed circumstances Ω and their respective expressions are listed in Table 1. We include all variables of EU-SILC containing information about the respondent's characteristics at birth and their living conditions during childhood. Descriptive statistics of circumstance variables are reported in Supplementary Material S.5.¹¹

The unit of observation is the individual and the outcome of interest is equivalized disposable household income. We obtain the latter by dividing household disposable income with the square root of household size. Reported incomes refer to the year preceding the survey wave, i.e. 2010 in the case of our empirical application. In line with the literature we focus on equivalized household income as it provides the closest income analogue to consumption possibilities and general economic well-being. Inequality statistics tend to be heavily influenced by outliers (Cowell and Victoria-Feser, 1996); therefore we adopt a standard winsorization method according to which we set all non-positive incomes to 1 and scale back all incomes exceeding the 99.5th percentile of the country-specific income distribution to this lower threshold. Our analysis is focused on the working age population. Therefore, we restrict the sample to respondents aged between 30 and 60. To assure the representativeness of our inequality of opportunity estimates we use individual cross-sectional weights when calculating $I(\hat{y}^C)$.

Table 2 shows considerable heterogeneity in income distributions across Europe. While the average households in Norway and Switzerland obtained incomes above €40,000 in 2010, the average household income in Romania, Bulgaria and Lithuania did not exceed the €5,000 mark. Lowest inequality prevails in Norway, Iceland and Denmark, all of which have Gini coeffi-

¹¹In contrast to some existing work we do not consider age as a circumstance (among others Checchi et al., 2016). This choice is motivated by the fact that cross-sectional income disparities across age groups even out across the life-cycle of individuals. In Supplementary Material S.7, we provide robustness analyses based on income distributions that are residualized from variation across age groups. Our conclusions remain unaffected.

TABLE 1 – List of Circumstances

<p>1. Respondent's sex:</p> <ul style="list-style-type: none"> - Male - Female <p>2. Respondent's country of birth:</p> <ul style="list-style-type: none"> - Respondent's present country of residence - European country - Non-European country <p>3. Presence of parents at home*:</p> <ul style="list-style-type: none"> - Both present - Only mother - Only father - Without parents - Lived in a private household without any parent <p>4. Number of adults (aged 18 or more) in respondent's household*</p> <p>5. Number of working adults (aged 18 or more) in respondent's household*</p> <p>6. Number of children (under 18) in respondent's household*</p> <p>7. Father's/mother's country of birth and citizenship:</p> <ul style="list-style-type: none"> - Born/citizen of the respondent's present country of residence - Born/citizen of another EU-27 country - Born/citizen of another European country - Born/citizen of a country outside Europe <p>8. Father's/mother's education (based on the International Standard Classification of Education 1997 [ISCED-97]):</p> <ul style="list-style-type: none"> - Unknown father/mother - Illiterate - Low (0-2 ISCED-97) - Medium (3-4 ISCED-97) 	<ul style="list-style-type: none"> - High (5-6 ISCED-97) <p>9. Father's/mother's occupational status*:</p> <ul style="list-style-type: none"> - Unknown or dead father/mother - Employed - Self-employed - Unemployed - Retired - House worker - Other inactive <p>10. Father's/mother's main occupation (based on the International Standard Classification of Occupations, published by the International Labour Office [ISCO-08]):</p> <ul style="list-style-type: none"> - Managers (I-01) - Professionals (I-02) - Technicians (I-03) - Clerical support workers (I-04) - Service and sales workers (I-05 and 10) - Skilled agricultural, forestry and fishery workers (I-06) - Craft and related trades workers (I-07) - Plant and machine operators, and assemblers (I-08) - Elementary occupations (I-09) - Armed forces occupations (I-00) - Father/mother did not work, was unknown or was dead <p>11. Managerial position of father/mother*:</p> <ul style="list-style-type: none"> - Supervisory - Non-supervisory <p>12. Tenancy status of the house in which the respondent was living*:</p> <ul style="list-style-type: none"> - Owned - Not owned
--	--

Data: EU-SILC 2011 cross-sectional (rev.5, June 2015).

Note: This table lists the circumstance variables available in EU-SILC 2011. Questions marked with * refer to the time when the respondent was 14 years old. Item 7 (11) is missing for Slovenia (Finland).

coefficients below 0.230. At the other end of the spectrum we find Latvia and Lithuania with Gini coefficients above 0.340.

4.2 Benchmark Methods

We compare trees and forests to three benchmark estimation methods from the extant literature.

TABLE 2 – Summary Statistics

Country	N	Equivalent Disposable Household Income in €		
		μ	σ	Gini
Austria	6,220	25,538	13,408	0.267
Belgium	6,011	23,314	10,769	0.247
Bulgaria	7,146	3,698	2,457	0.331
Croatia	6,945	6,631	3,764	0.306
Cyprus	4,589	21,074	11,554	0.278
Czech Republic	8,711	9,036	4,610	0.253
Denmark	5,795	32,471	14,422	0.227
Estonia	5,338	6,924	4,364	0.330
France	11,078	24,320	14,695	0.287
Germany	12,683	22,862	12,468	0.284
Greece	6,184	13,184	8,887	0.334
Hungary	13,330	5,305	2,830	0.275
Iceland	3,682	21,562	9,290	0.221
Ireland	4,318	24,882	14,078	0.295
Italy	21,070	18,774	11,348	0.314
Latvia	6,423	5,339	3,751	0.362
Lithuania	5,403	4,774	3,116	0.344
Luxembourg	6,765	37,948	19,412	0.270
Malta	4,701	13,058	6,758	0.272
Netherlands	11,411	24,322	11,452	0.243
Norway	5,026	42,265	16,679	0.206
Poland	15,545	6,087	3,837	0.316
Portugal	5,899	10,796	7,354	0.333
Romania	7,820	2,527	1,612	0.336
Slovakia	6,779	7,309	3,509	0.256
Slovenia	13,183	13,373	5,896	0.234
Spain	15,481	17,088	10,684	0.328
Sweden	6,599	25,098	11,157	0.237
Switzerland	7,583	42,844	23,877	0.278
United Kingdom	7,391	22,768	15,164	0.319

Data: EU-SILC 2011 cross-sectional (rev.5, June 2015).

Note: This table provides summary statistics by country. N indicates the total number of observations. The last three columns summarize the distribution of equivalized disposable household income: mean (μ), standard deviation (σ), and Gini coefficient.

First, we draw on the parametric approach as proposed by Bourguignon et al. (2007) and Ferreira and Gignoux (2011). In line with equation (4), estimates are obtained by a Mincerian regression of log income on the following circumstances: educational attainment of mother and father (5 categories each), father's occupation (11 categories), area of birth (3 categories), and tenancy status of the household at age 14 (2 categories). The prediction model includes 22 parameters.

Second, we draw on the non-parametric approach as proposed by Checchi and Peragine (2010). In line with equation (6), non-parametric estimates are obtained by calculating average outcomes in non-overlapping circumstance types. Types are homogeneous with respect to educational attainment of the highest educated parent (5 categories), fathers' occupation (4 categories), and migration status (2 categories).¹² The prediction model includes 40 parameters.

¹²To minimize the frequency of sparsely populated types we divert from the occupational list given in Table 1 by

Third, we draw on the latent class approach as proposed by Li Donni et al. (2015). We use the union of circumstances used in the parametric and non-parametric approach from which the algorithm infers the appropriate number of unobserved types in the data by minimizing BIC.

Do these specification choices serve for a fair assessment of these benchmark methods? As outlined in section 2, model specification in the (non-)parametric approach is a discretionary choice of the researcher; therefore there are many different specifications that could be used for the benchmarking. To make the comparison non-arbitrary, we anchor our comparison on model specifications of existing studies. The specification of the parametric approach is inspired by Palomino et al. (2019). We divert from their specification by excluding gender (due to our focus on disposable household income) and retrospective information on the financial situation during childhood (due to potential recall bias) from the list of circumstances. In comparison, our prediction model (22 parameters) is more parsimonious than the model in Palomino et al. (2019, 24 parameters). The specification of the non-parametric approach is inspired by Checchi et al. (2016). We divert from their specification by excluding gender (due to our focus on disposable household income) and age (due to its interpretation as a proxy for life-cycle effects) from the list of circumstances. In comparison, our prediction model (40 parameters) is more parsimonious than the model in Checchi et al. (2016, 96 parameters). As outlined in section 2, model specification in latent class analysis is data driven. We therefore do not need to specify the model itself but commit to a model selection criterion. We anchor our comparison on the study of Li Donni et al. (2015) who use BIC to select the number of latent classes to be estimated.

4.3 *Simulation*

We begin our analysis with a simulation exercise. The simulation allows us to assess the properties of different estimation approaches while maintaining control over the true DGP. As a consequence we can i) assess the prediction accuracy by decomposing MSE^{Test} into its variance and bias components, and ii) assess the resulting bias in inequality of opportunity estimates.

re-coding occupations into the following categories: high-skilled non-manual (I-01–I-03), low-skilled non-manual (I-04–I-05 and I-10), skilled manual and elementary occupation (I-06–I-09), and unemployed/unknown/dead.

In general, simulation results are sensitive to assumptions about the true DGP and sample sizes. To make the simulation relevant to the context of our empirical analysis, we choose DGPs that are anchored in the benchmark methods presented in section 4.3 and choose sample sizes to broadly cover the range of country samples in EU-SILC. We note that our additional simulation choices are conservative. First, we construct a simulation sample without missing data points. As a consequence—and in contrast to actual empirical applications—parametric and non-parametric approaches do not suffer from data reductions through list-wise deletion. Second, we restrict circumstances used by trees and forests to the union of circumstances used in the (non-)parametric approach. As a consequence—and in contrast to actual empirical applications—we deprive data-driven approaches from the advantage of using all available circumstance information in the data.

We impose three DGPs that are summarized in Table 3. The *parametric DGP* and *non-parametric DGP* correspond to the estimation models outlined in section 4.2. They present a challenging test for data-driven estimation methods since the latter have to compete against fixed specifications (parametric, non-parametric) that correspond to the ground truth. In addition, we specify a *mixed DGP* that integrates features of both the parametric and the non-parametric DGP. This is a more realistic scenario since it is plausible to assume that researchers devise fixed specifications without prior knowledge of the true DGP. We estimate all three models on the full sample of EU-SILC while list-wise deleting observations with missing information ($N = 197,565$). In turn, we retain the predictions from these estimations and add a disturbance term with $\mathcal{N}(0, 2000)$.¹³ Thus, we obtain three variables that define the distribution of income for the purpose of this simulation.

Next we specify five sample sizes that broadly cover the range of effective country sample sizes observed in EU-SILC (see Supplementary Table S.1): $N \in \{1,000; 2,000; 4,000; 8,000; 16,000\}$. For each sample size, we draw one test set of size $N^H = 1/3N$ and 50 training sets of size $N^{-H} = 2/3N$. Thus, for each observation in the test sets we obtain 50 predictions per combination of DGP and estimation approach. Based on these predictions we calculate two statistics: i) expected MSE^{Test} to assess out-of-sample prediction accuracy (James et al., 2013), and ii) the expected absolute difference between inequality of opportunity estimates and the true level of

¹³We choose a variance term small enough such that $y \in \mathbb{R}_{++}$.

TABLE 3 – Summary of Data Generating Processes

	Parametric	Non-Parametric	Mixed
Outcome	$\ln(y)$	y	$\ln(y)$
Parameters	22	40	18
Circumstances	Education Father Education Mother Occupation Father Birth Area Tenancy Status	Education Parents Occupation Father Mig. Background	Education Parents Occupation Father Mig. Background Tenancy Status
Non-Linearity	None	Full Interaction	All Circ. w/ Mig. Background (2 Levels)
ϵ	$\mathcal{N}(0, 2000)$	$\mathcal{N}(0, 2000)$	$\mathcal{N}(0, 2000)$

inequality of opportunity.

Figure 2 displays the results of our simulation. In its lower part, each panel describes expected MSE^{Test} per combination of DGP, estimation approach and sample size. Since we know the true DGP we can decompose MSE^{Test} into variance and expected bias.¹⁴ In its upper part, each panel describes the corresponding absolute bias in inequality of opportunity estimates on an inverse scale. The absolute bias is calculated as the expected absolute difference between inequality of opportunity estimates and the true level of inequality of opportunity as a percentage share of the latter.

The simulation results are in line with statistical theory. First, if fixed estimation approaches (parametric, non-parametric) invoke the true DGP, expected bias is zero and MSE^{Test} is driven by its variance component only. Second, with increasing N , the bias component of MSE^{Test} remains constant for fixed specifications (parametric, non-parametric) and decreases for data-driven approaches (LCA, trees, forests). Third, with increasing N the variance component of MSE^{Test} decreases for all combinations of DGPs and estimation approaches. Fourth, forests tend to have lower variance than trees—in our simulation this is true in 93% of all cases.

Furthermore, Figure 2 illustrates that the size of MSE^{Test} , and therefore bias in inequality of opportunity estimates, varies with sample size for all estimation methods. However, the sources of this bias vary across estimation methods. For example, forests incur downward bias in

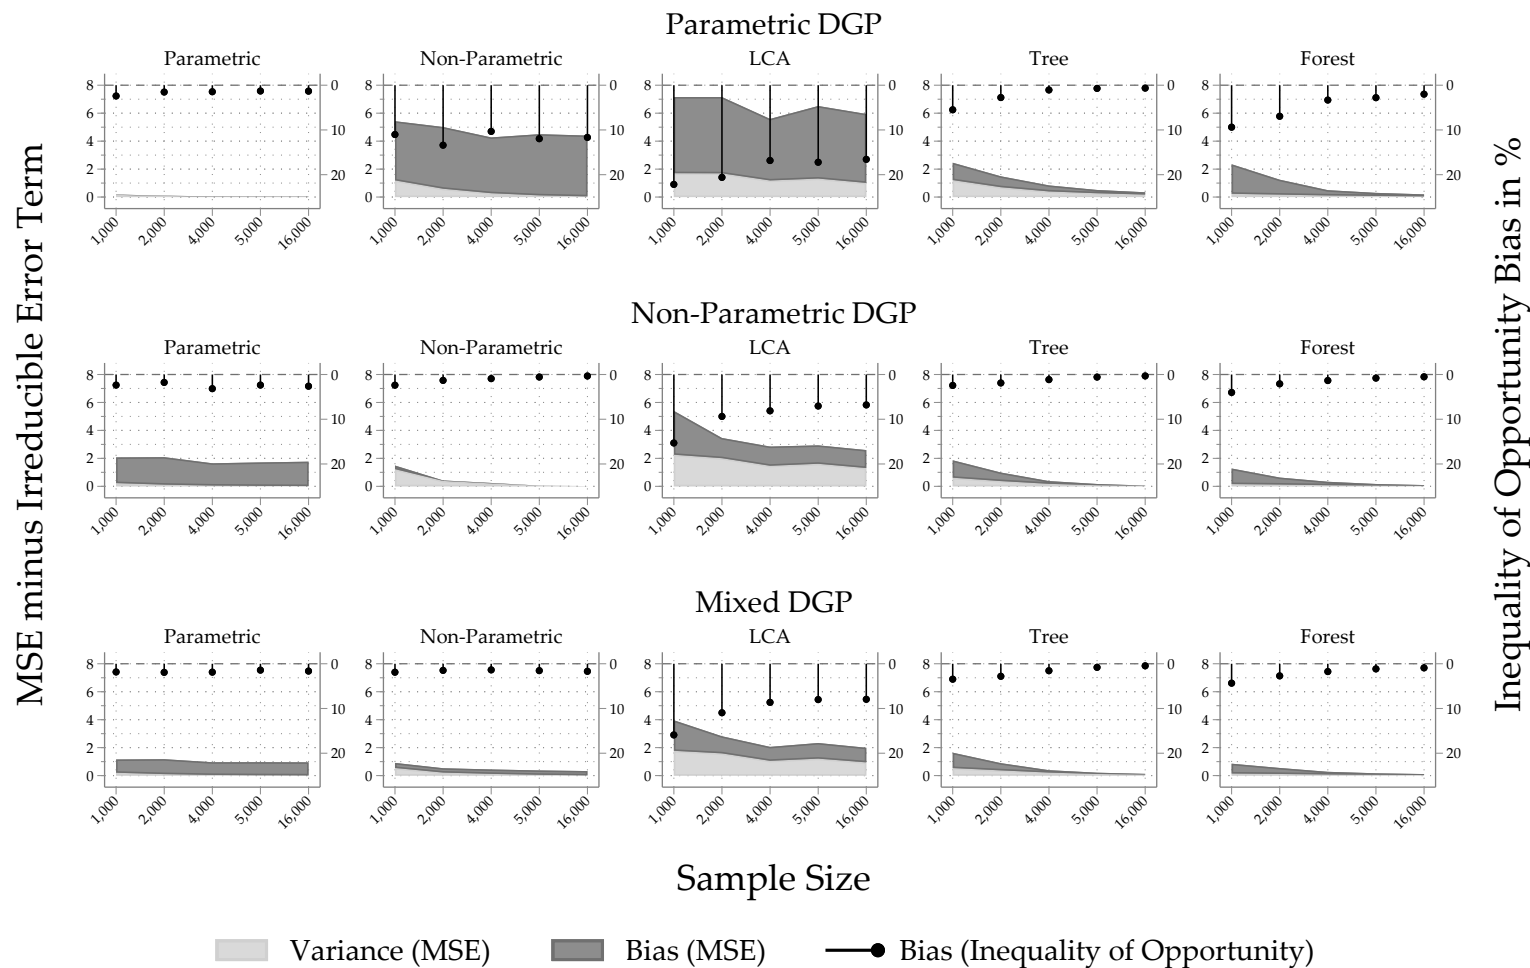
¹⁴See also Supplementary Material S.2 for an illustration of the variance-bias decomposition. The irreducible error term is uninformative for differences in MSE^{Test} because $\text{Var}(\epsilon) = 2,000^2$ is constant across specifications. Therefore, we only present evidence on the variance and the bias component of MSE^{Test} .

small samples as their algorithm prevents the detection of relevant splits. To the contrary, fixed specifications incur upward bias in small samples as the underlying parameters are noisily estimated (although unbiased in expectation). The crucial question is whether this sensitivity is larger for ML than for traditional econometric methods which is a case-specific and empirical question.

Under the reasonable assumption that researchers do not know the true DGP, forests clearly dominate all other estimation approaches in terms of expected MSE^{Test} . This result holds both in comparison to fixed estimation approaches (parametric, non-parametric) and in comparison to LCA as an alternative data-driven estimation approach. The results for trees are also persuasive, however, they have a weaker performance than forests when samples are small. Even in the unlikely case that researchers were to specify (non-)parametric models correctly, trees and forests quickly converge to the test error of the fixed model that invokes the true DGP. The simulation results furthermore highlight the close correspondence between MSE^{Test} and expected bias in inequality of opportunity estimates: the higher MSE^{Test} , the stronger inequality of opportunity estimates diverge from the ground truth.

In summary, the simulation results support the use of regression trees and forests. They flexibly approximate the true DGP. Thereby, they outperform fixed estimation approaches (parametric, non-parametric) and alternative data-driven estimation approaches (LCA) in terms of the expected MSE^{Test} which itself is tightly linked to expected bias in inequality of opportunity estimates.

FIGURE 2 – Simulation Results



Data: EU-SILC 2011 cross-sectional (rev.5, June 2015).

Note: This figure shows expected MSE^{Test} and expected bias in inequality of opportunity estimates from the simulation exercise. Each row corresponds to one data generating process (see Table 3). Each column corresponds to one estimation method (see sections 3.1, 3.2, 4.2). We multiply MSE^{Test} by 1×10^{-6} and deduct the irreducible error term. Inequality of opportunity is measured by the Gini coefficient of the counterfactual distribution \hat{y}^C . We measure expected bias in inequality of opportunity by the average absolute difference between inequality of opportunity estimates and the true level of inequality of opportunity as specified by the data generating process.

4.4 Cross-Country Comparison

We now turn to a cross-country comparison based on actual data. First, we calculate MSE^{Test} to assess the prediction accuracy of different estimation approaches. Second, we calculate inequality of opportunity estimates. In contrast to the simulation exercise, we do not know the true DGP and we cannot assess bias in inequality of opportunity estimates by comparison to the ground truth. Therefore, we assess bias in inequality of opportunity estimates by comparing estimation approaches against the method with the highest prediction accuracy, i.e. the method yielding the lowest MSE^{Test} .


Prediction Accuracy. Figure 3 compares MSE^{Test} across countries and estimation approaches. For each method, MSE^{Test} is presented in differences relative to random forests. By differencing across methods, we provide a close analogue to the simulation exercise in section 4.3: we omit the irreducible error term from the comparison, and relative MSE^{Test} is driven by variance and bias components, only. For better visual clarity, we again scale MSE^{Test} by 1×10^{-6} . Relative $\text{MSE}^{\text{Test}} > 0$ indicates poorer prediction accuracy in comparison to random forests.

Random forests outperform all other methods in terms of prediction accuracy. On average, the parametric approach yields test errors that exceed random forests by 8.4 (7.8%) (Figure 3, Panel [a]). Somewhat smaller average shortfalls in prediction accuracy are observed for non-parametric (Figure 3, Panel [b]) and latent class models (Figure 3, Panel [c]). Averages across countries, however, mask considerable heterogeneity. For example, the relative test error of parametric estimates for Eastern European countries like Slovenia or Czech Republic are close to zero. To the contrary, relative test errors of parametric estimates for Sweden, Luxembourg, and Switzerland diverge significantly from the forest benchmark.

Conditional inference trees are closest to the test error rate of forests ($\text{MSE}^{\text{Test}} = 2.8$ [2.2%]). Yet, they also fall short of the performance of forests due to higher variance, imposing non-linearity, and omitting less relevant circumstances (see section 3.2).

We conclude that among all considered methods, conditional inference forests deliver the high-

FIGURE 3 – Comparison of MSE^{Test} by Method



Data: EU-SILC 2011 cross-sectional (rev.5, June 2015).

Note: This figure shows differences of MSE^{Test} from different estimation approaches relative to random forests. For all methods, we multiply MSE^{Test} by 1×10^{-6} . Values $>$ ($<$)0 indicate worse (better) out-of-sample prediction accuracy than random forests. Vertical lines indicates unweighted cross-country averages. Point estimates and associated standard errors are listed in Supplementary Table S.1.

est out-of-sample prediction accuracy.¹⁵ Hence, relative to random forests, other methods underutilize or overutilize the information contained in Ω which in expectation will lead to bias

¹⁵In Supplementary Table S.3 we show that the overwhelming majority of detected differences are statistically significant at conventional levels.

in inequality of opportunity estimates.

Inequality of Opportunity Estimates. Figure 4 displays inequality of opportunity estimates across countries and estimation approaches. In each panel we plot inequality of opportunity estimates for a particular method, as well as the associated differences to estimates from forests. We emphasize that results from forests cannot be interpreted as the truth. However, since forests yield the lowest test error among all considered methods, they provide the best *approximation of the true DGP in a given data environment*. Therefore, they are a useful benchmark to assess bias of other estimation methods.¹⁶


Panel (a) shows estimates from the parametric approach. In our country sample the chosen model specification for the parametric approach tends to overstate inequality of opportunity relative to forests which is the method providing the lowest expected bias in comparison to the true DGP. For 21 out of 31 countries the inequality of opportunity estimates are higher than the results from forests. Most pronounced overstatements are observed in countries that are typically considered high-opportunity societies. For example, forests classify Sweden and the Netherlands as societies offering high equality of opportunity. To the contrary, the parametric estimate would rank them at similar levels like Germany and France.

Panel (b) shows estimates from the non-parametric approach. The overall pattern is more heterogeneous than for the parametric approach. While overstatements prevail in countries that are typically considered as high-opportunity societies, there are 20 out of 31 countries for which the non-parametric estimate falls short of the forest estimate. These countries are clustered in the lower part of the equal-opportunity ranking. For example, forests classify Italy at a worse position than most countries in Europe. To the contrary, the non-parametric estimate would elevate Italy towards the mid-field close to Sweden.

We highlight that any resemblance between forests and (non-)parametric estimation approaches

¹⁶It is important to keep this relative interpretation of "bias" in mind. We compare method-specific estimates to the best estimate of inequality of opportunity in a given data environment. In light of all methods lacking information on unobserved circumstances, methods that are upward biased in this comparison may potentially be closer to the ground truth than our reference estimate. Such conclusions, however, are purely speculative and can neither be verified nor falsified without knowledge of the ground truth (see also Supplementary Material S.2 for a thorough explanation). Therefore, another interpretation of forests is that they provide the reliable maximum lower bound estimate of inequality of opportunity in a given data environment.

FIGURE 4 – Comparison of Inequality of Opportunity Estimates by Method



Data: EU-SILC 2011 cross-sectional (rev.5, June 2015).

Note: This figure shows inequality of opportunity estimates from different estimation methods relative to forests. Inequality of opportunity is measured by the Gini coefficient of the counterfactual distribution \hat{y}^C . Point estimates and associated standard errors are listed in Supplementary Table S.2.

should be interpreted as a luck of the draw rather than a property inherent to the estimation approach. Under alternative plausible model specifications, estimates from the both approaches may diverge much more strongly than under the specifications adopted in this work. This property of fixed model specifications is apparent from the simulation results in section 4.3.

Panel (c) shows estimates from the latent class approach. In our country sample, LCA tends to understate inequality of opportunity relative to forests. For 25 out of 31 countries the LCA estimate falls short of its forest-based analogue. LCA chooses rather coarse type partitions. Therefore, understatements are clustered in the lower tail of the equal-opportunity ranking, i.e. in societies in which many circumstances co-produce the outcome of interest. To the contrary, in high-opportunity societies the parsimonious models chosen by LCA tend to replicate the results from forests reasonably well.

Panel (d) shows that trees and forests tend to produce similar results. The correlation between point estimates is high (0.98). In contrast to all other approaches there is no general tendency to over- or underestimate inequality of opportunity relative to forests.

Finally, detected differences between the benchmark estimation approaches and forests persist when estimating equality of opportunity in a pooled European sample.¹⁷ For example, the parametric approach overestimates inequality of opportunity relative to forests, whereas LCA yields lower estimates than forests.

Robustness to Differences in Sample Size. Effective sample sizes differ by estimation method and country (Supplementary Table S.1). First, samples for the benchmark methods (parametric, non-parametric, LCA) are reduced as they rely on list-wise deletion in case of missing circumstance information. These reductions can be sizable and exceed 50% in 6 countries of our sample (Denmark, Iceland, Netherlands, Norway, Slovenia, Sweden). Second, even when accounting for missing information the largest country sample in EU-SILC (Italy, $N = 21,070$) is almost seven times as large as the smallest country sample (Iceland, $N = 3,682$). Therefore, we perform two robustness analyses.

First, we recompute inequality of opportunity after completing missing data through multiple imputation (Schafer, 1999).¹⁸ As a consequence, we can compare inequality of opportunity esti-

¹⁷Note that we do not include country of residence as a circumstance. We acknowledge that country of residence is congruent with country of birth for most individuals. Therefore, it could be used as a proxy circumstance (Milanovic, 2015). However, our foremost concern is a methodological comparison of estimation approaches in different data environments. Therefore, we prefer to keep the set of circumstances comparable to our within-country estimates.

¹⁸List-wise deletion yields unbiased parameter estimates if data is missing completely at random (MCAR). Mul-

Accepted Article

mates across methods on the same effective sample size per country. Supplementary Figure S.1 shows a decrease in inequality of opportunity estimates relative to forests for all benchmark methods (parametric, non-parametric, LCA). This result is in line with the intuition that upward biases decrease as sample sizes grow relative to the number of model parameters. To the contrary, the patterns for trees and forests remain unaffected since they handle missing values by default through surrogate splits.

Second, we recompute inequality of opportunity while reducing sample sizes across countries to the smallest common denominator. As a consequence, we can compare inequality of opportunity estimates across countries on the same effective sample size. Supplementary Figure S.2 shows that point estimates and country rankings differ for all benchmark methods (parametric, non-parametric, LCA). Furthermore, also trees show some variability as sample sizes decrease. To the contrary, point estimates and country rankings of forests are unaffected by harmonization in sample sizes across countries. On the one hand, these results highlight that the high variance of trees may lead to suboptimal results in some applications and that researchers should give preference to forest estimates where possible. On the other hand, these results bolster confidence that opportunity rankings of forests are not an artifact of cross-country variation in sample sizes.¹⁹

Comparison to Existing Literature. We have shown that benchmark methods from the existing literature yield markedly different estimates of inequality of opportunity relative to the method for which we expect the lowest bias. These differences are manifested in both point estimates and country rankings. Therefore, these methods may be misleading in two related dimensions. First, they may mis-classify European societies regarding their need for opportunity equalizing policy interventions. Second, researchers and policymakers in search of best practices to devise opportunity equalizing policy interventions may turn to the wrong country examples. In the following we will assess the extent to which such concerns are reflected in the

multiple imputation weakens this assumption by assuming that data is missing at random (MAR), i.e. that missing data is random conditional on observed variables.

¹⁹We perform a similar exercise on the pooled sample: we re-estimate our results for the pooled sample on increasingly smaller fractions of the total sample. In Supplementary Figure S.3 we again show that benchmark methods (parametric, non-parametric, LCA) and trees are sensitive to changes in sample size when fractions become small. To the contrary, forests again emerge as the method that is most robust in small samples.

extant literature on inequality of opportunity in Europe.

We proceed in two steps. First, we assess whether existing literature on inequality of opportunity in Europe is consistent, i.e. whether it yields similar opportunity rankings across European societies. If the literature were consistent, researcher discretion in model selection would be irrelevant for conclusions about inequality of opportunity in Europe. Second, we assess whether existing literature on inequality of opportunity in Europe conforms with evidence on the intergenerational income elasticity (IGE). The IGE is a commonly used proxy statistic for equality of opportunity that is based on data links across generations. The IGE provides a suitable benchmark since it can be interpreted as an ex-ante utilitarian measure of inequality of opportunity (see footnote 5) and it is often based on richer (administrative) panel data. If there was conformity, current estimation approaches would yield opportunity rankings that are strongly in line with common priors about mobility in European societies. We answer both questions by calculating correlations in opportunity rankings across i) existing studies on inequality of opportunity,²⁰ ii) existing consensus estimates of the IGE,²¹ and iii) inequality of opportunity estimates from our preferred methods—regression trees and forests.

Panel (a) of Table 4 suggests that existing literature on inequality of opportunity in Europe is not consistent. Rank correlations as low as 0.09 indicate strong heterogeneity in country rankings. On the one hand, all studies that inform this comparison have a very high degree of harmonization in relevant dimensions: estimates were derived from the same underlying data source (EU-SILC), refer to a similar age group (25-60), and summarize counterfactual distributions \hat{y}^C by the same inequality metric (mean log deviation). On the other hand, all studies specify different prediction functions to estimate inequality of opportunity. This suggests that discretionary choices with respect to model specifications may be a major force behind inconclusive evidence in the inequality of opportunity literature. We cannot fully rule out the possibility

²⁰We focus on published studies estimating ex-ante measures of inequality of opportunity on the 2011 wave of EU-SILC. Further studies that do not meet both criteria include Andreoli and Fusco (2019) and Carranza (2020). Furthermore, we do not include Brzezinski (2020) since he derives estimates based on the methods proposed in this paper.

²¹We focus on IGE estimates based on actual data linkages across generations and exclude IGE estimates based on two-sample instrumental variable estimators to mitigate distortions through measurement error. Estimates are extracted from Stuhler (2018) and Carmichael et al. (2020). Jointly both studies contain the following subset of our country sample: Denmark, Finland, France, Germany, Italy, Netherlands, Norway, Sweden, Spain, and the United Kingdom.

that differences in income concept definitions, i.e. individual income (Checchi et al., 2016) vs. household income (Palomino et al., 2019; Suárez Álvarez and López Menéndez, 2021), may also contribute to observed lack of consistency. However, as we will detail in the next paragraph, regardless of their income definition all of these studies are only moderately correlated with IGE rankings that are calculated with respect to individual incomes. This pattern does not support an alternative explanation based on differences in income definitions (see Panel [b] of Table 4).

TABLE 4 – Rank Correlations of Existing Studies

	Existing Studies			This Paper	
	Checchi et al. (2016)	Palomino et al. (2019)	Suarez et al. (2021)	Tree	Forest
<i>Panel (a): Equality of Opportunity (23 countries)</i>					
Tree	.	.	.	1.000	.
Forest	.	.	.	0.984	1.000
Checchi et al. (2016)	1.000	.	.	0.363	0.345
Palomino et al. (2019)	0.281	1.000	.	0.882	0.859
Suarez et al. (2021)	0.090	0.855	1.000	0.756	0.757
<i>Panel (b): Intergenerational Elasticity (10 countries)</i>					
Stuhler (2018) & Carmichael et al. (2020)	0.535	0.657	0.444	0.900	0.887

Data: EU-SILC 2011 cross-sectional (rev.5, June 2015).

Note: This table shows country rank correlations in inequality of opportunity estimates across existing studies. Panel (a) is based on the intersection of countries included in this paper, Palomino et al. (2019), Checchi et al. (2016), and Suárez Álvarez and López Menéndez (2021) (23 countries). All ranks are calculated from the mean log deviation of the counterfactual distribution \hat{y}^C . Panel (b) is based on the intersection of countries included in this paper, Palomino et al. (2019), Checchi et al. (2016), and Suárez Álvarez and López Menéndez (2021), and the union of Stuhler (2018) and Carmichael et al. (2020) (10 countries). Ranks in Stuhler (2018) and Carmichael et al. (2020) are calculated from consensus estimates of the intergenerational earnings elasticity (IGE). All rank correlations are based on Spearman's ρ .

In Panel (b) of Table 4 we test for conformity of opportunity rankings with the IGE literature. Inequality of opportunity rankings of existing studies are only moderately correlated with IGE rankings. In fact, various findings contradict comparative evidence on the IGE (Carmichael et al., 2020; Bratberg et al., 2017). For example, Palomino et al. (2019) and Suárez Álvarez and López Menéndez (2021) find inequality of opportunity in Germany to be at par with the Nordic countries. Checchi et al. (2016) find the Netherlands to be in the lower part of the opportunity

Accepted Article

ranking. To the contrary, rankings based on trees and forests strongly increase conformity with IGE estimates and therefore yield results that are more strongly in line with common priors about mobility in Europe.

We conclude that regression trees and forests foster consistency in the inequality of opportunity literature by reducing researcher discretion and increase conformity with evidence from the neighboring IGE literature. Both findings further bolster confidence in the ability of trees and forests to make reliable distinctions among high and low opportunity societies in Europe.

5 CONCLUSION

In this paper we propose the use of conditional inference trees and forests to estimate inequality of opportunity. Both estimation approaches minimize arbitrary model selection by the researcher while trading off downward and upward biases in inequality of opportunity estimates.

Conditional inference forests outperform all methods considered in this paper in terms of their out-of-sample prediction accuracy. This observation is valid both for simulated data generating processes and representative survey data from 31 European countries. Hence, within a given data environment they provide estimates of inequality of opportunity that have the lowest expected bias. Conditional inference trees closely mirror forests in terms of their out-of-sample prediction accuracy and their inequality of opportunity estimates. Hence, they provide a fair first-order approximation to the least biased inequality of opportunity estimates. Nevertheless, researchers should be conscious that trees may yield suboptimal results in applications with smaller samples.

We note that the improvements of our preferred methods are conditional on a given data environment. As a consequence, they do not address two major challenges of the existing literature: bias due to unobserved circumstance information, and bias due to small sample sizes. These challenges exist independent of the chosen estimation technique and can only be overcome through the availability of improved data sources in the future.

Accepted Article

Next to their advantages, we acknowledge two potential drawbacks of our preferred methods for empirical research. First, (non-)parametric estimation approaches can be estimated by OLS—one of the workhorse estimation methods in economics and other social sciences. To the contrary, machine learning tools may require some upfront investment of applied researchers to familiarize themselves with these methods. However, as evidenced by the large volume of recent review articles, machine learning methods are increasingly integrated into the statistical toolkit of economists (Varian, 2014; Mullainathan and Spiess, 2017; Athey, 2018). Therefore, we expect this drawback to vanish over time. Second, trees and forests are computationally more costly than predictions via OLS regressions. However, in our empirical application trees approach the computation times of the (non-)parametric approach.²² Therefore, time-constrained researchers who are willing to settle for a fair first-order approximation of the least biased method may consider using trees instead of forests.

The development of machine learning algorithms and their integration into the analytical toolkit of economists is a dynamic process. Finding the best machine learning algorithm for inequality of opportunity estimations is a methodological horse race that eventually will lead to some method outperforming the ones employed in this work. Therefore, the main contribution of this work should be understood as paving the way for new methods that are able to handle the intricacies of model selection for inequality of opportunity estimations. A particularly interesting extension may be the application of local linear forests that outperform more traditional forest algorithms in their ability to capture the linear impact of predictor variables (Friedberg et al., 2020).

Finally, we restricted ourselves to ex-ante utilitarian measures of inequality of opportunity. The exploration of these algorithms for other measurement approaches in the inequality of opportunity literature provides another interesting avenue for future research (Lefranc et al., 2009; Kanbur and Snell, 2018; Pistoiesi, 2009; Brunori and Neidhöfer, 2021).

²²The simulation of section 4.3 has the following computation times: 0.5 min (parametric), 0.5 min (non-parametric), 121.4 min (LCA), 2.1 min (trees), 2,816.2 min (forests). These computation times are based on a machine with a AMD Ryzen 7 4700U Processor (8 cores) and 16 GB RAM working memory.

References

- ALESINA, Alberto, Stefanie STANTCHEVA, and Edoardo TESO (2018). "Intergenerational Mobility and Preferences for Redistribution". *American Economic Review* 108 (2), pp. 521–554.
- ANDREOLI, Francesco and Alessio FUSCO (2019). "Robust cross-country analysis of inequality of opportunity". *Economics Letters* 182, pp. 86–89.
- ARNESON, Richard J. (2018). "Four Conceptions of equal opportunity". *Economic Journal* 128 (612), F152–F173.
- ATHEY, Susan (2018). "The Impact of Machine Learning on Economics". *The Economics of Artificial Intelligence: An Agenda*. Ed. by Ajay K. AGRAWAL, Joshua GANS, and Avi GOLDFARB. Chicago: University of Chicago Press. Chap. 21.
- BIAU, Gérard and Erwan SCORNET (2016). "A random forest guided tour". *Test* 25 (2), pp. 197–227.
- BJÖRKLUND, Anders, Markus JÄNTTI, and John E. ROEMER (2012). "Equality of opportunity and the distribution of long-run income in Sweden". *Social Choice and Welfare* 39 (2-3), pp. 675–696.
- BLACKBURN, McKinley L. (2007). "Estimating wage differentials without logarithms". *Labour Economics* 14 (1), pp. 73–98.
- BLAU, Francine D. and Lawrence M. KAHN (2017). "The Gender Wage Gap: Extent, Trends, and Explanations". *Journal of Economic Literature* 55 (3), pp. 789–865.
- BLUNDELL, Jack and Erling RISA (2019). "Income and family background: Are we using the right models?" *mimeo*.
- BOURGUIGNON, François, Francisco H. G FERREIRA, and Marta MENÉNDEZ (2007). "Inequality of Opportunity in Brazil". *Review of Income and Wealth* 53 (4), pp. 585–618.
- BRATBERG, Espen, Jonathan DAVIS, Bhashkar MAZUMDER, Martin NYBOM, Daniel D. SCHNITZLEIN, and Kjell VAAGE (2017). "A Comparison of Intergenerational Mobility Curves in Germany, Norway, Sweden, and the US". *Scandinavian Journal of Economics* 119 (1), pp. 72–101.
- BREIMAN, Leo (2001). "Random Forests". *Machine Learning* 45 (1), pp. 5–32.
- BREIMAN, Leo, Jerome FRIEDMAN, Charles J. STONE, and R.A. OLSHEN (1984). *Classification and Regression Trees*. Belmont: Taylor & Francis.

- BRUNORI, Paolo and Guido NEIDHÖFER (2021). "The Evolution of Inequality of Opportunity in Germany: A Machine Learning Approach". *Review of Income and Wealth*, Forthcoming.
- BRUNORI, Paolo, Vito PERAGINE, and Laura SERLENGA (2019). "Upward and downward bias when measuring inequality of opportunity". *Social Choice and Welfare* 52, pp. 635–661.
- BRZEZINSKI, Michal (2020). "The evolution of inequality of opportunity in Europe". *Applied Economics Letters* 27 (4), pp. 262–266.
- CAPPELEN, Alexander W., Astri Drange HOLE, Erik Ø. SØRENSEN, and Bertil TUNGODDEN (2007). "The Pluralism of Fairness Ideals: An Experimental Approach". *American Economic Review* 97 (3), pp. 818–827.
- CARMICHAEL, Fiona, Christian K. DARKO, Marco G. ERCOLANI, Ceren OZGEN, and W. Stanley SIEBERT (2020). "Evidence on intergenerational income transmission using complete Dutch population data". *Economics Letters* 189, p. 108996.
- CARRANZA, Rafael (2020). "Upper and lower bound estimates of inequality of opportunity: A cross-national comparison for Europe". *ECINEQ Working Paper Series* 511.
- CHECCHI, Daniele and Vito PERAGINE (2010). "Inequality of opportunity in Italy". *Journal of Economic Inequality* 8 (4), pp. 429–450.
- CHECCHI, Daniele, Vito PERAGINE, and Laura SERLENGA (2016). "Inequality of Opportunity in Europe: Is There a Role for Institutions?" *Inequality: Causes and Consequences*. Ed. by Lorenzo CAPPELLARI, Solomon W. POLACHEK, and Konstantinos TATSIRAMOS. Vol. 43. Emerald Insight. Chap. 1, pp. 1–44.
- CHETTY, Raj, Nathaniel HENDREN, Patrick KLINE, and Emmanuel SAEZ (2014a). "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States". *Quarterly Journal of Economics* 129 (4), pp. 1553–1623.
- CHETTY, Raj, Nathaniel HENDREN, Patrick KLINE, Emmanuel SAEZ, and Nicholas TURNER (2014b). "Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility". *American Economic Review* 104 (5), pp. 141–47.
- CHETTY, Raj, Nathaniel HENDREN, Frina LIN, Jeremy MAJEROVITZ, and Benjamin SCUDERI (2016). "Childhood Environment and Gender Gaps in Adulthood". *American Economic Review* 106 (5), pp. 282–88.
- CORAK, Miles (2013). "Income Inequality, Equality of Opportunity, and Intergenerational Mobility". *Journal of Economic Perspectives* 27 (3), pp. 79–102.

- COWELL, Frank A. (2016). "Inequality and Poverty Measures". *Oxford Handbook of Well-Being and Public Policy*. Ed. by Matthew D. ADLER and Mark FLEURBAEY. Oxford: Oxford University Press. Chap. 4, pp. 82–125.
- COWELL, Frank A. and Maria-Pia VICTORIA-FESER (1996). "Robustness Properties of Inequality Measures". *Econometrica* 64 (1), pp. 77–101.
- DAHL, Gordon B. and Lance LOCHNER (2012). "The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit". *American Economic Review* 102 (5), pp. 1927–56.
- FERREIRA, Francisco H. G and Jérémie GIGNOUX (2011). "The Measurement of Inequality of Opportunity: Theory and an Application to Latin America". *Review of Income and Wealth* 57 (4), pp. 622–657.
- FRIEDBERG, Rina, Julie TIBSHIRANI, Susan ATHEY, and Stefan WAGER (2020). "Local Linear Forests". *Journal of Computational and Graphical Statistics* 30 (2), pp. 503–517.
- FRIEDMAN, Jerome, Trevor HASTIE, and Robert TIBSHIRANI (2009). *The elements of statistical learning*. New York: Springer.
- HOTHORN, Torsten, Kurt HORNIK, and Achim ZEILEIS (2006). "Unbiased Recursive Partitioning: A Conditional Inference Framework". *Journal of Computational and Graphical Statistics* 15 (3), pp. 651–674.
- HUFE, Paul, Andreas PEICHL, John E. ROEMER, and Martin UNGERER (2017). "Inequality of Income Acquisition: The Role of Childhood Circumstances". *Social Choice and Welfare* 143 (3-4), pp. 499–544.
- JAMES, Gareth, Daniela WITTEN, Trevor HASTIE, and Robert TIBSHIRANI (2013). *An Introduction to Statistical Learning with Applications in R*. New York: Springer.
- KANBUR, Ravi and Andy SNELL (2018). "Inequality Measures as Tests of Fairness". *Economic Journal* Forthcoming.
- KREISMAN, Daniel and Marcos A. RANGEL (2015). "On the Blurring of the Color Line: Wages and Employment for Black Males of Different Skin Tones". *Review of Economics and Statistics* 97 (1), pp. 1–13.
- LANZA, Stephanie T., Xianming TAN, and Bethany C. BRAY (2013). "Latent Class Analysis With Distal Outcomes: A Flexible Model-Based Approach". *Structural Equation Modelling* 20 (1), pp. 1–26.

- LEFRANC, Arnaud, Nicolas PISTOLESI, and Alain TRANNOY (2009). "Equality of opportunity and luck: Definitions and testable conditions, with an application to income in France". *Journal of Public Economics* 93 (11–12), pp. 1189–1207.
- LI DONNI, Paolo, Juan Gabriel RODRÍGUEZ, and Pedro ROSA DIAS (2015). "Empirical definition of social types in the analysis of inequality of opportunity: a latent classes approach". *Social Choice and Welfare* 44 (3), pp. 673–701.
- MILANOVIC, Branko (2015). "Global Inequality of Opportunity: How Much of Our Income is Determined by Where We Live?" *Review of Economics and Statistics* 97 (2), pp. 452–460.
- MORGAN, James N. and John A. SONQUIST (1963). "Problems in the Analysis of Survey Data, and a Proposal". *Journal of the American Statistical Association* 58 (302), pp. 415–434.
- MULLAINATHAN, Sendhil and Jann SPIESS (2017). "Machine Learning: An Applied Econometric Approach". *Journal of Economic Perspectives* 31 (2), pp. 87–106.
- PALOMINO, Juan C., Gustavo A. MARRERO, and Juan G. RODRÍGUEZ (2019). "Channels of inequality of opportunity: The role of education and occupation in Europe". *Social Indicators Research* 143, pp. 1045–1074.
- PISTOLESI, Nicolas (2009). "Inequality of opportunity in the land of opportunities, 1968–2001". *Journal of Economic Inequality* 7 (4), pp. 411–433.
- RAMOS, Xavier and Dirk VAN DE GAER (2016). "Empirical Approaches to Inequality of Opportunity: Principles, Measures, and Evidence". *Journal of Economic Surveys* 30 (5), pp. 855–883.
- ROEMER, John E. (1998). *Equality of Opportunity*. Cambridge: Harvard University Press.
- ROEMER, John E. and Alain TRANNOY (2015). "Equality of Opportunity". *Handbook of Income Distribution*. Ed. by Anthony B. ATKINSON and François BOURGUIGNON. Vol. 2A. Amsterdam: Elsevier. Chap. 4, pp. 217–300.
- SCHAFFER, Joseph L (1999). "Multiple imputation: a primer". *Statistical Methods in Medical Research* 8 (1), pp. 3–15.
- STUHLER, Jan (2018). *A Review of Intergenerational Mobility and its Drivers*. Luxembourg: Publications Office of the European Union.
- SUÁREZ ÁLVAREZ, Ana and Ana Jesús LÓPEZ MENÉNDEZ (2021). "Dynamics of inequality and opportunities within European countries". *Bulletin of Economic Research*, Forthcoming.

VAN DE GAER, Dirk (1993). "Equality of Opportunity and Investment in Human Capital". PhD thesis. University of Leuven.

VARIAN, Hal R. (2014). "Big Data: New Tricks for Econometrics". *Journal of Economic Perspectives* 28 (2), pp. 3–27.