Social Indicators Research (2026) 181:65
https://doi.org/10.1007/s11205-025-03788-3

ORIGINAL RESEARCH

An Atlas of Educational Inequality in Italy: Outcomes, Disparities and Opportunities

Paolo Brunori[1,2] · Emanuele Fedeli[3] · Moris Triventi[3]

Received: 1 April 2025 / Accepted: 12 December 2025

Abstract

We present an in-depth analysis of educational inequality in Italy, focusing on disparities in educational outcomes and opportunities across different socio-economic, gender, and migration backgrounds. Leveraging administrative longitudinal data, we construct a dataset of 386 small geographical areas with a sufficient sample size to assess the extent to which key ascriptive characteristics predict the mathematical achievement of Italian students in the 5th grade of primary school. Our findings highlight a substantial influence of ascriptive characteristics on students' educational attainment, able to correctly predict out-of-sample up to 20% of the variability despite the relatively small sample size. We show significant geographical variation that previous studies, based on larger geographical aggregations, were unable to observe comprehensively. Additionally, we identify a weak yet negative trade-off between equality and average attainment, which is more pronounced in southern areas, where higher achievement is associated with greater variance and a stronger influence of ascriptive characteristics. Among the predictors, we find that mother's education plays a predominant role in most of the country.

Keywords Inequality · Opportunity · Machine learning · Education · Mathematics · INVALSI

This research was supported by Fondazione Cariplo under the project EDUFAIR – Mapping and Tackling Unfair Educational Inequality (Project Code 2024-1575), led by Moris Triventi, and funded within the INEQUALITIES RESEARCH 2024.

✉ Paolo Brunori
p.brunori@lse.ac.uk

Emanuele Fedeli
emanuele.fedeli@unimi.it

Moris Triventi
moris.triventi@unimi.it

1 London School of Economics, London, UK

2 University of Florence, Florence, Italy

3 University of Milan, Milan, Italy

Published online: 30 January 2026

1 Introduction

The persistent link between individuals' ascriptive characteristics – such as gender, socio-economic and migration background – and their educational attainment raises significant concerns for educational systems, both in terms of equity and efficiency. From an equity standpoint, a society where students' educational outcomes are heavily influenced by their family background or other ascribed characteristics is often seen as fundamentally unjust, as it undermines the ideal of equal opportunity (Rawls, 1971; Roemer, 1998; Swift, 2005). This situation can be criticized since it perpetuates inequality and hinders intergenerational social mobility (Corak, 2013; Breen & Müller, 2020). Moreover, the widely accepted principle of meritocracy suggests that individuals should have access to education and occupational opportunities based on their abilities and efforts rather than their inherited characteristics (Breen & Goldthorpe, 2001). However, empirical research has consistently shown that socioeconomic background, ethnic origin, gender and other individual circumstances strongly affects educational success (Boudon, 1974; Erikson & Jonsson, 1996; Blossfeld et al., 2016; Ferreira & Giugnoux, 2014; Biagi et al., 2022; Liu et al., 2022).

A society characterized by high levels of inequality, which penalizes competent and hard-working students from disadvantaged backgrounds, also misallocates human capital investment. When access to high-quality education is determined by individual circumstances on which individuals bear no responsibility rather than ability (Roemer, 1998), economies suffer from underutilization of talent, leading to lower productivity and innovation (Heckman, 2006; Brunello & Checchi, 2007). From an efficiency perspective, such inequalities are also detrimental to the effective use of human capital, which ultimately impacts economic development and social cohesion (Becker, 1964; Hanushek & Woessmann, 2008)[1].

While educational attainment has traditionally been the primary focus of sociological research, the growing role of human capital in modern economies has shifted attention towards more direct measures of student competencies (Heckman, Stixrud, & Urzua, 2006; Hanushek & Woessmann, 2011). Standardized test scores in key subject domains – such as mathematics and reading – allow researchers to study educational inequalities by capturing differences in cognitive skills and academic preparedness, which are crucial for later labor market success (Jencks et al., 1972; Murnane, Willett, & Levy, 1995) and life-chances in general (OECD 2010).

Since "equal opportunity" is a term understood in various ways, ranging from the absence of discrimination to meritocracy and equality, it is important to clarify the scope of our analysis. We remain agnostic about the normative implications of equal opportunity and instead adopt a pragmatic definition. We define equal educational opportunity (EEOp hereafter) as a condition in which ascriptive characteristics do not predict educational outcomes. Therefore, when inherited characteristics, such as gender, place of birth, and key aspects of socioeconomic background, are found to predict (out-of-sample) inequality in educational outcomes, we consider this variability as inequality of educational opportunities (IEOp hereafter).

[1] Note that this does not imply that an educational system that is more able to eliminate advantages arising from ascriptive characteristics will necessarily produce higher average results. As we will discuss when commenting on the results, redistribution does have costs in terms of resources, and an efficiency-equity trade-off may emerge in practice.

A substantial body of research has documented cross-national differences in the IEOp among children from diverse social backgrounds, both in terms of educational transitions and academic performance (Breen & Jonsson, 2005; Jackson, 2013; van de Werfhorst, 2018; Blossfeld et al., 2016). Moreover, several works documented the existence of significant achievement gaps related to migration background and ethnic origin (Alba et al., 2011; Borgna & Contini, 2014), as well as gender, but with heterogeneous patterns across subjects (Buchmann, DiPrete & McDaniel 2008; Tsai, Smith, & Hauser 2018).

Despite extensive research on educational inequalities, several important limitations persist, particularly in comparative studies. First, much of the existing literature has focused on cross-national differences in IEOp, often overlooking within-country heterogeneity (Jackson, 2013; OECD, 2018). While cross-national analyses provide valuable insights into broad institutional patterns, they tend to treat countries as homogeneous units, failing to capture substantial territorial disparities in educational outcomes within national borders (Agasisti & Cordero-Ferrera, 2013). This is particularly problematic in countries like Italy, where regional and local inequalities in educational performance are well-documented but often neglected in large-scale international comparisons (Bratti, Checchi, & Filippin, 2007).

Second, most studies have relied on single indicators of individual circumstances or focus on isolated predictors of student achievement. While these indicators are undoubtedly important, they fail to capture the complex interplay of multiple ascriptive characteristics, such as the joint influence of mother's and father's background, migration status, and gender. This omission leads to a partial understanding of IEOp, as it does not account for how different social circumstances interact in shaping educational opportunities (Ferreira & Gignoux, 2014; Marrero, Biagi et al., 2022; Palomino & Sicilia, 2024).

Third, the dominance of parametric regression models in prior research has constrained the ability to flexibly model interactions between variables capturing individual circumstances (Brunori, Hufe & Mahler, 2023). Traditional approaches often impose additive effects and assume linear relationships, limiting the ability to detect more complex and context-dependent inequalities.

The aim of this work is to develop a comprehensive measure of IEOp and to assess to what extent it varies across territories within Italy, a country characterized by substantial heterogeneity in socioeconomic and cultural resources across geographical areas (Felice 2019; Istat 2022). Our research questions are as follows: What is the level of IEOp and to what extent does it vary across Italian micro-areas? Is there a trade-off between achievement levels and IEOp? Which are the most important drivers of IEOp and how this varies across micro-areas?

Our study aims to provide a fine-grained analysis of inequality of IEOp across Italian micro-areas (a mix of medium and large municipalities and aggregations of small municipalities), addressing key gaps in previous research. Specifically, we examine how the influence of students' individual circumstances on student mathematics achievement varies geographically, moving beyond broad national or regional comparisons. To do so, we adopt a holistic approach that considers multiple ascriptive characteristics—including parental education, occupational class, migration background, and gender—rather than treating them in isolation. We focus primarily on mathematics achievement, as it is more strongly linked to long-term educational and labour market outcomes and tends to exhibit greater variability between students than other domains, thereby offering a more sensitive lens for detecting geographical patterns of inequality (Hanushek & Woessmann, 2008). However, we in replicate all analysis also for language and we comment key differences between the two outcomes.

Methodologically, we employ advanced machine learning techniques, particularly conditional inference random forest (Hothorn et al., 2006), to capture complex interactions between individual circumstances and educational outcomes in a data-driven and flexible manner. We select random forests as they outperform both a simpler and more traditional OLS-based approaches in the majortity of the areas. The use of regression-based algorithm was tested both including and excluding pairwise interactions of variables and by attempting an improvement of their predictive ability by shrinkage (James et al., 2013). An additional advantage of random forests is that, unlike traditional regression models, they are able to quantify patterns of IEOp taking into account complex interactions of circumstances by aggregating hundreds of conditional inference regression trees. Using INVALSI-SNV data on 5th-grade students from 2012 to 2022, we leverage the dataset's large scale (around 3 million cases) to map territorial variations in IEOp. Through a fine-grained geographical analysis, a multidimensional conceptualization of individual circumstances, and an advanced machine learning-based methodology, we thus contribute to a more comprehensive understanding of how inequality of educational opportunity varies across micro-areas and student categories.

The paper is organized as follows: in the next section, we discuss key contributions in the empirical literature, by focusing on studies that analysed (a) educational achievement inequality; (b) geographical/local variation in intergenerational mobility and educational performance; (c) the relationship between educational achievement levels and inequality. In the third section, we present the data, variables and methods used, while the fourth section presents the empirical results. At the end, in the fifth section we discuss the main findings in relation to previous studies and draw some conclusions.

2 Literature Review

2.1 Empirical Evidence on Educational Achievement Inequality

In what follows, we primarily review empirical studies conducted in the Italian context, while also drawing on comparative research that includes Italy, to situate our analysis within both national and cross-national perspectives. To better understand the persistent disparities in educational outcomes, two important lines of research have emerged, both highlighting the systemic inequalities present in the educational system. The first line of inquiry focuses on the IEOp, examining how gender, social and ethnic origins significantly shape children's competencies in key subjects. The second line of research investigates geographical inequality in student achievement, uncovering large regional and provincial disparities in educational outcomes, particularly within countries like Italy. This research emphasizes how differences in regional resources, local educational systems, and socio-economic contexts contribute to uneven educational standards, with students in northern areas typically outperforming those in the south.

In the first line of research, studies on inequality of educational opportunity consistently demonstrate that gender, social and ethnic origin are key factors shaping children's academic performance, both in terms of average competencies and across the entire distribution of test scores. For instance, various studies, both in comparative perspective (Schnepf, 2007; Borgna, 2014; Marrero et al. 2024) and in Italy (Azzolini and Schnell, 2012; Triventi et al., 2022) found that children from ethnic minority backgrounds perform significantly worse than their native peers, even after accounting for socio-economic status. Similarly, several studies emphasize that social class and parental education remains a strong predictor of performance in core subjects such as math and reading (Pensiero, Giancola & Barone 2019; Giambona and Porcu 2015; Costanzo and Desimoni 2007; Biagi et al. 2022). These studies highlight that inequalities are systemic and deeply rooted in both social and educational structures, often reflecting broader issues of social stratification and exclusion. Further extending these findings, various studies show that these inequalities are not confined to early childhood or primary school but persist throughout secondary education and even into young adulthood (Dammrich and Triventi, 2018; Skopek & Passaretta, 2021).

Gender differences in academic achievement, particularly in standardized test scores, have been widely documented across various national contexts. While girls generally outperform boys in reading and verbal skills, boys tend to achieve higher scores in mathematics and, in some cases, science-related subjects. These patterns, however, are not uniform across countries and education systems, reflecting the role of cultural, institutional, and socio-economic factors in shaping gender disparities (Buchmann, DiPrete, & McDaniel, 2008; Tsai, Smith, & Hauser, 2018). Italy follows the general international trend, with girls significantly outperforming boys in reading and boys holding an advantage in mathematics, although the magnitude of the gender gap in math is particularly notable (Contini, Di Tommaso, & Mendolia, 2017). Evidence from INVALSI national standardized assessments and PISA data suggests that the gender gap in mathematics in Italy is larger than the OECD average, particularly in Southern regions, where socio-economic inequalities may amplify gender disparities.

2.2 The Relevance of Geographical Areas in Creating Educational and Social Mobility Opportunities

Geographical differences play a crucial role in shaping opportunities for educational attainment and social mobility. A growing body of research highlights that intergenerational mobility—the extent to which children's socioeconomic status differs from that of their parents—varies significantly within countries, suggesting that where one grows up can be just as influential as individual or family characteristics in determining life chances. These findings challenge traditional views of mobility as a uniform national phenomenon, emphasizing instead the importance of local institutional, economic, and social structures in fostering or constraining mobility.

Studies on intergenerational mobility in Britain, the United States, and across Europe consistently show substantial spatial variation in mobility rates. In the United States, Chetty et al. (2014) demonstrate that intergenerational mobility varies sharply across different cities and states. Key factors associated with higher mobility include school quality, social capital, income equality, and family stability. These differences have led to the idea of 'place effects', where the local environment systematically influences an individual's long-term socioeconomic outcomes. Berger and Engzell (2019) further show that these patterns are not uniquely American but share similarities with historical mobility trends observed in European countries, suggesting deep-rooted institutional and policy-driven determinants of regional disparities in mobility. In Britain, Breen and In (2024) document striking regional disparities, with mobility being higher in areas with greater economic opportunities, better educational resources, and more inclusive labor markets. These findings resonate with the well-documented ‘geography of opportunity’ framework, which emphasizes that access to quality education, social networks, and economic opportunities is unevenly distributed across regions.

Italy provides a particularly striking example of spatial inequality in social mobility. Acciari, Polo, and Violante (2022) document large interregional differences in mobility rates, with the North generally offering greater opportunities for upward mobility compared to the South. Their analysis highlights that children born in wealthier northern regions have a much higher probability of surpassing their parents’ income and educational status than those born in the economically disadvantaged South, where labor market rigidities, lower educational investments, and weaker social infrastructures limit mobility. This North-South divide in Italy reflects long-standing structural disparities. The educational system plays a crucial role in reinforcing these gaps, as students in the South tend to attend lower-performing schools with fewer resources and limited access to higher education pathways. Additionally, labor market conditions differ significantly, with higher youth unemployment and weaker job placement services in the South contributing to lower returns on education. Checchi and Peragine (2005) examine regional disparities in inequality of opportunity in Italy, emphasizing how differences in economic and educational conditions across regions contribute to unequal life chances. Using a framework that distinguishes between inequality due to circumstances (e.g., family background, region of birth) and inequality due to effort, they find that regional disparities significantly shape labor market outcomes. Their analysis highlights that inequality of opportunity is particularly pronounced in Southern Italy, where structural disadvantages limit social mobility, reinforcing persistent economic and educational gaps between the North and the South.

Looking more specifically at educational outcomes, previous works have clearly highlighted the existence of a North-South gradient in student competencies levels (Bratti, Checchi & Filippin 2007; Agasisti & Cordero-Ferrera 2013), and geographical heterogeneity in educational standards (Argentin & Triventi 2015).These studies consistently find that in the more economically developed North average performance in standardized tests and educational standards are higher than in Southern areas. Less is known about whether there is significant geographical variation in educational inequalities and whether individual circumstances play a different role in generating educational opportunities in different contexts. One exception is the recent study by Triventi & Fedeli (2025), who show the existence of substantial provincial heterogeneity of social background inequalities in accessing the academic track in high school. They suggest that geographical disparities in educational inequality can arise due to differences in school resources, teacher quality, economic conditions, and local educational policies. Furthermore, Daniele (2021) shows that relative poverty—not just absolute deprivation—plays a key role in shaping educational outcomes across geographical areas in Italy, as students from low-income backgrounds face additional social and psychological barriers to achievement.

While these studies illuminate the existence of geographical patterns in educational performance, they mostly focus on the North-South divide, overlooking the potential heterogeneity within these broad macro-areas. Moreover, they do not examine whether IEOp—defined as the portion of outcome variability attributable to individual circumstances beyond students' control—varies across territories, nor do they analyze the relative contribution of specific stratification factors. Addressing these gaps is a key objective of our study.

2.3 Is There a Trade-off Between Levels of Educational Achievement and inequality?

One of the central debates in educational policy concerns the trade-off between equity (or equality) and effectiveness (or performance levels) in student achievement. This debate revolves around whether policies that promote educational equity—reducing achievement gaps between students of different socioeconomic backgrounds—come at the cost of overall academic performance, or whether it is possible to design systems that achieve both equity and high effectiveness (Ferrer-Esteban, 2016; Parker et al., 2021).

The trade-off between equity and effectiveness is often conceptualized in terms of resource allocation, school-level policies, and institutional structures (Ferraro & Põder, 2018; Parker et al., 2021). Some scholars argue that policies aimed at reducing educational inequalities, such as comprehensive schooling and greater public investment in disadvantaged schools, may lower overall performance by redistributing resources away from high-achieving students or by limiting selection mechanisms that allow the best students to thrive. Others contend that reducing educational disparities enhances overall efficiency, as greater inclusivity and equal opportunities foster a broader pool of well-educated individuals, benefiting society as a whole (Van de Werfhorst & Mijs, 2010).

Empirical evidence on this trade-off remains mixed. Cross-national studies indicate that education systems that emphasize selection and tracking tend to produce higher levels of student performance but also greater inequality in outcomes (Parker et al., 2018). Conversely, systems that emphasize equal educational opportunities, such as those in Scandinavian countries, tend to achieve smaller performance gaps. While their average levels of performance are usually rather high, sometimes they struggle with maintaining high levels of excellence at the top of the achievement distribution. Ferraro and Põder (2018) analyze school-level policies and find that while more inclusive and equitable policies can improve average achievement among disadvantaged students, they may also lead to lower overall efficiency, especially in competitive environments. At the national and regional levels, research in developing and middle-income countries suggests that increasing public expenditure on disadvantaged students can improve equity without necessarily lowering overall performance. Gershberg and Schuermann (2001) provide evidence from Mexico showing that well-targeted investments in educational resources for low-income students improve both equity and efficiency, challenging the notion of an inevitable trade-off.

While much of the existing research has focused on cross-national comparisons, an important question is whether similar patterns emerge within the same country, across regions, cities, or even neighborhoods. Studies on regional disparities in educational achievement suggest that within-country variation in student performance and equity often mirrors the patterns found at the cross-national level (Ning et al., 2016; Daniele, 2021). For instance, within Italy, significant disparities exist between northern and southern regions in terms of both average performance and educational inequality (Acciari, Polo & Violante, 2022). The North tends to perform better overall, but also exhibits higher levels of stratification, whereas the South has lower performance levels but smaller within-region inequalities. Similarly, Checchi and Peragine (2005) argue that regional disparities in opportunity play a crucial role in shaping educational achievement, indicating that place-based factors contribute to both equity and effectiveness.

3 Analytical Strategy

3.1 Data

We use data from INVALSI-SNV, which provide information on competencies and socio-demographic characteristics across key school grades for the whole population of students. In this work, we compile data from seven academic years (2012–2018) to create a student population panel dataset, tracking 3,468 million students enrolled in the 5th grade of primary school—around 500,000 students per academic year. Our analytical sample consists of 2,473,514 students, with a 27% missing rate over seven academic years.

We retain all observations with valid outcome data. For categorical parental background variables with missing values, we introduce a distinct ‚missing' category for each variable. This approach preserves the full analytical sample while allowing the model to capture potential patterns associated with missingness. Descriptive statistics on the extent of missing data are reported in the Online Appendix (Table A2).

We pool INVALSI data for 5th grade students from the 2012–2018 cohorts into a single analytical dataset. This choice is primarily motivated by the need to ensure an adequate sample size within each LMA—set at a minimum threshold of 3,000 students—to allow robust estimation of machine learning models and facilitate out-of-sample prediction. Pooling also enables us to construct fine-grained geographical estimates that would not be reliable if based on a single year's data (see below). This approach implicitly assumes that the structure of educational inequalities within LMAs remains relatively stable over the pooled period. To assess the plausibility of this assumption, we conducted a test to assess whether merging more than one year results in systematically different estimates (reported in the last section of the Online Appendix).

3.1.1 Geographical Variables

By leveraging the large size of the INVALSI dataset, we are able to disaggregate the analysis by LMAs and LMA hinterlands, thereby gaining a more fine-grained perspective on territorial heterogeneity in educational inequality. LMAs, our primary geographical unit of analysis, are functional territories defined by the Italian National Statistical Institute using commuting flow data, grouping municipalities into coherent socio-economic spaces where most residents both live and work. Compared to administrative boundaries (e.g., provinces or regions), LMAs better capture the socio-economic and institutional environments in which schools operate, including local labour markets, commuting patterns, and shared public services. This is particularly relevant for educational inequality, as school catchment areas and student flows often extend beyond individual municipalities but remain embedded within local economic and social systems (Istat, 2014).

More aggregate units, such as macro-areas or regions, risk masking important intra-regional heterogeneity, while smaller clusters of municipalities based only on geographical proximity may not reflect the functional socio-economic linkages that shape educational opportunities. LMAs thus provide an analytically meaningful compromise: they are large enough to ensure adequate sample sizes for robust statistical modelling, yet fine-grained enough to detect spatial variation in inequality that would be invisible at broader scales.

Starting from over 5,000 municipalities with at least one primary school, we applied aggregation rules to ensure each area includes a sufficient number of observations for robust statistical modelling. These rules account for student population size, distinguish between stand-alone municipalities and their hinterlands, and merge smaller LMAs where necessary. This process results in 386 analytical areas that better reflect local socio-economic contexts and schooling markets than traditional administrative boundaries such as provinces. A detailed description of the aggregation procedure is provided in the Online Appendix.

To assess whether the geographical variations in the phenomena of interest vary mostly across or within more traditional territorial classifications used in the literature, we also consider macro-area as an additional geographical-level variable. It distinguishes between North-West, North-East Center, South and Isles. In some analyses, for parsimony, it has recoded into three categories (North, Center and South).

3.1.2 Individual Variables

The outcome variable in this analysis represents academic performance in mathematics for fifth-grade students, as assessed annually by INVALSI. The predictors include fundamental socio-demographic characteristics such as gender, migration background (natives, first-generation, and second-generation migrants), parental education level, and parental occupation (see Online Appendix Table A1 for summary statistics). Gender is coded as 0 for boys and 1 for girls. Migration background is coded as 0 for native students, 1 for second-generation migrants, and 2 for first-generation migrants. Parental education levels are recorded separately for mothers and fathers, with detailed information on each educational credential, including primary education, lower secondary education, upper secondary education, post-secondary qualifications, tertiary education. Parental occupations are classified into the following categories: Unemployed, Housewife, Retired, Worker, Service Worker/Cooperative Member, Teacher, Employee, Career Military Officer, Salaried Professional, Non-Commissioned Military Officer, Self-Employed Worker (e.g., merchant, direct farmer, craftsman, mechanic), Executive, University Professor, and Civil Servant.

4 Methods

4.1 Measures of Predictability

We are interested in mapping how mathematics scores are predicted by ascriptive characteristics in small areas of Italy. This empirical exercise can be understood from two distinct points of view: a descriptive perspective and a normative one. Predictability can inform us about specific conditions associated with particularly low performance, signaling the need for and facilitating targeted intervention. Moreover, a description of unexpectedly high performance for certain students in specific areas can shed light on virtuous situations and best practices, benefiting the broader policy debate. From a normative point of view, the fact that certain ascriptive characteristics predict school performance may signal a failure to provide equal educational opportunities and the persistence of educational inequality across generations. From a methodological perspective, our mapping exercise is similar in spirit to the Fragile Families Challenge (Salganik et al., 2019). In this competition dozens of research teams applied machine learning algorithms to predict six key outcomes[2] at age 15, based on thousands of predictors observed between ages 0 and 9, leveraging data from the 'Future of Families and Child Wellbeing Study' (FFCWS) project. Of course, our study is much narrower in scope, focusing only on mathematics outcomes in grade 5, and having access to a limited set of predictors. However, our maps are broader in terms of population coverage, while FFCWS sampled only children of unmarried couples, we do consider all children in the INVALSI data with a coverage close to 100%.

Consistent with the suggestion by Salganik et al. (2019), we define "predictability" as the proportion of variability in the variable of interest, y, that can be correctly predicted out-of-sample. This is obtained by training the learning algorithm on a sample comprising 70% of the observations (*training* set), then predicting the mathematical score on a hold-out sample consisting of the remaining 30% of the observations (*test* set) and evaluating the model's accuracy. Our measure of predictability is the out-of-sample R-squared (R^2_{OOS}):

$$R^2_{OOS} = 1 - \frac{\sum_{i\in Test}\left(y_i - \widehat{y}_i\right)^2}{\sum_{i\in Training}\left(y_i - \overline{y}_{Training}\right)^2}$$

Where $R^2_{OOS} \leq 1$.

R^2_{OOS} provides a relative measure of accuracy with respect two references: a negative R^2_{OOS} signals a predictive ability of the model lower than predicting using the mean in the training set, $R^2_{OOS} = 1$ signals perfect accuracy in the test set, which implies that the model predicts with no error out-of-sample. Note that this measure is an out-of-sample version of the so-called relative ex-ante inequality of opportunity, first introduced by van de Gaer (1993) and widely used in normative economics to quantify inequality of opportunity. For example, see Checchi and Peragine (2010) and Ferreira and Gignoux (2011) for applications to income distributions, Biagi et al. (2022) for an application to educational attainments and Ferreira and Gignoux (2014) for an analysis of educational outcomes, more closely related to our context (PISA standardized test scores).

To benchmark our analysis, consider that Salganik et al. (2020) report a predictability of 0.19 for Grade Point Average (GPA) at age 15 using the best-performing algorithm out of 160 participating in the competition. This value reflects several factors relevant to our analysis: the number of observable predictors available, the strength of the relationship between the outcome and these predictors, and the sample size (a larger training set provides more degrees of freedom to detect robust relationships between predictors and the dependent variable). In each of these aspects, the FFCWS data surpass those used in our analysis. The dataset contains 12,942 variables, while many may be irrelevant for predicting educational outcomes, a substantial portion captures key environmental factors and may relate to genetic endowment. Additionally, the training set consisted of 4,242 observations, a sample size larger than what is available for most small areas considered in our study. Given these differences, expecting a predictability score far above 0.19 in our analysis would be

[2] (1) child grade point average (GPA), (2) child grit, (3) household eviction, (4) household material hardship, (5) caregiver layoff, and (6) caregiver participation in job training.

unreasonable, even while recognizing that we observe characteristics and outcomes at the same point in time.

The fact that the ability to predict is monotonically increasing with the sample size explains also why we fixed a common sample size for all areas (n^*). When a municipality for a given year and grade does not have sufficient size, we merge its sample with the sample of other municipalities based on geographical proximity. When a municipality j has sample size larger than n^*, estimates are the average of the estimates obtained on a number of draws of n^* observations randomly selected without repetition from the original population (n_j). The number of iteration is inversely proportional to the sample. We adopt a rule of thumb based on the ratio between n^* and the population size to guarantee that, in each municipality, the probability to not include any observation to be smaller than 1%:

$$k = 1 + \frac{\log(0.01)}{\log\left(\frac{n_j - n^*}{n^*}\right)}$$

Where k is the number of samples.

Predictability as share of total variability, R^2_{OOS}, is our first variable of interest. However, this indicator by normalizing predicted inequality by total inequality is a misleading indicator of the absolute level of variability predicted by regressors. For this reason, we complement this value with a second measure of predictability that takes into account total explained variability: $IEOp = \sum_{i \in Test} (y_i - \widehat{y}_i)^2 \times R^2_{OOS}$.

Note that this measure is again analogous to an out-of-sample version of the ex-ante inequality of opportunity measure introduced by van de Gaer (1993) in its absolute version.

4.2 Algorithms

To evaluate predictability, we have estimated a series of conditional inference random forests, as proposed by Hothorn et al. (2006), one for each LMA.

Random forests obtain predictions by estimating a large number of regression trees to explain score variability. A regression tree is a simple algorithm that partitions the predictor space into non-overlapping regions and makes predictions by averaging the outcomes within each region. The partitioning is defined through recursive binary splits: the space is first divided into two regions, each of which is then subdivided into two, and so on, until a stopping rule is met. This recursive splitting is particularly well-suited for modeling data-generating processes in which the interaction of factors, rather than the factors themselves, drives variability in the dependent variable. Each split can be interpreted as introducing an interaction term into the model. However, regression trees, because of their interaction-based structure, tend to suffer from high variance. By estimating hundreds of such trees and aggregating their predictions, random forests provide an extremely effective approach for improving out-of-sample predictive performance (see again James et al. (2013) for a general discussion and Hothorn et al. (2006) for the specific property of conditional inference random forests).

Conditional inference random forests are tuned via 10-fold cross validation by testing different critical Bonferroni adjusted p-value to allow split (0, 0.05, 0.1, 0.25, 0.4) and for different number of input variables randomly sampled as candidates at each node [1–4].

Each forest is made of 500 trees. We chose to rely on random forest since it guarantees flexibility and parsimony, and it is the one that best performs on average in predicting educational score at LMA level, compared to two alternative models: (1) A benchmark OLS regression; (2) a regularized regression of a model including (and excluding) all pairwise interactions (LASSO) (see Friedman et al., 2008), in which shrinkage is performed by 10-fold cross validation.[3]

Random forest is the best-performing method for predicting out-of-sample data. Specifically, it produces the highest R-squared out-of-sample in 156 areas $(R^2_{OOS} = 0.0943)$. The second-best method, LASSO regularization of a standard OLS, performs best in 144 areas. A simple OLS achieves the highest performance in 80 areas, while a LASSO applied to an OLS model with all pairwise interactions performs best in only 12 areas.

The mean absolute difference between the predictions from the two best methods (LASSO and random forest) is 0.0099, exceeding 4% in only four areas. This suggests that using complex data-driven algorithms only marginally improves our ability to predict standardized educational outcomes in this specific setting. This could be due to specific generation process behind IEOp in Italy, in which interactions between ascriptive characteristics might play a limited role in predicting student achievement vis-à-vis standard additive effects. It could also be that the relatively limited difference is due to the fact we rely on a modest number of variables, which are all categorical. It could be that machine learning algorithms display more substantial advantages in terms of out-of-sample predictability in contexts with more complex data generating process, higher number of circumstances and more quantitative measures of individual circumstances. Moreover, the sample size set for the analysis (3,000 observations) also tends to favour algorithms characterized by low variance and relatively higher bias.

Moreover, R^2_{OOS} ranges from just above 1% to over 20%, but on average, they are about half of what is obtained with FFCWS data. To maximize comparability, in what follows we present all results based on the random forest model. We chose this model not only because it performs best overall, but also because its tree-aggregation structure provides a straightforward way to quantify the relative importance of predictors while accounting for complex interactions among them (see James et al. (2013) for a discussion). Results from other algorithms are available upon request.

4.3 Relative Importance and Partial Dependency Plots

Predictability is an interesting measure in itself; however, policymakers may be also interested in understanding the importance of specific factors in predicting variability in standardized mathematics test scores. We use two complementary statistical tools to shed light on the relative role of different ascriptive characteristics.

The first measure we use is the relative importance of each variable, estimated through a permutation approach. This method computes importance by randomly permuting values within covariates associated with the variable of interest. Consider a variable X_j by permuting randomly its values, X_j becomes necessarily orthogonal to the dependent variable. When all remaining (unpermuted) predictors are used to predict the dependent variable

[3] In the Online Appendix we show how the R^2_{OOS} based on random forests compare with R^2_{OOS} estimated using LASSO and a standard OLS.

alongside the permuted variable, the prediction accuracy necessarily weakens. We quantify relative importance of variable X_j by measuring the reduction in accuracy before and after permuting it, as explained in detail by Strobl et al. (2007).

This relative importance measure is determined by two factors: the extent to which variability in X_j is associated with variability in the dependent variable and [2] the extent to which X_j varies within the sample. To illustrate this, consider that we are assessing the impact of migration background. Suppose being a migrant is associated with a substantial penalty in a particular municipality. However, if the proportion of migrants in the sample is very small, the ability to predict math scores based on migration background will also be low. After all, if nearly all students were born in Italy, knowing their migration background would provide little predictive power. This would result on a low relative importance for the migration background variable. Conversely, if the proportion of migrants in the sample were close to 50%, even a moderate penalty associated with migration status could result in a substantial relative importance.

To avoid mixing the two concerns: the relative importance of a feature Vs. the magnitude of the penalty/premium associated with a given characteristic we complement the inspection of the relative importance measure with Partial Dependency Plots. Partial Dependence Plots (PDPs), originally introduced by Friedman (2001), are visual tools designed to help interpret machine learning outputs. They illustrate how changes in a specific predictor variable affect the predicted outcome while holding the influence of other variables constant.

The partial dependence function at a particular feature value, for example, having a mother in a high-skilled occupation, represents the average predicted outcome if all data points were forced to assume that feature value. In other words, what would be the average score if all students had a mother in a high-skilled occupation? This counterfactual exercise is implemented while keeping the distribution of all other features constant. The algorithm constructs a hypothetical distribution in which all mothers are white-collar workers, even though, in reality, it may be unrealistic for someone to have no formal education yet hold a high-skilled occupation.

PDPs can be used to detect areas in which being affected by a particular circumstance value is predictive of a particularly low/high performance. The value is not influenced by the share of individuals characterized with the particular value and therefore complement relative importance. However, PDPs should be interpreted with caution, particularly in the presence of highly correlated predictors or when analyzing categories or values that appear in very few observations.

5 Empirical Findings

5.1 Levels and Geographical Distribution of IEOp

Figures 1 and 2 display four choropleth maps of Italy summarizing key dimensions of educational performance and inequality across local labor market areas (LMAs). Graph A shows average INVALSI scores in mathematics at the LMA level, while Graph B presents the total variance in student performance. Graphs C and D focus on the inequality of educational opportunity (IEOp), respectively displaying its absolute magnitude (Graph C) and its relative contribution to overall variance (Graph D). The latter measure captures the

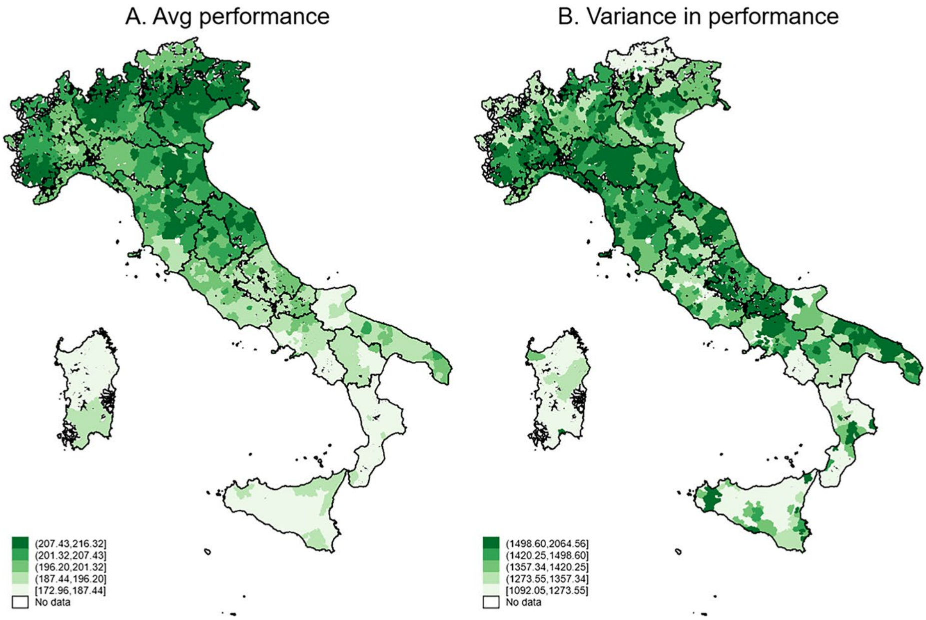

Fig. 1 Average performance (graph **A**) and total variance (graph **B**) in mathematics achievement scores

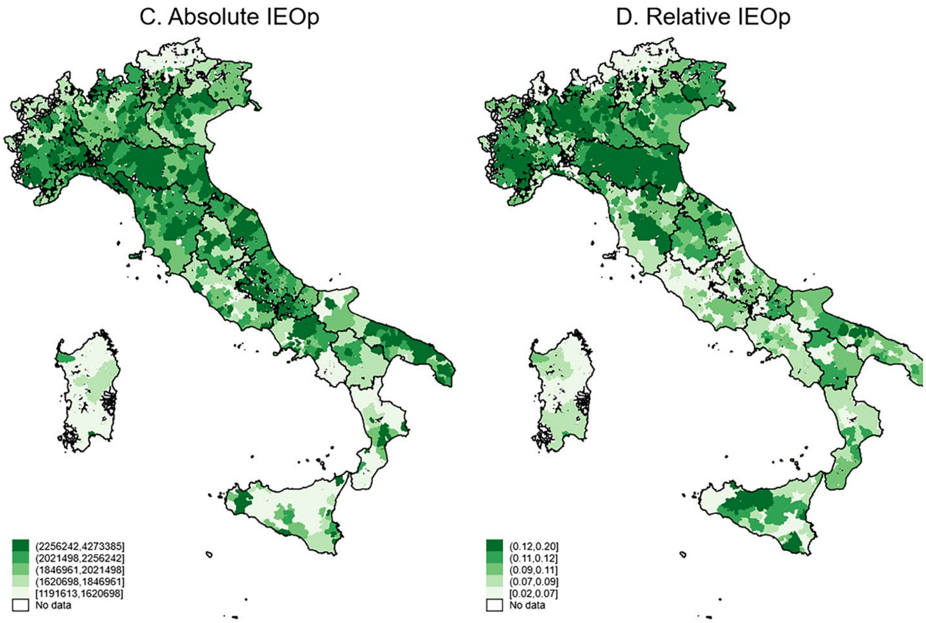

Fig. 2 Absolute (graph **C**) and relative IEOp (graph **D**) in mathematics achievement scores

proportion of performance variation predicted out-of-sample by individual circumstances, accounting for both their main effects and interactions.

To complement the spatial visualization, we conducted descriptive analyses and one-way ANOVA to decompose the variance of these four indicators across and within Italian geographical macro-areas.[4]

The average mathematics score across LMAs is approximately 198, with values ranging from 172 (Acerra, Campania) to 216 (Rest of the Province of Udine). The difference between the 90th and 10th percentiles (P90–P10) amounts to 25 points. A clear geographic gradient emerges: the average score is 204 in the North, 200 in the Center, and 188 in the South and Islands. One-way ANOVA results indicate that the five macro-areas explain about 64% of the variance in average performance across LMAs, with the remaining 36% attributable to variation within macro-areas. Notably, some regions exhibit relatively homogeneous performance across LMAs—especially in the South (e.g., Calabria) and Islands (e.g., Sicily, Sardinia)—while several Northern regions (Lombardy, Piedmont, Veneto, Tuscany) display substantial internal heterogeneity.

The map of total variance in performance (Graph B) captures within-LMA variability in student outcomes. We observe considerable heterogeneity: the LMA with the lowest variance has nearly half the variability of the one with the highest (1091 vs. 2069). In this case, within macro-area variation accounts for a larger share (92%) of the overall variance, which is also visually reflected in the map's more scattered color patterns within regions, with the notable exception of the Sardinian LMAs, which appear more internally consistent.

Turning to the measurement of IEOp, we present two indicators: the absolute amount of variance in performance explained by individual circumstances (Graph C), and the relative share of total variance attributable to these circumstances (Graph D). The correlation between the two measures across LMAs is relatively low ($r \approx 0.15$), underscoring the importance of considering both perspectives.

The absolute IEOp values show large territorial disparities. The LMA with the highest absolute IEOp (Avellino, Campania) exhibits more than three times the value of the lowest one (Rest of the Province of Palermo, excluding Palermo municipality). Interestingly, absolute IEOp tends to be somewhat lower in the South compared to the North and Center. When considering the relative IEOp, the North registers a slightly higher average share (12%) than the Center and South (both around 8%). However, the LMA-level variation is substantial, ranging from 1.6% (Formia, Lazio) to 20.4% (Novara, Piedmont). Looking at Graph D, we observe that some regions are internally homogeneous (e.g., Emilia-Romagna in the North, Sardinia among the Islands), while others—such as Lombardy (6%–19%) and Puglia (2%–20%)—display wide internal disparities in how strongly individual background predicts mathematics performance.

Overall, both absolute and relative IEOp display strong intra-regional heterogeneity. One-way ANOVA shows that 93% of the variation in absolute IEOp and 71% of the variation in relative IEOp occur within macro-areas. These findings underscore the importance of adopting a fine-grained geographical lens when analyzing educational inequality, as aggregate statistics at the macro-area level risk obscuring substantial local disparities.

[4] We have performed a similar analysis using degree of geographical marginality, relying on the proportion of individuals living in peripheral and ultra-peripheral areas using the official SNAI definition. The analysis reveals the IEOp variation across LMAs is not explained by this dimension of territorial marginality.

5.2 Is There a trade-off Between Levels of Educational Achievement and Inequality?

We now examine whether there is a systematic relationship between average achievement levels and inequality indicators at the LMA level. Specifically, we ask whether a trade-off exists between effectiveness (high average scores) and equity (low inequality), or whether it is possible for LMAs to combine strong performance with limited influence of ascriptive characteristics on achievement.

Figure 3 presents descriptive evidence addressing this question. Graph A plots average mathematics scores (x-axis) against total performance variance (y-axis), while Graph B relates average performance to absolute IEOp. We find positive correlations between average performance and both total variance ($r=0.36$) and absolute IEOp ($r=0.34$). The correlation with relative IEOp is even larger, at 0.49 (not shown). These findings suggest the presence of a trade-off: LMAs with higher average scores also tend to have greater performance dispersion and stronger associations between achievement and individual circumstances.

However, given Italy's pronounced regional disparities in socioeconomic development, school infrastructure, and cultural traditions, we investigate whether this trade-off is consistent across macro-areas. Graphs C and D in Fig. 3 present binned scatterplots—where observations are grouped into bins by average score to reduce noise—separately for the North, Center, and South. These plots show the fitted regression lines of the relationship between average performance and (C) variance in achievement, and (D) absolute IEOp, within each macro-area.

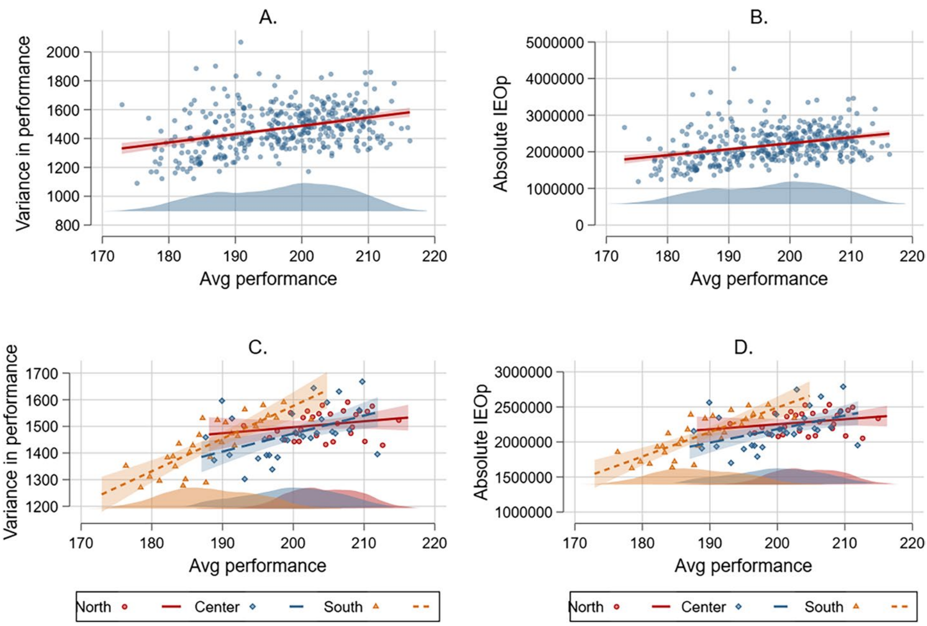

Fig. 3 Scatterplots with the relationship between average performance and total variance (graph **A**) and absolute IEOp (graph **B**); and relationship between average performance and total variance across macro-area (graph **C**) and absolute IEOp across macro-area (graph **D**)

The results reveal important regional differences. In the South, the trade-off is pronounced, with correlations between average performance and both variance ($r=0.44$) and absolute IEOp ($r=0.42$) being strongly positive. Moreover, the densities plotted at the bottom of the figure indicates that southern regions tend to have lower average score. A similar, though weaker, pattern is observed in the Center (r=around 0.35 for both correlations). However, the pattern in the North is much less evident, since the correlation between average performance and both total variance and absolute IEOp is smaller ($r=0.10$). This suggests that in areas with better infrastructure and more efficient educational system, higher inequality of opportunity is not a necessary condition to obtain higher average performance. In other words, the trade-off between effectiveness and equity is not a universal phenomenon and may be moderated by broader contextual factors such as school quality, governance, and social capital.

5.3 The Structure of IEOp Across Labor Market Areas

In this section, we focus on the role of individual ascriptive characteristics in predicting students' performance in mathematics at the end of lower secondary education (Grade 8), when students are approximately 14 years old. As outlined in the methods section, we rely on two complementary approaches to investigate this question: variable importance and partial dependency analyses.

First, we examine the relative importance of each predictor in explaining standardized test scores. This measure captures the average contribution of each factor to the overall prediction, taking into account both the magnitude of score differences across groups and the relative size of each group in the population. Second, we use partial dependence plots, which show the average association between a single variable and student performance, holding all other variables constant. These plots are conceptually similar to average partial effects at the mean in traditional regression models, and they help interpret the direction and magnitude of performance gaps associated with each variable.

Figure 4 shows the most predictive circumstances in each LMA. Mother's education appears as the most predictive variable in most parts of the country. Meanwhile, father's education emerges as the most relevant predictor in several internal areas. Similarly, mother's occupation is the most predictive variable in some areas in the North and, to a lesser extent, in central Italy and Sicily.

A distinct pattern emerges for migration, which is the most important variable in the internal areas of the North (but not in the eastern or western parts of the North) and in a few cases in the center-north. This is consistent with the fact that a higher prevalence of migrants (higher variability of the predictor) inevitably tend to be a better predictor of a relatively uniform predictor (such as in the case of a LMA where the vast majority are natives. A notable result concerns father's occupation, which turns out to be the most predictive variable only in the extreme North, specifically in the South Tyrol and Trentino regions.

5.4 Group Gaps and Predictive Importance

In Fig. 5 we turn to PDP and we compare predicted performance across selected groups. For gender, we contrast boys and girls; for parental education, we compare students whose parents hold a tertiary degree with those whose parents have only completed lower secondary

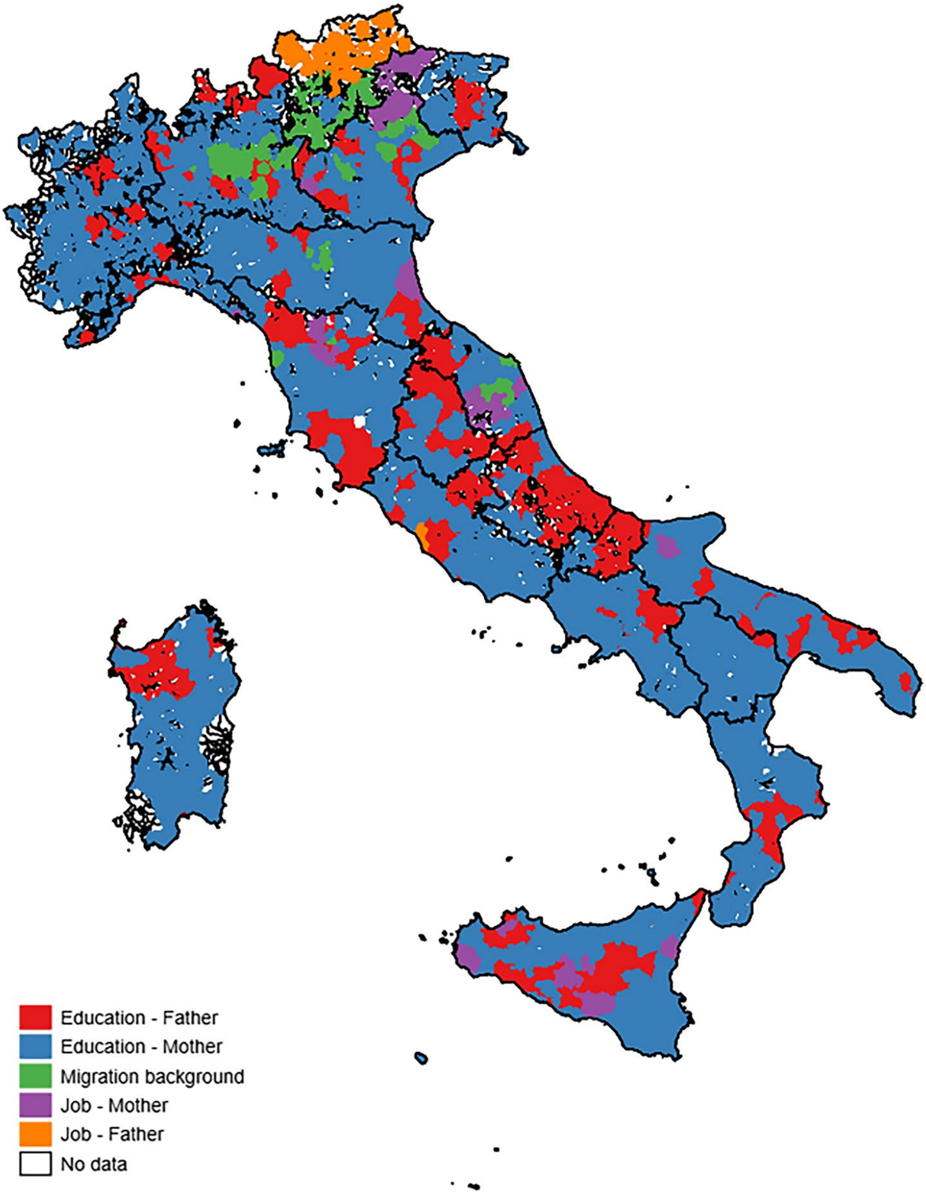

Fig. 4 Most important individual circumstance across Labour Market Areas (LMAs). Note that gender is not the most important circumstance in any area

education. In terms of occupational background, we contrast students from middle or upper-class (white-collar) backgrounds with those from working-class families. Finally, to capture the potential heterogeneity among students with a migration background, we examine both first- and second-generation immigrants in comparison to native students.

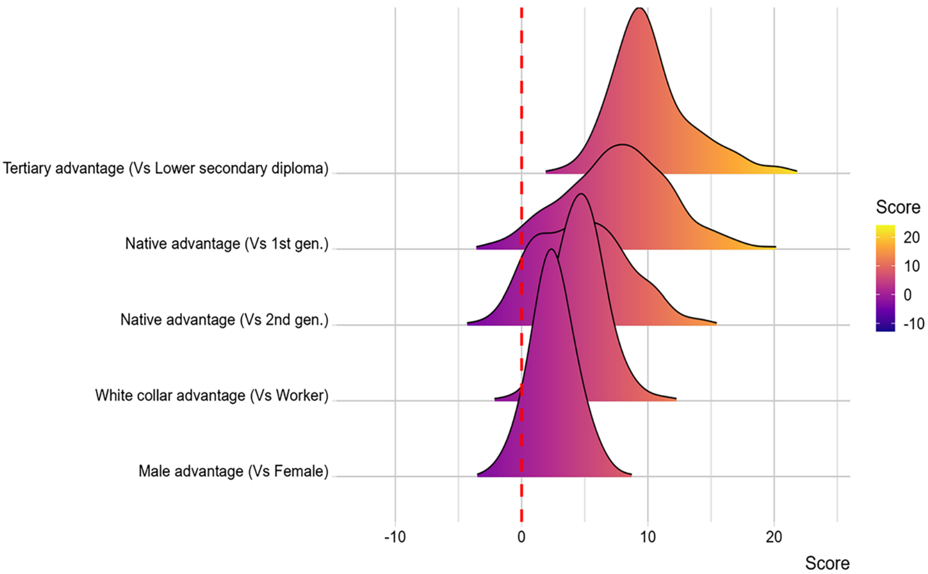

Fig. 5 Distributions of partial dependencies of selected individual circumstances across Labour Market Areas (LMAs)

The observed score gaps are consistent with established findings in the literature. Boys slightly outperform girls in mathematics, natives perform better than students with a migration background, and students with more advantaged family backgrounds—both in terms of education and occupation—tend to score higher than their peers from less advantaged backgrounds. On average across LMAs (unweighted), the largest performance gap is associated with parental education: students with tertiary-educated parents outperform those with lower-secondary-educated parents by approximately 10 points. This is followed by the gap between natives and first-generation immigrants (7.7 points). The gaps between natives and second-generation students and between middle-class and working-class students are smaller and roughly comparable (around 4.7–4.8 points). The gender gap is the smallest in magnitude but notably consistent.

5.5 Geographical Patterns and Regional Heterogeneity

Figure 6 displays choropleth maps of Italy illustrating the geographic distribution of these gaps. The gender gap shows a distinct spatial pattern: it tends to be larger in the North and Center than in the South, where boys' advantage in mathematics is more modest. One-way ANOVA indicates that 77% of the variation in the gender gap occurs within macro-areas, highlighting significant territorial heterogeneity beyond the macro-areas divides. Interestingly, in 27 LMAs, the gender gap slightly favors girls. In most of these cases, this reflects a lower average performance of boys (often around 180 points), rather than exceptionally high performance among girls.

A similar North–South divide emerges for performance gaps between natives and students with a migration background. In the South, the performance disadvantage for immigrant-origin students is approximately half that observed in the North. This is largely due

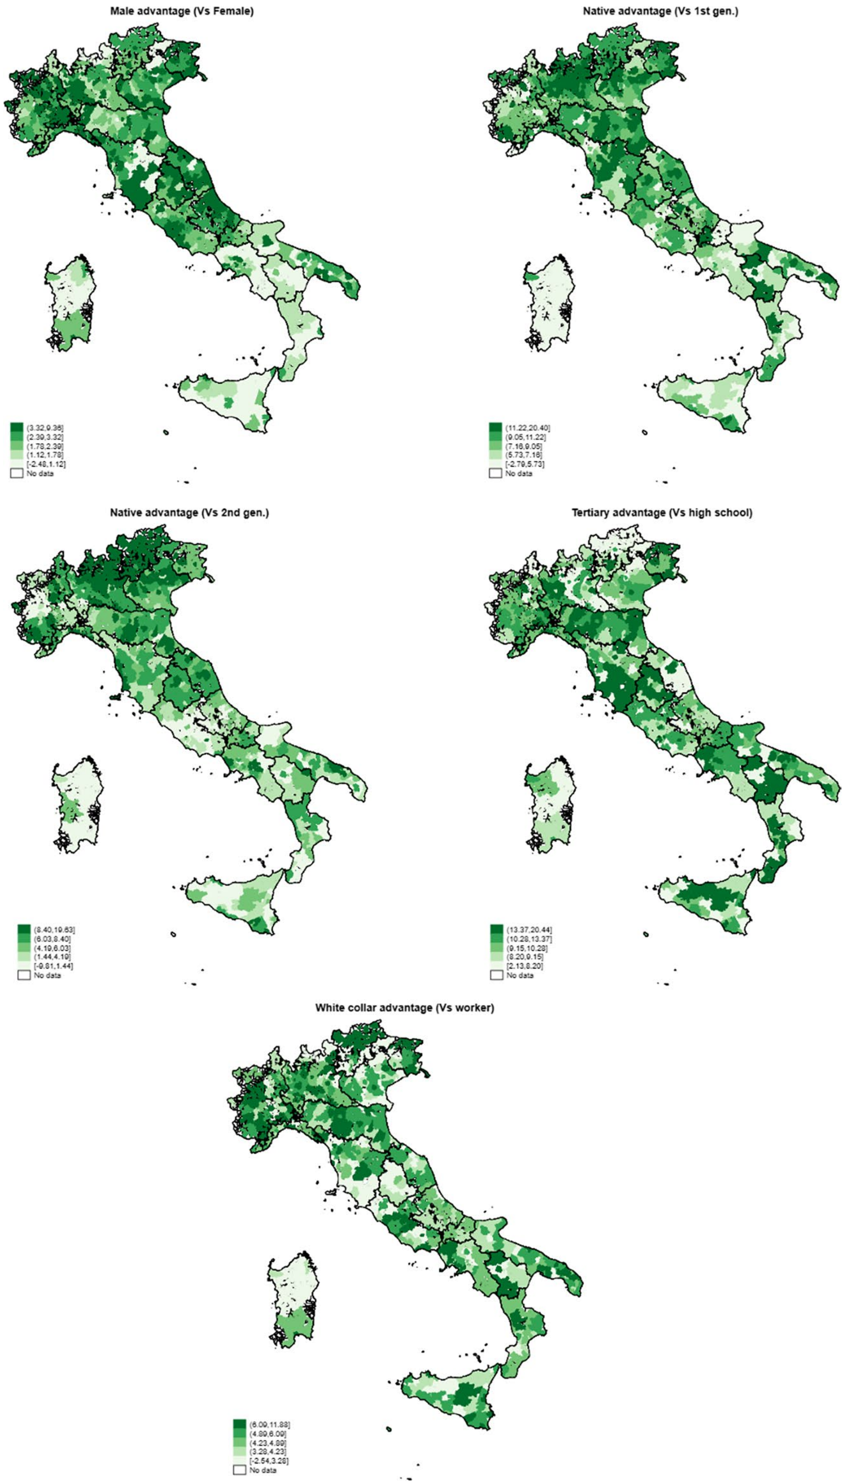

Fig. 6 Choropleth maps of selected individual circumstances across Labour Market Areas (LMAs)

to higher average performance among native students in the North, while the scores of immigrant-origin students are more similar across the country. The share of variance in these gaps explained by macro-areas ranges from 27% for second-generation immigrants to 30% for first-generation immigrants, suggesting moderate regional structuring of these inequalities and possible important drivers at the local level.

When turning to socioeconomic background, both in terms of parental education and occupational class, we observe a different pattern. The average gaps across macro-areas are more similar, and no clear geographical gradient emerges. Nevertheless, territorial heterogeneity is substantial. According to one-way ANOVA, only 2–3% of the variance in these gaps is between macro-areas, while 97–98% occurs within them. When using the twenty regions as group factor, the within-region variation is estimated to be still very high between 85 and 92%. This is clearly reflected in the maps: in most of Italy's 20 regions, the spatial distribution of the score gaps appears scattered and heterogeneous, with significant variation even among neighboring LMAs.

Large and diverse regions—such as Lombardy and Emilia-Romagna in the North, or Puglia and Sicily in the South—exemplify such internal variability. These findings highlight the limitations of prior research that relied on highly aggregated geographical data (typically at the macro-area or regional level), often based on PISA. By leveraging fine-grained local data, we uncover a more nuanced geography of educational inequality in Italy, which had largely been masked in previous studies (Fig. 6).

All the evidence described above for mathematics has also been estimated for language test scores. Without discussing all the results—which are reported in more detail in the Online Appendix—we briefly highlight a few key differences that emerge. First of all predictability is substatially higher for language: R^2_{OOS}=0.135 (R^2_{OOS}= 0.0943 for mathematics). A second substantial difference lies in the relative importance of different ascriptive characteristics. Gender accounts for over 10% of total predictability in language outcomes, compared to less than 3% in mathematics. Similarly, migration background has greater predictive power for language than for mathematics (17% vs. 10%), with some extreme cases—such as in the metropolitan area western of Florence and in some LMAs between the east of Lombardy and west of Veneto—where it explains more than 50% of the predicted variability. Differently, parental education is less predictive in this case. As shown in Fig. 7, which reports the correlation of mathematics and language IEOp in both absolute (top) and relative (bottom) terms, there is a substantial correlation between the two outcomes. In LMAs with high predicted variance in mathematics, the predicted variance in language is also high in both absolute and relative terms. Interestingly, areas performing particularly poorly in absolute terms are concentrated in the South, whereas areas with a high share of predicted variability for both subjects are more frequent in the North (due to a lower denominator). By contrast, the few LMAs performing poorly in terms of IEOp in education but not as poorly in terms of mathematics IEOp are found in the Center. Finally, while IEOp and average score are positively correlated (r=0.37) the same trade-off is not found for total variability and average score that are in fact negatively correlated (r=−0.22).

Therefore, while the language suggests two different patterns — IEOp positively correlated with average performance, but total variance negatively correlated with it — in mathematical terms both indicate a trade-off between achievement and equality. In high-performing areas, total variability in outcomes shrinks, but the reduced variance is more strongly structured by individual circumstances, raising IEOp.

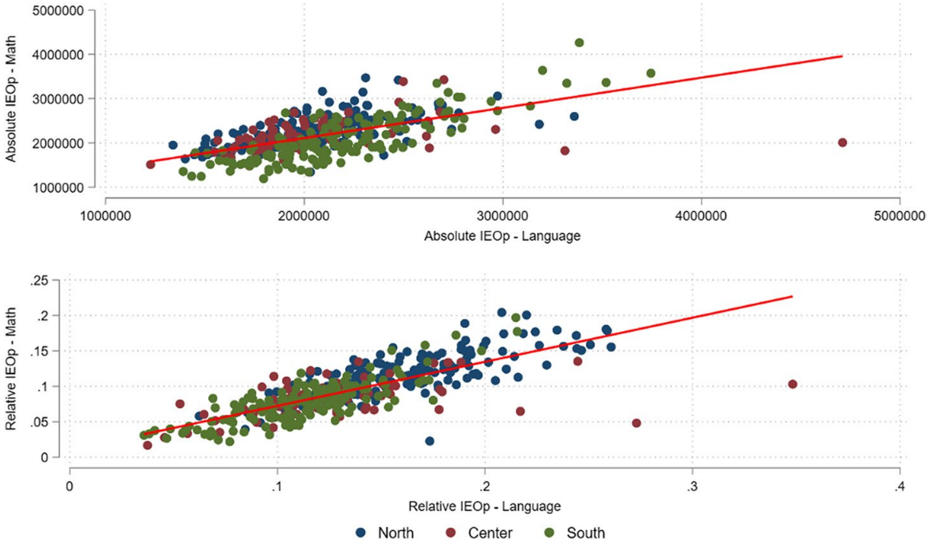

Fig. 7 Correlation of math and language measure of IEOp in absolute (top) and relative (bottom) terms

6 Discussion and Conclusions

Educational systems that allow individual circumstances—such as parental background, migration status, and gender—to shape students' academic trajectories pose critical challenges to both equity and efficiency. While a considerable body of research has addressed the inequality of educational opportunity (IEOp) across countries, there remains limited understanding of its spatial variation within countries. This is particularly the case in Italy, a country characterized by longstanding regional disparities in economic development, institutional quality, and school performance. Our study addresses this gap by mapping the geography of IEOp across Italian micro-areas and investigating its structure and drivers using advanced machine learning methods and detailed standardized test data.

Our results lead to four main conclusions. First, we find substantial spatial heterogeneity in both average student performance and IEOp across local labor market areas (LMAs). While national and regional-level analyses have traditionally emphasized the North–South divide, our fine-grained analysis reveals that a significant share of the variation in educational inequalities occurs within macro-areas and even within regions. For example, 93% of the variation in absolute IEOp and 71% of the variation in relative IEOp are found within macro-areas, suggesting that regional averages can mask important local disparities.

Second, we document a moderate but positive correlation between average performance and inequality, suggesting a potential trade-off between equity and effectiveness. LMAs with higher average mathematics scores tend to have greater variability in performance and stronger associations between achievement and individual circumstances. However, this pattern is not universal. In Northern Italy, we find cases of LMAs that achieve both high average performance and low inequality, indicating that the trade-off is not structurally inevitable and may be mitigated by favorable contextual conditions.

Third, our analysis of predictive importance and partial dependence shows that parental education is the strongest predictor of students' performance, followed by migration background and parental occupation, while gender differences are relatively modest and slightly more homogeneous across macro-areas. However, the influence of these factors varies considerably across LMAs. Notably, while the gap between students from different socioeconomic backgrounds is fairly homogeneous, gaps linked to migration background are more geographically dispersed, sometimes even reversing in direction. An interesting pattern emerging from our analysis is that in several Southern LMAs, mother's occupation is the most predictive individual circumstance for mathematics achievement. One possible explanation relates to the markedly lower female labor force participation in Southern Italy (Istat, 2022). In contexts where fewer women are employed, maternal occupational status may become a sharper indicator of socio-economic differentiation, reflecting both household economic resources and broader social capital. This could amplify its predictive power relative to other parental characteristics in these areas.

Another remarkable finding is that the gender gap in mathematics achievement, favouring boys, is more pronounced in Northern and Central LMAs, while it is comparatively smaller in the South. One possible explanation is that the higher overall achievement levels in the North and Centre may allow underlying gender differences to emerge more clearly, whereas in lower-performing contexts (often in the South) general performance constraints affect both genders similarly, compressing the gap. This is consistent with findings from cross-national research showing that gender gaps in mathematics tend to be larger in high-performing education systems (Guiso et al., 2008; Fryer & Levitt, 2010). A second, complementary interpretation points to subject-specific stereotyping and gendered expectations in more academically competitive environments (OECD, 2015). In the North and Centre, where schools are more selective and competitive, cultural norms and tracking patterns may reinforce boys' advantage in mathematics through differential encouragement, self-concept, and participation in STEM-related activities (Breda et al., 2020).

Fourth, our results emphasize the multidimensional nature of IEOp. Rather than being driven by a single factor, inequality arises from the complex interplay of multiple individual characteristics whose effects are context-dependent. Our methodological approach, which captures interaction effects through aggregation of multiple regression trees into random forest, allows us to move beyond additive models and uncover nuanced local dynamics.

Our study confirms several findings from prior research, such as the importance of parental education and migration background in shaping academic achievement (e.g., Pensiero et al., 2019; Borgna & Contini, 2014). We also corroborate the North–South divide in student performance (Bratti et al., 2007; Agasisti & Cordero-Ferrera, 2013), and extend it by showing how IEOp follows more complex spatial patterns.

Compared to earlier studies that rely on parametric regression models or highly aggregated geographical units, our work offers a more detailed and flexible account of inequality. By leveraging machine learning techniques on nearly complete population data, we capture more accurately the predictive power of social origin and its variation across small areas. Moreover, our analysis aligns with the growing literature on "geographies of opportunity" and intra-national disparities in mobility (Chetty et al., 2014; Acciari et al., 2022), bringing these insights into the domain of educational performance. Compared to the findings of Salganik et al. (2020) in the Fragile Families Challenge—where the best-performing algorithm predicted Grade Point Average (GPA) at age 15 with an R^2 of approximately 0.19

using thousands of predictors collected over a decade—our models achieve only slightly lower levels of predictability in many LMAs despite relying on a far more limited set of individual-level variables and observing students at a single point in time. This contrast underscores both the robust association between a few key ascriptive characteristics and educational outcomes in the Italian context, and the potential gains in explanatory power that could be achieved by incorporating richer contextual or longitudinal data.

These findings have several implications for policy. First, national education policies should recognize the high degree of territorial fragmentation in educational inequality. While macro-level initiatives are essential, they need to be complemented by place-based interventions that account for local circumstances. Policies tailored to specific LMAs could better address the unique constellations of disadvantage that affect student outcomes.

Second, the existence of high-performing, low-inequality LMAs—particularly in the North—suggests that equity and effectiveness can be jointly pursued under favorable institutional and social conditions. Investigating the characteristics of these areas – also with qualitative methods such as in-depth interviews and ethnographic work - could inform best practices for reducing inequality without compromising achievement. Third, the differential impact of migration background across areas points to the importance of local integration policies, school climate, and cultural sensitivity. Educational systems must move beyond one-size-fits-all approaches and acknowledge the diversity of student experiences and needs. In this direction, initiatives such as those promoted by *Con i Bambini*—the implementing agency of Italy's Fund to Combat Educational Poverty—represent promising examples of locally grounded interventions. These programs finance educational projects that are co-designed and implemented by schools, municipalities, third-sector organizations, and local communities, with the goal of addressing the specific educational challenges and social vulnerabilities present in each area. By fostering inter-institutional collaboration and adapting interventions to the territorial context, they embody a bottom-up approach that complements national education policy and enhances its local effectiveness.

While our study offers several innovations, it also has limitations. First, although our models consider a broad set of individual circumstances, we are constrained by the availability of variables in the INVALSI dataset. Key dimensions such as household income, religion, disability are not available, potentially limiting our understanding of the stratification of educational opportunities and providing lower-bounded estimates of IEOp. Second, our focus is on performance at a specific point in time (5th Grade mathematics), which provides only a snapshot of educational inequality. Future research could extend this to longitudinal data to explore how IEOp evolves over time and affects later transitions in the educational system.

Several avenues for future work emerge from our findings. First, similar methodologies could be applied to other educational stages to assess whether patterns of IEOp persist or shift along the school career. While our study period predates the COVID-19 pandemic, which may have altered the spatial distribution of educational inequalities, our methodological framework could be readily applied to more recent post-pandemic data to investigate such potential changes in more recent years.

Second, researchers could integrate school-level and neighborhood-level data to better understand the contextual factors driving spatial variation in IEOp. Third, comparative studies applying this framework across countries with similarly decentralized or regionally heterogeneous education systems could help uncover broader institutional determinants of territorial inequality.

In sum, this study highlights that IEOp is not only a national or regional issue but also a profoundly local one. By applying a machine learning approach to large-scale administrative data, we reveal how ascriptive characteristics shape student achievement in context-specific ways, with striking variations across Italian micro-areas. These findings call for a more nuanced, data-informed approach to educational policy—one that combines national goals with local solutions and places the reduction of inequality at the heart of educational reform.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11205-025-03788-3.

Funding This research was supported by Fondazione Cariplo under the project EDUFAIR – Mapping and Tackling Unfair Educational Inequality (Project Code 2024-1575), led by Moris Triventi, and funded within the INEQUALITIES RESEARCH 2024.

Data Availability Data availability is subject to official requests to INVALSI, as outlined in the 'Regolamento per l'accesso ai dati, ai metadati e alle relative banche dati.' Researchers must first access the 'Area Dati' section on the INVALSI website to identify the relevant dataset. A formal request, including the necessary forms must then be submitted to the Servizio Statistico at uff.statistico@invalsi.it Upon validation, the data, devoid of identifiable information, will be provided. Additionally, researchers must later submit bibliographic details of any publications using INVALSI data to biblioteca@invalsi.it.

Declarations

Conflict of Interest The authors, Paolo Brunori, Emanuele Fedeli, and Moris Triventi, declare no conflicts of interest.

References

Acciari, P., Polo, A., & Violante, G. L. (2022). And yet it moves: Intergenerational mobility in Italy. *American Economic Journal: Applied Economics, 14*(3), 118–163.

Alba, R., Sloan, J., & Sperling, J. (2011). The integration imperative: The children of low-status immigrants in the schools of wealthy societies. *Annual Review of Sociology*, *37*(1), 395–415.

Agasisti, T., & Cordero-Ferrera, J. M. (2013). Educational disparities across regions: A multilevel analysis for Italy and Spain. *Journal of Policy Modeling, 35*(6), 1079–1102.

Berger, T., & Engzell, P. (2019). American geography of opportunity reveals European origins. *Proceedings of the National Academy of Sciences, 116*(13), 6045–6050.

Biagi, F., Palmisano, F., & Peragine, V. (2022). Inequality of opportunity in tertiary education in Europe. *Research in Higher Education*, *63*, 514–565.

Blossfeld, H. P., Buchholz, S., Skopek, J., & Triventi, M. (Eds.). (2016). *Models of secondary education and social inequality: An international comparison*. Edward Elgar Publishing.

Borgna, C., & Contini, D. (2014). Migrant achievement penalties in Western Europe: Do educational systems matter? *European Sociological Review, 30*(5), 670–683.

Bratti, M., Checchi, D., & Filippin, A. (2007). Geographical Differences in Italian Students' mathematical Competencies: Evidence From Pisa 2003. *Giornale degli economisti e Annali di economia*, 299–333.

Breda, T., Jouini, E., & Napp, C. (2020). Societal inequalities amplify gender gaps in math. *Science*, *368*(6493), 1225–1228. https://doi.org/10.1126/science.aaz9892

Breen, R., & Goldthorpe, J. H. (2001). Class, mobility, and merit: The experience of two British birth cohorts. *European Sociological Review*, *17*(2), 81–101.

Breen, R., & Jonsson, J. O. (2005). Inequality of opportunity in comparative perspective: Recent research on educational attainment and social mobility. *Annual Review of Sociology, 31*, 223–243.

Breen, R., & Müller, W. (2020). Social mobility in Europe: New approaches and findings. *Annual Review of Sociology, 46*, 1–21.

Brunello, G., & Checchi, D. (2007). Does school tracking affect equality of opportunity? New international evidence. *Economic Policy*, *22*(52), 781–861.

Brunori, P., Hufe, P., & Mahler, D. (2023). The roots of inequality: Estimating inequality of opportunity from regression trees and forests. *The Scandinavian Journal of Economics*, *125*(4), 900–932.

Buchmann, C., DiPrete, T. A., & McDaniel, A. (2008). Gender inequalities in education. *Annual Review of Sociology, 34*(1), 319–337.

Chetty, R., Hendren, N., Kline, P., Saez, E., & Turner, N. (2014). Is the united States still a land of opportunity? Recent trends in intergenerational mobility. *American Economic Review*, *104*(5), 141–147.

Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives*, *27*(3), 79–102.

Daniele, V. (2021). Socioeconomic inequality and regional disparities in educational achievement: The role of relative poverty. *Intelligence, 84*, Article 101515.

Erikson, R., & Jonsson, J. O. (Eds.). (1996). *Can education be equalized? The Swedish case in comparative perspective*. Westview.

Felice, E. (2019). Regional disparities in Italy over the long run: The interplay of human and social capital. *Journal of Economic Surveys*, *33*(3), 729–751.

Ferraro, S., & Põder, K. (2018). School-level policies and the efficiency and equity trade-off in education. *Journal of Policy Modeling*, *40*(5), 1022–1037.

Ferrer-Esteban, G. (2016). Trade-off between effectiveness and equity? An analysis of social sorting between classrooms and between schools. *Comparative Education Review*, *60*(1), 151–183.

Ferreira, F. H. G., & Gignoux, J. (2014). The measurement of educational inequality: Achievement and opportunity. *The World Bank Economic Review*, *28*(2), 210–246.

Fryer, R. G., & Levitt, S. D. (2010). An empirical analysis of the gender gap in mathematics. *American Economic Journal: Applied Economics, 2*(2), 210–240. https://doi.org/10.1257/app.2.2.210

Gershberg, A. I., & Schuermann, T. (2001). The efficiency–equity trade-off of schooling outcomes: Public education expenditures and welfare in Mexico. *Economics of Education Review, 20*(1), 27–40.

Guiso, L., Monte, F., Sapienza, P., & Zingales, L. (2008). Culture, gender, and math. *Science*, *320*(5880), 1164–1165. https://doi.org/10.1126/science.1154094

Hanushek, E. A., & Woessmann, L. (2008). The role of cognitive skills in economic development. *Journal of Economic Literature, 46*(3), 607–668.

Hanushek, E. A., & Woessmann, L. (2011). The economics of international differences in educational achievement. *Handbook of the Economics of Education, 3*, 89–200.

Heckman, J. J., Stixrud, J., & Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics, 24*(3), 411–482.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning* (1st ed.). Springer.

Istat (2014). I sistemi locali del lavoro 2011. Statistiche report. Roma. 17 dicembre 2014.

Istat (2022). *Rapporto annuale Sulla situazione Del Paese*. Istituto Nazionale di Statistica (ISTAT).

Liu, J., Peng, P., Zhao, B., & Luo, L. (2022). Socioeconomic status and academic achievement in primary and secondary education: A meta-analytic review. *Educational Psychology Review*, *34*(4), 2867–2896.

Marrero, G. A., Palomino, J. C., & Sicilia, G. (2024). Inequality of opportunity in educational achievement in Western Europe: Contributors and channels. *Journal of Economic Inequality, 22*(2), 383–410.

Murnane, R. J., Willett, J. B., & Levy, F. (1995). The growing importance of cognitive skills in wage determination. *Review of Economics and Statistics, 77*(2), 251–266.

OECD (2010). The High Cost of Low Educational Performance: The Long-Run Economic Impact of Improving PISA Outcomes. OECD Publishing. https://doi.org/10.1787/9789264077485-en

OECD (2015). The ABC of gender equality in education: Aptitude, behaviour, confidence. OECD Publishing. https://doi.org/10.1787/9789264229945-en

Parker, P. D., Marsh, H. W., Jerrim, J. P., Guo, J., & Dicke, T. (2018). Inequity and excellence in academic performance: Evidence from 27 countries. *American Educational Research Journal, 55*(4), 836–858.

Rawls, J. (1971). *A theory of justice*. Harvard University Press.

Roemer, J. E. (1998). *Equality of opportunity*. Harvard University Press.
Schnepf, S. (2007). Immigrants' educational disadvantage: An examination across ten countries and three surveys. *Journal of Population Economics, 20*(3), 527–545.
Shavit, Y., & Blossfeld, H. P. (Eds.). (1993). *Persistent inequality: Changing educational attainment in thirteen countries*. Westview.
Swift, A. (2005). *Justice, luck, and the family*. Princeton University Press.
Triventi, M., & Fedeli, E. (2025). The geography of social background inequalities in educational opportunities: A within-country study. *European Societies*, 1–32. https://doi.org/10.1162/euso.a.35
Tsai, S. L., Smith, M. L., & Hauser, R. M. (2018). Gender gaps in student academic achievement and inequality. *Research in the sociology of education* (pp. 181–218). Emerald Publishing Limited.
Van de Werfhorst, H. G., & Mijs, J. J. (2010). Achievement inequality and the institutional structure of educational systems: A comparative perspective. Annual review of sociology, 36(1), 407–428.