

Decomposing Learning Inequalities in East Africa: How Much Does Sorting Matter?

Paul Anand, Jere R. Behrman, Hai-Anh H. Dang, and Sam Jones

Abstract

Inequalities in learning opportunities arise from both household- and school-related factors. Although these factors are unlikely to be independent, few studies have considered the extent to which sorting between schools and households might aggravate educational inequalities. To fill this gap, this article presents a novel variance decomposition, which is then applied to data from over one million children from East Africa. Results indicate that sorting accounts for around 8 percent of the test-score variance, similar in magnitude to the contribution of differences in school quality. Empirical simulations of steady-state educational inequalities reveal that policies to mitigate sorting could substantially reduce educational inequalities over the long run.

JEL classification: I21, I24, O15

Keywords: Africa, education, inequality of opportunity, sorting, variance decomposition

1. Introduction

Children in low- and middle-income countries face a wide range of educational and developmental challenges (Pritchett 2013; Black et al. 2017; Lu et al. 2020). Even so, average national conditions often mask considerable within-country variation, evidenced by large gaps in educational achievement along dimensions typically associated with household income inequalities (e.g., Watkins 2012; Behrman 2020). Inequalities associated with educational achievement feed into subsequent labor market inequalities, as well as longer-term differences in social mobility. And while the presence of such educational inequalities has motivated increasing calls for investments in public services, espoused by the UN's Sustainable

Paul Anand is Research Associate at the Department of Social Policy and Intervention, Oxford University, Oxford, UK and Professor at the Open University, Milton Keynes, UK; his email address is paul.anand@open.ac.uk. Jere R. Behrman is the William R. Kenan, Jr. Professor of Economics at the University of Pennsylvania, Philadelphia, USA; his email address is jbehrman@sas.upenn.edu. Hai-Anh H. Dang is a senior economist in the Data Production and Methods Unit, Development Data Group, World Bank, Washington DC, USA; he is also affiliated with the International School, Vietnam National University, Hanoi, the Center for Analysis and Forecasting, Vietnam's Academy of Social Sciences, GLO, Indiana University, USA, and Institut zur Zukunft der Arbeit (IZA); his email address is hdang@worldbank.org. Sam Jones (corresponding author) is a Research Fellow with UNU-WIDER in Maputo, Mozambique; his email address is jones@wider.unu.edu. The authors thank the editor, referees, Jed Friedman, Ha Nguyen, Ron Smith, and participants at the Centre for the Study of African Economies 2019 Conference for helpful comments on earlier versions. We are grateful to the UK Department of International Development for funding assistance through its Strategic Research Program (SRP) program. The study was also prepared within the UNU-WIDER project on "Social mobility in the Global South – concepts, measures, and determinants." Previous versions were published as World Bank Policy Research Paper #8622 and UNU-WIDER Working Paper 2019/110. A supplementary online appendix is available with this article at *The World Bank Economic Review* website.

Development Goals to “leave no child behind,” the drivers of educational inequalities are still only partly understood, particularly in lower-income countries.

This article focuses on educational inequality arising from one potentially important yet understudied source: sorting between schools and households (pupils). While a vast literature attests that both household and school factors are key determinants of educational achievement (e.g., [Coleman 1966](#); [Björklund and Salvanes 2011](#); [Hanushek and Rivkin 2012](#)), these factors are unlikely to be independent. Various processes, such as residential segregation ([Hanushek and Yilmaz 2007](#); [Benabou 1993](#)) or endogenous selection of teachers into schools ([Pop-Eleches and Urquiola 2013](#); [Hanushek et al. 2005](#)), have been connected to differences in the quality of public goods. Sorting between households and schools—defined generically as the covariance in quality-adjusted inputs—may thus accentuate inequalities in educational opportunities. Where material, sorting may not only confound precise estimation of the unique contribution of individual factors of educational outcome, such as that of teachers or households, but it would justify enhanced attention to how *existing* inputs are distributed.

To date, the vast bulk of research on educational sorting pertains to richer countries. For instance, various studies have shown that children from disadvantaged families disproportionately attend lower-quality schools ([Hsieh and Urquiola 2006](#); [Nechyba 2006](#)). While this may reflect a range of factors, including peer effects, the evidence suggests teachers serving schools in socioeconomically deprived areas tend to be less experienced or less qualified ([Jackson 2009](#); [Clotfelter, Ladd, and Vigdor 2011](#); [Sass et al. 2012](#); [OECD 2018](#)). For the United States, [Jackson \(2013\)](#) shows that the quality of the match between teachers and children is economically important, representing around two-thirds the explanatory power of teacher effects with respect to student learning outcomes. Even so, the overall significance of these processes remains unsettled. [Kremer \(1997\)](#) argues that eliminating neighborhood segregation would decrease long-run educational inequality in the United States by less than 2 percent. In contrast, [Fernández and Rogerson \(2001\)](#) develop a model where enhanced sorting can have much larger effects on inequality (also [Fernandez 2003](#)), and data from Canada suggest that elimination of sorting either by home language or by parental education could reduce test-score variance by as much as 40 percent in some subjects, at least in locations where school segregation is substantial ([Friesen and Krauth 2007](#)).

In developing countries, indirect evidence indicates that educational sorting may often aggravate educational inequalities. A small literature supports the conjecture that children from disadvantaged backgrounds tend to receive lower-quality schooling inputs. For Guatemala, [McEwan and Trowbridge \(2007\)](#) decomposes the large achievement gap between indigenous and nonindigenous children, arguing this is predominately attributable to schooling attributes not family background. There are similar findings for other Latin American countries ([McEwan 2004](#)), but Vietnam stands out as providing a more equitable treatment of disadvantaged children ([Glewwe, Krutikova, and Rolleston 2017](#)).

Estimating the contribution of sorting to variation in educational outcomes presents a number of methodological challenges, especially where the factors of interest are incompletely observed. To address these, our first contribution—applicable in various domains—is to develop a new two-way fixed-effects variance decomposition procedure, which has some advantages over existing methods, especially in low-income country contexts where data are limited. Second, this procedure is applied to a large dataset on educational achievement, comprising over one million children in three East African countries. This analysis yields some of the first estimates of the impact of household–school sorting on educational inequalities, as well as the magnitude of the stand-alone variance contributions of household and school factors. Finally, we simulate the long-term gains for educational equality were the magnitude of sorting to be reduced.

A preliminary review of the East African data reveals a very substantial degree of spatial segregation in age-adjusted educational outcomes across communities: i.e., location matters for educational achievement. But even after controlling for this form of spatial clustering, the contribution of sorting of children across schools within communities is positive, accounting for around 8 percent of the total test-score

variance. As such, almost a fifth of the joint variation in test scores due to systematic supraindividual circumstances—i.e., due to schools, communities, and households—can be attributed to sorting. Further analysis of heterogeneity suggests that the same sorting contribution tends to be larger among families at both the top and bottom ends of the socioeconomic distribution and among those sending their children to private schools, as well as in specific locations.

Our simulations show that for the average district in the region, the steady-state level of educational inequality would fall by around 15 percent if sorting between households and schools were fully eliminated and up to 30 percent if sorting across communities were also cut to zero. Roughly speaking, this would be equivalent to cutting the magnitude of the intergenerational persistence of education by more than half. Moreover, in a significant number of districts, the reduction in inequality from eliminating all forms of sorting would be over 40 percent of the total variance. Together, our results suggest that to be effective, educational policies in lower-income countries must counter the consequences of household sorting.

2. Drivers of Educational Inequality

At a most general level, the drivers of educational inequality encompass any factor that produces differences in learning outcomes across children. However, following [Roemer \(1996\)](#) who advocates for a focus on equality of opportunity as the most universally supported conception of justice in advanced societies, we narrow our attention to inequality of educational opportunity (IEO), defined as those factors beyond the influence of individual children that affect learning outcomes.

Without attempting to cover the vast literature concerning the different sources of IEO (for discussion see [Reynolds and Teddlie 2000](#); [Gamoran and Long 2007](#); [Raffo et al. 2007](#); [Watkins 2012](#)), as a prelude to the remainder of the article, [table 1](#) provides a selective sample of existing studies that shed light on the contribution of supraindividual factors to inequalities in test scores, where the studies were chosen to cover the predominant range of methods and (to the extent possible) consider developing country contexts. As [Ferreira and Gignoux \(2014\)](#) note, the variance is generally the preferred metric of dispersion in such accounting procedures. This reflects two attractive properties: the variance is ordinaly invariant to standardization procedures often used to express test scores on a comparable scale, and the linear additive nature of the variance makes it straightforward to isolate the contributions of individual factors ([Shorrocks 1982](#)) and permits subgroup decomposability ([Chakravarty 2001](#)).

The studies in the table are classified according to four distinct methodological approaches. Panel (i) highlights studies that focus on a single main factor, such as households or schools. In most cases, these studies recognize that observed dimensions of the chosen factor (e.g., parental education or books per pupil) are partial, necessitating use of random (RE) or fixed-effects (FE) procedures that treat the specified factors as unobserved. Panel (ii) highlights two recent studies that take the fixed-effects approach further. Namely, based on a first-stage regression containing a single main unobserved fixed effect (classrooms), the authors are able to further partition each estimated effect into two or more subcomponents, focusing in particular on the “permanent” contribution of individual teachers. The studies in panel (iii) use mixed-effects estimators, which represent an alternative and popular means to estimate unobserved effects at multiple levels. But this necessitates making strict assumptions about the correlation structure of the unobserved factors, including their mutual independence.

The studies in panel (iv) relax this latter assumption, but do so by combining one unobserved factor (estimated via fixed effects) and proxying the other factor by observed covariates (OE). In so doing, they are able to estimate the covariance between the two factors (the sorting contribution; see below). For instance, in his analysis of rich test-score data from North Carolina, [Mansfield \(2015\)](#) finds that student background (which includes family and individual characteristics, as well as past performance)

Table 1. Selection of Literature on Drivers of IEO

Study	Context	Outcome	Factor(s)			
			Scope ^a	Spec. ^b	Var. share ^c	
(i) Mazumder (2011) Ferreira and Gignoux (2014) Pritchett and Viarengo (2015) Anger and Schnitzlein (2017)	United States	Reading comprehension	—	Family background	RE	0.46
	Multi-country	PISA scores	—	Given circumstances	OE	0.35
	Multi-country	PISA scores	—	Schools	FE	0.28
	Germany	Word fluency test	—	Family background	RE	0.60
(ii) Buhl-Wiggers et al. (2018) Bau and Das (2020)	Uganda	EGRA scores	(1)	Teachers	FE	0.06
	Pakistan	LEAPS scores	(2)	Classrooms	FE	0.04
			(1)	Teachers	FE	0.02
(iii) Nicoletti and Rabe (2013) Troncoso, Pampaka, and Olsen (2016) Dang and Glewwe (2018)	UK	Key Stage 2 tests	(2)	Classrooms	FE	0.08
	Chile	SIMCE math tests	(3)	Schools	FE	0.03
			(1)	Family background	RE	0.50
Portugal	PISA reading scores	(2)	Neighborhoods	RE	0.10	
		(1)	Classrooms	RE	0.09	
		(2)	Schools	RE	0.37	
(iv) Carneiro (2008) Mansfield (2015)	Vietnam	Std. Grade 5 test scores	(3)	Local authorities	RE	0.07
			(1)	Schools	RE	0.45
United States	Std. test scores	(2)	Provinces	RE	0.13	
		(1)	Family background	OE	0.10	
			(2)	Schools	FE	0.23
			(1)	Student background	OE	0.60
			(2)	School/teacher quality	FE	0.06

Source: Authors' elaboration.

Note: ^a'Scope' indicates the (different) factors (levels) included in the analysis, where numbers (1)-(3) denote the number of separate effects included. ^b'Spec.' indicates how these factors are specified in the model, where OE are observed effects, RE random effects, and FE fixed effects. ^c'Var. share' is the proportion of variance in the outcome attributable to the corresponding factor.

accounts for 60 percent of the variation in outcomes, while overall school quality (including teachers) accounts for about 6 percent and the covariance of these factors is also positive and similar in magnitude to the school effect at 4 percent of the variance.

Our classification underlines some of the persistent methodological challenges plaguing identification of the drivers of IEO. First, treating a factor as unobserved requires multiple observations for each element (index) of each factor. Data limitations thus explain why measures of the contribution of family background to IEO in developing countries have predominantly relied on observed proxies (Ferreira and Gignoux 2014). To our knowledge, in such contexts sibling correlations have only been estimated for income and educational attainment (e.g., Dahan and Gaviria 2001; Louw, Van Der Berg, and Yu 2007; Emran and Shilpi 2015), not learning outcomes. Second, while scholars have largely converged on value-added models to identify the causal effects of specific inputs within education production functions (see Todd and Wolpin 2003), these models cannot easily incorporate multiple high-dimensional unobserved factors. Under conventional OLS estimation methods, combining a lagged outcome with fixed effects introduces bias (Nickell 1981), which will only be negligible if a large number of observations (at least 20) are available to identify each unobserved effect. And in mixed-effects estimators, strict exogeneity of the (initial) unobserved effects must be assumed for value-added specifications to be consistent.

Third, from the perspective of estimating the contribution of sorting, existing studies must either rule this out by assumption or reduce one of the factors to its observed components, implying that only a partial view of this factor and its (co)variance contribution can be captured. For this reason, such models have rarely been used either to consider questions of sorting or to estimate the relative importance of different

components of IEO. Historically, technical challenges in accurate estimation of multiple unobserved fixed effects have limited the development of more flexible models. Early on, [Abowd, Creecy, and Kramarz \(2002\)](#) showed that the joint contribution of two-way fixed effects can be obtained fairly easily (for implementation in Stata see [Correia 2017](#)). The drawback is that where the individual effects themselves are of direct interest, their pairwise correlation tends to be biased downwards. [Andrews et al. \(2008\)](#) demonstrate this is driven by a quasi-mechanical relation, whereby if one factor (e.g., household effects) is overestimated then on average the other factor (e.g., schools) is underestimated. And, as [Jochmans and Weidner \(2019\)](#) prove, this bias is larger where there is limited mobility or overlap of units across different effects, the limiting case being where their bipartite graph is disconnected.

Development of two-way fixed-effects estimators has focused predominantly on bias-corrections in the context of the contribution of workers and firms to the variance of log-earnings. [Bonhomme et al. \(2020\)](#), for instance, show that the variance contribution due to the sorting of workers across firms increases from around zero to 30 percent when bias corrections are applied. Even so, not only do these methods remain work in progress but also they demand rich (longitudinal) data, with a very large connected set of factors. In our case, this would require observations of members of the same household across multiple different schools, including connections that span different geographic locations. This kind of data is rarely available even in high-income countries (for an exception, see [Thiemann 2018](#)), motivating an alternative approach to identify sorting, addressed in the next section.

Before doing so, a comment on the findings summarized in [table 1](#) is warranted. While it is inappropriate to compare effect sizes across studies, due to differences in tests, methods, and populations, it is nonetheless clear that variation in school (teacher) quality constitutes just one part of overall IEO. As also summarized by [Hanushek and Rivkin \(2012\)](#) for the United States, less than a quarter of the variance in test scores can be attributed to dispersion in teacher quality. In contrast, as stressed early on by the Coleman Report ([1966](#)), the home environment seems consistently of much greater importance, often accounting for more than half of the overall test-score variance (see also [Björklund and Salvanes 2011](#)). And this broad pattern appears to extend to lower-income countries, although previous studies have not covered both types of factors simultaneously.

3. Decomposition Framework

3.1. Preliminaries

To fill the methodological gap in the literature, we now propose a flexible decomposition procedure. This has the advantage of treating two separate effects as unobserved (here, households and schools): it does not impose any a priori structure on their correlation structure, and it also does not depend on the (biased) factor covariance identified from a two-way fixed-effects estimator. To begin, consider a simple (unconditional) factor model:

$$t_{ijk} = f(h_j, s_k) + e_{ijk}, \quad (1)$$

where t is a measure of educational achievement (e.g., test scores), and indexes i , j , and k refer to individual children, families, and schools respectively. The factor h_j represents the contribution of all factors shared by children in the same household to test scores, including common genetic endowments, and s_k represents the contribution of the given school to their learning. In both cases we assume these measures are comprehensive, in the sense that they reflect cumulative inputs until now. Finally, e_{ijk} represents the remaining individual or idiosyncratic variation, which we assume is orthogonal to the household effects and school effects: $E(e_{ijk} | s_k, h_j) = 0$.

To make this framework tractable for an empirical analysis of inequality, some structure must be placed on the IEO term $f(\cdot)$. With respect to the levels expression in [equation \(1\)](#), we adopt a simple additive

Table 2. Summary of Alternative Test-Score Data-Generating Processes

Model	Score level	Score variance	Description	
1	Restricted linear	$b_j + s_k + e_{ijk}$	$\sigma_b^2 + \sigma_s^2 + \sigma_e^2$	Independent households & schools
2	Household upper bound	$b_j + \gamma \bar{b}_{jk} + v_k + e_{ijk}$	$(1 + \gamma)^2 \sigma_b^2 + \sigma_v^2 + \sigma_e^2$	Household effects partly absorb school effects
3	School upper bound	$s_k + \theta \bar{s}_{kj} + \omega_j + e_{ijk}$	$(1 + \theta)^2 \sigma_s^2 + \sigma_\omega^2 + \sigma_e^2$	School effects partly absorb household effects
4	Unrestricted linear	$b_j + s_k + e_{ijk}$	$\sigma_b^2 + \sigma_s^2 + 2\sigma_{bs} + \sigma_e^2$	Correlated household & school factors

Source: Authors' elaboration.

Note: Variances in rows 3 and 4 assume $\sigma_b^2 \equiv \text{Var}(b_j) \approx \text{Var}(\bar{b}_{jk})$ and $\text{Var}(s_k) \approx \text{Var}(\bar{s}_{kj})$.

linear model. However, without further assumptions, this specification does not pin down the (co)variance structure. To make this point clear, [table 2](#) describes four alternative cases, where each row invokes specific assumptions about the level *and* variance of *t*. Row 1 invokes a zero-covariance assumption, corresponding to the most restrictive random-effects models described earlier: e.g., [table 1](#), panel (iii). However, this effectively rules out a range of sorting processes, including residential segregation, school choice (by parents), and teacher allocation rules.

Following [Solon, Page, and Duncan \(2000\)](#) and others, a straightforward way to loosen this restriction is to regress the outcome on a set of household dummies only. As per row 2 of [table 2](#), this will capture both the stand-alone household effect and the effect of all unobserved variables correlated with this factor (see also [Raaum, Salvanes, and Sørensen 2006](#); [Björklund and Salvanes 2011](#)). To see this, imagine part of the contribution of schools to learning reflects a direct causal effect of (average) constituent households, such as when households make direct financial or time commitments to school functioning. This kind of mechanism is also suggested by versions of cream-skimming models, where average peer quality in a school (or class) is driven by household characteristics, which in turn directly influences individual achievement ([Walsh 2009](#)). So, assuming *s* can be partitioned into components that are oblique or parallel to *b* plus an orthogonal component *v*, with its own variance σ_v^2 , we have

$$s_k = \gamma \frac{1}{N_k} \sum_{j|K=k} b_{jK} + v_k = \gamma \bar{b}_{jk} + v_k, \tag{2}$$

$$E(\bar{b}'_{jk} v_k) = 0.$$

The corollary is given in [table 2](#), row 3, where household effects are assumed to be (partial) reflections of given school effects, plus an orthogonal component (ω). In both these cases, the observed covariance between household and school effects is attributed wholly to one of the factors, meaning there is no remaining covariance by construction. These two models thus directly provide upper bounds on the variance components of interest, while lower bounds can be calculated in a two-step procedure (e.g., [Raaum, Salvanes, and Sørensen 2006](#); [Bau and Das 2020](#))—in the first stage a one-way fixed-effects model is estimated; the second effect is extracted from the combined residuals—e.g., $v_k = \frac{1}{N_k} \sum_{j|K=k} (t_{ijK} - b_j - \gamma \bar{b}_{jK})$.

3.2. Methods

Interpreting upper-bound variance contributions as gross effects (including sorting) and lower bounds for the same factor as net effects (excluding sorting), a natural approximation to the overall contribution of the factor covariance is just their arithmetic difference. This is similar in spirit to approaches used to tighten estimated teacher-effect variance estimates (as in [McCaffrey et al. 2009](#); [Hoffmann and Oreopoulos 2009](#)).¹ Formally, appendix S1 of the supplementary online appendix shows how to

1 Thanks to an anonymous referee for highlighting this approach.

reexpress both the household and school lower-bound variance contributions in terms of their “true” effects: $\sigma_\omega^2 = (1 - \rho_{hs}^2)\sigma_h^2$ and $\sigma_v^2 = (1 - \rho_{hs}^2)\sigma_s^2$, where ρ_{hs} is the pairwise factor correlation.

Combining these expressions yields the following statement for the difference in the household upper and lower bounds:

$$(1 + \gamma)^2 \sigma_h^2 - \sigma_\omega^2 = \underbrace{2\rho_{hs}\sigma_h\sigma_s}_{\text{sorting}} + \underbrace{\rho_{hs}^2(\sigma_h^2 + \sigma_s^2)}_{\text{bias}}. \tag{3}$$

This states that the previously defined range for the household variance contribution (on the left-hand side) represents a strict upper bound on the sorting term. The bias is positive and increasing in the factor correlation, reflecting the fact that the two lower-bound contributions subtracted from the test-score variance only exactly offset the “true” contributions associated with each individual factor in the special case of zero covariance. Due to the equality of the upper-bound household and school models (table 2, rows 2 and 3) it follows that the same expression gives the difference in the school upper and lower bounds.

The advantage of this approximation of the sorting term is that it is straightforward to implement. However, further elaboration reveals a complete decomposition is also feasible. Applying previous expressions, rewrite the unrestricted linear model as

$$\sigma_t^2 - \sigma_e^2 = \sigma_h^2 + \sigma_s^2 + 2\Sigma_{sh} = \frac{\sigma_\omega^2 + \sigma_v^2 + 2\rho_{hs}\sigma_\omega\sigma_v}{1 - \rho_{hs}^2},$$

which, rearranging, yields a quadratic in ρ :

$$(\rho_{hs}^2 - 1)(\sigma_t^2 - \sigma_e^2) + 2\rho_{hs}\sigma_\omega\sigma_v + (\sigma_\omega^2 + \sigma_v^2) = 0. \tag{4}$$

It follows that once three specific quantities are known—namely, the two lower-bound (uncorrelated) variance components and the total variance jointly attributable to the two latent factors ($\sigma_t^2 - \sigma_e^2$)—then the unknown correlation coefficient can be obtained as the root of equation (4). Where there is doubt as to the relevant sign of the pairwise correlation, the approximate bounds can provide helpful guidance: in principle, the upper- and lower-bound difference could be small or even negative (e.g., for $-1 \leq \gamma < 0$). And as summarized in table S4.1 in the supplementary online appendix, any of the variance decomposition components of interest, including the contribution of sorting as well as the upper- and lower-bound terms, can be calculated from three primitives alone: namely, σ_ω^2 , σ_v^2 , ρ_{hs} .

3.3. Implementation

The ultimate objective of the variance decomposition is to estimate the three primitives. As already suggested, however, upper- and lower-bound estimates of household and school effects can be derived using separate one-way fixed-effect models. These directly permit calculation of the sorting approximation (equation (3)), and the correlation parameter can be derived in an additional step. Details of this procedure, which we refer to as indirect fixed effects (FEi), are set out in the supplementary online appendix S2.

A drawback of approaches using high-dimensional fixed effects is that estimates of latent factors generally include measurement error, meaning the corresponding raw-variance shares will be upwards biased (Koedel, Mihaly, and Rockoff 2015). To address this, we use empirical Bayes shrinkage, which involves adjusting each estimated effect towards a common prior by a factor proportional to the estimated noise-to-signal ratio in the original estimates. Following Stanek, Well, and Ockene (1999), we shrink each estimated fixed effect (e.g., \hat{h}_j) towards a global mean as follows:

$$\tilde{h}_j = \bar{h}_j + \psi(\hat{h}_j - \bar{h}_j)$$

with shrinkage factor $0 < \psi = \frac{\sigma_b^2}{\sigma_b^2 + \sigma_\epsilon^2/N_j} < 1$.

Here, N_j is the effective degrees of freedom available to estimate each of the j effects, σ_b^2 is the raw variance of the estimated household fixed effect, σ_ϵ^2 is the estimated overall residual variance, and \bar{h}_j is the sample fixed-effect mean, typically zero under conventional normalization restrictions.

A complementary approach to dealing with the presence of sorting is suggested by [Altonji and Mansfield \(2018\)](#) (hereafter, AM18). Motivated by a concern that sorting on unobserved variables may bias estimates of the effects of group-level inputs, such as schools, these authors show that where households with different characteristics also have differing effective preferences (willingness-to-pay) for underlying school or community amenities, group averages of observed household or pupil characteristics can be employed as control functions. Inclusion of these generated variables, henceforth referred to as sorting proxies, in a regression model thus permits identification of a lower bound on the unique (variance) contribution of school factors. Here, the specification of interest becomes

$$t_{ijkl} = b'_{O,j}\beta + \bar{h}'_{O,k}\gamma + [h_{jl}] + [s_{kl}] + e_{ijkl}, \quad (5)$$

where $b_{O,j}$ are observed household-level covariates, and $\bar{h}_{O,k}$ are the set of sorting proxies derived directly from the former. Treating the terms in brackets as (uncorrelated) household and school random effects, the specification can be estimated using a linear mixed-effects estimator (denoted MLM; as used by AM18).

An advantage of this procedure, which is useful for triangulation, is that it relies on a single estimation equation. A drawback is that the random effects are assumed to be orthogonal to the set of included covariates, implying [equation \(5\)](#) may remain misspecified.² Furthermore, while the variances of the random effects are estimated as model parameters, their associated best-linear-unbiased predictors (BLUPs or conditional modes) typically do not share the same variance–covariance structure as the estimated population-level moments (see [Morris 2002](#)). Consequently, the BLUPs may not be reliable for the purposes of investigating systematic patterns in the random effects (e.g., subgroup heterogeneity).³ Consequently, as an additional validation procedure, we extend the AM18 approach to incorporate fixed- as opposed to random-effects terms. As described in supplementary online appendix S2, this involves modifying the indirect estimation procedure whereby the sorting proxies are used directly to account for variation in test scores before estimation of the (one-way) household and school fixed effects. This procedure is denoted FED.

Before proceeding to the application, it merits emphasis that under all three methods we seek to estimate the (upper- and) lower-bound variance components, not the “true” components (σ_b^2 , σ_s^2). Furthermore, note that in the absence of longitudinal data, both households and schools are only encountered in the same locations and, thus, may share a common community (neighborhood) component. To capture this explicitly, we elaborate the household and school effects as follows:

$$z_{jkl} = c_l + h_{jl} + s_{kl}, \quad (6)$$

where z represents the *joint* contribution of schools and households in community l , and c is the location-specific effect. Admittedly, the nature and magnitude of these latter effects remains controversial (e.g., [Oreopoulos 2003](#)). However, there is growing evidence that the quality of local community environments can have a material influence on child development trajectories ([Chetty, Hendren, and Katz 2016](#); [Black et al. 2017](#); [Richter et al. 2017](#); [Chetty and Hendren 2018](#)). Furthermore, ignoring the contribution of the latter term would mean they are simply absorbed by the upper-bound household or school-effect

- 2 For households, a “tighter” lower-bound variance share is obtained by adding the estimated variance of the fitted observed component to the variance of the random component.
- 3 Estimation of high-dimensional mixed-effects models is also highly computationally intensive and prone to convergence problems. In the present case we use the `lme4` package in R for estimation ([Bates 2010](#)).

estimates, in turn muddying the interpretation of the variance components. Consequently, in the empirical implementation we treat community effects as a separate factor. In doing so, we estimate community effects first, in keeping with upper-bound fixed-effects models, in turn imposing the assumption that the estimated household/school effects are orthogonal to the community effects (see supplementary online appendix S2 for details). As such, our estimates of sorting should be net of any selection of households or schools into different communities, meaning the estimated contribution of sorting (defined above) should be interpreted in a narrow sense: i.e., it refers to processes of sorting between schools within a specific location.⁴ Similarly, the household and school variance components will capture the average contribution (variation) of these factors within communities.

4. Data

Since 2010, the Uwezo initiative has undertaken large-scale household-based surveys of academic achievement in Kenya, mainland Tanzania, and Uganda.⁵ The surveys target children residing in households aged between the official starting-school age and 16 and are representative at both national and district levels. Excluding the initial surveys, five rounds of the Uwezo surveys are publicly available (2011–2015) and used here. For each household, the surveys collected information covering general characteristics, as well as the demographic and educational details of resident children (e.g., age, gender, whether or not attending school). Also, all children of school age were individually administered a set of basic oral literacy and numeracy tests, which were tailored to each country and varied by survey round based on a common template, to reflect competencies stipulated in the national curriculum at the grade 2 level.

Importantly, the Uwezo tests are not adapted to the children's ages or their completed level of schooling. Given that they focus on basic competencies, there are strong age-related differences, which affect both the level and variance of scores between age cohorts. Consequently, we exclude children above 14 years old since they tend to perform at the upper end of the tests (and show much lower variance). Our analytical sample also excludes observations that can be perfectly predicted using either household or school fixed effects: i.e., singletons have been removed. [Table 3](#) summarizes alternative metrics of learning for this sample, pooling data from all countries and rounds. The first part shows the raw literacy and numeracy test scores, described in detail in [Jones et al. \(2014\)](#). The literacy tests refer to the national languages of instruction in which pupils are tested at the end of primary school: i.e., English and Kiswahili in Kenya and Tanzania, and just English in Uganda. For consistency, we then take the main language as the overall literacy score, which is combined with the score in math to give the overall mean.

Simple averages of scores on different tests may provide poor guidance to underlying levels of ability, particularly where tests are composed of questions of variable difficulty ([Bond and Lang 2013](#)). Thus, we use two other approaches to combine the information on math and literacy into an overall test score. First, we take the first principal component of the suite of tests answered by each child (denoted PCA); second, in line with recommendations in the literature, we apply a graded response IRT model to the same tests, where the empirical Bayes mean of the latent trait is taken as the metric of achievement (ability).⁶

- 4 Note, however, that unlike in many advanced countries, migration driven by school choice is not common in the East Africa region, at least at young ages. Not only are barriers to internal migration generally recognized as high ([Hirvonen 2016](#)), but the drivers of migration are predominantly of an economic nature and migrants are typically younger adults (without children) (see [Beegle, De Weerd, and Dercon 2011](#)).
- 5 The approach adopted by Uwezo has been inspired by exercises carried out in India by the Assessment Survey Evaluation Research Centre (ASER). For further details and comparison to other regional assessments see [Uwezo \(2012\)](#) and [Jones et al. \(2014\)](#).
- 6 Further details of both these procedures are available on request. Due to differences in test forms between countries and across years, we estimate the IRT models separately by country and survey round. In a small number of cases, the graded response model did not converge, and instead we used a partial credit model.

Table 3. Summary of Alternative Metrics of Learning (Test Scores)

Outcomes	Mean	SD	Pairwise correlations						
			(a)	(b)	(c)	(d)	(e)	(f)	(g)
(–) Raw Kiswahili score	2.83	1.94							
(–) Raw English score	2.09	1.85							
(a) Raw literacy score	2.39	1.89	1.00						
(b) Raw math score	3.72	2.39	0.75	1.00					
(c) Mean score (raw)	2.98	1.97	0.87	0.96	1.00				
(d) PCA score (std.)	0.00	1.00	0.90	0.90	0.95	1.00			
(e) IRT score (std.)	0.00	1.00	0.92	0.82	0.91	0.97	1.00		
(f) Mean score (age std.)	0.00	1.00	0.74	0.76	0.80	0.84	0.82	1.00	
(g) PCA score (age std.)	0.00	1.00	0.76	0.75	0.79	0.84	0.83	0.99	1.00
(h) IRT score (age std.)	–0.00	1.00	0.78	0.67	0.76	0.81	0.86	0.96	0.96

Source: Authors' elaboration.

Note: SD is the standard deviation; mean scores is the arithmetic mean of math and language scores; PCA score is estimated as the first principal component of math and language scores; IRT score is estimated via graded response models; "std." indicates scores are standardized by country and year; "age std." indicates scores are standardized by age, country, and year; all survey rounds are pooled ($N = 869, 030$).

Since neither the PCA or IRT scores have a natural scale, for each country and survey round we standardize them to have a mean of 0 and standard deviation of 1. These scores are shown in the second part of [table 3](#). Notably, they are extremely highly correlated with one another and with the overall mean (correlation >0.90).

Since the same Uwezo test forms are administered to children of all ages in each household, there is likely to be remaining between-cohort variation in test scores. From the present perspective, however, this can be considered unwanted noise (see [Mazumder 2008](#)). Thus, to ensure the score standard deviations are constant across cohorts over time, we further standardize the overall scores by age group within each country and survey year. These scores are shown in the final part of [table 3](#). They continue to be highly mutually correlated and, reflecting their attractive theoretical properties and widespread usage, we focus on the IRT scores henceforth.

Table S4.2 in the supplementary online appendix further reports regional means and standard deviations of the raw and IRT test scores for the analytical sample. The first column reports the simple means of the raw competency tests (ordinal scores), column (2) reports the IRT scores standardized by country and round, but not age, and column (3) reports the final measures, including age standardization. As can be seen, movement from the second to the third metric constitutes a simple monotone transformation. Also, as per the methodological discussion of the Methods section, the rank position of each region according to its test-score variance is largely preserved, regardless of the transformation applied.

To implement the decomposition procedures, household and school indexes must be defined. The former is trivial: unique indexes are ascribed to all siblings in the same household (in each year).⁷ The school effects are less straightforward. In the present data, limited information about schools is provided. Nonetheless, we know the kind of school attended (none, public, or private) and whether or not children attend the specified main public school in their local community (catchment area). Consequently, for each enumeration area, we categorize children into one of four school categories: (1) those not attending school; (2) those attending the main local public school, which we denote as "matched"; (3) those attending other public schools; and (4) those attending private schools.

The advantage of our definition of school effects is that, within each household, children can be mapped to different school effects: i.e., the school and household effects are crossed. The main downside is that

7 The Uwezo surveys are cross sectional in nature and no attempt is made to track the same children over time.

for the last two school categories we do not identify which specific school they attend. Further, the first category is composed of a mix of children that either never enrolled or dropped out. Recognizing that estimates of the school effects for these three groups are likely to be noisy averages of multiple underlying effects, unless otherwise indicated we present our main results for the matched group only. In doing so, however, we first run estimations on the full sample (e.g., to ensure household effects are identified correctly), but then report variance decompositions for the chosen subgroup.⁸ Full results are also presented for comparative purposes.

Further descriptive statistics for the dataset are reported in table S4.3 in the supplementary online appendix. This shows the number of unique children (i), households (j), and schools (k) covered in the dataset. Additionally, the table reports summary statistics, including average child characteristics and schooling status indicators (those out of school, the shares attending the specific matched public schools, and those attending private schools). Overall, these indicate the sample is comprehensive and balanced (by age and gender).⁹ It also reveals there are systematic differences in schooling among countries as well as across regions within each country: e.g., in all countries enrolment rates differ substantially across regions (e.g., from 80.0 to 96.5 percent in Kenya), as do the shares of children attending private school.

5. Results

5.1. Spatial Clustering

Motivated by our interest in sorting, we begin with a preliminary review of the degree of spatial clustering or segregation: i.e., the extent to which similar types of children are found in the same schools or communities. To do so, we compare the correlation in educational outcomes between individuals within different aggregate units, which is equivalent to the proportion of the variance attributable to the between-group structure of the data (for similar exercises see [Fryer and Levitt 2004](#); [Friesen and Krauth 2007, 2010](#); [Lindahl 2011](#); [Emran and Shilpi 2015](#)). The magnitude of the correlation across members of the same group, and how quickly this declines as we move to higher levels of aggregation, thus indicates the extent to which variables are spatially segregated (clustered). For instance, if local communities only contained households with exactly the same socioeconomic status, then the proportion of variance in socioeconomic status accounted for by communities would be equal to that accounted for by households, indicating a very high degree of clustering.

Results from this exercise are reported in [table 4](#), covering a range of variables and using the full sample of children. As might be expected, assuming child gender is approximately random, the between-group variance share accounted for by households is extremely low. For the remaining variables, however, clustering by households, schools, or communities is much higher. For instance, more than two-thirds of the variation in access to clean water is accounted for both by schools and by communities. Overall, a little more than half of the total variation in aggregate socioeconomic status in the region (SES) is attributable to (average) differences between distinct communities, implying substantial levels of residential clustering or economic segregation.

Turning to the educational outcomes at the bottom of the table, we observe somewhat lower but nonetheless substantial magnitudes of clustering across communities and schools, in part reflecting variation in schooling status between children within households.¹⁰ In terms of achievement on the Uwezo tests as measured by the age-adjusted IRT scores, the correlation between siblings is almost 50 percent, which is highly comparable to magnitudes found in a range of other countries (see Decomposition Framework).

8 If we were to restrict the sample before running the regression models, then we would have no crossed effects: i.e., in a given community, we would only observe children attending a single school.

9 Average ages are higher in Tanzania as the starting school age is 7, compared to 6 in the other countries.

10 Of course, enrolment is fully accounted for by school-level variation, but there is within-household variation in enrolment.

Table 4. Percent of Between-Group Variance Attributable to Alternative Grouping Structures

Group level →		Hhld.	Sch.	Vill.	District	Region	Ratios	
Variable		(a)	(b)	(c)	(d)	(e)	(b)/(a)	(c)/(b)
Female	0.5	3.0	3.0	1.7	0.1	0.0	1.01	0.56
Hhld. has electricity	0.2	100.0	54.0	47.5	19.1	7.4	0.54	0.88
Hhld. has clean water	0.4	100.0	70.5	66.9	13.4	2.8	0.70	0.95
Hhld. owns phone	0.7	100.0	38.9	33.0	10.3	5.4	0.39	0.85
Mother's education	4.4	100.0	52.7	46.4	13.6	6.7	0.53	0.88
Aggregate SES index	−0.1	100.0	58.6	51.8	22.5	9.4	0.59	0.88
Enrolled	0.9	40.6	100.0	27.7	8.2	3.2	2.46	0.28
Highest grade	−0.0	50.2	42.9	31.1	11.8	5.2	0.85	0.72
Math achievement (IRT)	0.0	42.6	31.9	25.4	7.3	2.9	0.75	0.80
Literacy achievement (IRT)	0.0	45.0	34.4	28.2	10.7	5.0	0.76	0.82
Overall achievement (IRT)	−0.0	48.8	37.5	30.2	10.6	4.6	0.77	0.81

Source: Authors' elaboration.

Note: Cells in columns (a)–(e) report the correlation between children within the same grouping unit (e.g., households, schools, etc.), which is the adjusted R^2 from a one-way fixed-effects model, without covariates; for highest grade, this correlation is calculated conditioning on child age; the final two columns report ratios; SES is socio-economic status; IRT refers to item response theory metrics.

More notably, however, the correlation between pupils attending the same schools, as well as between children in the same communities, is only moderately smaller. Nearly 40 percent of the overall variation in achievement is accounted for by the schooling effects, and 30 percent between distinct residential locations. And although the definition of communities is not equivalent across studies, these latter magnitudes appear to be a factor larger than encountered in developed countries. That is, the degree of clustering of both economic and educational outcomes in East Africa appears to be large.

5.2. Variance Bounds

Moving to the variance decomposition, [table 5](#) presents results from the upper-bound models for households (HUB) and schools (SUB), estimated using separate one-way fixed-effects estimators (the FE_i approach) as described in supplementary online appendix S2. Panel (a) focuses attention on the subsample of individuals attending the matched public school, and panel (b) refers to the full sample. To account for observed individual-level covariates, such as gender, grade-for-age, birth order, and whether they have ever attended school, these variables are added to the factor models, and their combined variance contribution is calculated manually from the corresponding estimates, denoted the “individual” effect. Given that both these and the community-level effects are estimated *before* extracting the separate contributions of the unobserved factors, their variance shares are fixed across the two models: i.e., following the section Decomposition Framework, the difference between the household and school models only concerns where the factor covariance is attributed. Both absolute and relative variance contributions are reported in the table, where the former is reported in standard deviation units ($\times 100$). Standard errors are calculated via a clustered bootstrap procedure.

Three immediate findings merit comment. First, our models are able to account for about half of the variation in test scores. This is broadly consistent with studies elsewhere (see below), but confirms most variation in outcomes occurs between children in the same household. Second, as expected following [table 4](#), while the community effect is sizeable, both the household and school upper-bound variance contributions are large. For the full sample these represent 32 percent and 20 percent of the total variance, respectively; however, in the matched sample, the school effect contribution drops to 13 percent. This gap implies almost half of the variation associated with schooling relates to differences across types of schools, including not being enrolled, a point to which we return below. Even so, the variation between matched

Table 5. Variance Partitions from Household and School Upper-Bound Models, FEi Procedure

	Absolute (in SD units)		Relative (in percent)	
	HUB	SUB	HUB	SUB
<i>(a) Matched subsample</i>				
Individual		20.6 (0.2)		4.8 (0.1)
Household	53.4 (0.4)	41.8 (0.3)	32.1 (0.4)	19.6 (0.3)
School	9.2 (0.3)	34.5 (0.3)	0.9 (0.1)	13.4 (0.3)
Community		33.0 (1.0)		12.3 (0.7)
Residual		66.5 (0.4)		49.9 (0.5)
Sorting (approx.)		33.3 (0.4)		12.5 (0.3)
<i>(b) Full sample</i>				
Individual		21.3 (0.2)		4.5 (0.1)
Household	56.4 (0.7)	41.1 (0.4)	31.8 (0.6)	16.9 (0.5)
School	22.7 (0.8)	44.8 (1.0)	5.1 (0.3)	20.0 (0.7)
Community		36.0 (1.8)		13.0 (1.1)
Residual		67.5 (0.5)		45.5 (0.8)
Sorting (approx.)		38.6 (0.8)		14.9 (0.5)

Source: Authors' elaboration.

Note: HUB is the household upper-bound model (table 2, row 2); SUB is the school upper-bound model (table 2, row 3); sorting is the upper-bound approximation to the factor covariance contribution from equation (3); individual, community, and residual components do not vary across HUB and SUB models; standard errors (in parentheses) are calculated via a clustered bootstrap procedure; SD units are standard deviation units; FEi refers to indirect fixed effects.

public schools remains sizeable: a 1 standard deviation increase in school quality corresponds to a 0.34 standard deviation improvement in test scores.

Last, the (upper-bound) approximation to the sorting component is highly material. For the matched sample this equals 13 percent of the variance, or 16 percent for the full sample. These magnitudes are far from trivial, being just marginally smaller than the upper-bound estimates for the school effects or about half the upper-bound household effect contribution. According to the Uwezo data, an additional year of schooling is associated with an approximate 0.18 standard deviation increase in a child's test score *ceteris paribus*. Against this benchmark, a shift of 1 standard deviation in either the quality of (public) schools or degree of sorting would be roughly equivalent to about two years of schooling. However, since this estimate of the sorting term is expected to be biased upwards, we now turn to results from the full decomposition.

5.3. Complete Decomposition

Using these lower-bound estimates, table S4.4 in the supplementary online appendix reports estimates for the three primitives required to identify the within-community sorting component. It also sets out the corresponding estimates from the two other empirical methods, based on AM18 (see Methods; also

Table 6. Complete Variance Decomposition, All East Africa

	Absolute (in SD units)			Relative (in percent)		
	FEi	MLM	FEd	FEi	MLM	FEd
<i>(a) Matched subsample</i>						
Individual	20.6 (0.2)	24.5 (0.3)	20.6 (0.2)	4.8 (0.1)	6.8 (0.2)	4.8 (0.1)
Household	48.1 (0.3)	41.2 (0.4)	44.1 (0.3)	26.1 (0.4)	19.1 (0.3)	21.9 (0.3)
School	10.6 (0.3)	24.5 (0.2)	20.6 (0.4)	1.3 (0.1)	6.8 (0.1)	4.8 (0.2)
Sorting	22.5 (0.3)	19.4 (0.8)	23.8 (0.3)	5.7 (0.1)	4.3 (0.4)	6.4 (0.2)
Community	33.0 (1.0)	38.3 (0.5)	33.0 (1.0)	12.3 (0.7)	16.5 (0.4)	12.3 (0.7)
Residual	66.5 (0.4)	64.3 (0.6)	66.5 (0.4)	49.9 (0.5)	46.5 (0.6)	49.9 (0.5)
<i>(b) Full sample</i>						
Individual	21.3 (0.2)	25.0 (0.3)	21.3 (0.2)	4.5 (0.1)	6.3 (0.2)	4.5 (0.1)
Household	45.6 (0.4)	43.4 (0.5)	44.4 (0.4)	20.8 (0.4)	18.8 (0.3)	19.7 (0.4)
School	25.1 (0.8)	31.2 (0.5)	28.2 (0.7)	6.3 (0.4)	9.8 (0.3)	7.9 (0.3)
Sorting	31.4 (0.7)	29.5 (1.3)	30.5 (0.7)	9.9 (0.3)	8.7 (0.6)	9.3 (0.3)
Community	36.0 (1.8)	38.1 (0.6)	36.0 (1.8)	13.0 (1.1)	14.5 (0.4)	13.0 (1.1)
Residual	67.5 (0.5)	64.7 (0.7)	67.5 (0.5)	45.5 (0.8)	41.9 (0.9)	45.5 (0.8)

Source: Authors' elaboration.

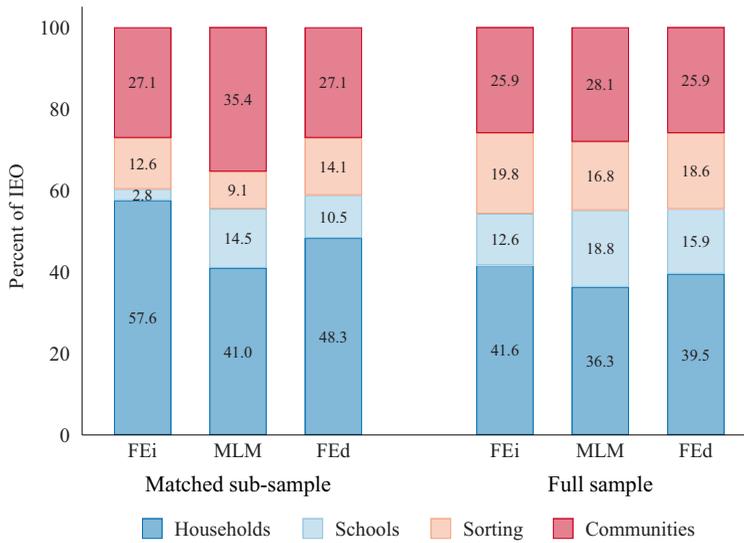
Note: Columns indicate the outcome transformation and decomposition method, abbreviated as FEi = indirect fixed effects, MLM = mixed-effects linear model, and FEd = direct fixed effects; "absolute" refers to standard deviation units ($\times 100$); "relative" is as a percentage of the total variance; standard errors (in parentheses) calculated via a clustered bootstrap procedure; matched subsample refers to children attending the main local public school in each location; SD units are standard deviation units.

supplementary online appendix S2): namely, the MLM and the FEd approaches.¹¹ Together these form the basis for the complete decomposition, the results of which are reported in table 6 for both the matched and full samples. The relative contributions of the different constitutive components to total IEO, defined as the overall variance contribution of systematic components above the level of the individual, are further illustrated in figure 1.

As can be seen from the primitives, the three decomposition methods yield highly consistent results and, hence, the different methods provide similar estimates for the various final variance components. Despite some differences in precise magnitudes, discussed further below, this suggests the present results are robust to the method chosen (e.g., mixed versus fixed effects) and, more specifically, that our simple indirect approach (FEi) performs very adequately. In the full sample, we find a positive correlation between the household and school effects, ranging between 0.32 and 0.43. Consequently, the derived

11 For the MLM results, variance components are calculated from the estimated BLUPs of the random effects plus their standard error, which match the population second moments. Furthermore, the estimated BLUPs of the random effects show material cross correlation, particularly with the community effect. As a result, in order to ensure comparability with the fixed-effects procedures in which the school and household primitives are defined as uncorrelated factors, we use only the component of the estimated school and household BLUPs that are orthogonal to the community component. Further details available on request.

Figure 1. Variance Components (in Percent of Inequality of Opportunities), Alternative Methods, All East Africa



Source: Authors' elaboration.

Note: Bars indicate the percentage contribution of each factor to inequality of educational opportunity (IEO), which is the variance accounted for by systematic supra-individual factors (households, schools, communities plus their covariance); matched subsample refers to children attending the main local public school in each location; abbreviations for methods: FEi is indirect fixed effects, MLM is the mixed-effects model, and FEi is direct fixed effects (FEi).

within-community sorting component, shown in table 6, remains substantial and is often similar in magnitude to the stand-alone contribution of schools. And although the estimates for this component are all lower than the earlier approximation (as expected), they nonetheless confirm that within-community sorting makes a nontrivial contribution to the total variance in test scores, equal to around 8 percent.

Five further points stand out. First, as before, the primary difference between the matched and full sample estimates refers to the contribution of schools. And, as this variance component feeds into the sorting component, the latter also remains somewhat smaller in the matched sample. The main implication is that sorting of children from distinct types of households across different types of schools, as well as at the extensive margin of schooling, is a material aspect of the overall covariance between household and school effects.

Second, differences between the three empirical methods appear to derive from two main sources. On the one hand, some of the differences for the correlation and sorting components reflect the sensitivity of calculations based on equation (4). Recall we estimate ρ_{hs} from the household- and school-effect lower bounds. Thus, even minor differences in either of these components, such as estimates for school effects under the FEi procedure, influence the estimates for the correlation term.

On the other hand, the mixed- and fixed-effects approaches treat the community effects in a slightly different fashion. In the latter case, community effects are removed from the joint fixed effect (households plus schools) in the first of the series of orthogonal projections (supplementary online appendix S2). As such, and in similar fashion to the upper-bound models discussed in Decomposition Framework, the estimated community variance contribution will capture *both* the direct effect of shared community-level factors on outcomes (e.g., via environmental conditions) plus any covariance with either household or school factors. Thus, any possible sorting of households (or teachers) across different communities will be absorbed by the community effect.¹² In the mixed-effects approach, all terms are estimated simultaneously,

12 Since we do not observe the same households (or teachers) across different communities, we are not able to distinguish between the direct and sorting effects associated with residential location.

meaning that a strict distinction between within- and between-community components is not enforced *ex ante*. Although we implement a correction for this to obtain the uncorrelated school and household components (see above), this methodological difference is likely to explain some of the disparities in our estimates of the community effect and the uncorrelated variance components. However, overall, the qualitative differences in the results from the different approaches are trivial.

Third, in absolute and relative terms, the stand-alone contribution of schools might be deemed moderate, representing no more than 10 percent of the total variance and even less for the matched sample (7 percent). However, these magnitudes are not small when compared with school-effects estimates in the literature (see [table 1](#)). Also, AM18 estimate that both school *and* community effects jointly explain around 5 percent of the variation in high-school graduation rates (lower bounds) in the United States. Here, schools alone account for roughly the same amount of variation in achievement, and combined with community effects the joint contribution is nearly 20 percent ([table 6](#)). Similarly, the joint contribution of schools and communities taken from the school upper-bound estimates ([table 5](#)), which are generally more directly comparable to estimates found elsewhere (e.g., [Freeman and Viarengo 2014](#)) are substantial. Our estimates for East Africa place this joint effect at around 25 percent, which is similar to the average school-effect variance contribution calculated by [Pritchett and Viarengo \(2015\)](#) for a range of countries. Thus, while the present results show some similarities with findings elsewhere, the material contribution of residential location within the joint school/community effect stands out.

Fourth, as partly indicated, while all sources of inequality are substantial, systematic environmental and institutional factors that lie at levels above the household clearly merit consideration in the assessment of IEO. As shown in [figure 1](#), the combination of schools, communities, and sorting is typically greater than the household contribution alone, accounting for around half of the total IEO. Not only does this reinforce our interest in quantifying multiple sources of IEO simultaneously, but also, assuming that household-based inequalities are generally persistent, it suggests the scope for policy measures to reduce such inequalities are substantial.

Last, we recognize the residual component remains large, accounting for roughly 50 percent of the test-score variance. On the one hand, this may well reflect the limitations of our explanatory variables. For instance, we do not measure the extent to which children within the same household receive different parental inputs, nor do we capture differences in the quality of teachers within each school. These factors will effectively remain in the residual, thereby implying we may be underestimating the magnitude of IEO, which underlines the importance of the research agenda on measuring teacher effects (e.g., [Buhl-Wiggers et al. 2018](#); [Bau and Das 2020](#)) and variable household inputs (e.g., [Das et al. 2013](#)). On the other hand, a large residual is not inconsistent with substantial differences in ability and effort, as well as other idiosyncratic factors at the individual level. Indeed, studies elsewhere suggest these effects are likely to be highly material (see [Troncoso, Pampaka, and Olsen 2016](#), [table 1](#)). For instance, on the PISA tests, [Zamarro, Hitt, and Mendez \(2019\)](#) find that around a third of the cross-country variation in scores might be explained by variation in effort across countries. Put differently, a large residual confirms that not all outcome inequality necessarily reflects opportunity inequalities.

5.4. Heterogeneity

Averages can often hide substantial heterogeneity. Thus, to look behind the pooled results, we take advantage of the properties of the variance and calculate the decomposition for specific subgroups. These results are presented in [table 7](#), where we stratify individuals along various dimensions: gender, age group, schooling effect category, and household SES quintile. For simplicity, and given the broad consistency across the empirical procedures, hereafter we only present results for the direct fixed-effects method. This is chosen

Table 7. Subgroup Variance Decomposition (in Percent of Total Variance), Full Sample

Group	Level	Indiv. σ_x^2	Hhld σ_h^2	Sch. σ_s^2	Sort $2\Sigma_{bs}$	Comm. σ_c^2	Resid. σ_e^2	Total σ_t^2	Correl. ρ_{bs}
Female	No	4.4 (0.1)	19.7 (0.4)	8.1 (0.3)	9.2 (0.3)	12.8 (1.0)	45.7 (0.7)	100.0 —	0.37 (0.01)
	Yes	4.4 (0.1)	19.8 (0.4)	7.8 (0.3)	9.4 (0.3)	13.1 (1.0)	45.4 (0.8)	100.0 —	0.38 (0.01)
Age group	6–9	4.5 (0.2)	20.5 (0.6)	9.2 (0.5)	10.0 (0.4)	14.0 (1.1)	41.8 (1.2)	100.0 —	0.36 (0.01)
	10–13	4.5 (0.2)	19.2 (0.4)	6.9 (0.2)	8.7 (0.3)	12.1 (1.1)	48.6 (0.7)	100.0 —	0.38 (0.01)
SES tercile	Low	4.5 (0.1)	20.4 (0.3)	9.2 (0.5)	9.6 (0.5)	11.1 (0.5)	45.1 (1.0)	100.0 —	0.35 (0.01)
	Medium	4.8 (0.1)	20.0 (0.3)	8.0 (0.3)	8.3 (0.4)	11.3 (0.6)	47.5 (0.7)	100.0 —	0.33 (0.01)
	High	5.0 (0.2)	18.3 (0.5)	7.5 (0.3)	8.0 (0.4)	14.3 (1.4)	46.9 (0.6)	100.0 —	0.34 (0.01)
Sch. cat.	Unenrolled	4.2 (0.3)	19.1 (0.5)	12.1 (0.4)	11.0 (0.7)	15.3 (0.9)	38.3 (1.0)	100.0 —	0.36 (0.02)
	Matched	4.8 (0.1)	21.9 (0.3)	4.8 (0.2)	6.4 (0.2)	12.3 (0.6)	49.9 (0.6)	100.0 —	0.31 (0.01)
	Other public	4.7 (0.1)	19.9 (0.5)	5.8 (0.1)	7.0 (0.2)	13.3 (1.0)	49.3 (0.6)	100.0 —	0.33 (0.01)
	Private	5.0 (0.1)	18.4 (0.7)	6.7 (0.6)	6.4 (0.7)	14.1 (2.1)	49.4 (0.6)	100.0 —	0.29 (0.01)
All		4.5 (0.1)	19.7 (0.4)	7.9 (0.3)	9.3 (0.3)	13.0 (1.1)	45.5 (0.8)	100.0 —	0.37 (0.01)

Source: Authors' elaboration.

Note: The table reports direct fixed-effects variance decomposition estimates (FEd), based on the unrestricted linear model representation; standard errors (in parentheses) are calculated via a clustered bootstrap procedure; SES is socio-economic status.

as it gives variance component estimates that lie mostly in between those of the two other methods (e.g., see table S4.4 in the supplementary online appendix).¹³

Focusing on the main insights, there large differences in the variance components over the subgroups.¹⁴ Nonetheless, the four explained variance components above the level of the individual are all generally larger for younger children as well as among children who are *not* in school. Recall that the school effects treat all children within each community who are not enrolled as a separate category. Thus, the variance decomposition for the various categories of enrolled children captures variation between schools within each category, while results for the subgroup of nonenrolled children captures variation in environmental effects outside the schooling system. In this sense, as the correlation coefficient and the school effects are largest in the nonenrolled category, this suggests schools do play a role in reducing inequality of opportunity. At the same time, among the three school-type categories, the contribution of variation between

13 For reference, table S4.5 in the supplementary online appendix presents the identical subgroup decomposition in absolute magnitudes, and tables S4.6 and S4.7 present the same relative and absolute results based on the mixed-effects method, which provides the largest estimates of the sorting component among the three methods.

14 Statistically speaking, based on a nonparametric medians test, across all variance components the differences between groups are generally significant at the 1 percent level, except for between boys and girls (details on request); however, the relative magnitudes of these differences are generally within ± 10 percent of the overall effect.

schools is moderately larger in the private sector. While this may reflect measurement error, it may also reflect differences between lower- and higher-cost private schools, where these segments serve children of quite different backgrounds.

A related form of heterogeneity runs along spatial lines, which is particularly relevant from the perspective of pursuing geographically targeted interventions. Here, an advantage of the Uwezo data is that we can run the same analysis at the district level. Figure S5.1 in the supplementary online appendix plots the cumulative empirical district-level distributions of the relative variance shares for the four main IEO components taken from the FE estimates, by country. These show substantial variations across all components within each country, but country-specific differences are also evident. The structure of (relative) IEO appears most distinct in Tanzania, where the contributions of communities, sorting, and households tend to be much larger (on average). Even so, within countries, the variance contributions of the different components all display a wide range in comparison to the mean.

Leveraging this granular data, [table 8](#) presents results from a seemingly unrelated regression of the five absolute variance components (for the full sample) on a set of explanatory variables, which reflect the composition of the sample in each district. Recognizing that lower-level components of IEO are likely to influence higher-level components, we structure the system recursively. Five results stand out. First, while we make no claims these results are causal, it is evident that the variance components do vary systematically across districts: i.e., around half of the variation in component magnitudes or more can be attributed to the included covariates. Second, gender differences are notable. Higher female shares are strongly associated with lower contributions of both individual and sorting effects, suggesting educational investments for male children are a typical focus of active choice. Indeed, the data suggest boys are more likely to drop out or never attend school.

Third, further reinforcing previous findings, school status matters. The regression estimates employ a baseline category of children attending a public school in Kenya. Thus, higher shares of children not attending school (either never enrolled or dropped out) are associated with much larger school effects. So, attending school significantly compensates for other location-based differences in learning. Also, a greater share of children attending private school is also associated with a moderately larger school effect, indicative of significant differences in quality between public and private schools.

Fourth, a simple metric of school competition or choice, defined as 1 minus the normalized Herfindahl index calculated from the share of children attending each of the three different school type categories (out of all enrolled children), is also relevant. In particular, greater school choice is associated with larger household and schooling effects, but lower sorting. This may reflect the point that districts with greater diversity in household types demand greater school choice and (naturally) a greater variation in schools is associated with larger dispersion in school quality. Even so, at a minimum, this reinforces the conclusion that schools do matter. However, the negative association with sorting would point to other mechanisms at play: e.g., contrary to a standard cream-skimming hypothesis, greater school choice would enable those households with preferences for higher-quality education to choose accordingly, regardless of their own “type.” So, competition may not be inequality enhancing.

Fifth, socioeconomic inequality *within* each district, as measured by the standard deviation of the SES variable, remains material. Not only is there a clear positive association between the magnitude of the household effect and the extent of local inequality (as expected), but also districts with greater inequality appear to display larger school and sorting contributions. So, holding other conditions fixed, including the share of private schools and school choice, less economically equal districts show greater variation in school quality and more acute sorting. This is consistent with the idea that, even among public schools, poorer communities generally have access to poorer quality schools.

5.5. Long-Run Inequalities

As a final exercise we follow the spirit of [Kremer \(1997\)](#) and consider the implications of sorting for long-run (steady-state) educational inequalities. The key idea here is that in the presence of positive educational

Table 8. Systematic Patterns in Absolute Variance Components across Districts, Based on Full Sample Results from the FEd Procedure

	Effect (SD units)				
	Indiv.	Hhld.	Sch.	Sorting	Vill.
Age of children	0.00 (0.01)	-0.08*** (0.02)	-0.08*** (0.02)	0.01 (0.01)	-0.06*** (0.02)
Female share	-0.12*** (0.04)	0.15 (0.10)	-0.12 (0.10)	-0.14** (0.06)	-0.08 (0.09)
Never attended sch.	0.05*** (0.01)	0.07* (0.04)	0.68*** (0.04)	-0.03 (0.03)	-0.00 (0.05)
No longer in sch.	0.17*** (0.06)	0.29* (0.17)	0.59*** (0.18)	-0.22** (0.10)	-0.41*** (0.16)
Attends matched sch.	-0.01 (0.01)	0.03 (0.03)	0.06 (0.04)	-0.09*** (0.02)	0.03 (0.03)
Attends private sch.	0.05*** (0.01)	-0.07*** (0.03)	0.13*** (0.03)	-0.02 (0.02)	-0.03 (0.02)
School competition	-0.01 (0.01)	0.09*** (0.02)	0.11*** (0.02)	-0.08*** (0.01)	0.02 (0.02)
SES (level)	-0.01*** (0.00)	-0.04*** (0.01)	-0.04*** (0.01)	-0.00 (0.00)	0.00 (0.01)
SES (st. dev.)	0.01 (0.01)	0.05*** (0.02)	0.05*** (0.02)	0.03*** (0.01)	0.02 (0.02)
Tanzania	-0.04*** (0.00)	0.07*** (0.01)	0.04*** (0.01)	-0.01* (0.01)	0.05*** (0.01)
Uganda	-0.04*** (0.00)	0.02*** (0.01)	-0.00 (0.01)	-0.03*** (0.00)	0.01 (0.01)
Individual effect		0.13 (0.14)	0.33** (0.15)	-0.47*** (0.08)	-0.05 (0.13)
Household effect			0.20*** (0.05)	0.31*** (0.03)	-0.01 (0.05)
School effect				0.56*** (0.03)	-0.30*** (0.06)
Sorting effect					1.01*** (0.07)
Constant	0.23*** (0.01)	0.37*** (0.04)	-0.01 (0.05)	0.16*** (0.03)	0.04 (0.04)
R ²	0.73	0.47	0.80	0.87	0.67
RMSE	0.01	0.03	0.03	0.02	0.03

Source: Authors' elaboration.

Note: Dependent variables, indicated in the columns and ordered hierarchically, are absolute variance contributions in standard deviation (SD) units; all explanatory variables are district-level means (proportions for dummy variables) or standard deviations; system estimated by a (triangular) seemingly unrelated regression; standard errors in parentheses; FEd is direct fixed effects; RMSE is root mean square error; SES is socio-economic status. Significance levels are * 10 percent, ** 5 percent, *** 1 percent.

sorting, children from more advantaged families receive a double benefit that can be expected to exacerbate inequalities over time. Such children benefit both from family circumstances and from higher-quality schooling, which has a multiplicative effect across generations.

In order to quantify the potential magnitude of these effects, we move to a dynamic or intergenerational setting. Furthermore, we adopt the simple formulation that household effects (*h*) are directly proportional to parental achievement, reflecting both parental capacity to support learning and other characteristics that flow from their educational level. Indexing generations by *g*, and using previous notation, these

Table 9. Summary of Long-Run Inequality Simulations, All Districts

Scenario →	Fall in inequality (in percent)			Fall in $\bar{\delta}$ (in percent)
	(I)	(II)	(III)	(III)
10th percentile	0.12	0.07	0.19	0.30
Median	0.15	0.15	0.29	0.60
90th percentile	0.19	0.20	0.38	0.88
Mean	0.15	0.14	0.28	0.59

Source: Authors' elaboration.

Note: In the first three columns the cells report summary statistics of the percentage reduction in steady-state inequality of educational attainment associated with the indicated scenario; in the final column we report the fall in the intergenerational persistence parameter equivalent to the results under scenario (III); simulations are undertaken at the district level.

assumptions imply $h_{jg} \approx h_{j0} + \delta t_{j, g-1}$, which in turn means

$$\delta^2 = \frac{\sigma_{h_g}^2}{\sigma_{t_{g-1}}^2} = \frac{\sigma_{\omega_g}^2}{(1 - \rho_{bs,g}^2)\sigma_{t_{g-1}}^2} = \frac{\bar{\delta}^2}{1 - \rho_{bs,g}^2}, \quad (7)$$

where $\bar{\delta} = \sigma_{\omega_g}^2 / \sigma_{t_{g-1}}^2$ is the sorting-invariant or raw intergenerational persistence parameter. From this, we can then solve for the long-run steady-state level of educational inequality:

$$\sigma_t^2 = \frac{\sigma_v^2 + 2\bar{\delta}\rho\sigma_v\sigma_t + (\sigma_c^2 + \sigma_e^2)(1 - \rho^2)}{(1 - \rho^2)(1 - \bar{\delta}^2)}, \quad (8)$$

where σ_c^2 is the upper-bound variance contribution associated with residential location, which we treat as exogenous for simplicity (see supplementary online appendix S1 for the derivation).

To simulate this model, a value for the raw persistence parameter needs to be chosen. To do so, we opt for a data-driven approach and make the assumption that currently observed inequality is approximately at the steady state (i.e., $\sigma_t^2 = \sigma_{\bar{t}}^2$), allowing us to solve for $\bar{\delta}$ using equation (8) based on previously estimated values of the primitives. This gives us the magnitude of the persistence parameter that would maintain achievement inequality at current levels. Fixing $\bar{\delta}$ at this value also means that simulations based on changes to (other) primitives in the model can be interpreted directly with reference to the current magnitude of achievement inequality. Figure S5.2 in the supplementary online appendix plots the resulting district-level cumulative distributions of the inequality-constant persistence parameter, which range from 0.34 to 0.60; the corresponding pooled overall coefficient is 0.48. These magnitudes are broadly consistent with intergenerational regression coefficients in educational attainment found elsewhere.¹⁵ In other words, our baseline assumption is likely to be conservative.

With the constant-inequality persistence parameter in hand, we now use equation (8) to simulate the steady-state level of inequality under alternative assumptions for other inputs. We consider three scenarios: (I) we set the within-community sorting correlation coefficient ρ to zero, under which the previous equation simplifies to $\sigma_t^2 = (\sigma_v^2 + \sigma_c^2 + \sigma_e^2) / (1 - \bar{\delta}^2)$; (II) we remove the variance contribution due to community effects, setting it to zero; and (III) we combine the assumptions of scenarios (I) and (II) simultaneously. Looking across all districts, table 9 summarizes the results and figure S5.3 in the supplementary online appendix shows them visually. On average, eliminating within-community sorting leads

15 For instance, Hertz et al. (2007) estimate an average persistence parameter of 0.42 for a range of mostly developed countries. However, existing estimates for (low income) developing countries typically point to somewhat larger values. For instance, Emran and Shilpi (2015) find a persistence parameter of around 0.50 in India, and in a sample of African countries, Azomahou and Yitbarek (2016) find an average persistence of 0.66. These values are generally larger than our own estimates, implying a higher steady-state level of educational inequality than observed today.

steady-state inequality to fall by about 15 percent relative to the baseline. As per [equation \(8\)](#), this underlines that eliminating educational sorting would lead to a substantially larger reduction in long-run inequality than might be suggested by the relative variance contribution of sorting alone (e.g., [figure S5.1](#) in the supplementary online appendix). Indeed, the relative variance contribution due to sorting is roughly only half the expected proportional decline in inequality if sorting were to be reduced.

Eliminating within-community sorting does not address segregation of either households or teachers (schools) by residential location, as captured by the community effect. Our earlier results indicated this term was material, potentially reflecting significant barriers to internal mobility in the East African context. Exogenously cutting out this component, as per scenario (II), generates a similar reduction in inequality compared to the first scenario (14 percent on average), but with more dispersion. Combining these two assumptions, as in scenario (III), equates to a fall in long-run inequality by about one-third, on average. To put this magnitude of reduction in context, it is equivalent to a very substantial fall in the intergenerational persistence of educational achievement. The final column of [table 9](#) (also [figure S5.3](#), panel IV in the supplementary online appendix) reports the percentage reduction in the persistence parameter required to match the effect of scenario (III): namely, eliminating both within-community sorting and between-community segregation. On average, the equivalent reduction is around 60 percent, implying a mean persistence parameter of 0.20 (versus 0.48 in the baseline).

6. Conclusions

There is continued debate regarding the extent to which sorting exacerbates inequalities, including in educational attainment and achievement. We contribute to this debate by setting out and implementing a novel variance decomposition, identifying sorting as the contribution of the covariance between household and school effects to the variation in educational outcomes. We quantify the contribution of sorting to learning inequalities among over one million children in three East African countries: Kenya, mainland Tanzania, and Uganda.

To estimate the contribution of sorting, we first propose a simple bounding procedure based on separate fixed-effects estimators, where sorting is given by the difference between the upper- and lower-bound contributions of households (or schools). Second, we show how the same estimates can be combined to back out the factor correlation coefficient, from which the covariance is derived straightforwardly. To triangulate our results based on fixed-effects estimators, we further adopt the approach of AM18 ([Altonji and Mansfield 2018](#)), who show how group means of observed factors can absorb confounding effects of sorting on unobserved characteristics. This helps isolate the independent contributions of the factors of interest, from which lower-bound variance contributions can be estimated using mixed-effects estimators. Last, we also extend the AM18 approach to allow for fixed as opposed to random effects, which is both computationally more tractable and better suited to exploring subgroup heterogeneity.

Empirically, all three estimation procedures for the complete decomposition indicate positive sorting of pupils across schools. We estimate that sorting accounts for around 8 percent of the total test-score variance and almost a fifth of the joint variation in test scores due to schools, communities, and households. The sorting contribution tends to be larger among families with children not enrolled in school, among those sending their children to private schools, where there is less school competition or in locations with greater socioeconomic inequality. We further find the stand-alone contributions of schools and households to the variance are around 10 percent and 20 percent respectively, where the former is in keeping with the order of magnitude of institutional-level effects found elsewhere on the continent (e.g., [Bau and Das 2020](#)).

To explore the implications of these results, we conduct simulations of how learning inequality evolves over time. These show that for the average district in the region, the steady-state level of educational inequality would fall by around 15 percent if within-community sorting between schools and households

were to be fully eliminated. Moreover, in a number of districts, the reduction in inequality from eliminating this form of sorting would be over 20 percent. If between-community sorting were also cut to zero, this would have an equivalent effect of cutting the intergenerational persistence of educational attainment by more than half. Overall, this suggests that policies that take sorting into account, and even proactively seek to counteract it by allocating additional resources (e.g., better teachers) to schools in disadvantaged locations, merit serious consideration.

References

- Abowd, J. M., R. H. Creecy, and F. Kramarz. 2002. "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data." Longitudinal Employer-Household Dynamics Technical Papers 2002-06, Center for Economic Studies, US Census Bureau, Washington, DC, USA.
- Altonji, J. G., and R. K. Mansfield. 2018. "Estimating Group Effects Using Averages of Observables to Control for Sorting on Unobservables: School and Neighborhood Effects." *American Economic Review* 108 (10): 2902–46.
- Andrews, M. J., L. Gill, T. Schank, and R. Upward. 2008. "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171 (3): 673–97.
- Anger, S., and D. D. Schnitzlein. 2017. "Cognitive Skills, Non-Cognitive Skills, and Family Background: Evidence from Sibling Correlations." *Journal of Population Economics* 30 (2): 591–620.
- Azomahou, T. T., and E. A. Yitbarek. 2016. "Intergenerational Education Mobility in Africa: Has Progress Been Inclusive?". Policy Research Working Paper 7843, World Bank, Washington, DC, USA.
- Bates, D. M. 2010. *lme4: Mixed-Effects Modeling with R*. New York: Springer.
- Bau, N., and J. Das. 2020. "Teacher Value Added in a Low-Income Country." *American Economic Journal: Economic Policy* 12 (1): 62–96.
- Beegle, K., J. De Weerdt, and S. Dercon. 2011. "Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey." *Review of Economics and Statistics* 93 (3): 1010–33.
- Behrman, J. 2020. "Social Mobility and Human Capital in Low- and Middle-Income Countries." In *Social Mobility in Developing Countries: Concepts, Methods, and Determinants*, edited by A. K. V. Iversen and K. Sen. Oxford: Oxford University Press.
- Benabou, R. 1993. "Workings of a City: Location, Education, and Production." *Quarterly Journal of Economics* 108 (3): 619–52.
- Björklund, A., and K. G. Salvanes. 2011. "Education and Family Background: Mechanisms and Policies." In *Handbook of the Economics of Education*, edited by E. Hanushek, S. Machin and L. Wößmann, Volume 3, Chapter 3, 201–47. Amsterdam: Elsevier.
- Black, M. M., S. P. Walker, L. C. Fernald, C. T. Andersen, A. M. DiGirolamo, C. Lu, and D. C. McCoy et al. 2017. "Early Childhood Development Coming of Age: Science through the Life Course." *Lancet* 389 (10064): 77–90.
- Bond, T. N., and K. Lang. 2013. "The Evolution of the Black-White Test Score Gap in Grades K–3: The Fragility of Results." *Review of Economics and Statistics* 95 (5): 1468–79.
- Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler. 2020. "How Much Should We Trust Estimates of Firm Effects and Worker Sorting?" NBER Working Paper 27368, National Bureau of Economic Research, Cambridge, MA, USA.
- Buhl-Wiggers, J., J. T. Kerwin, J. A. Smith, and R. Thornton. 2018. "Teacher Effectiveness in Africa." Working Paper S-89238-UGA-1, International Growth Center, London, UK.
- Carneiro, P. 2008. "Equality of Opportunity and Educational Achievement in Portugal." *Portuguese Economic Journal* 7 (1): 17–41.
- Chakravarty, S. R. 2001. "The Variance as a Subgroup Decomposable Measure of Inequality." *Social Indicators Research* 53 (1): 79–95.
- Chetty, R., and N. Hendren. 2018. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *Quarterly Journal of Economics* 133 (3): 1107–62.
- Chetty, R., N. Hendren, and L. F. Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106 (4): 855–902.

- Clotfelter, C. T., H. F. Ladd, and J. L. Vigdor. 2011. "Teacher Mobility, School Segregation, and Pay-Based Policies to Level the Playing Field." *Education Finance and Policy* 6 (3): 399–438.
- Coleman, J. S. 1966. *Equality of Educational Opportunity*. US Department of Health, Education, and Welfare, Office of Education, Washington DC, USA.
- Correia, S. 2017. "reghdfe: Stata Module for Linear and Instrumental-Variable/GMM Regression Absorbing Multiple Levels of Fixed Effects." Technical Report Statistical Software Components s457874, Boston College Department of Economics, Boston, USA.
- Dahan, M., and A. Gaviria. 2001. "Sibling Correlations and Intergenerational Mobility in Latin America." *Economic Development and Cultural Change* 49 (3): 537–54.
- Dang, H.-A., and P. Glewwe. 2018. "Well Begun, but Aiming Higher: A Review of Vietnam's Education Trends in the Past 20 Years and Emerging Challenges." *Journal of Development Studies* 54 (7): 1171–95.
- Das, J., S. Dercon, J. Habyarimana, P. Krishnan, K. Muralidharan, and V. Sundararaman. 2013. "School Inputs, Household Substitution, and Test Scores." *American Economic Journal: Applied Economics* 5 (2): 29–57.
- Emran, M. S., and F. Shilpi. 2015. "Gender, Geography, and Generations: Intergenerational Educational Mobility in Post-Reform India." *World Development* 72: 362–80.
- Fernandez, R. 2003. "Sorting, Education, and Inequality." In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, edited by M. Dewatripont, L. P. Hansen and S. J. Turnovsky, Chapter 1, 1–40. Cambridge: Cambridge University Press.
- Fernández, R., and R. Rogerson. 2001. "Sorting and Long-Run Inequality." *Quarterly Journal of Economics* 116 (4): 1305–41.
- Ferreira, F. H., and J. Gignoux. 2014. "The Measurement of Educational Inequality: Achievement and Opportunity." *World Bank Economic Review* 28 (2): 210–46.
- Freeman, R. B., and M. Viarengo. 2014. "School and Family Effects on Educational Outcomes across Countries." *Economic Policy* 29 (79): 395–446.
- Friesen, J., and B. Krauth. 2007. "Sorting and Inequality in Canadian Schools." *Journal of Public Economics* 91 (11–12): 2185–212.
- . 2010. "Sorting, Peers, and Achievement of Aboriginal Students in British Columbia." *Canadian Journal of Economics/Revue canadienne d'économique* 43 (4): 1273–301.
- Fryer, R. G., and S. D. Levitt. 2004. "Understanding the Black-White Test Score Gap in the First Two Years of School." *Review of Economics and Statistics* 86 (2): 447–64.
- Gamoran, A., and D. A. Long. 2007. "Equality of Educational Opportunity: A 40 Year Retrospective." In *International Studies in Educational Inequality, Theory and Policy*, edited by R. Teese, S. Lamb, M. Duru-Bellat and S. Helme, 23–47. Springer Netherlands.
- Glewwe, P., S. Krutikova, and C. Rolleston. 2017. "Do Schools Reinforce or Reduce Learning Gaps between Advantaged and Disadvantaged Students? Evidence from Vietnam and Peru." *Economic Development and Cultural Change* 65 (4): 699–739.
- Hanushek, E. A., J. F. Kain, D. M. O'Brien, and S. G. Rivkin. 2005. "The Market for Teacher Quality." NBER Working Papers 11154, National Bureau of Economic Research, Cambridge, MA, USA.
- Hanushek, E. A., and S. G. Rivkin. 2012. "The Distribution of Teacher Quality and Implications for Policy." *Annual Review of Economics* 4 (1): 131–57.
- Hanushek, E., and K. Yilmaz. 2007. "The Complementarity of Tiebout and Alonso." *Journal of Housing Economics* 16 (2): 243–61.
- Hertz, T., T. Jayasundera, P. Piraino, S. Selcuk, N. Smith, and A. Verashchagina. 2007. "The Inheritance of Educational Inequality: International Comparisons and Fifty-Year Trends." *BE Journal of Economic Analysis & Policy* 7 (2): Article 10.
- Hirvonen, K. 2016. "Temperature Changes, Household Consumption, and Internal Migration: Evidence from Tanzania." *American Journal of Agricultural Economics* 98 (4): 1230–49.
- Hoffmann, F., and P. Oreopoulos. 2009. "Professor Qualities and Student Achievement." *Review of Economics and Statistics* 91 (1): 83–92.
- Hsieh, C.-T., and M. Urquiola. 2006. "The Effects of Generalized School Choice on Achievement and Stratification: Evidence from Chile's Voucher Program." *Journal of Public Economics* 90 (8): 1477–503.

- Jackson, C. K. 2009. "Student Demographics, Teacher Sorting, and Teacher Quality: Evidence from the End of School Desegregation." *Journal of Labor Economics* 27 (2): 213–56.
- . 2013. "Match Quality, Worker Productivity, and Worker Mobility: Direct Evidence from Teachers." *Review of Economics and Statistics* 95 (4): 1096–1116.
- Jochmans, K., and M. Weidner. 2019. "Fixed-Effect Regressions on Network Data." *Econometrica* 87 (5): 1543–60.
- Jones, S., Y. Schipper, S. Ruto, and R. Rajani. 2014. "Can Your Child Read and Count? Measuring Learning Outcomes in East Africa." *Journal of African Economies* 23 (5): 643–672.
- Koedel, C., K. Mihaly, and J. E. Rockoff. 2015. "Value-Added Modeling: A Review." *Economics of Education Review* 47: 180–95.
- Kremer, M. 1997. "How Much Does Sorting Increase Inequality?" *Quarterly Journal of Economics* 112 (1): 115–39.
- Lindahl, L. 2011. "A Comparison of Family and Neighborhood Effects on Grades, Test Scores, Educational Attainment and Income – Evidence from Sweden." *Journal of Economic Inequality* 9 (2): 207–26.
- Louw, M., S. Van Der Berg, and D. Yu. 2007. "Convergence of a Kind: Educational Attainment and Intergenerational Social Mobility in South Africa." *South African Journal of Economics* 75 (3): 548–71.
- Lu, C., J. Cuartas, G. Fink, D. McCoy, K. Liu, Z. Li, and B. Daelmans et al. 2020. "Inequalities in Early Childhood Care and Development in Low/Middle-Income Countries: 2010–2018." *BMJ Global Health* 5 (2): e002314.
- Mansfield, R. K. 2015. "Teacher Quality and Student Inequality." *Journal of Labor Economics* 33 (3): 751–88.
- Mazumder, B. 2008. "Sibling Similarities and Economic Inequality in the US." *Journal of Population Economics* 21 (3): 685–701.
- . 2011. "Family and Community Influences on Health and Socioeconomic Status: Sibling Correlations over the Life Course." *B.E. Journal of Economic Analysis & Policy* 11 (3): Article 1.
- McCaffrey, D. F., T. R. Sass, J. Lockwood, and K. Mihaly. 2009. "The Intertemporal Variability of Teacher Effect Estimates." *Education Finance and Policy* 4 (4): 572–606.
- McEwan, P. J. 2004. "The Indigenous Test Score Gap in Bolivia and Chile." *Economic Development and Cultural Change* 53 (1): 157–90.
- McEwan, P. J., and M. Trowbridge. 2007. "The Achievement of Indigenous Students in Guatemalan Primary Schools." *International Journal of Educational Development* 27 (1): 61–76.
- Morris, J. S. 2002. "The BLUPs Are Not "Best" When It Comes to Bootstrapping." *Statistics & Probability Letters* 56 (4): 425–30.
- Nechyba, T. J. 2006. "Income and Peer Quality Sorting in Public and Private Schools." In *Handbook of the Economics of Education*, edited by E. A. Hanushek and F. Welch, Volume 2, Chapter 22, pages 1327–68. Amsterdam: Elsevier.
- Nickell, S. J. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica* 49 (6): 1417–26.
- Nicoletti, C., and B. Rabe. 2013. "Inequality in Pupils' Test Scores: How Much Do Family, Sibling Type and Neighbourhood Matter?" *Economica* 80 (318): 197–218.
- OECD. 2018. *Effective Teacher Policies*. Organization for Economic Cooperation and Development. Paris: OECD.
- Oreopoulos, P. 2003. "The Long-Run Consequences of Living in a Poor Neighborhood." *Quarterly Journal of Economics* 118 (4): 1533–75.
- Pop-Eleches, C., and M. Urquiola. 2013. "Going to a Better School: Effects and Behavioral Responses." *American Economic Review* 103 (4): 1289–1324.
- Pritchett, L. 2013. *The Rebirth of Education: Schooling Ain't Learning*. Washington, DC: Center for Global Development.
- Pritchett, L., and M. Viarengo. 2015. "Does Public Sector Control Reduce Variance in School Quality?" *Education Economics* 23 (5): 557–76.
- Raaum, O., K. G. Salvanes, and E. Ø. Sørensen. 2006. "The Neighbourhood Is Not What It Used to Be." *Economic Journal* 116 (508): 200–22.
- Raffo, C., D. Dyson, H. M. Gunter, D. Hall, and A. Kalambouka. 2007. *Education and Poverty: A Critical Review of Theory, Policy and Practice*. York, UK: Joseph Rowntree Foundation.
- Reynolds, D., and C. Teddlie. 2000. *The International Handbook of School Effectiveness Research*. London: Routledge.
- Richter, L. M., B. Daelmans, J. Lombardi, J. Heymann, F. L. Boo, J. R. Behrman, and C. Lu et al. 2017. "Investing in the Foundation of Sustainable Development: Pathways to Scale Up for Early Childhood Development." *Lancet* 389 (10064): 103–18.

- Roemer, J. E. 1996. *Theories of Distributive Justice*. Cambridge MA: Harvard University Press.
- Sass, T. R., J. Hannaway, Z. Xu, D. N. Figlio, and L. Feng. 2012. "Value Added of Teachers in High-Poverty Schools and Lower Poverty Schools." *Journal of Urban Economics* 72 (2): 104–22.
- Shorrocks, A. F. 1982. "Inequality Decomposition by Factor Components." *Econometrica* 50 (1): 193–211.
- Solon, G., M. E. Page, and G. J. Duncan. 2000. "Correlations between Neighboring Children in Their Subsequent Educational Attainment." *Review of Economics and Statistics* 82 (3): 383–92.
- Stanek, E. J., A. Well, and I. Ockene. 1999. "Why Not Routinely Use Best Linear Unbiased Predictors (BLUPs) as Estimates of Cholesterol, Per Cent Fat from Kcal and Physical Activity?" *Statistics in Medicine* 18 (21): 2943–59.
- Thiemann, P. 2018. "Inequality in Education Outcomes: The Role of Sorting among Students, Teachers, and Schools." Unpublished manuscript, Lund University. https://ekstern.filer.uib.no/svf/Econ%20web/Petra%20Thiemann_teacher-sorting-2018-10-16.pdf.
- Todd, P. E., and K. I. Wolpin. 2003. "On the Specification and Estimation of the Production Function for Cognitive Achievement." *Economic Journal* 113 (485): F3–F33.
- Troncoso, P., M. Pampaka, and W. Olsen. 2016. "Beyond Traditional School Value-Added Models: A Multilevel Analysis of Complex School Effects in Chile." *School Effectiveness and School Improvement* 27 (3): 293–314.
- Uwezo. 2012. "Are Our Children Learning? Literacy and Numeracy across East Africa." Technical report, Uwezo, Dar Es Salaam, Tanzania.
- Walsh, P. 2009. "Effects of School Choice on the Margin: The Cream Is Already Skimmed." *Economics of Education Review* 28 (2): 227–36.
- Watkins, K. 2012. "The Power of Circumstance: A New Approach to Measuring Education Inequality." Center for Universal Education Working Paper 5, Brookings Institute, Washington DC, USA.
- Zamarro, G., C. Hitt, and I. Mendez. 2019. "When Students Don't Care: Reexamining International Differences in Achievement and Student Effort." *Journal of Human Capital* 13 (4): 519–52.