

LWS

Working Paper Series

No. 32

Inheritances and Wealth Inequality: a Machine Learning Approach

Pedro Salas-Rojo, Juan Gabriel Rodríguez

December 2020



CROSS-NATIONAL
DATA CENTER
in Luxembourg

Luxembourg Income Study (LIS), asbl

Inheritances and Wealth Inequality: a Machine Learning Approach¹

Pedro Salas-Rojo

Universidad Complutense de Madrid (Spain), ICAE and EQUALITAS
Tel: +34 91 3942626. E-mail: pedsalas@ucm.es

Juan Gabriel Rodríguez

Universidad Complutense de Madrid (Spain), ICAE, EQUALITAS and CEDESOG
Tel: +34 91 3942516. E-mail: juangabr@ucm.es

Abstract

This paper explores how the inheritances received influence the distribution of wealth (financial, non-financial and total) in four developed –but substantially different– countries: the United States, Canada, Italy and Spain. Following the inequality of opportunity literature, we first group individuals into types based on the inheritances received. Then, we estimate the between-types wealth inequality to approximate the part of overall wealth inequality explained by inheritances. After showing that traditional approaches lead to non-robust and arbitrary results, we apply Machine Learning methods to overcome this limitation. Among the available computing methods, we observe that the *random forests* is the most precise algorithm. By using this technique, we find that inheritances explain more than 65% of wealth inequality (Gini coefficient) in the US and Spain, and more than 40% in Italy and Canada. Finally, for the US and Italy, given the availability of parental education, we also include this circumstance in the analysis and study its interaction with inheritances. It is observed that the effect of inheritances is more prominent at the middle of the wealth distribution, while parental education is more important for the asset-poor.

JEL Codes: C60, D31, D63, G51.

Keywords: Wealth inequality; inheritances; Machine Learning; inequality of opportunity; parental education.

¹ We acknowledge useful comments and suggestions from LIS Datacenter seminar participants, Juan César Palomino, Paolo Brunori, Amaia Palencia, Alberto Palacios, Pablo de la Vega, Jordi Domenech, Josep Espasa and Piotr Paradowski. This paper has received financial support from the Ministerio de Economía y Competitividad of Spain under project ECO2016-76506-C4-1-R and Comunidad de Madrid (Spain) under project H2019/HUM-5793-OPINBI-CM. Salas-Rojo also received financial support from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement 730998, InGRID-2 – Integrating Research Infrastructure for European Expertise on Inclusive Growth from Data to Policy and the Ministerio de Ciencia, Innovación y Universidades through FPU grant 17/735. The usual disclaimer applies.

1. Introduction

Wealth inequality is on the rise. Since the publication of *Capital in the Twenty-First Century* (Piketty, 2014), the literature has followed different approaches to explain the causes of this persisting trend. For instance, Lusardi et al. (2017) find that those with better financial knowledge extract higher profits from their investments. Others, such as Alstadsaeter et al. (2017), claim that wealth evasion through tax havens rises the net return to capital of the very rich, allowing them to accumulate more wealth. Meanwhile, Zucman (2019) blames the fall in progressive taxation, which has hindered the effect of distributional policies. Surprisingly, there is still a large dissent on the role that intergenerational wealth transmission through inheritances plays on shaping the wealth distribution. This paper contributes to the matter by employing Machine Learning (ML) techniques that enlighten how the bequests received affect the individual opportunities to accumulate wealth.

A wide stream of the literature has found that inheritances are to be blamed for a relevant part of wealth inequality. Piketty (2011), Piketty and Zucman (2015) and Alvaredo et al. (2017) provide empirical evidence on the rising shares of wealth accumulated at the top of the distribution, where inheritances are the vehicle driven by the very rich to channel wealth through generations. In the same vein, Fessler and Schurz (2018) claim that intergenerational wealth transfers are the largest contributors to net wealth inequality. Other authors, such as De Nardi and Yang (2016), Adermon et al. (2018), Palomino et al. (2020) and Nolan et al. (2020) find similar results after applying different econometric techniques and agent-based models.

However, other empirical evidence puts into question the reliability of these results. Wolff and Gittleman (2014) for the United States, Crawford and Hood (2016) and Karagiannaki (2017) for the United Kingdom, and Elinder et al. (2018) for Sweden coincide on the equalizing net effect that inheritances have upon the wealth distribution. They argue that, as inheritances are more equally distributed than wealth, its intergenerational transmission actually provokes a net decrease in overall inequality. In the same vein, Boserup et al. (2016) find that, despite inheritances increase absolute inequality, this effect is not reflected in relative inequality measures calculated at the top of the distribution.

Given this lack of consensus, we propose a different approach based on the Inequality of Opportunity literature. Following Roemer (1993) and Van de Gaer (1993) any economic outcome such as wealth, income or health status is the result of the interaction between two sets of factors. On the one hand, exogenous factors beyond the individuals' control, such as sex, parental education, race or the inheritances received. They are called *circumstances* because they are out of individuals' responsibility. By considering different combinations of these factors we can divide the population into a set of mutually exclusive and exhaustive groups called *types*. On the other hand, the remaining factors are considered to be endogenous, as they are within the

individuals' set of choices. It is the case, for instance, of the number of hours worked or the nutritional habits. They are called *efforts*.

Both sets of factors decompose overall inequality in two terms: inequality of opportunity (IOp, the component attributed to circumstances) and inequality of effort (IE). The first component is found to be undesirable for social justice matters, as this “unfair” inequality cannot be palliated by individual choices (Rawls, 1971; Sen, 1980). In addition, the literature has found IOp to have a negative effect on economic growth, because hindering individual opportunities for education and work causes a misallocation of talent (Marrero and Rodríguez, 2013 and 2019; Bradbury and Triest, 2016; Carranza, 2020).

In this paper we consider inheritances and gifts to be circumstances whose transmission is independent from the behavior of the recipient.² As such, this variable can be used to divide the population into types, allowing us to estimate IOp as the between-type component of overall wealth inequality. The idea behind this proposal is that the IOp component measures the influence of circumstances (in our case, bequests) on the final distribution of wealth. Thus, if overall wealth inequality was independent from the inheritances and gifts received, there should be no dispersion between types. Otherwise, bequests would have a role on the shape of the observed distribution of wealth.

Wealth is, by definition, a stock variable accumulated through life, so it is affected by life-cycle and behavioral dynamics. Clearly, analyzing how inheritances affect consumption, the saving paths or capitalization would require behavioral and agent-based models (see De Nardi, 2015). However, in this paper, these effects are already included in the final wealth measures under consideration, because wealth surveyed at a certain point of life reflects all past decisions and exogenous dynamics. Thus, once we control for sex and age, the estimated types allow us to measure the aggregate impact of inheritances on the wealth distribution.

Bearing all this in mind, and before the estimation of the between-type inequality component is conducted, an important issue must be solved. The definition of types, essential to calculate the between-type inequality component, is straightforward and easy for categorical circumstances such as parental education or sex. However, being the value of the inheritances received a non-linear continuous circumstance, generating types from this variable is difficult. Indeed, our first result shows that traditional empirical approaches to measure IOp, such as the ex-ante parametric method (Ferreira and Guignoux, 2011), lead to arbitrary and non-robust IOp estimates when ad hoc discretizations of inheritances are applied.

To solve this limitation and generate more objective estimates, we propose the use of Machine Learning (ML) algorithms. These computing techniques extract statistical information from the

² Descendants' behavior might have an influence on the bequests received, but this case is probably impossible to discern with the data available, so we assume that this case is negligible.

data, such as their distribution or correlation between variables, limiting the biases introduced by the researcher. To date, ML has been used in the IOp framework to select the relevant circumstances among a wide range of candidates (Brunori et al. 2019, Brunori et al. 2020, Brunori and Neidhöfer, 2020).³ Here, we propose three ML techniques to discretize a (non-linear) continuous circumstance, like the inheritances received, generating statistically meaningful types. First, we present *ChiMerge* (Kerber, 1992), a computing technique that is exclusively based on the distribution of the variable under consideration (inheritances and gifts received). Second, we employ *Conditional Inference Trees* (Hothorn et al. 2006) which also consider the distribution of the outcome variable (wealth), and *Random Forests* (Strobl et al. 2007) as the bootstrapped version of the tree algorithms.

When comparing these methods, we find that random forests provide the most robust and objective measures of wealth IOp. This algorithm performs a non-arbitrary discretization based on the relationship between the wealth and the inheritances distributions, also controlling for the non linearities inherent to the wealth data. By means of these methods, we analyze how the inheritances received condition the opportunities to accumulate total, financial and non-financial wealth.

Our proposal is applied to the United States (US), Canada, Italy and Spain. The data comes from the Luxembourg Wealth Survey (LWS), which is one of the most comprehensive databases containing information on wealth. In the wealth inequality context, these four developed economies present some differences on welfare and fiscal systems that justify their study. On the one hand, both south European economies have highly developed welfare state systems based on the public provision of services and social security schemes. However, they differ on the fiscal treatment to inheritances and wealth. Italy has a national tax (*Imposta sulle Successioni e Donazioni*) ranging between 4% and 8% of the total amount inherited, and a wealth tax on different assets held out of the country (*Imposta sul valore degli immobile situati all'estero*, and *Imposta sul valore delle Attività Finanziarie detenute all'estero*), while Spain has a heterogeneous tax on inheritances (*Impuesto de Sucesiones y Donaciones*) whose precise application depends on the regional administrations, and also an specific wealth tax (*Impuesto sobre el Patrimonio*). On the other hand, both North American economies, Canada and the US, present important differences. In the US, the public sector does not guarantee free access neither to sanitation nor tertiary education, while the private initiative is prevalent in pension schemes. Interestingly, their inheritances tax scheme mixes the Italian and the Spanish models, with a national tax modified through conditional deductions with some States also implementing their own taxes and benefits schemes (*Estate Tax*). Meanwhile, Canada has a welfare state system more similar to the

³ A precedent of the definition of types based on statistical information can be found in Li Donni et al. (2015).

European countries, but has not a specific inheritance tax. None of these two countries has an specific wealth tax.

Because the literature has shown that parental education is a circumstance highly related to wealth inequality (Adermon et al., 2018; Palomino et al., 2020), we also consider this variable in addition to inheritances. Despite data limitations restrict our analysis for parental education to Italy and the US, the differences across these two countries make the analysis of particular interest. While the educational system, including tertiary education, is free in Italy, access to university in the US is often expensive, being unfordable for many individuals. Thus, the access to higher levels of education might probably be strongly dependent on the opportunities defined by previous generations.

We find inheritances to explain around 68% of overall wealth inequality (measured by the Gini coefficient) in the US and Spain, with ratios reaching 74.96% (US) or 65.15% (Spain) of financial wealth and 66.57% (US) or 76.43% (Spain) of non-financial wealth. Lower rates are found in Canada and Italy, where around 40% of overall wealth inequality is explained by this circumstance, rising to 56.98% (Canada) and 43.94% (Italy) in financial wealth, and descending to 36.57% (Canada) or 38.28% (Italy) in non-financial wealth. Thus, the intergenerational transmission of wealth is found to severely condition the overall wealth distribution. In particular, we find that financial wealth inequality is more affected by inheritances than non-financial wealth, maybe because they act as security nets that allow individuals to face riskier investments.

Including the parental education in the analysis provides a deeper insight. The share of wealth inequality explained by inheritances and parental education does not vary in the US, but in Italy the ratios rise to 52.44% for total wealth, 61.63% for financial wealth and 51.51% for non-financial wealth. This may suggest that the effect of both circumstances can be partially overlapped, particularly in the US. To disentangle both effects, we perform a Shapley value decomposition (Sastre and Trannoy, 2002; Shorrocks, 2013). Around 20% of any wealth distribution in the US can be actually explained by parental education, being the rest attributed to inheritances. For Italy, more than 30% of financial wealth inequality is explained by parental education, descending the share actually attributed to inheritances to around 25%. In this country, total wealth inequality and non-financial wealth inequality are evenly explained by both circumstances.

Finally, we explore whether the effect of the inheritances received is homogeneous along the wealth distribution. To do so, we modify the relative weight given to the observations at different parts of that distribution by using the Single Parameter Gini index (Donaldson and Weymark, 1980; Yitzhaki, 1983) for several parameters of inequality aversion. We find that, in the four countries analyzed, the inheritances received are more important for the middle and upper tail individuals, fostering their opportunities to accumulate higher levels of wealth. The asset poor

receive small inheritances, so there are other factors determining inequalities across them. In particular, we show that parental education for the US and Italy is highly important for the opportunities of people at the bottom of the wealth distribution.

The article is structured as follows. Section 2 introduces the wealth IOp theoretical framework and the ML algorithms employed in the paper. Section 3 describes the LWS database and the adjustments applied to the data. In Section 4 we present the main results for inheritances and parental education while in Section 5 we explore the effect of these two circumstances along the wealth distribution. Finally, Section 6 highlights our main conclusions.

2. IOp measurement and Machine Learnings techniques

In this section we first present the IOp framework to measure the influence of inheritances on wealth inequality, and then justify and explain the implementation of Machine Learning techniques to overcome some limitations of traditional IOp approaches.

2.1. The inequality of opportunity approach

Consider a population of discrete individuals indexed by $i \in \{1, \dots, N\}$ and a variable w characterizing our economic outcome of interest, wealth, whose distribution is a function of the set of circumstances faced by the individual, C_i , and the amount of effort exerted, e_i , such that $w_i = f(C_i, e_i)$. Circumstances are defined as a finite discrete vector of J elements and are assumed to be exogenous because they cannot be affected by individual's choices. Simultaneously, we consider effort to be a continuous variable that depends on both, personal decisions and circumstances, such that individual wealth can be rewritten as $w_i = f[C_i, e_i(C_i)]$.

Then, the population is divided into M exhaustive and mutually exclusive groups, called types, $\Pi = \{T_1, \dots, T_M\}$, such that all individuals belonging to the same type T_m share the same circumstances: $T_1 \cup T_2 \cup \dots \cup T_M = \{1, \dots, N\}$, $T_r \cap T_s = \emptyset, \forall r, s$, and $C_i = C_j, \forall i, j | i, j \in T_m, \forall m$.

Because wealth is a continuous variable, a simple way to assess the effect of circumstances (in our case, the inheritances received) on overall wealth inequality is to compare the density function of w across types. Then, circumstances would have no role on the final distribution of wealth if:

$$\int w | T_m d_{T_m} = \int w | T_s d_{T_s}, \quad \forall m, s | T_m \in \Pi, T_s \in \Pi, \quad (1)$$

where subscripts m and s indicate two different types.

Individual circumstances are irrelevant to explain wealth inequality if the distribution of wealth across types is the same. In this case the individual's outcome is independent of her circumstances. Otherwise, individual non-responsibility factors are important and contribute to shape the final

distribution of wealth. The problem with this approach is that, in general, comparisons between distributions do not fulfill the stochastic dominance property (Atkinson, 1970), because distributions can be significantly different and yet cross each other (so no comparison can be made). To avoid this problem and break potential ties, a common alternative is to focus on a particular moment, typically the mean. Following the ex-ante approach in Van de Gaer (1993) circumstances have no role on the final distribution of wealth if:

$$\bar{w}_s = \bar{w}_m, \quad \forall m, s | T_m, T_s \in \Pi, \quad (2)$$

where the mean wealth in a type, $\bar{w}$, is the expected wealth value that an individual can get conditioned on her type.⁴ The ex-ante method requires the calculation of a counterfactual smoothed distribution $\hat{w}$ assigning to each observation the mean wealth of her type. Applying an inequality measure $I(\cdot)$ to this counterfactual distribution gives us the part of overall inequality attributed to the set of observed circumstances (absolute IOp). Dividing this measure over total inequality reveals the share of overall inequality explained by the set of circumstances (relative IOp):

$$Iop\ Abs = I(\{\hat{w}\}) \quad (3)$$

$$Iop\ Rel = \frac{I(\{\hat{w}\})}{I(\{w\})} \quad (4)$$

For our analysis we use the Gini coefficient but, for robustness, we also employ the Mean Logarithmic Deviation (MLD). The IOp literature has generally leaned towards the MLD, mainly because it is the only additively and path independent decomposable inequality index (Foster and Shneyerov, 2000). However, we base our main analysis on the Gini coefficient because it is the most extended measure on wealth inequality (Cowell and Van Kerm, 2015), so our results can easily be compared with the related literature. Moreover, the Gini index has recently been applied in the IOp framework by Lefranc et al. (2008), Brunori et al. (2019) and Cabrera et al. (2020) among others, highlighting some of their advantages. First, IOp indices measure inequality based on the distribution of means across types (equation (2)), so extreme values are by construction smoothed. However, the MLD is more sensitive to extreme values than the Gini coefficient, so the transformation of the original distribution into the smoothed distribution has a larger effect on the former than in the latter inequality index. Consequently, estimates based on the MLD are likely to under-estimate wealth IOp. Second, as we want to calculate the size of IOp caused by inheritances, we are not particularly interested in decomposing overall inequality into IOp and IE, so the Gini index being not additively decomposable is not an issue.

⁴ A different method, the ex-post approach, compares the average outcome across types for individuals who have exerted the same *degree of effort*, i.e., who belong to the same quantile of the type distribution. We have not followed this approach because it requires a strong assumption on the definition of effort. Moreover, as explained in Brunori et al (2020), the adaptation of some ML techniques to the ex-ante method is straightforward, while that is not the case for the ex-post approach.

The main problem we face to analyze the effect of inheritances on the wealth distribution –through the IOp framework– lies on the construction of the smoothed distribution $\hat{w}$. The literature provides several ex-ante methods, but for comparability reasons we lean towards the standard method, so we consider the parametric approach proposed by Ferreira and Guignoux (2011) (see also Bourguignon et al., 2007).⁵ This method exploits the reduced form of a loglinear OLS regression in which the natural logarithm of the outcome variable is regressed against the vector of available circumstances C_i :

$$\ln(w_i) = \alpha + \varphi C_i + \varepsilon_i. \quad (5)$$

Assuming that the estimates of φ are satisfactory, we obtain the smoothed vector $\hat{w}$ by fitting the parameters of equation (5):

$$\bar{w}_i = \exp[\hat{\alpha} + \hat{\varphi} C_i] \quad (6)$$

Notice that when the set of circumstances is defined with discrete or categorical variables, the creation of types is straightforward. But when one circumstance is continuous and highly non-linear, such as the value of the inheritances received (see below), we face a problem. Including the circumstance without any previous discretization would produce one type for each possible value of that variable, potentially leading to an over fitted regression and producing biased estimates (Brunori et al. 2019). On the contrary, discretizing that variable under the ad-hoc criteria of the researcher might lead to non-robust results. For instance, taking the quintiles of the distribution of inheritances as cutting points can generate rather different smoothed vectors $\hat{w}$, as non-linearities at the top tail might affect differently the imputed mean values. This can lead to quite different between-type inequality measures, impeding accurate comparisons across countries or precise policy assessments. To palliate this problem, we propose the use of Machine Learning techniques that base the discretization of the continuous circumstance on the statistical properties of the wealth distribution.

2.2. Machine Learning and Inequality of Opportunity

We call Machine Learning to the stream of computing techniques that take information from the data, identify patterns and make statistical decisions with the smallest possible human intervention. The main idea of these algorithms lies on the motto “let the data talk”, avoiding potential biases introduced by researchers such as the selection of variables, the level of significance to obtain statistical inference or the way in which we discretize continuous variables. Indeed, the discretization problem previously described is not exclusive of the IOp literature, as it is highlighted by the use of ML techniques in many divergent fields, such as biomedicine

⁵ Nevertheless, we have replicated the whole analysis for the non-parametric approach proposed by Checchi and Peragine (2010), and all conclusions are similar.

(Lutsgarten et al., 2011), genetics (Gallo et al., 2016), the stock market dynamics (Lalithendra and Prasad, 2018) and the price of gold (Banerjee et al., 2019).

To the best of our knowledge, this paper is the first attempt to employ Machine Learning methods to correct the discretization problems in the literature on wealth inequality. For this reason, we believe it can be useful to present and compare unsupervised and supervised techniques of ML discretizing algorithms (Varian, 2014; Mullainathan and Spiess, 2017; Athey and Imbens, 2019). On the one hand, unsupervised algorithms consider information streaming only from the variable to be discretized, in our case, the value of the inheritances received. For this task, we employ Chi-based methods, defined as merging processes that start with a pre-defined partition (in our case, each value of the continuous variable), removing the cutting points by joining the adjacent intervals until some statistical criteria is fulfilled. On the other hand, supervised algorithms take into account the relation of the variable to be discretized with the distribution of the dependent (in our case, wealth) and other independent variables (for instance, parental education). In this context, we use trees and forests as splitting algorithms that subsequently divide the continuous variables into discrete categories. Being trees and forests supervised merging methods, they present a difference that justifies their separate implementation. While trees are run just once over the data, forests can be thought as their bootstrapped version which corrects for potential inconsistencies attributed to the data structure. As the underlying mechanism of forests cannot be explained without understanding that of trees, we present both.

Chi-based Algorithms

The *ChiMerge* algorithm was first proposed in Kerber (1992), and its main idea consists on testing whether each interval of the objective continuous variable, also called *class*, is independent from its adjacent intervals. If both intervals are statistically dependent, they are merged; if they are shown to be statistically independent, the algorithm considers that they belong to different categories, maintaining their separation and discretizing the continuous variable.

The merging process is effectuated by repeating two consecutive steps. First, the algorithm computes the Chi-squared independence test, getting a χ^2 value for each pair of adjacent intervals (see the Technical Appendix). Second, if the null hypothesis of dependence is not rejected, this algorithm merges the pair of adjacent intervals with the lowest χ^2 value. The algorithm ends when the null hypothesis of dependence is rejected for a given level of significance alpha, discretizing the continuous variable by setting the correspondent class as cut point. After this process, we define the statistically-based types and apply the ex-ante parametric method previously described, now with a partition based on the distribution of the independent variable.

For robustness, we have also checked that our results do not vary after employing the Chi2 algorithm, usually considered to be a robust version of the ChiMerge method (Liu and Setiono,

1995). This technique tests the dependence of the intervals with a remarkably low level of significance ($\alpha = 0.5$). If the null hypothesis of dependence cannot be rejected at that α , the algorithm repeats the first step, descending the level of significance by 1%, to $\alpha=0.49$. This process is repeated until some selected α (0.01) rejects the null hypothesis.

Trees

Tree algorithms divide a dataset on exhaustive and mutually exclusive groups of observations, based on sequential and hierarchical decisions given some statistical criteria. Once all partitions are performed, the algorithm imputes to each observation the average value of the objective dependent variable considered, conditioned to the group it belongs. Its application to the IOp framework is straightforward, as these methods directly generate types based on the distribution of the dependent variable across the observed circumstances and, then, assign to each individual the mean value of the dependent variable conditioned on her type.

From the family of tree-based methods we select those based on conditional inference algorithms, which have the advantage of not being biased towards splitting only continuous variables, as other tree-based techniques (Hothorn et al. 2006). Furthermore, they have been used in the IOp framework to select the relevant set of circumstances (Brunori et al. 2019; Brunori et al. 2020; Brunori and Neidhöfer, 2020).

The functioning of the algorithm can be explained in three consecutive steps:

1. First, it performs a t -test on the global null hypothesis of independence for each circumstance considered, at some α -value of significance.

$$H_0^C = D(w|C) = D(w) \quad (7)$$

For each circumstance, the algorithm provides a p -statistic, which needs to be adjusted to avoid type I errors. In this paper we apply the Bonferroni correction, quite common for multiple hypothesis testing:⁶

$$p_{adj} = 1 - (1 - p)^p \quad (8)$$

The algorithm selects the circumstance with the lowest p -adjusted value, i.e., the one with the strongest association to the dependent variable w . If $p_{adj} > \alpha$, the algorithm stops. Otherwise, it continues by setting the selected circumstance as a splitting variable.

2. Once we know that a circumstance is a splitting variable, conditional inference trees decide the cutting points. Note that when the variable is binary this step is trivial, as there is

⁶ Apart from the Bonferroni correction, for robustness, we also followed Genz (1992) and checked our results with a Montecarlo adjustment with 10000 iterations.

only one way to be split. But, when the variable is continuous, the algorithm needs to test all potential partitions. Thus, consider:

$$w_z = \{w_i: C_i < x\} \quad (9)$$

$$w_{-z} = \{w_i: C_i \geq x\}, \quad (10)$$

where x defines each possible value in the continuous variable, and z the resulting subsamples. For every x , tree algorithms test the discrepancy between both subsamples, applying a difference-in-means t -test and obtaining an associated p -value. Finally, the algorithm selects the splitting point delivered by the smallest p -value.

3. The whole process is repeated for each resulting subsample until the null hypothesis of independence cannot be rejected.⁷

The main shortcoming of tree-based algorithms is that their results are highly dependent on several factors. One of them is the selected alpha that rejects or accepts the null hypothesis defined in equation (7). To reduce this problem and following the spirit of ML techniques, we use an endogenously tuned alpha, obtained by the application of K-fold Cross-Validation (see the Technical Appendix). This endogenous alpha eliminates the external judgement of the researcher on setting a particular level of critical significance. Nevertheless, and following the canonical stream in econometrics, we test the robustness of our results by also presenting those for 0.05 and 0.01.

Another relevant factor for these algorithms is the data structure, i.e., the number of variables included, their correlation or their distribution (Friedman et al., 2009). Imagine the case of two highly correlated circumstances, where one delivers a slightly lower p -value than the other. In that case, the other circumstance might disappear from the splitting process, despite being almost as important as the selected one. A similar situation would be found when deciding the splitting point in step 2, in which the distribution leads to two similar cutting points. As a result, predictions inferred directly from trees might be fairly sensitive to alterations in the data structure. Tree methods usually perform well in-sample, but using them for out-of-sample inferences may bear reasonable doubts, so we need to include a more complete technique in the analysis to reinforce the robustness of our results.

⁷ Indeed, every time we test the null hypothesis of independence in the first step we are actually testing equality of opportunity, following the same ideas expressed in equation (1) and (2). Rejecting independence implies that the distribution of an outcome variable is significantly conditioned by a certain circumstance, also rejecting the existence of equality of opportunity (Brunori et al., 2020).

Forests

Conditional inference forests implement bootstrapping ideas into the ML framework.⁸ The algorithm generates a certain number of conditional inference trees, averaging their results. The repeated extraction of subsamples guarantees the independence of each tree, so each one provides different estimates. In the end, the law of large numbers smooths the discrepancies between the constructed trees, providing a distribution consistent with the out-of-sample reality (See Schlosser et al., 2019, for a discussion).

Each tree inside the random forest grows following the same three-step structure previously explained. Nevertheless, this algorithm bears some particularities. As related in Brunori et al. (2020), three factors determine how these forests grow. First, to control for the exclusion of highly correlated independent variables, each tree is generally grown after a random selection of circumstances. However, in our case this is not a problem, since we only have the value of the inheritances received and, when available, parental education. Second, we need to consider the number of trees grown in each forest. For robustness, we have checked the results for 100, 200 and 500 trees. Finally, same as for conditional inference trees, we apply the method not only for the endogenously tuned alpha, but also for values of alpha equal to 0.05 and 0.01.

4. Database and adjustments

The data comes from the Luxembourg Wealth Study (LWS) Database, provided by the Luxembourg Income Survey (LIS) cross-national data center.⁹ From the available set of countries, we present results for Canada (2016), Italy (2014), Spain (2014), and the US (2016), as they are those with the most extensive data on inheritances received.¹⁰ Another advantage of this selection of countries, as said in the Introduction, is that they present quite different welfare and fiscal systems.

Our analysis is based on three different dependent variables. First, *non-financial wealth*, defined as the combined market value of all real estate, non-housing and non-current assets owned by household members. Second, *financial wealth*, expressing the market value of all financial

⁸ Despite random forests and bootstrapping ideas are similar, they differ on a relevant aspect: the sampling process in random forests is performed without replacement because it can lead to biased results, as described in Strobl et al. (2007) and Strobl et al. (2009).

⁹ LIS is a non-profit organization whose main mission consists on acquiring datasets and harmonizing them, easing cross-country comparisons. All the relevant information of the institution and the data can be found at <https://www.lisdatacenter.org/>.

¹⁰ We also tried to include other countries, such as Austria, Norway or the United Kingdom. However, the limited number of valid observations with inheritances (less than 10% of the total sample) led us to inaccurate results.

investments, deposits, cash and the rest of liquid assets reported by the household. Finally, *total wealth*, defined as the sum of the two previous wealth concepts. Distinguishing between the first two wealth variables provides an insightful view on the unequal opportunities that people face to invest in certain assets, while analyzing total wealth leads to a more comprehensive global inequality analysis. All monetary variables are expressed in dollars of 2011, as we use the PPP conversion rates provided by the LIS.

Throughout the whole paper, the unit of analysis is the household, as this is the level at which wealth variables are usually reported. Nevertheless, for circumstances and other particular variables in our study, such as age and gender, we take the values reported by the household head. To make all countries comparable in terms of age, we restrict our final sample to individuals aged between 35 and 80 years old.

Analyzing stock variables such as wealth requires some previous adjustments. First, given the well-known right skewness of wealth distributions, we transform the wealth variable by taking natural logarithms (recall equation (5)).¹¹ Second, the necessity of using scales of equivalence to work with households of different size is unclear. In fact, there is still a heat debate about this issue in the wealth inequality literature (Cowell and Van Kerm, 2015). We consider a commonly used scale of equivalence, the square root of the number of household members (Buhmann et al. 1988; Bover, 2010), but all results and robustness checks have also been calculated without equivalence of scale adjustments, with no meaningful variations found.

Third, disposable wealth in the household might be negative or smaller than that reported by absolute gross measures, as debts can also be held. However, we are interested on the effect of inheritances on wealth accumulation, so incorporating debts to wealth measures might harden the interpretation of our results. Accumulating big debts generally requires big collaterals, which are often provided by individuals' parents.¹² If net worth was employed, individuals with big debts might be located on the wealth distribution at the same position as those with smaller debt and wealth levels, thus entangling the actual effect of inheritances and opportunities. Moreover, while the Gini index can deal with negative and zero values, the MLD cannot. Employing different subsamples would make inaccurate our robustness analysis, and excluding the observations with

¹¹ The logarithmic transformation has an obvious problem when the respondent reports to have zero gross wealth. To solve this problem, we add one monetary unit to all observations before taking logarithms. In addition, we have also checked results with the hyperbolic sine transformation which is typically used to reduce heteroscedasticity and the non-normality of the error term. The results did not change significantly.

¹² Parental assets can be used as collateral for mortgages, easing individuals' chances to acquire their own dwelling, and also to pay high university fees, easing their descendants' human capital accumulation. Unfortunately, we have no data on the current wealth status of the parents, nor the expectations of future bequests, to deal with this issue in our analysis.

negative values would artificially provoke a sample selection bias.¹³ Bearing all this in mind, we think that gross wealth measures can be more helpful to analyze the effect of inheritances on the wealth distribution. Nevertheless, we have also calculated the results for net wealth (using the Gini coefficient) by excluding and maintaining the negative observations, confirming that estimates only differ by 2-3% from those presented in this paper.¹⁴

Finally, being a woman and having a given age are also relevant circumstances. Both factors are, by definition, beyond the individual's control and strongly associated to the wealth distribution. For instance, the gender wealth gap in Europe varies from a lower bound set around 27% in Slovakia to the upper bound set around 48% in Greece, finding countries such as France (44%), Austria (45%), and Germany (47%) in between (Sierminska et al. 2010; Schneebaum et al., 2018). In addition, wealth is by nature strongly related to life cycle dynamics. Therefore, to compare households whose heads differ in their gender and age we must adjust our dependent variable.

Following Palomino et al. (2020), the adjustment is three-fold. First, we control by the gender of the household head. Second, we center the logarithm of wealth at the age of 65 which is, on average, the moment in which people retire and start de-investing. Small changes on that centering point, ranging 63-67 years old, have also been applied without finding remarkable differences on the resulting wealth distributions. Third, we consider the possible interaction between both factors, age and gender. To this end, the following regression is proposed:¹⁵

$$\ln(W_i) = \alpha + \beta F_i + \sum_{n=1}^4 \gamma_n (A_i - 65)^n + \sum_{n=1}^4 \delta_n F_i (A_i - 65)^n + \varepsilon_i, \quad (11)$$

where the dummy variable F_i is 1 when the household heads is a woman and A_i expresses the age of the household head. The forth-degree specification represents the life-cycle non-linear dynamics on wealth, as suggested in Solon (1992) (see also Palomino et al., 2018).

Table 1 deploys the summary statistics for age and gender. Overall, the mean age ranges between 50 and 60, with standard deviations surrounding the 16-17 years. Moreover, we observe that our samples are evenly distributed by gender.

¹³ Excluding negative values in net wealth reduces our sample by around 5% in Canada, Italy and Spain and 8% in the US.

¹⁴ The results for net wealth are available from the authors upon request.

¹⁵ The results for these regressions across countries and wealth measures are available from the authors upon request.

Table 1. Summary statistics for the discrete variables.

CANADA (N=3,627)		
Variable	Mean	Sd
Age	50.93	19.67
Gender	0.51	0.50
ITALY (N=4,142)		
Variable	Mean	Sd
Age	62.76	17.36
Gender	0.56	0.39
SPAIN (N=4,792)		
Variable	Mean	Sd
Age	55.99	16.44
Gender	0.52	0.33
US (N=3,325)		
Variable	Mean	Sd
Age	54.69	18.30
Gender	0.52	0.39

Note: Sd represents the standard deviation and the dummy variable gender is 1 for women.

Adjusted wealth ($W_{adj,i}$) is obtained after extracting the estimated coefficients from equation (11) as follows:

$$\ln(W_{adj,i}) = \ln(W_i) - \hat{\beta}F_i - \sum_{n=1}^4 \hat{\gamma}_n(A_i - 65)^n - \sum_{n=1}^4 \hat{\delta}_nF_i(A_i - 65)^n. \quad (12)$$

Table 2 presents the summary statistics for the three wealth dependent variables after the adjustment, our main circumstance of analysis (inheritances) and parental education. US households are, on average, the wealthiest, but the high value of the standard deviation highlights large inequality levels of wealth. Indeed, we find the US to be the most unequal country for the three wealth variables under consideration, reaching 80.3, 91.6 and 82.2 Gini points for total, financial and non-financial wealth, respectively. Italy and Spain have Gini coefficients for total and non-financial wealth at around 60 points, while Canada seems to be lying somewhere in between the European and the US economies.

Table 2. Summary statistics of the continuous variables (after adjustments).

CANADA (N=3,627)				
Variable	Mean	Sd	Gini	MLD
Total	379,048	605,423	70.66	1.51
Financial	72,068	268,456	83.70	2.61
Non-financial	306,962	554,043	74.90	1.91
Inheritances	46,938	216,652	92.26	1.38
ITALY (N=4,142)				
Variable	Mean	Sd	Gini	MLD
Total	272,602	510,148	59.00	0.89
Financial	31,354	177,073	73.96	2.21
Non-financial	241,248	491,171	60.61	1.01
Inheritances	19,346	139,091	93.89	1.03
SPAIN (N=4,792)				
Variable	Mean	Sd	Gini	MLD
Total	303,548	494,644	59.24	0.98
Financial	46,718	216,145	84.13	2.47
Non-financial	256,830	506,785	60.20	1.33
Inheritances	34,790	186,521	88.55	1.25
US (N=3,325)				
Variable	Mean	Sd	Gini	MLD
Total	1,697,203	1,104,777	80.28	1.85
Financial	426,507	653,076	91.60	3.24
Non-financial	1,270,696	1,127,253	82.17	2.44
Inheritances	9,348	302,239	95.24	1.55

Note: Sd represents the standard deviation. All monetary values are expressed in dollars of 2011, after using the LIS PPP adjustment. The categories of parental education are: 1 (Illiterates), 2 (Basic studies), 3 (Basic secondary), 4 (Upper secondary), 5 (University).

Bequests are not orthogonal to other circumstances. For instance, parental education has been shown to be a good proxy of ascendants' social status which reflects other aspects of wealth accumulation, such as human and social capital intergenerational transmission (Nordblom and Ohlsson, 2010, Adermon et al., 2018; Palomino et al., 2020). To perform the most comprehensive possible analysis and disentangle the potential overlapping effects between both circumstances, inheritances and parental education, we also consider the latter into the analysis. Unfortunately, as said in the Introduction, parental education is only available for Italy and the US. For these two countries, we first use only the inheritances received as a circumstance, so the results can be compared with those obtained for Spain and Canada. Then, we also include the highest parental education level achieved by any parent. Table 3 shows that, on average, U.S. citizens have more educated parents than their Italian counterparts. A Shapley value decomposition permits us to

estimate the contribution of each circumstance to overall inequality and size the potential overlapping term.¹⁶

Table 3. Summary statistics of parental education.

	Mean	Sd
Italy (N=4142)	2.14	0.99
U.S. (N=3325)	3.46	1.01

Note: Sd represents the standard deviation and the categories of parental education are: 1 (Illiterates), 2 (Basic studies), 3 (Basic secondary), 4 (Upper secondary), 5 (University).

5. Empirical results

This section employs the IOp framework to explore how inheritances affect wealth inequality in the US, Canada, Italy and Spain. To begin with, we estimate IOp following the Ferreira and Guignoux (2011) methodology after applying several ad-hoc discretizations over the value of inheritances received. These results are shown to be arbitrary because each partition of the continuous circumstance (inheritances) provides quite different estimates of wealth IOp. This empirical fact confirms the convenience of implementing ML techniques to generate non ad hoc types. For this reason, the role of inheritances is estimated later with the ML techniques explained in Section 2.

Table 4 deploys absolute and relative IOp measures with several ad-hoc discretizing points over the value of the inheritances received. First, we split the variable at 0\$, generating a type with individuals who have not inherited anything, and another with those who have inherited any positive bequest. Second, we generate three types: one for those who have not inherited at all and, for those who inherit, we divide the subsample by the median value of the bequests received. Third, we generate four types, dividing those who inherit by terciles. Finally, we consider the case for which it is only relevant to receive a big amount of wealth. We split the variable in two: those above the 75th percentile of inheritances, and the rest. Many other different discretionary cuts have also been checked, only to confirm their sizable differences and arbitrariness.

¹⁶ The Shapley value decomposition is the only decomposition method that solves the tension between marginality and additivity (See Sastre and Trannoy (2002), Rodríguez (2004) and Shorrocks, 2013).

Table 4. IOp measures with ad-hoc discretizations (Gini).

CANADA	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Financial	Non-Financial	Total	Financial	Non-Financial
> 0\$	19.68	26.78	17.56	27.85%	32.00%	23.44%
Median	23.13	36.05	19.92	32.73%	43.07%	26.60%
Terciles	29.80	44.52	26.50	42.17%	53.19%	35.38%
P75	24.24	38.91	22.20	34.31%	46.49%	29.64%
ITALY	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Financial	Non-Financial	Total	Financial	Non-Financial
> 0\$	17.14	16.84	17.81	29.05%	22.77%	29.38%
Median	17.43	15.80	18.92	29.54%	21.36%	31.22%
Terciles	22.68	23.19	24.14	38.44%	31.35%	39.83%
P75	13.92	18.04	14.85	23.59%	24.39%	24.50%
SPAIN	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Financial	Non-Financial	Total	Financial	Non-Financial
> 0\$	26.11	27.59	33.23	44.07%	32.79%	55.20%
Median	25.00	21.73	32.24	42.20%	25.83%	53.55%
Terciles	35.43	33.30	43.58	59.81%	39.58%	72.39%
P75	15.23	13.29	19.39	25.71%	15.80%	32.21%
US	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Financial	Non-Financial	Total	Financial	Non-Financial
> 0\$	26.74	32.10	33.75	33,31%	35,04%	41,07%
Median	28.12	39.18	29.94	35,03%	42,77%	36,44%
Terciles	37.92	49.27	41.89	47,23%	53,79%	50,98%
P75	25.76	39.44	21.03	32,09%	43,06%	25,59%

Note: IOp measures calculated with the ex-ante parametric Ferreira and Guignoux (2011) approach, with different ad-hoc discretizations over the continuous circumstance (inheritances).

Results in Table 4 show that IOp estimates are quite sensible to the researcher's decision on how the different types should be generated. For instance, relative non-financial wealth IOp in Spain range between 32.21% and 72.39% of overall inequality, with remarkable differences prevailing in the other countries and wealth definitions. Therefore, it is observed that applying subjective ad-hoc criteria leads to arbitrary wealth IOp estimates, hindering accurate cross-country comparisons. Indeed, when analyzing absolute financial IOp, taking the median cutting point sets Italy as the country with more equality of opportunities (15.80 Gini points), followed by Spain (21.73), Canada (36.05) and, finally, the US (39.18). Nevertheless, when we set the third quartile

as the discretizing point, Spain shows the lowest absolute IOp, followed by Italy, Canada and the US.

Our results also highlight an intrinsic characteristic of the ex-ante approach. By construction, including more types leads to higher IOp estimates, as it artificially generates more inequality (Ramos and Van de Gaer, 2016). Considering continuous non-linear circumstances deepens into this problem, as every additional cut increases the number of types. This explains why the discretization based on terciles is higher than that based on the median value. The implications are particularly worrying for policy assessment, as these estimates can easily provide downward or upward biased wealth IOp measures, directly provoked by the arbitrary criteria of the researcher. Indeed, the problem lies, precisely, on when should the cutting process stop, as including too many types would provoke upward biased estimates and overfitting problems, as just a few observations would be attributed to each type.

Both limitations call for a more objective method to estimate the impact of inheritances on wealth inequality. Overcoming these methodological problems, Table 5 presents and compares wealth IOp estimated with objective unsupervised (ChiMerge) and supervised (Conditional inference trees and forests) ML techniques.

ChiMerge provides estimates of wealth IOp after performing statistically justified discretizations on the inheritances received.¹⁷ According to them, in Canada, around 35% of total wealth inequality can be attributed to inheritances, while for financial the ratio rises to 41.63%, descending to 28.87% for non-financial wealth. In Italy, the ratios ascend to 35.49% in total wealth, 29.73% in financial wealth and 36.53% non-financial, while in Spain, they are much higher: almost 53.75% for total wealth, 39.31% for financial and 65.93% for non-financial. Finally, in the US, we find up to 38.85% of total wealth inequality can be explained by the inheritances received, while the ratios ascend to 40.62% in financial and 45.56% in non-financial wealth.

The non-arbitrary types based on ChiMerge algorithms can still be improved. As stated in Brunori et al. (2020), it is convenient to also take into account the distribution of the dependent objective variable. Consequently, we prefer results obtained through supervised methods, although unsupervised techniques can be used to perform discretizations when no information from other related variables is advised.

¹⁷ ChiMerge and Chi2 algorithms provide the same results. With an alpha of 1%, both algorithms cut the value of the inheritances received at 18,250\$ in Canada, at 70,356\$ in Italy, in Spain at 17,263\$, and at 37,924\$ in the US. Types are then generated upon those discretizing points. Other alphas varied those cuts points always by less than 1000\$, with no meaningful differences on the results.

Table 5. IOp measures with ML techniques (only inheritances, Gini).

CANADA	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	24.46	34.01	21.62	34,62%	40,63%	28,87%
Tree	29.57	49.35	26.95	41,85%	58,96%	35,98%
Forest	29.59	47.69	27.39	41,88%	56,98%	36,57%
ITALY	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	20.94	21.99	22.14	35,49%	29,73%	36,53%
Tree	21.70	31.44	23.75	36,78%	42,51%	39,18%
Forest	22.01	32.50	23.20	37,31%	43,94%	38,28%
SPAIN	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	31.84	33.07	39.69	53,75%	39,31%	65,93%
Tree	41.31	55.26	45.91	69,73%	65,68%	76,26%
Forest	40.77	54.81	46.01	68,82%	65,15%	76,43%
US	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	31.19	37.21	37.44	38,85%	40,62%	45,56%
Tree	56.50	69.74	56.51	70,38%	76,14%	68,77%
Forest	55.06	68.66	54.70	68,58%	74,96%	66,57%

Note: ChiMerge measures are calculated applying the Ferreira and Gignoux (2011) approach. Trees and Forests have endogenously tuned their respective alphas following the K-fold Cross Validation technique explained in the Technical Appendix. Forests are calculated after 500 replications.

Table 5 shows that trees and forests supervised algorithms provide similar measures of IOp.¹⁸ However, random forests with endogenously tuned alphas are more objective and totally based on statistical tests. Apart from taking into consideration the distribution of the dependent variable of interest, they correct for the particular structure of the data that might bias tree-based estimates. Consequently, they are set as our preferred specifications, and all commentaries on the results will be based on their estimates.

In line with Adermon et al (2018), Piketty (2011, 2014), Piketty and Zucman (2015) and Alvaredo et al (2017) our results point towards inheritances as a relevant vehicle for intergenerational transmission of wealth disparities. Thus, inheritances alone explain more than 68% to total wealth

¹⁸ All extra information concerning the algorithm results, such as the tree-plots, is available upon request.

inequality in both, the US and Spain, descending the ratios to 41.9% and 37.3% in Canada and Italy, respectively. When we look at financial wealth inequality, the contribution ascends in Canada (57%), Italy (43.9%) and the US (75%), while it descends to 65.2% in Spain. Finally, the inheritances received explain around 36.6% of non-financial inequality in Canada, 38.3% in Italy, 76.4% in Spain and 66.6% in the US. Table A2 in the Appendix shows the results for exogenous alphas of 0.01 and 0.05, while Table A3 replicates the complete analysis for the MLD index.¹⁹

Comparing different wealth definitions shows that, in general, financial wealth is more affected by bequests received than non-financial wealth. The former assets are more risky and volatile than non-financial, mainly because they are much more liquid (Jordá et al., 2019). In our context, we suggest that inheritances may act as insurances or safety nets. After receiving a bequest, individuals could be more prone to face risky investments than those who rely on their own savings, improving their opportunities to access the financial markets. Nevertheless, more factors must be at place since the share of non-financial wealth explained by inheritances is higher than that of financial wealth for Spain.

As said, another important circumstance that is not orthogonal to the inheritances received is parental education. Indeed, as shown by Palomino et al. (2020), more educated parents have, in general, higher income and saving levels, thus being able to bequeath more. In addition, the literature has repeatedly found that individual's education, occupation and income are highly correlated to parental education (Cabrera et al., 2020). Provided that the LWS data includes information on parental education for Italy and the US, we incorporate this variable in our study so, after dealing with overlapping, we obtain a more refined measure on the impact of inheritances on wealth inequality for these two countries.

Table 6 presents wealth IOp estimates for Italy and the US using two circumstances: the inheritances and the parental education.²⁰ Tables A4 and A5 in the Appendix present the rest of IOp indexes that confirm the robustness achieved with the proposed ML techniques, with the MLD estimates being again always below those obtained with the Gini index.

¹⁹ As expected, the MLD index provides much lower estimates of wealth IOp. In this respect, Palomino et al. (2020) recently studied the contribution of intergenerational transfers and socioeconomic background to wealth inequality in France, Spain, the UK and the US, ranging their estimates of the gross inheritance contribution measured with the MLD between 32.8% in the UK and 39.3% in France.

²⁰ Wealth IOp estimate with the parental education alone can be found in Table A7, in the row correspondent to $v=2$. The following section deeps into this Table and its interpretation.


Table 6. IOp measures with inheritances and parental education (Gini).

Italy	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	32.71	44.27	32.55	55,44%	59,86%	53,70%
Tree	31.71	45.77	31.14	53,75%	61,88%	51,38%
Forest	30.94	45.58	31.22	52,44%	61,63%	51,51%
US	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	41.99	50.79	48.48	52,30%	55,45%	59,00%
Tree	62.15	74.62	53.20	77,42%	81,46%	64,74%
Forest	55.65	68.83	53.55	69,32%	75,14%	65,17%

Note: ChiMerge measures are calculated applying the Ferreira and Gignoux (2011) approach. Trees and Forests have endogenously tuned their respective alphas following the K-fold Cross Validation technique explained in the Technical Appendix. Forests are calculated after 500 replications.

In Italy, including parental education as circumstance increases absolute and relative IOp measures. Using the Gini index, more than 52% of total inequality can be explained with both factors, reaching up to 61.63% in financial and 51.51% in non-financial wealth inequality. However, the US estimates do not significantly vary when compared to those presented in Table 5, showing that more information available on individual's background does not necessarily lead to higher contributions. Interestingly, deepening into the algorithm results, we have found that in the first step of tree construction, parental education is barely selected as a meaningful circumstance in the US. This result should not lead us to claim that parental education is not relevant to accumulate wealth in the US, nor that its contribution in Italy is the gross difference with respect to the IOp estimates presented in Table 5, as there might be some overlapping. The inheritances received are, potentially, a comprehensive variable that collects many different aspects of individual's background. Considering this possibility, Figure 1 plots the relative contribution of each circumstance to overall inequality, calculated with the Shapley value decomposition over the random forest indexes calculated in Table 6.

Figure 1. Shapley Value Decomposition of wealth IOp.



Note: Shapley value decomposition applied over the relative wealth IOp measures obtained in Table 6.

In Italy, around 26% of total and non-financial inequality is explained by the distribution of bequests, being the rest attributed to parental education. Indeed, this circumstance explains up to 35% of financial wealth inequality, while around 26% is finally attributed to the inheritances received. In the US, around 22% of overall inequality is actually explained by the parental education, being the distribution of inheritances able to explain 47% of total, 52% of financial and 43% of non-financial wealth inequality. Our results show that parents do not exclusively affect the individual opportunities to accumulate wealth through the bequest transmission itself, as other factors such as human capital transmission may also be important.

In the next section we study whether the effect of inheritances is homogeneous across the wealth distribution or, on the contrary, it depends on the part of the wealth distribution that is considered.

6. Inheritances and Parental Education

The distributions of wealth and inheritances are highly non-linear, so it is pertinent to check whether their relationship is homogeneous along the wealth distribution or, rather, more accentuated at the tails or the middle (see Rodríguez, 2008). For this task, we expand our analysis by using the Single-Parameter Gini (S-Gini) proposed in Donaldson and Weymark (1980) and Yitzhaki (1983).

S-Gini indexes are born from the idea that the canonical Gini index weights equally all parts of the wealth distribution. Thus, they introduce a vector of weights, defined by a parameter v of inequality aversion which modifies the weight given to different percentile positions, q . In

particular, as v increases, the delivered weights give more (less) importance to the lower (upper) part of the distribution. Formally, the S-Gini index is:

$$I_{S-Gini}(F; v) = 1 - v[v - 1] \int_0^1 [1 - q]^{v-2} L(F; q) dq \quad (13)$$

where L is the Lorenz curve, and $v > 1$ is an inequality aversion parameter. For the particular case $v = 2$, equation (13) provides the traditional Gini index.

Now, by applying this index to the counterfactual smoothed distribution $\hat{w}$, we obtain an absolute measure of IOp for an inequality aversion parameter v , $Iop(v)$. Interestingly, we might find three possible cases when the value of v is increased. First, if there were no significant changes in the estimates, the circumstances under consideration would have a homogeneous effect along the wealth distribution. Second, if absolute IOp increased with the inequality aversion parameter v , we could safely say that the effect of circumstances is higher at the lower tail of the wealth distribution. Finally, if absolute IOp decreased when the inequality aversion parameter rose, the effect of circumstances would be more intense at the upper tail of the wealth distribution.


Before continuing, we remark that we do not include relative IOp in these results, because its interpretation can be misleading: changes provoked in relative measures could not only be caused by absolute IOp, but also by overall inequality. As a result, the calculated variation in the relative IOp measures would not exclusively be explained by the heterogeneous effect of the circumstances at different parts of the wealth distribution but, rather, by their interaction with overall inequality.

We estimate the effect of circumstances along the wealth distribution for $v \in [2, 5]$. The precise values are included in Tables A6 and A7 in the Appendix, while Figure 2 plots the evolution of $Iop(v)$ for the inheritances received. It is clearly observed that the higher the weight on the lower tail of the distribution, the smaller the value of $Iop(v)$, with this index being close to zero for $v = 5$ in all countries and wealth measures. Hence, our findings show a persistent fact: the inheritances received lose importance when we focus at the lower tail of the wealth distribution, i.e., the significant contribution to wealth inequality shown in Tables 5 and 6 is mainly explained by their effect on the opportunities of the middle class and, particularly, the wealthy people.

The same analysis is replicated in Figure 3 for inheritances and parental education in Italy and the US. In the first country, $Iop(v)$ declines again, although at a lower rate than in Figure 2. Thus, $Iop(v)$ in Italy descends from 22.06 to 11.80 in total wealth, from 39.10 to 34.05 in financial wealth and from 20.97 to 10.33 in non-financial wealth. On the contrary, $Iop(v)$ rise in the US for total wealth (from 34.74 to 37.27) and financial wealth (from 42.42 to 48.89), while it remains stable for non-financial wealth (from 36.85 to 34.55). This result suggests that, for those located

at the lower part of the wealth distribution, parental education gains importance to determine the individual opportunities to accumulate wealth.²¹


Figure 2. S-Gini wealth IOP with inheritances.



Note: Absolute IOP measures with different inequality aversion parameters calculated with random forests with endogenously tuned alphas.

²¹ In line with Lusardi et al (2017), parental education in the US seems to be a relevant factor explaining financial wealth inequality, particularly at the bottom of the wealth distribution. A successful participation in the financial markets usually requires some skills highly affected by human capital. See also Cabrera et al (2020).

Figure 3. S-Gini wealth IOp with inheritances and parental education.



Note: Absolute IOp measures with different inequality aversion parameters calculated with random forests with endogenously tuned alphas.

Employing S-Gini indexes highlight the heterogeneous transmission of opportunities at different parts of the distribution, otherwise hidden under other traditional approaches. On the one hand, as inheritances are typically concentrated at the upper tail of the wealth distribution (Nolan et al, 2020), diminishing the weight on this part of the distribution makes the effect of inheritances on inequality to disappear. This fact together, with the large non-linearities of both, wealth and inheritances, might explain why inheritances are so relevant to explain wealth inequality, as they would intensively foster the individual opportunities of the wealthy. On the other hand, as inheritances at the lower tail of the wealth distribution are rather small, the individual opportunities of those at this part of the distribution are conditioned by other factors such as parental education.

7. Conclusions

In this paper we estimate the impact of the inheritances received on wealth inequality for four developed western economies: Canada, Italy, Spain and the US. Following Palomino et al. (2020) we estimate this impact by the between-types component of total wealth inequality. The idea is simple: after controlling by age and sex, and computing the type distributions based on the inheritances received, there should not exist any dispersion between types. Otherwise, overall wealth inequality is conditioned by the bequest distribution.

Unfortunately, the traditional definition of types in the literature (Ferreira and Guignoux, 2011) for a continuous circumstance like inheritances, requires a discretization process that, when left to the researchers' criteria, leads to arbitrary results. Depending on how the bequests are cut to define the types, the between-types component of overall inequality changes its size, impeding a

correct cross-country comparison or policy advising. To overcome this limitation, we propose the application of Machine Learning techniques which objectivize the discretization process, statistically justifying the cuts on the inheritances variable. Among all available methods, we select the *endogenously tuned random forests* as the approach that provides the most objective estimates of the impact of inheritances on wealth inequality.

Our results show that a remarkable share of overall wealth inequality is explained by the inheritances received. Particularly, this variable alone explains up to 43.94% of financial wealth inequality in Italy, 56.98% in Canada, 65.15% in Spain and 74.96% in the US. In the case of non-financial wealth, inheritances explain 36.57% overall inequality in Canada, 38.28% in Italy, 66.57% in the US and a significant 76.43% in Spain. On aggregate, almost 41.88% (Canada), 37.31% (Italy), 68.82% (Spain) and 68.58 (the US) of total wealth inequality can be assigned to this circumstance.

If we include parental education as an additional circumstance, both circumstances explain up to 61.63% in financial wealth, 51.51% in non-financial wealth and 52.44% of total wealth inequality in Italy. For the US, including parental education does not change our previous estimates. The literature has already shown that these two circumstances are not orthogonal (Adermon, 2018; Palomino et al., 2020), so we check for the existence of overlapping in their effects. For this task, we perform a Shapley value decomposition and find that, in the US, around two-thirds of wealth IOp can be attributed to the inheritances received, while the rest is actually explained by parental education. In Italy, the shares attributed to both circumstances, inheritances and parental education, are similar.

Finally, due to the remarkable non-linearities of the wealth and inheritance distributions, we apply the S-Gini index for several parameters of inequality aversion. The higher the aversion parameter, the larger the weight on the lower tail of the wealth distribution. We find that, the more we focus on the asset-poor, the inheritances received are less important to explain their opportunities to acquire wealth, while the opposite happens with parental education. That is, both circumstances, inheritances and parental education, have a heterogeneous effect along the wealth distribution, although these effects influence different tails of the distribution. Therefore, policies having a bearing on inheritances and on parental education are complementary (not substitutive) for the reduction of wealth inequality.

References

- Adermon, A., Lindahl, M., Waldenström, D.: Intergenerational wealth mobility and the role of inheritance: evidence from multiple generations. *The Eco Jou.* 128, 482 -513. (2018)
- Alstadsaeter, A., Johannesen, N., Zucman, G.: Tax evasion and inequality. NBER Working Paper, no 23772. (2017)
- Alvaredo, F., Chancel, L., Piketty, T., Saez, E. Zucman., G.: Global inequality dynamics: New findings from WID.world. *Ame. Eco. Rev.* 107, 404-409. (2017)
- Atkinson, A.: On the Measurement of Inequality. *Jou. of Eco. The.*, 2, 244-263. (1970).
- Athey, S., Imbens, G.: Machine Learning methods economists should know about. *Ann. Rev. of Eco.* 11, 685-725. (2019)
- Banerjee, D., Ghosal, A., Mukherjee, I.: Prediction of gold price movement using discretization procedure. *Comp. Int. in Dat. Min.* 1, 345-356. (2019)
- Boserup, S. H., Kopczuk, W., Kreiner, C. T.: The role of bequests in shaping wealth inequality: Evidence from Danish wealth records. *Ame. Eco. Rev.* 106, 656-61. (2016)
- Bover, O.: Wealth inequality and household structure: U.S. vs Spain. *Rev. of Inc. and Wea.* 56, 259-290. (2010)
- Bourguignon, F., Ferreira, F. H., Menéndez, M.: Inequality of opportunity in Brazil. *Rev. of Inc. and Wea.* 53, 585-618. (2007)
- Bradbury, K., Triest, R. K.: Inequality of opportunity and aggregate economic performance. *The Rus. Sag. Fou. Jou. of the Soc. Sci.* 2, 178-201. (2016)
- Brunori, P., Hufe, P., Mahler, D. G.: The roots of inequality: Estimating inequality of opportunity from regression trees. World Bank Group Policy Research Working Paper no 8349. (2020)
- Brunori, P. and Neidhöfer, G.: The evolution of inequality of opportunity in Germany: A machine learning approach. ZEW Discussion Papers, no 20-013. (2020)
- Brunori, P., Peragine, V., Serlenga, L.: Upward and downward bias when measuring inequality of opportunity. *Soc. Cho. and Wel.* 52, 635-661. (2019)
- Buhmann, B., Rainwater, L., Schmaus, G., Smeeding, T. M.: Equivalence scales, well-being, inequality and poverty: Sensitivity estimates across ten countries using the Luxembourg Income Study Database. *Rev. of Inc. and Wea.* 34, 115-142. (1988)
- Cabrera, L., Marrero, G. A., Rodríguez, J. G., Salas-Rojo, P.: Inequality of opportunity in Spain: New insights from new data. *Rev. of Pub. Eco.* (*forthcoming*). (2020)
- Carranza, R.: Inequality of outcomes, inequality of opportunity, and economic growth. ECINEQ Working Paper no 534. (2020)
- Cecchi, D., Peragine, V.: Inequality of opportunity in Italy. *The Jou. of Eco. Inc.* 8, 429-450. (2010)
- Crawford, R., Hood, A.: Lifetime receipt of inheritances and the distribution of wealth in England. *Fis. Stu.* 37, 55-75. (2016)
- Cowell, F., Van Kerm, P.: Wealth inequality: A survey. *Jou. of Eco. Sur.* 29, 671-710. (2015)

- De Nardi, M.: Quantitative models of wealth inequality: A survey. NBER Working Paper no 21106. (2015)
- De Nardi, M., Yang, F.: Wealth inequality, family background and estate taxation. *Jou. of Mon. Eco.* 77, 130-145. (2016)
- Donaldson, D., Weymark, J. A.: A single parameter generalization of the Gini indices of inequality. *Jou. of Eco. The.* 22, 67-86. (1980)
- Elinder, M., Erixson, O., Waldenström, D.: Inheritance and wealth inequality: Evidence from population registers. *Jou. of Pub. Eco.* 165, 17-30. (2018)
- Ferreira, F. H. G., Guignoux, J.: The measurement of inequality of opportunity: Theory and application to Latin America. *Rev. of Inc. and Wea.* 57, 622-657. (2011)
- Fessler, P., Schurz, M.: Private wealth across European countries: The role of income, inheritance and the welfare state. *Jou. of Hum. Dev. and Cap.* 19, 1-29. (2018)
- Fushiki, T.: Estimation of prediction error by using K-Fold Cross Validation. *Sta. and Com.* 21, 137-146. (2011)
- Friedman, F., Hastie, T., Tibshirani, R.: *The elements of Statistical Learning*. Springer, Berlin. (2009)
- Foster, J. E., Shneyorov, A.: Path independent inequality measures. *Jou. of Eco. The.* 91, 199-222. (2000)
- Gallo, C. A., Cecchirini, R. L., Carballido, J. A., Micheletto, S., Ponzoni, I.: Discretization of gene expression data revised. *Bri. Bio.* 17, 758-770. (2016)
- Genz, A.: Numerical computation of multivariate normal probabilities. *Jou. of Com. and Gra. Sta.* 1, 141-149. (1992)
- Holmes, D. E., Jain, L. C.: *Data mining: Foundations and intelligent paradigms: Statistical bayesian, time series and other theoretical aspects*. Springer, Berlin. (2012)
- Hothorn, T., Hornik, K., Zeileis, A.: Unbiased recursive partitioning: A conditional inference framework. *Jou. of Com. and Gra. Sta.* 15, 651-674. (2006)
- Jordá, O., Knoll, K., Kuvshinov, D., Schularick, M., Taylor, A. M.: The rate of return on everything, 1870-2015. *The Qua. Jou. of Eco.* 134, 1225-1298. (2019)
- Karagiannaki, E.: The impact of inheritance on the distribution of wealth: Evidence from Great Britain. *Rev. of Inc. and Wea.* 63, 394-408. (2017)
- Kerber, R.: Chimerge: Discretization of numeric attributes. *Pro. of the Ten. Nat. Con. on Art. Int.* 1, 123-128. (1992)
- Lalithendra, N. V., Prasad, G. S.: Support vector machine in the anticipation of currency markets. *Int. Jou. of Eng. and Tec.* (2018)
- Lefranc, A., Pistolesi, N. and Trannoy A.: Inequality of opportunities vs. Inequality of outcomes: Are Western societies all alike?, *Rev. of Inc. and Wea.* 54, 513 -546. (2008)
- Li Donni, P., Rodríguez, J. G., Rosa Días, P.: Empirical definition of social types in the analysis of inequality of opportunity: A latent classes approach. *Soc. Cho. and Wel.* 44, 673-701. (2015)
- Liu, H., Setiono, R.: Chi2: Feature selection and discretization of numeric attributes. *Pro. of Sev. Int. Con. on Too. Art. Int.* 1, 388-391. (1995)

- Lusardi, A., Michaud, P. C., Mitchell, O. S.: Optimal financial knowledge and wealth inequality. *Jou. of Pol. Eco.* 125, 431-477. (2017)
- Lutsgarten, J. L., Visweswaran, S., Gopalakrishnan, V.: Application of an efficient bayesian discretization method to biomedical data. *Bio.* 12, 1-15. (2011)
- Luxembourg Wealth Study (LWS) Database, <http://www.lisdatacenter.org> (Italy, Canada, U.S., Spain; access 03/03/2020 – 30/06/2020). Luxembourg: LIS.
- Marrero, G. A. and Rodríguez, J. G.: Inequality of Opportunity and Growth. *Jou. of Dev. Eco.* 104, 107-122. (2013)
- Marrero, G. A., Rodríguez, J. G.: Inequality and growth: The cholesterol hypothesis. ECINEQ Working Paper no 501. (2019)
- Mullainathan, S., Spiess, J.: Machine Learning: An applied econometric approach. *Jou. of Eco. Pers.* 31, 87-106. (2017)
- Nolan, B., Palomino, J.C., Van Kerm, P. V. and Morelli, S.: The Intergenerational Transmission of Wealth in Rich Countries. Centre for Studies in Economics and Finances Working Papers, no. 578. (2020)
- Nordblom K., Ohlsson, H.: Bequests, gifts and education: Links between intergenerational transfers. *Emp. Eco.* 40, 343-358. (2010)
- Palomino, J. G., Marrero, G. A., Rodríguez, J.G.: One size doesn't fit all: a quantile analysis of intergenerational income mobility in the U.S. (1980-2010). *The Jou. of Eco. Ine.*, 16, 347-367. (2018)
- Palomino, J. G., Marrero, G. A., Nolan, B., Rodríguez, J.G.: Inheritances, family background and wealth inequality. INET Oxford Working Paper No. 2020-15. (2020)
- Piketty, T.: On the long-run evolution of inheritance: France 1820-2050. *The Qua. Jou. of Eco.* 126, 1041-1131. (2011)
- Piketty, T.: *Capital in the Twenty-First Century*. Harvard University Press, Cambridge. (2014)
- Piketty, T., Zucman, G.: Wealth and inheritance in the long run. *Han. of Inc. Dis.* 2, 1303-1368. (2015)
- Ramos, X., Van de Gaer, D.: Approaches to inequality of opportunity: Principles, measures and evidence. *Jou. of Eco. Sur.* 30, 855-883. (2016)
- Rawls, J.: *A theory of justice*. Harvard University Press, Cambridge. (1971)
- Rodríguez, J. G.: Partial equality of opportunity orderings. *Soc. Cho. and Wel.* 31, 435-456. (2008)
- Rodríguez, J. D., Pérez, A., Lozano, J. A.: Sensitivity analysis of K-Fold Cross Validation in prediction error estimation. *Tran. on Pat. Ana. and Mac. Int.* 32, 569-575. (2009)
- Roemer, J. E.: A Pragmatic approach to responsibility for the egalitarian planner. *Phi. and Pub. Aff.* 20, 146-166. (1993)
- Sastre, M. and Trannoy, A.: Shapley Inequality Decomposition by Factor Components: Some Methodological Issues. *Jou. of Eco.* 77, 51-89. (2002)

- Schlosser, L., Hothorn, T., Stauffer, R., Zeileis, A.: Distributional regression forests for probabilistic precipitation forecasting in complex terrain. *The Ann. of App. Sta.* 13, 1564-1589. (2019)
- Schneebaum, A., Rehm, M., Mader, K., Hollan, K.: The gender wealth gap across European countries. *Rev. of Inc. and Wea.* 64, 295-331. (2018)
- Sen, A.: Equality of what?. *The Tan. Lec. on Hum. Val.*, University of Utah Press, Salt Lake City. (1980)
- Shorrocks, A. F.: Decomposition procedures for distributional analysis: A unified framework based on the Shapley value. *The Jou. of Eco. Ine.* 11, 99-126. (2013)
- Sierminska, E. M., Frick, J. R., Grabka, M., M.: Examining the gender wealth gap. *Oxf. Eco. Pap.* 62, 669-690. (2010)
- Solon, G.: Intergenerational income mobility in the United States. *Ame. Eco. Rev.* 82, 393-408. (1992)
- Strobl, C., Boulesteix, A. L., Zeileis, A. and Hothorn, T.: Bias in Random Forest Variable Importance Measures : Illustrations, Sources and a Solution. *Bio.* 8, 25-46. (2007)
- Strobl C., Malley, J. and Tutz, G.: An Introduction to Recursive Partitioning: Rationale, Application and Characteristics of Classification and Regression Trees, Bagging and Random Forests. *Psy. Met.* 14, 323-348. (2009)
- Van de Gaer, D.: Equality of Opportunity and Investment in Human Capital. PhD Dissertation, Catholic University of Leuven. (1993)
- Varian, H. R.: Big Data: New Tricks for Econometrics. *Jou. of Eco. Per.* 28, 3-28. (2014)
- Wolff, E. N. and Gittleman, M.: Inheritances and the Distribution of Wealth or Whatever Happened to the Great Inheritance Boom?. *The Jou. of Eco. Ine.* 12, 439-468. (2014)
- Yitzhaki, S.: On An Extension of the Gini Inequality Index. *Int. Eco. Rev.* 24, 617-628. (1983)
- Zucman, G.: *The Triumph of Injustice : How the Rich Dodge Taxes and How to Make Them Pay.* Norton and Company, New York. (2019)

Appendix

Table A.1: IOp with ad-hoc discretizations (MLD).

CANADA	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Total	Total	Total	Financial	Non-Financial
> 0\$	0.08	0.14	0.06	5.30%	5.36%	3.14%
Median	0.13	0.29	0.10	8.61%	11.11%	5.24%
Terciles	0.18	0.39	0.15	11.92%	14.94%	7.85%
P75	0.17	0.38	0.15	11.26%	14.56%	7.85%
ITALY	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Total	Total	Total	Financial	Non-Financial
> 0\$	0.07	0.06	0.07	7.87%	2.71%	6.93%
Median	0.09	0.07	0.10	10.11%	3.17%	9.90%
Terciles	0.12	0.13	0.14	13.48%	5.88%	13.86%
P75	0.08	0.13	0.10	8.99%	5.88%	9.90%
SPAIN	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Total	Total	Total	Financial	Non-Financial
> 0\$	0.14	0.16	0.24	14.29%	6.48%	18.05%
Median	0.15	0.12	0.24	15.31%	4.86%	18.05%
Terciles	0.22	0.20	0.33	22.45%	8.10%	24.81%
P75	0.10	0.08	0.15	10.20%	3.24%	11.28%
US	Absolute IOp			Relative IOp		
Partition of inheritances	Total	Total	Total	Total	Financial	Non-Financial
> 0\$	0.15	0.21	0.24	8.11%	6.48%	9.84%
Median	0.20	0.36	0.22	10.81%	11.11%	9.02%
Terciles	0.28	0.46	0.32	15.14%	14.20%	13.11%
P75	0.23	0.43	0.18	12.43%	13.27%	7.38%

Note: IOp measures calculated with the ex-ante parametric Ferreira and Guignoux (2011) approach, with different ad-hoc discretizations over the continuous circumstance (inheritances).

Table A.2: IOp with Trees and Forests, alphas at 1% and 5% (Gini).

CANADA	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
Tree 1%	28.16	42.72	26.95	39.85%	51.04%	35.98%
Tree 5%	29.57	49.35	26.95	41.85%	58.96%	35.98%
Forest 1%	29.61	48.59	27.47	41.90%	58.05%	36.68%
Forest 5%	29.59	48.66	27.44	41.88%	58.14%	36.64%
ITALY	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
Tree 1%	21.71	31.45	22.76	36.80%	42.52%	37.55%
Tree 5%	21.71	31.45	23.75	36.80%	42.52%	39.18%
Forest 1%	22.02	32.47	23.32	37.32%	43.90%	38.48%
Forest 5%	22.02	32.62	23.20	37.32%	44.10%	38.28%
SPAIN	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
Tree 1%	41.37	55.68	46.08	69.83%	66.18%	76.54%
Tree 5%	41.31	55.27	45.91	69.73%	65.70%	76.26%
Forest 1%	40.73	54.95	45.97	68.75%	65.32%	76.36%
Forest 5%	40.74	54.87	45.98	68.77%	65.22%	76.38%
US	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
Tree 1%	55.84	69.74	54.74	69.56%	76.14%	66.62%
Tree 5%	56.50	69.74	56.51	70.38%	76.14%	68.77%
Forest 1%	55.11	68.43	54.72	68.65%	74.71%	66.59%
Forest 5%	54.87	68.82	54.83	68.35%	75.13%	66.73%

Note: IOp measures using trees and random forests. Alphas are exogenously set to 1% and 5%. Forests are calculated after 500 replications.

Table A.3: IOp with ChiMerge, Trees and Forests, alphas endogenous at 1% and 5% (MLD).

CANADA	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	0,12	0,23	0,10	7.95%	8.81%	5.24%
Tree 1%	0,16	0,36	0,15	10.60%	13.79%	7.85%
Tree 5%	0,18	0,44	0,15	11.92%	16.86%	7.85%
Tree end	0,18	0,44	0,15	11.92%	16.86%	7.85%
Forest 1%	0,18	0,42	0,15	11.92%	16.09%	7.85%
Forest 5%	0,18	0,43	0,15	11.92%	16.48%	7.85%
Forest end.	0,17	0,42	0,15	11.26%	16.09%	7.85%
ITALY	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	0,11	0,12	0,12	12.36%	5.43%	11.88%
Tree 1%	0,12	0,21	0,14	13.48%	9.50%	13.86%
Tree 5%	0,12	0,21	0,14	13.48%	9.50%	13.86%
Tree end	0,12	0,21	0,14	13.48%	9.50%	13.86%
Forest 1%	0,12	0,24	0,13	13.48%	10.86%	12.87%
Forest 5%	0,12	0,24	0,13	13.48%	10.86%	12.87%
Forest end.	0,12	0,22	0,13	13.48%	9.95%	12.87%
SPAIN	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	0,17	0,16	0,28	17.35%	6.48%	21.05%
Tree 1%	0,29	0,55	0,36	29.59%	22.27%	27.07%
Tree 5%	0,28	0,55	0,36	28.57%	22.27%	27.07%
Tree end	0,29	0,55	0,36	29.59%	22.27%	27.07%
Forest 1%	0,29	0,53	0,35	29.59%	21.46%	26.32%
Forest 5%	0,29	0,53	0,36	29.59%	21.46%	27.07%
Forest end.	0,28	0,53	0,35	28.57%	21.46%	26.32%
US	Absolute IOp			Relative IOp		
ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
ChiMerge	0,19	0,27	0,27	10.27%	8.33%	11.07%
Tree 1%	0,63	0,99	0,61	34.05%	30.56%	25.00%
Tree 5%	0,65	0,99	0,61	35.14%	30.56%	25.00%
Tree end	0,65	0,99	0,61	35.14%	30.56%	25.00%
Forest 1%	0,59	0,93	0,56	31.89%	28.70%	22.95%
Forest 5%	0,59	0,94	0,56	31.89%	29.01%	22.95%
Forest end.	0,60	0,94	0,56	32.43%	29.01%	22.95%

Note: IOp measures calculated with several Machine Learning techniques. ChiMerge measures are calculated applying the Ferreira and Gignoux (2011). Trees and Forests have endogenously tuned their respective alphas following the K-fold Cross Validation technique explained in the Technical Appendix. Other alphas are exogenously set to 1% and 5%. Forests are calculated after 500 replications.

Table A.4: IOp estimates with the inheritances and parental education (Italy).

Italy	Absolute IOp			Relative IOp		
Partition and ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
Gini						
> 0\$	30.43	42.15	29.81	51.58%	56.99%	49.18%
Median	31.18	41.64	31.21	52.85%	56.30%	51.49%
Terciles	34.02	44.90	34.17	57.66%	60.71%	56.38%
P75	29.44	43.83	28.89	49.90%	59.26%	47.67%
Tree 1%	30.73	43.20	30.71	52.08%	58.41%	50.67%
Tree 5%	31.71	45.77	31.13	53.75%	61.88%	51.36%
Forest 1%	31.11	46.39	31.37	52.73%	62.72%	51.76%
Forest 5%	31.12	46.39	31.35	52.75%	62.72%	51.72%
MLD						
> 0\$	0.15	0.31	0.14	16.85%	14.03%	13.86%
Median	0,16	0,31	0,17	17.98%	14.03%	16.83%
Terciles	0,20	0,35	0,20	22.47%	15.84%	19.80%
P75	0.15	0.35	0.16	16.85%	15.84%	15.84%
ChiMerge	0,18	0,35	0,18	20.22%	15.84%	17.82%
Tree 1%	0,17	0,38	0,17	19.10%	17.19%	16.83%
Tree 5%	0,18	0,41	0,17	20.22%	18.55%	16.83%
Tree end	0.18	0.41	0.17	20.22%	18.55%	16.83%
Forest 1%	0,17	0,41	0,17	19.10%	18.55%	16.83%
Forest 5%	0,17	0,41	0,17	19.10%	18.55%	16.83%
Forest end	0.17	0.41	0.17	19.10%	18.55%	16.83%

Note: IOp measures calculated with several Machine Learning techniques. ChiMerge measures are calculated applying the Ferreira and Gignoux (2011). Trees and Forests have endogenously tuned their respective alphas following the K-fold Cross Validation technique explained in the Technical Appendix. Other alphas are exogenously set to 1% and 5%. Forests are calculated after 500 replications.

Table A.5: IOp estimates with the inheritances and parental education (the US).

US	Absolute IOp			Relative IOp		
Partition and ML method	Total	Financial	Non-Financial	Total	Financial	Non-Financial
Gini						
> 0\$	38.51	46.91	46.09	47.97%	51.21%	56.09%
Median	42.07	54.37	45.98	52.40%	59.36%	55.96%
Terciles	46.61	58.60	51.18	58.06%	63.97%	62.29%
P75	42.52	56.19	42.24	52.96%	61.34%	51.41%
Tree 1%	61.85	73.53	53.20	77.04%	80.27%	64.74%
Tree 5%	62.14	74.62	53.20	77.40%	81.46%	64.74%
Forest 1%	55.54	69.82	53.99	69.18%	76.22%	65.71%
Forest 5%	55.46	69.10	53.71	69.08%	75.44%	65.36%
MLD						
> 0\$	0.26	0.45	0.38	14.05%	13.89%	15.57%
Median	0,33	0,59	0,38	17.84%	18.21%	15.57%
Terciles	0,39	0,67	0,46	21.08%	20.68%	18.85%
P75	0.35	0.66	0.34	18.92%	20.37%	13.93%
ChiMerge	0,31	0,51	0,42	16.76%	15.74%	17.21%
Tree 1%	0,73	1,18	0,57	39.46%	36.42%	23.36%
Tree 5%	0,74	1,21	0,57	40.00%	37.35%	23.36%
Tree end	0,74	1,21	0,57	40.00%	37.35%	23.36%
Forest 1%	0,61	1,05	0,60	32.97%	32.41%	24.59%
Forest 5%	0,58	1,02	0,56	31.35%	31.48%	22.95%
Forest end.	0,61	1,04	0,55	14.05%	13.89%	15.57%

Note: IOp measures calculated with several Machine Learning techniques. ChiMerge measures are calculated applying the Ferreira and Gignoux (2011). Trees and Forests have endogenously tuned their respective alphas following the K-fold Cross Validation technique explained in the Technical Appendix. Other alphas are exogenously set to 1% and 5%. Forests are calculated after 500 replications.

Table A.6: Single Parameter Gini Absolute IOp with inheritances.

Canada	Total	Financial	Non-financial
v=2	29.59	47.69	27.39
v=2.5	14.79	25.81	13.95
v=3	6.93	12.76	6.86
v=3.5	3.28	6.26	3.55
v=4	1.59	3.14	2.07
v=4.5	0.80	1.62	1.44
v=5	0.42	0.86	1.19
Spain	Total	Financial	Non-financial
v=2	40.77	54.81	46.01
v=2.5	29.70	42.10	35.98
v=3	20.66	30.63	26.72
v=3.5	13.98	21.59	19.12
v=4	9.33	15.01	13.39
v=4.5	6.20	10.44	9.27
v=5	4.12	7.35	6.39
Italy	Total	Financial	Non-financial
v=2	22.01	32.50	23.20
v=2.5	8.09	14.27	8.49
v=3	3.16	6.06	3.29
v=3.5	1.28	2.61	1.32
v=4	0.53	1.17	0.54
v=4.5	0.22	0.58	0.22
v=5	0.09	0.34	0.09
US	Total	Financial	Non-financial
v=2	55.06	68.66	54.70
v=2.5	28.43	40.70	32.55
v=3	12.46	18.71	16.62
v=3.5	5.33	7.95	8.03
v=4	2.40	3.65	3.90
v=4.5	1.21	2.03	1.98
v=5	0.73	1.45	1.10

Note: IOp measures calculated with random forests with 500 replications and an alpha endogenously tuned following the K-fold Cross Validation technique explained in the Technical Appendix.

Table A.7: Single Parameter Gini Absolute IOp with inheritances and parental education.

Italy	Parental Education			Parental Education and Inheritances		
	Total	Financial	Non-financial	Total	Financial	Non-financial
v=2	22.06	39.10	20.97	30.94	45.58	31.22
v=2.5	18.24	37.56	17.17	20.42	41.36	19.88
v=3	15.71	36.04	14.57	15.71	38.88	14.85
v=3.5	14.10	34.99	12.86	13.32	35.89	12.25
v=4	13.07	34.42	11.73	11.90	34.78	10.68
v=4.5	12.35	34.17	10.94	10.87	33.92	9.59
v=5	11.80	34.05	10.33	10.06	32.91	8.71
US	Parental Education			Parental Education and Inheritances		
	Total	Financial	Non-financial	Total	Financial	Non-financial
v=2	34.74	42.42	36.85	55.65	68.83	53.55
v=2.5	36.75	45.85	37.42	45.06	56.26	46.90
v=3	37.73	47.90	37.15	40.06	50.75	42.56
v=3.5	38.05	48.96	36.50	37.70	48.65	39.44
v=4	37.97	49.34	35.77	36.31	47.51	37.24
v=4.5	37.67	49.26	35.10	35.34	46.57	35.79
v=5	37.27	48.89	34.55	34.63	45.68	34.91

Note: IOp measures calculated with random forests with 500 replications and an alpha endogenously tuned following the K-fold Cross Validation technique explained in the Technical Appendix.

Technical Appendix

In this Appendix we offer a detailed explanation of the Chi-test and the K -fold cross validation methods.

The Chi-Test Method

The null hypothesis of this test can be stated as “The relative cumulative distribution in two adjacent intervals are realizations of the same underlying distribution”, while, the alternative hypothesis says that “The relative cumulative distribution in two adjacent intervals are realizations of different underlying distributions” (Holmes and Jain, 2012). The test is performed using the basic Pearson’s χ^2 test:

$$D = \sum_{l=1}^m \sum_{b=1}^c \left(\frac{N_{lb} - E_{lb}}{E_{lb}} \right)^2 \quad (T.A. 1)$$

$$E_{lb} = \frac{R_l C_b}{N} \quad (T.A. 2)$$

$$R_l = \sum_{b=1}^c N_{lb} \quad (T.A. 3)$$

$$C_b = \sum_{l=1}^2 N_{lb}, \quad (T.A. 4)$$

$$N = \sum_{b=1}^c C_b \quad (T.A. 5)$$

where $m = 2$, as there are 2 intervals being compared, and c is the total number of classes. N_{lb} is the number of observations in the l^{th} interval and b^{th} class, while E^{lb} is the expected frequency of N_{lb} . Note that R_l is the number of observations in the l^{th} interval, C_b is the number of observations in the b^{th} class, and N is the total number of observations,.

In short, this test sums the relative squared differences between the expected and the observed class occurrences in both intervals potentially combined. This deviation is assumed to be distributed as a χ^2 with $f = c - 1$ degrees of freedom. Then, we compare the obtained statistic D with the critical value of χ^2 , given a predefined level of significance. If D is smaller than the critical value, we do not reject the null hypothesis and combine the two intervals under scrutiny. If D is higher than the selected critical value, we reject the null hypothesis and leave the cut point.

The K-fold Cross Validation Method

Imposing a certain alpha level, such as 0.01 or 0.05, can bias the results, as they are exogenous econometric conventions that might not allow us to collect all information gathered in the data (Brunori et al. 2020). To solve this issue and test the accuracy of our results we implement the so-called *K*-fold Cross Validation method, one of the most popular ML techniques (Rodríguez et al., 2009; Fushiki, 2011).

First, the algorithm divides the whole sample into *K* subsamples, also called folds. The optimal number of possible folds depends on the dataset. To make sure that we always have enough degrees of freedom in every fold, we set $K = 5$. For robustness, we have also tested $K = 6$ and $K = 7$, without significant changes in the results.

The conditional inference tree algorithm is run on the union of $K - 1$ folds (training sample, m) for varying levels of alpha, while taking out the last k^{th} fold (validation sample, k). After that, we use the mean squared prediction error (MSPE) to evaluate the difference of the prediction in the training sample with respect to the validation sample:

$$MSPE_k(\alpha) = \sum \frac{N^m}{N^k} \sum \frac{1}{N^m} (w_i^k - \mu^m(\alpha))^2, \quad (T.A. 6)$$

where w_i^k is the output vector of the validation fold, and $\mu^m(\alpha)$ collects the predictions emanated from the training sample for a certain alpha level. Note that N^m and N^k are the sample sizes of the training and validation folds, respectively. This exercise is repeated four times more for the same alpha level, sequentially leaving out one *K*-fold at a time.

Finally, we construct the average mean squared error of the cross-validation process (MSECV) as the average of the five MSPEs:

$$MSECV(\alpha) = \frac{1}{5} \sum_{k=1}^5 MSPE_k(\alpha) \quad (T.A. 7)$$

After running the algorithm a set of possible alpha levels, we select the one that provides the smallest MSECV.