Spatial Inequality of Opportunity in West Africa

Luke Heath Milsom*

1st April, 2023

Abstract

Using a movers design I find that on average moving to a one standard deviation better location at birth increases a child's probability of completing primary school by 15 percentage points in Benin, Cameroon, and Mali. Place matters more for girls than boys and is correlated with local characteristics such as urbanization, access to electricity, and employment in agriculture. Religion is weakly correlated with average causal place effects but strongly correlated with the gender gap in the effect place has on educational attainment.

*Milsom: Oxford University, Department of Economics, Manor Road Building, luke.milsom@economics.ox.ac.uk. With thanks to: Ferdinand Rauch, Niclas Moneke, and Gabriel Ulyssea for their guidance, patience, and supervision; Hannah Zillessen, Sam Altmann, Binta Zahra Diop, Shihang Hou, Verena Wiedemann, Itzhak Rasooly, and the entire DPhil cohort at Oxford for many helpful comments and long discussions; Johannes Abeler, for his continued advice and assistance; Barbara Petrongolo, Simon Quinn, Ian Crawford, Isabelle Roland, Hamish Low, Abigail Adams-Prassl, Jessica Pan, Christian Dustmann, Nina Guyon, Tony Venables, and Dave Donaldson, for useful comments and conversations. I am also extremely grateful to the editor, Nathaniel Hendren, and three anonymous referees for their comments and suggestions which greatly improved the paper.

Since 1970 in Benin, Cameroon, and Mali, GDP per capita has doubled in real terms, child mortality has halved and the human development index has also doubled.[1] However, the rising tide has not lifted all boats. Inequality within countries, and across space is pervasive. More worryingly, this spatial inequality might indicate differing opportunities accruing to individuals growing up in different locations, rather than the result of sorting across space. Inequality of opportunity determined by where a child grows up is intrinsically undesirable, could interact with gender inequality, and may have negative consequences on aggregate development.

In this paper, I quantify the magnitude of spatial inequality of opportunity, analyze the correlates of local opportunity, and study it's consequences on gender inequality in Benin, Cameroon, and Mali. I use a movers design to estimate the causal effect of spending an additional year of childhood on primary school completion, in every locality[2]-year cell, both overall and by gender. I focus on primary education as it is the most salient margin of education in this setting (a third of people in my sample complete primary school whereas only 7% go on to complete secondary school) as is correlated with greater later life outcomes broadly defined such as not working in agriculture, better housing quality, and higher wages. The movers design focuses on a sub-sample consisting of those who move across localities, and leverages variation in the exact age of a child when a family moves to identify place effects.[3] This approach builds on the nascent *place effects* literature which aims to map the causal effect of growing up in different areas on child outcomes. In this, I follow seminal work by Chetty et al. [2016], Chetty and Hendren [2018a,b] and Chyn [2018] in the US context, as well as work by Damm and Dustmann [2014] in Denmark, Deutscher [2020] in Australia, Laliberté [2021] in Montreal, and Alesina et al. [2021] in Sub-Saharan Africa.

On average, children converge to the outcome of permanent residents in a given location at a rate of 2.7% per year of childhood spent in the location.[4] The validity of the movers design rests on the key identifying assumption that selection effects do not vary over the age

[1] Statistics based on author's calculations using data from the World Development Indicators, and the World Bank, Feenstra et al. [2015] and Prados de la Escosura [2021], between 1970 and 2015.

[2] In this paper, I refer to second administrative divisions (Communes in Benin, Departments in Cameroon, and Circles in Mali) as localities. These have a median population of 267,000 and can be thought of as similar in scale to commuter zones in the US.

[3] Intuitively, this strategy considers children born in the same period t and in the same location o, who move to the same destination location d, but do so at different ages. I then compare the outcomes of children who spend a larger proportion of their childhood in d vis-à-vis o to infer the causal effect of d relative to o on primary school completion under the assumption that other determinants of children's outcomes are unrelated to the age at which they move. I then make all such possible comparisons across localities to uncover the (relative) causal effect of spending an additional year of childhood in each location on the probability of completing primary education.

[4] This is consistent with results from Chetty and Hendren [2018a] in the US and Alesina et al. [2021] in Sub-Saharan Africa, both of which look at intergenerational mobility. Chetty and Hendren [2018a] find a convergence rate of about 4% and Alesina et al. [2021] find a rate between 1.5% and 3%.

of a child when they move.[5] I validate this assumption in five ways. First, I show that there is no evidence of varying selection effects on observables. Second, I demonstrate that results are robust to including household fixed effects, that is leveraging cross-sibling age variation. Third, I consider increasingly likely to be exogenous move events and show that results remain stable. Fourth, I perform an over-identification test using variation across gender in locality quality. Finally, I employ a placebo exercise and show that estimated place effects for those who move *after* they've completed primary education are not correlated with those found in my main specification.

Estimating locality-year effects, I find that opportunity is indeed unequally distributed across space. Moving to a location with one standard deviation greater opportunity at birth increases the probability of a child completing primary school by 15 percentage points. I analyze whether spatial inequality varies across countries or over time and find that all three countries exhibit high levels of inequality of opportunity across space, indicating that the aggregate results are not driven by a single country. Next, I characterize areas of high opportunity. Urbanization, electricity provision, sanitation, and housing quality are associated with higher observed primary completion rates and opportunity. Conversely, employment in agriculture is negatively associated with both observed primary completion rates and the causal effect of place. However, in all cases, less than half of the correlation between characteristics and observed completion rates operates through the causal effects channel.

A key characteristic, often associated with education completion, is the prevailing religion in a locality [Alesina et al., 2020, Panin et al., 2021]. I find that, although the proportion of individuals in a locality indicating Islam as their religion is on average strongly negatively associated with observed primary completion rates, religion impacts the causal effect of place to a much lesser degree. Indeed, religion stands out among the characteristics considered as one where a strong observational correlation doesn't translate into a significant impact on the causal effect place has on primary education completion.

Place could also play a role in mitigating or exacerbating prevailing gender inequalities.[6] To study the relationship between spatial and gender inequality of opportunity, I estimate separate overall convergence rates for boys and girls, as well as gender-specific locality-year causal place effects. I find that girls converge to the outcomes of permanent residents at a rate over twice that of boys, implying that place matters more for female than male outcomes. By correlating place characteristics with the difference between relative male and female causal place effects, I test whether place characteristics are relatively better for male

[5]This assumption allows households with better unobservables to move to better areas, but is violated if they systematically do so earlier or later depending on their child's age.

[6]Although overall primary completion rates have increased substantially over the sample period, gender gaps have remained persistently high: In Benin (1992-2013) increasing from 17 to 20 percentage points, in Cameroon (1976-2005) decreasing from 21 to 12 percentage points, and in Mali (1998-2009) increasing from 6 to 8 percentage points.

than female outcomes and therefore contribute to, or detract from, gender inequality of opportunity. Urbanization and electricity availability increase gender inequality of opportunity, but employment in agriculture reduces it. I also find that the average null correlation of local religion on causal place effects hides considerable gender heterogeneity.

This paper is most closely related to Chetty and Hendren [2018a] and Chetty and Hendren [2018b] who study the impact of place on income mobility in the US and estimate locality-specific effects, and Alesina et al. [2021] who study observational education mobility in Africa and estimate average place effects. I consider a different outcome, setting, and time frame from Chetty and Hendren [2018a,b] and go further than Alesina et al. [2021] by estimating place-specific effects over time, considering heterogeneity across countries, and studying the ramifications of inequality of opportunity on gender inequality. My main contributions relative to these papers are fourfold. First, I quantify the time-varying spatial inequality of opportunity in Benin, Cameroon, and Mali, and find that opportunity is unequally distributed in all countries. Second, I characterize place effects, describing what areas of high opportunity look like, and show that religion does not impact causal place effects on average. Third, I study the confluence of spatial and gender inequality and find that place has a potentially significant role to play in reducing gender inequality, but that characteristics associated with higher opportunity may also be associated with greater gender inequality. Fourth, I find that although religion is weakly correlated with average causal place effects, it is strongly correlated with the gender gap in the effect place has on educational attainment.

This paper proceeds as follows: Section 1 describes the data and empirical setting; section 2 estimates overall causal place effects; section 3 estimates locality-year specific place effects; section 4 characterizes place effects and considers the role of religion; section 5 studies heterogeneity by gender and place variables associated with greater(or lower) gender inequality. Finally, section 6 concludes.

1 Setting and data sources

1.1 Why focus on Benin, Cameroon and Mali

Benin, Cameroon, and Mali, whilst representing important countries to study in and of themselves, also provide a particularly interesting setting to examine how the location of upbringing affects primary education completion and thus later life outcomes. There are three main reasons for this. First, Benin, Cameroon, and Mali are developing countries aiming to grow rapidly in the near future, and ensuring equitable growth is key to individual well-being, as well as potentially aggregate success. Although exhibiting some similarities Benin, Cameroon, and Mali also have many differences. Thus enabling this paper to speak to issues of external validity and cross-country heterogeneity that could not be achieved by just focusing on one country. Second, in these countries, primary completion is a key measure of the opportunities available to individuals broadly defined as discussed in more detail in subsection 1.3. Third: data requirements. Benin, Cameroon, and Mali are the only countries in Sub-Saharan Africa that meet the high data requirements of this paper which are explained in more detail in subsection 1.2.

1.2 Data description

This paper uses data on 8 million observations from 1976 to 2013, across 164 localities from every available census from Benin, Cameroon, and Mali.[7] To estimate causal place effects using a movers design it's necessary to observe where an individual is when the census was taken, their previous location, their birth location, and how long they have resided in their current location.[8] In addition to migration information, I require data on primary completion,[9] age, and gender, at a fine and consistent geographic level over multiple cross-sectional waves.

Only three countries in Sub-Saharan Africa — Benin, Cameroon, and Mali fit the stringent data requirements. In Benin, I use censuses from 1992, 2002, and 2013, encompassing over 2.2m individual-level observations in 77 consistent localities (Communes). In Cameroon, I use censuses from 1976, 1987, and 2005 covering over 3.4m individual-level observations in 39 consistent localities (Departments). Finally, in Mali, I use censuses from 1998, and 2009 giving over 2.4m individual-level observations in 48 consistent localities (Circles). In total, this gives me a sample of over 8m individual-level observations across 164 localities and 444 locality-year cells. A benefit of using this data is that it is freely and publicly available from IPUMS International for anyone to access.

Using available variables on the current location, previous location, and birth location

[7]Data is accessed from IPUMS [2020] and consists of 10% samples, with thanks to the National Institute of Statistics and Economic Analysis in Benin, the Central Bureau of Census and Population Studies in Cameroon and the National Directorate of Statistics and Informatics in Mali, who provided the underlying data.

[8]In the censuses described, I can geo-locate individuals at the second administrative unit level. These localities have a median population of 267,000 overall samples. Benin's Communes have a median population of 103,000 with an inter-quartile range of 71,000 to 173,000. Mali's circles have a median population of 308,000 with an inter-quartile range of 197,000 to 520,000. Cameroon's departments have a median population of 456,000 with an inter-quartile range of 225,000 to 907,000. A broadly comparable geographical unit in the US would be commuter zones. To uncover causal place effects previous research [Chetty and Hendren, 2018a, Laliberté, 2021, Deutscher, 2020] has mainly relied on administrative data that is not typically publicly available. This data is necessary for these studies to observe full migration histories and to match child and parent outcomes over long time frames. Such rich data is not needed to estimate place effects, in the censuses I use, it's possible to discern migration histories from cross-sectional evidence.

[9]In Benin, Cameroon, and Mali, education is compulsory for the first six years between the ages of 6 and 11/12 which covers primary school.

of individuals, I define three groups: (i) Permanent residents whose birth location equals their current location. This group covers 81% of 14 to 18-year-olds in my sample. (ii) One-time movers whose birth location does not equal their current location but is their previous location. This group covers 13% of 14 to 18-year-olds in my sample. (iii) Multi-time movers whose birth location doesn't equal their current location or their previous location. This group covers 6% of 14 to 18-year-olds in my sample. Thus, for all but 6% of the sample, I can construct each 14 to 18-year-old individual's full location history, including the age and year of moving. In the online appendix, I discuss the differences and similarities between the three groups.

The data described above comes from the same source as that used in Alesina et al. [2021]. However, whereas Alesina et al. [2021] consider 27 counties in Africa for their main analysis, and 16 countries when they analyze place effects, I focus only on three: Benin, Cameroon, and Mali. This is because I have stricter criteria for inclusion in my sample. I only consider censuses that include information on previous location as well as birth and current location allowing me to trace out entire migration histories of one-time movers. To the best of my understanding Alesina et al. [2021] discern age-moved using only birth and current location coupled with information regarding time spent at the current location. Thus Alesina et al. [2021] measure previous location with error bundling in those who indeed have moved from their birth location with those who may have moved from some intermediate location (32% of all movers in my sample). As Alesina et al. [2021] show, moves are on average to better locations thus this has the potential to introduce non-classical measurement error into movers' change in location quality, and could bias estimates of the importance of place. In addition, I only include countries with multiple cross-sections to enable an analysis of place effects over time.

The Online Appendix provides some descriptive statistics from the census data. I show that the distribution over localities of primary completion rates in Benin, Cameroon, and Mali has shifted over time. There are considerable gains in Cameroon and Benin although intra-national variation remains large, while in Mali the gains have been smaller. Additionally, although most students complete primary education by age 12/13, (in all three countries primary school *should* be completed by this age) a significant minority complete it only at 14 or 15. In the online appendix, I also analyze migration patterns and find that migration in Benin, Cameroon, and Mali is not easily characterizable by to-primate migration, or rural-to-urban migration.

1.3 Education as a measure of opportunity across space in Benin, Cameroon, and Mali

Education is one of the most important dimensions of development[10] and for the majority of individuals in Benin, Cameroon, and Mali, primary education is the most salient margin of variation. In my sample, among those 19 or older 34% have completed primary education and only 7% secondary education. Schooling has also been found to have high Mincerian returns, Psacharopoulos and Patrinos [2018] estimate the return to an additional year of education to be 10.5% in Sub-Saharan Africa — the highest of any region in the world.[11]

Rather than relying on estimates of returns to education from other studies which consider different samples, I can indirectly estimate the returns to education in my data. Although I don't observe wages or wealth, I do observe detailed information about housing quality. In particular, I observe the material walls, floors, and roofs are made from, the level of sanitation, the availability of clean running water, and whether or not the household has access to electricity. From this, I create a composite housing quality index as the standardized first principle component over these variables. Returns to education in each locality are measured as the average housing quality those with primary education have above those without. Figure Ia displays the distribution of returns by country over localities. In each country returns to primary education are on average around half a standard deviation of housing quality. This high mean hides significant variation across localities within a country. In all three countries, there are some locations with very low returns to primary education while in others they are as large as a full standard deviation.

Primary completion is also highly correlated with employment prospects. Figures Ic and Id show the long differenced relationship between locality-level average primary completion and the proportion working in professional services or agriculture respectively. All slopes of the lines of best fit are significant at the 5% level. Both figures tell the same, albeit correlational, story that areas in which primary completion rates most increased saw the greatest transition away from agriculture and towards employment in the professional sector.

At the individual level, there is also evidence linking education and occupation. Figure Ib shows that the proportion of individuals working in agriculture decreases substantially with education completion. Primary education is a key route out of agricultural work (much

[10] To *achieve universal primary education* was one of the Millennium development goals [Bank, 2015] and this aim continues in the updated sustainable development goals [Sachs, 2012]. Millennium development goal 2: *achieve universal primary education* aimed to *ensure that, by 2015, children everywhere, boys and girls alike, will be able to complete a full course of primary schooling*. Sustainable development goal 4: *quality education* aims, among other things to *by 2030, ensure that all girls and boys complete free, equitable, and quality primary and secondary education leading to relevant and Goal-4 effective learning outcomes*.

[11] Similarly Psacharopoulos and Patrinos [2004], Barouni and Broecke [2014], and Schultz [2004] find evidence of high returns to education in Sub-Saharan Africa.

of which is in subsistence agriculture[12]). However, it shouldn't be taken as a panacea. Many individuals with primary education are still employed in agriculture.

Figure I Education is associated with positive later life outcomes

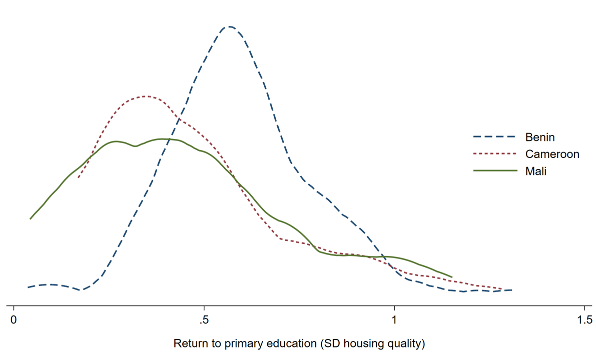

(a) Distribution of returns to education

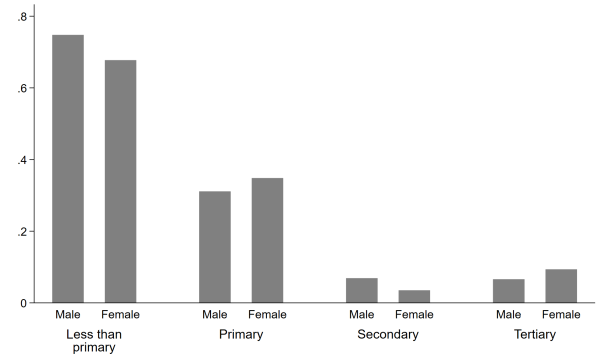

(b) Employment in Agriculture by education and sex

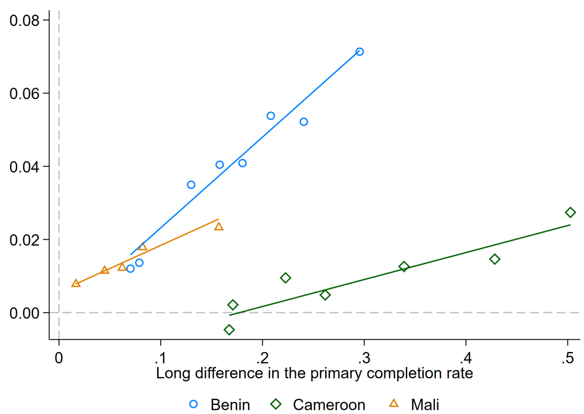

(c) Proportion working professionally

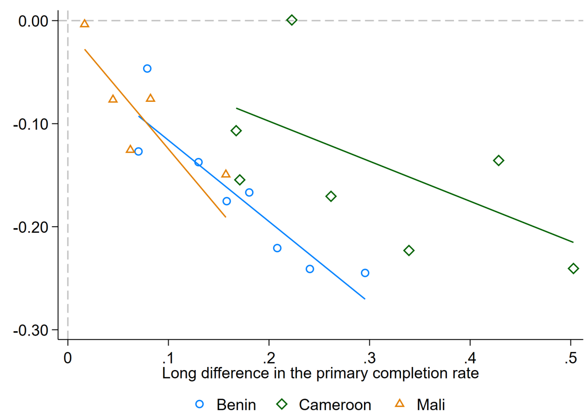

(d) Proportion working in agriculture

Note: Panel (a) shows the distribution of returns to primary education across localities for Benin, Cameroon, and Mali. Returns to primary education are measured as the average housing quality of those with primary education minus the average housing quality of those without primary education. Housing quality is measured as the standardized first principle component over observables describing housing quality including floor, wall, and roof material, sanitation provision, access to running water, and access to electricity. Panel (b) shows the proportion employed in agriculture by education and sex. In both figures, the sample is restricted to include only those of core working age, 25 to 55. Panels (c) and (d) display binscatter plots with corresponding linear regression lines of the association between the locality level change in primary completion on the x-axis vs the change in the proportion employed in a professional capacity in the left-hand panel and the change in the proportion employed in agriculture in the right-hand panel. Each line represents a different regression, the slope of each is significant at the 5% level with robust standard errors. The long difference for Benin is from 1979-2013, for Cameroon from 1976-2005, and for Mali from 1987-2012.

[12] Measures of the proportion of individuals employed in subsistence agriculture are by their very nature difficult to ascertain. Medina et al. [2017] estimates the size of the informal economy (part of which is subsistence agriculture) to be 20-40% of GDP in Cameroon and greater than 40% in Mali and Benin, however, a much larger proportion of the labor force will be employed in the informal economy due to its relatively lower productivity.

In sum, the evidence presented in this section suggests that completing primary education is associated with greater opportunity defined more broadly. Therefore, unequal access to primary educational opportunities across space will be taken as indicative of unequal access across space to opportunities more generally defined. Of course, there will be other dimensions of opportunity that are not captured by primary education completion and notions of inequality of opportunity that go beyond that which is across space. In this paper, I will focus on inequality of opportunity defined across space and in terms of primary education completion as this is a particularly salient margin of inequality of opportunity in this setting.

Primary school completion is also one of the most commonly available, and comparable variables across countries and decades. Due to its universality and ease of recognition, it is less likely to suffer from substantial measurement error. Primary school completion also has the advantage of being a historically determined variable that is observable in a cross-section. Most individuals have completed primary schooling by age 12 to 14, but it continues to be observed no matter what age the individual is at the time data is collected. This allows me to consider the impact of place *after* schooling has been determined constituting a natural placebo test, place cannot causally impact primary school completion after the event.

2 How unequally distributed is opportunity across space in Benin, Cameroon, and Mali?

Figure II shows that the observed distribution of primary education completion over space is highly unequal in Benin, Cameroon, and Mali. However, it's unclear from the raw variation to what extent these discrepancies reflect differing characteristics of individuals living in different locations or the underlying causal impact of place on primary completion rates. The distinction has normative as well as policy implications: inequality in causal effects is inequality of opportunity — simply by growing up in one location rather than another children are more or less likely to complete primary education.

To set ideas consider a simple place effects model.

$$y_i = \sum_{a=1}^{13} \mu_{c(i,a)} + \theta_i \tag{1}$$

Where y_i is a dummy variable taking the value one if an individual i has completed primary education, and $c(i,a)$ denotes the location c individual i resided in at age a. I split the determinants of y into two parts. First, θ_i represents the individual-specific effects such as raw ability which are orthogonal to location effects. Second, $\sum_{a=1}^{13} \mu_{c(i,a)}$ captures the sum of all place effects individual i has been exposed to during childhood, where place effects

Figure II Locality level observed primary completion rates

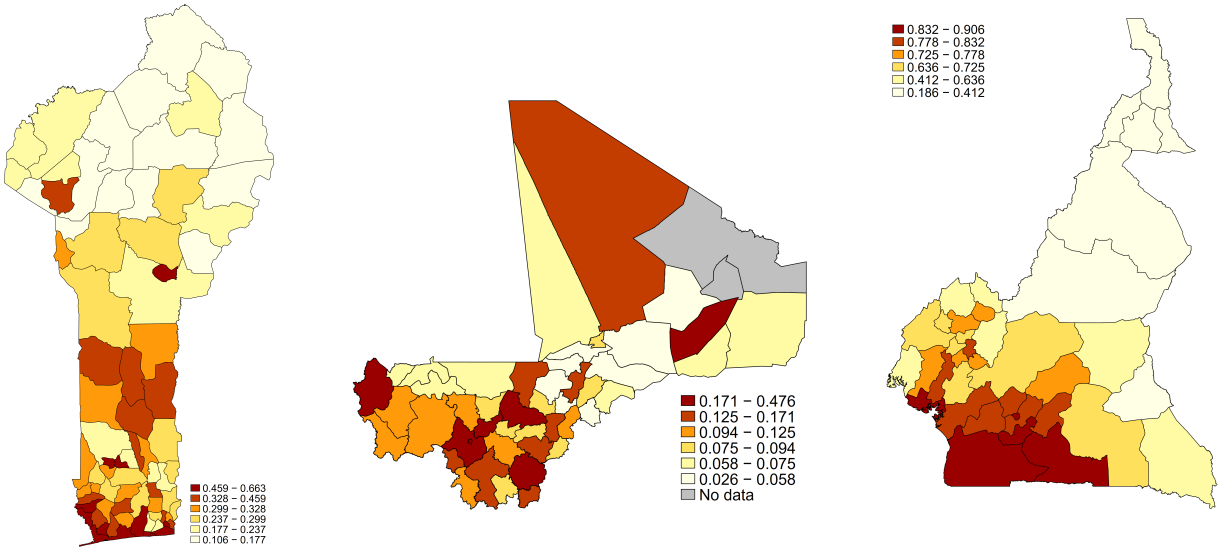

(a) Benin (2013) (b) Mali (2009) (c) Cameroon (2005)

Note: This figure shows the spatial distribution of observed primary education completion rates for those above the age of 12 in Benin, Cameroon, and Mali. Each figure has its own scale and corresponding legend where darker orange/red indicates higher completion rates. The data for Benin comes from the 2013 census, for Mali from the 2009 census, and for Cameroon from the 2005 census.

can be thought of as location-specific fixed effects μ_c. Intuitively, μ_c captures any locality-specific characteristic which may impact y, for example: weather, religio-cultural factors, or the prevailing local returns to education.

This formalization introduces two assumptions. First, I assume that place impacts each individual i in the same manner. I relax this when considering differing effects by gender, but in general, do not probe this assumption. Instead, results reflect the average over potentially heterogeneous effects. Second, I assume additive separability of effects in what is effectively a linear probability model. In the results presented below, I retain this assumption but replicate qualitatively and quantitatively similar results using probit, logit, or pseudo-Poisson approaches in the online appendix.

In the absence of random variation assigning children to locations, to estimate causal place effects I use a movers design building on work by Chetty and Hendren [2018a], Aaronson [1998], Weinhardt [2014], Finkelstein et al. [2016], Molitor [2018] and, Bronnenberg et al. [2012] among others. This approach uses observational variation in the outcomes of children that move across localities at different ages. If children who move to better areas[13] earlier have a higher probability of completing primary school, I conclude that the quality of a

[13]In general *better* can be defined on any dimension of interest, but In this setting, a *better* area is one with a higher primary completion rate among permanent residents, that is those who don't move from their birth location. I consider other definitions of location quality in the online appendix and find similar results.

location matters. However, individuals who move earlier to better places could also have better-unobserved characteristics. To disentangle causal place effects and selection I make, and interrogate, one key assumption: that selection effects do not vary with the age of the child at move.

2.1 Estimating the overall convergence rate

Formalizing the above identification strategy within the place effects framework, I estimate models of the form given in equation 2.

$$y_i = \rho_{os} + \alpha_m + \beta_m \cdot \Delta\bar{y}^p_{ods} + \varepsilon_i \tag{2}$$

Where i denotes an individual, o their origin (birth) location, d their destination (current) location, s their cohort and m the age at which they move. y_i is an indicator variable taking the value one if primary education has been completed. ρ_{os} is a set of origin by cohort fixed effects. α_m is a set of age-at-move fixed effects which flexibly absorbs movement costs. $\Delta\bar{y}^p_{ods} = \bar{y}^p_{ds} - \bar{y}^p_{os}$ is the difference in primary completion rates between the destination and origin locations' permanent residents who are in the same cohort as i and ε_i is an idiosyncratic error term.[14] This equation is estimated on a sample of one-time movers between the ages of 14 and 18. The decision to use a sample of 14 to 18-year-olds is the result of a balancing act, on one hand, I want to be able to match as many children as possible to their parents and can only do so if they cohabit, thus a younger sample is better. On the other hand, I want to capture children after their primary completion status has been determined, so later is better. β_m are the coefficients of interest and can be interpreted as the percentage point impact of moving at age m to a location with a 1 percentage point higher primary completion rate among permanent residents.

We can decompose β_m into the causal effect of place b_m and the selection effect of place δ_m as follows: $\beta_m = b_m + \delta_m$. b_m is the object of interest but cannot, in general, be separately identified from δ_m. To identify b_m, I make the additional assumption that selection effects do not vary over the age at move. This is a potentially strong assumption and, in sub-section 2.2 I relax it, and test its validity, finding considerable evidence in support and thus, for the main analysis, take it as given. Under the identifying assumption, I can estimate the causal effect of spending an additional year of one's childhood at age m in a one percentage point

[14] In the online appendix, I show that the distribution of the difference in location quality among one time movers, that is the distribution of $\Delta\bar{y}^p_{ods}$, is symmetric around a mean close to zero. Implying that it is not the case that all, or even a large majority, of moves, are to better locations. This gives some indicative evidence against large selection effects. I also show that male and female children are similarly likely to move to better or worse locations.

better location as $\beta_{m+1} - \beta_m = b_{m+1} - b_m$. As primary schooling *should*[15] be completed by age 12, it must be that $b_m = 0$ for $m > 12$ and so in this range, β_m estimates the selection effect, $\beta_m = \delta$ for all m greater than 12.

Figure III shows the results from estimating equation 2 where $\hat{\beta}_m$ is plotted against age at move m in grey with corresponding confidence intervals. Standard errors are two-way clustered at the birth locality and current locality level.[16] After age 12, $\hat{\beta}_m$ is constant, providing initial evidence in favor of the identifying assumption, as this shows selection effects are constant for all $m > 12$ when primary completion is already mainly determined. Between ages 1 and 12, I find a positive and increasing (with earlier ages) effect of place which is well approximated by a linear trend (as illustrated by the red overlay). This implies that where children grow up in Benin, Cameroon, and Mali, *does* exert a causal impact on their probability of completing primary education. Additionally, the linearity of the effect implies that there are no tipping points — or particularly important years [Cunha and Heckman, 2008]. Each additional year spent in a better location exerts roughly the same positive impact as any other before age 13.

Due to the linearity apparent in causal place effects, I consider the pre-13 slope as the single key parameter. This slope, of -0.027 (0.04), is interpreted as the rate of convergence. That is, for each additional year spent in a new location, how *fast* do movers converge to permanent residents in that location? As a robustness test, in the online appendix, I estimate a parametric analog to equation 2 allowing a different slope and intercept for children moving between the ages of 1 and 12 and those moving after age 12. Over various fixed effects specifications, this recovers a very similar convergence rate and a precisely estimated zero slope for post-12 moves.

Figure III is analogous to figure 4 in Chetty and Hendren [2018a] and figure 11 in Alesina et al. [2021]. It recovers remarkably similar results to these studies which consider different settings, the US in Chetty and Hendren [2018a], or different outcome variables in a larger sample in Alesina et al. [2021]. Chetty and Hendren [2018a] find a convergence rate of about 4% and Alesina et al. [2021] find a rate between 1.5% and 3%. Relative to these studies, the main contribution of this paper lies in estimating and analyzing locality-year specific place effects which I turn to in section 3. Analyzing locality-year specific place effects allows me to quantify and characterize the changing spatial distribution of opportunity and study the ramifications of unequal opportunity across space on gender equality.

[15] In each of the countries studied and over the whole sample period, primary school is expected to end at age 12. The online appendix documents that although many individuals complete primary school after age 12 by age 14 most who will complete have done so.

[16] In the online appendix I discuss inference in this setting and provide permutation-based inference results which affirm the conclusions reached here.

Figure III Plotting $\hat{\beta}$ from estimating equation 2

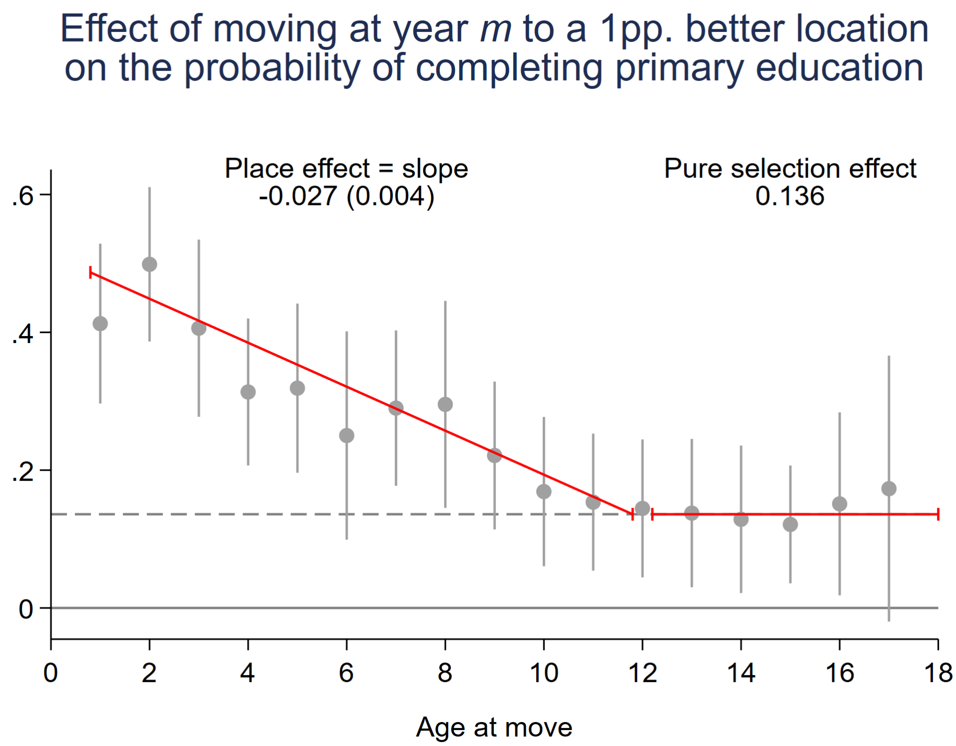

Note: This figure shows the estimated β_m coefficients from the regression: $y_i = \rho_{os} + \alpha_m + \beta_m \cdot \Delta\bar{y}^p_{ods} + \varepsilon_i$ estimated on a sample of one time movers between the ages of 14 and 18 for each available census in Benin, Cameroon, and Mali. Corresponding 95% confidence intervals are indicated by gray bars where standard errors have been clustered at the birth locality and current locality level. Superimposed on this in red are two lines of best fit with a break at age 12. The horizontal gray dashed line indicates the post 12 mean of $\hat{\beta}_m$.

2.2 Validity of the movers design

I perform five validation exercises of the identifying assumption that selection effects do not vary with the age of a child when the family decides to move. They are described in detail in the appendix B and summarised here.

First, there is no evidence of varying selection effects on mothers' observed primary education. Mothers with primary education are more likely to move to better areas, but they do not systematically do this to a greater or lesser extent with respect to the age of their child at move. Second, results are robust to including household fixed effects and therefore only considering cross-sibling age variation. By including household fixed effects, I shut down any time-invariant dimensions of possible selection effects, for example, socioeconomic bracket, the child's revealed pre-school potential academic ability or parents' attitudes towards education.

Third, I restrict the sample to only include increasingly likely to be exogenous move events and show that results remain stable. To do this, I define statistically increasingly unusually large out-move events from localities and restrict the sample to include only increasingly unlikely events — capturing an increasing proportion of exogenous moves. I find that the estimate of the causal effect of place is stable as I continue to restrict the sample, providing evidence for the identifying assumption.

Fourth, I perform an over-identification test using variation across gender in locality quality. This exercise utilizes differences in predictions from the place effects vs selection effects explanations. The place effects model postulates that individuals should converge to outcomes of permanent residents in their gender-specific sub-group in the destination location. Whereas, to the extent that households cannot observe gender-specific outcomes, selection effects should only operate at the aggregate level. By including own-gender-specific location quality (average primary completion rate of permanent residents with the same reported gender as individual i) and *other* gender quality in one regression, I can explicitly test the predictions of the two possible scenarios. If place effects are the driving force behind the observed effects, then individuals will only react to the outcomes of children of their own gender. Whereas, if selection effects were predominant, individuals will react equally to the outcomes of children of both sexes. When performing this test, the coefficient on own gender remains positive and significant while that on other gender attenuates to zero. Giving further evidence in support of the place effects explanation.

In sum, any omitted variable bias or selection effects would have to operate within families proportionally to exposure times, persist in the presence of moves induces by displacement shocks, and precisely replicate permanent residents' outcomes by gender in proportion to exposure times in addition to not being observed on mothers education. I find it unlikely that any qualitatively or quantitatively important threats to identification operate in this manner.

In the online appendix, I also discuss the generalisability of my results which are based on a sample of movers to the whole population. The assumption that place effects are the same for movers and stayers can be decomposed into two components. First, that movers and stayers do not systematically differ in a manner that interacts with place effects [Lagakos, 2020, Hamory et al., 2021]. Second, movers and stayers could be similar but due to the nature of having moved to a new location, movers may interact differently with their surroundings. In the online appendix, I describe and provide some evidence against both of these dimensions by considering increasingly likely to be exogenous moves and replacing origin fixed effects with destination fixed effects.

3 Estimating locality-year specific place effects

To estimate locality-year specific effects, I generalize the specification presented in equation 2 by estimating equation 3 following Chetty and Hendren [2018b]:

$$y_i = \alpha_{odt} + \mu_{lt} \cdot e_{il} + \varepsilon_i \tag{3}$$

Where $l \in \mathcal{L}$ is a location, $o \in \mathcal{L}$ is the birth location of individual i, $d \in \mathcal{L}$ is i's destination location and t is the sample year in which i is observed.[17] α_{odt} are origin by destination by sample year fixed effects and ε_i is an idiosyncratic error term. Finally, e_{il} is a variable equal to the years of childhood (1 to 13) i spends in location l and is given by equation 4.[18]

$$e_{il} = \begin{cases} 13 - m_i & \text{if } l = d(i) \\ m_i & \text{if } l = o(i) \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

The causal place effects, $\{\mu_{lt}\}$, can only be recovered in a relative sense as I can only compare across localities within a given census. Intuitively this approach considers all individuals who move from a given origin o to a given destination d in a given census year t and compares the outcomes of those who move earlier relative to later. If it is the case that those who moved from o to d earlier have better outcomes (are more likely to complete primary education) I conclude that $\mu_{dt} > \mu_{ot}$. I then combine information from all such comparisons to estimate each locality's causal place effect. By including origin-destination-year fixed effects, α_{odt}, I only use variation in the age of a child at move rather than comparing the outcomes of families that move to/ from different areas. μ_{lt} can therefore be interpreted as the causal effect of spending an additional year of childhood in location l in a given year t, on primary completion rates relative to the average location.

Figure IV shows the recovered causal place effects in each country-year sample. It shows considerable variation across space within each country. Across all censuses, I find that the average signal variance[19] of opportunity across localities is $\sigma^2_\mu = 1.2$ percentage points per year. This means that on average moving to a location with one standard deviation higher opportunity at birth causes a 15.0 percentage point (1.2pp. a year times 13 years of

[17] t is defined as the year the census is taken, however, it can be thought of as capturing place effects across the childhoods of the 14 to 18-year-olds in the sample at time t.

[18] This formulation assumes linearity of place effects, an assumption that is made based on the evidence provided in section 2.1.

[19] From the set of estimated place effects I can back out an estimate of $\sigma_{\hat{\mu}}$ but in actuality, we're interested in σ_μ, the signal standard deviation. As place effects are measured with error I can write $\hat{\mu}_{lt} = \mu_{lt} + \eta_{lt}$, and then find the signal standard deviation as follows: $\sigma^2_\mu = \sigma^2_{\hat{\mu}} - \sigma^2_\eta$. The standard deviation of the error, σ_η is estimated as the average standard error of $\hat{\mu}_{lt}$ denoted as s_{lt}. Then $\sigma_\eta = \sqrt{\mathbb{E}[s^2_{lt}]}$, and so an estimate of σ_μ can be recovered. Details of this procedure are presented in the appendix section C.

childhood) increase in the probability a child completes primary education. Figure IV shows significant variation in all country-year cells with the exception of Cameroon in 2005 which has a considerably tighter distribution. This is because, in contrast to the other countries and other years, spatial variation in observed primary completion rates in Cameroon in 2005 is relatively small due to a rightward shift in the distribution (and obvious truncation at 1). This doesn't mean that opportunity is equally distributed in Cameroon in 2005, in the online appendix I show that considerable variation remains in the causal effect of place on secondary school completion in Cameroon in 2005.

Figure IV Distribution of recovered causal place effects

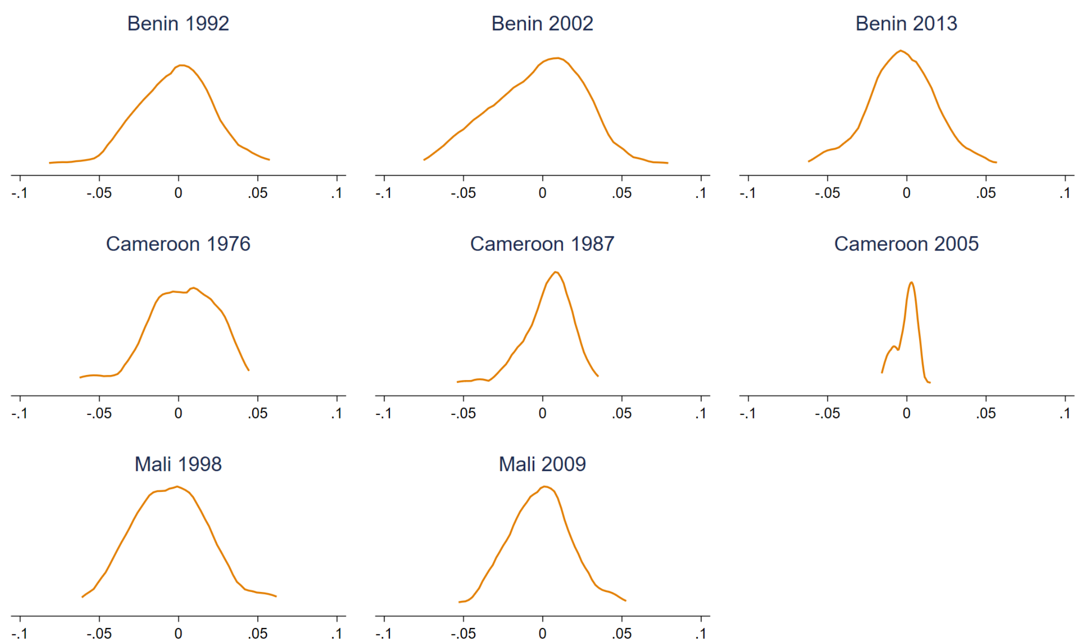

Note: This figure shows the distribution of recovered causal place effects in each census. Causal place effects are estimated from equation 3 and plotted relative to a 0 baseline. The x-axis is measured in yearly percentage-point effects, that is the impact of moving to a given location a year earlier on a child's probability of completing primary school.

Table I reports the total and signal variance of causal place effects in each country-year cell. I do not report estimates for Cameroon in 2005 for the reasons described above. Table I confirms what figure IV shows graphically, that the considerable average inequality of opportunity uncovered is not restricted to a specific country or specific time period. That being said, there are considerable differences in the signal variance of place effects across censuses. Table I also reports the within-country cross-census correlation. Although the average within-country cross-time correlation is positive, this is not the case for each consecutive period. In Benin locations with high place effects in 1992 on average had higher place effects in 2002 but locations with high place effects in 2002 had on average lower place effects in 2013. In Cameroon, the correlation between locations' place effects in 1976 and 1986 is high, whereas in Mali the relationship between 1998 and 2009 remains positive but considerably weaker. The online appendix provides maps of the resulting estimated spatial distribution of opportunity.

Table I Quantifying inequality of opportunity

	Benin			Cameroon		Mali	
	1992	2002	2013	1976	1986	1998	2009
Total Variance	4.80	7.01	3.76	3.94	2.67	5.25	2.51
Signal Variance	0.72	4.64	1.75	2.04	1.47	2.74	1.21
Correlation with previous census		0.29	-0.23		0.62		0.09

Note: This table reports statistics from causal place effects recovered from equation 3 across censuses, that is country-year pairs. The first row reports the total variance measured per year of exposure. The second row reports the signal variance measured per year of exposure. The third row reports the within-country cross-year correlation between estimated effects.

4 Characteristics of place effects

Previous work has considered various correlates of observed primary completion (or educational inter-generational mobility) rates in Sub-Saharan Africa (e.g. Majgaard and Mingat [2012], Alesina et al. [2020], Shafiq et al. [2019], Shifa and Leibbrandt [2022]). These correlations could reflect sorting or they may be predictive of a characteristics' correlation with the causal effect a place has on individuals. In this paper, I can ask whether such known relationships indicate that an area with such a characteristic is associated with more opportunity or simply attracts people more likely to complete primary education anyway. To do this, I use the estimated place effects, $\hat{\mu}_{lt}$, as the dependent variable in regressions with locality-level covariates on the right-hand side. Specifically, I decompose the observed location-year primary completion rate, $\bar{y}^p_{lt}$, into an estimated causal effect of place, $\hat{\mu}_{lt}$, and a selection effect, $\bar{\theta}_{lt}$.

$$\bar{y}^p_{lt} = \alpha_{census} + 13 \cdot \hat{\mu}_{lt} + \bar{\theta}_{lt} + v_{lt} \tag{5}$$

Where α_{census} are census fixed effects, and v_{lt} is an idiosyncratic error term. Consider some explanatory variables X_{lt} then suppose the following relationships hold where β_μ and β_θ are the objects of interest and $\varepsilon^1_{lt}, \varepsilon^2_{lt}, \varepsilon^3_{lt}$ are idiosyncractic error terms.

$$
\begin{aligned}
13 \cdot \hat{\mu}_{lt} &= \beta_\mu \cdot X_{lt} + \varepsilon^1_{lt} \\
\bar{\theta}_{lt} &= \beta_\theta \cdot X_{lt} + \varepsilon^2_{lt} \\
\upsilon_{lt} &= \varepsilon^3_{lt}
\end{aligned}
$$

As $\bar{\theta}_{lt}$ is not observed we cannot directly estimate β_θ. Instead, I combine the above terms with equation 5 and estimate β from equation 6.

$$
\bar{y}^p_{lt} = \alpha_{census} + \underbrace{(\beta_\mu + \beta_\theta)}_{\beta} \cdot X_{lt} + \underbrace{\varepsilon^1_{lt} + \varepsilon^2_{lt} + \varepsilon^3_{lt}}_{\varepsilon_{lt}} \tag{6}
$$

From this, I can then back out β_θ. Note that β_θ or β_μ can only be interpreted as capturing correlations. That is, a significant positive estimated $\hat{\beta}_\mu$ implies only that a characteristic of a location is correlated with higher local opportunity, not that it causes greater opportunity in a location. When estimating β_μ I use an outcome variable (the estimated causal place effects $\widehat{\mu}_{lt}$) which is observed with error. However, classical measurement error in the left hand side variable will not impact estimates of β_μ and thus is not a concern.

Figure VIII plots the impact of standardized locality characteristics on observed primary completion rates ($\hat{\beta}$ reported in orange) and causal place effects ($\hat{\beta}_\mu$ reported in blue). Greater urbanization, greater access to electricity, better housing quality, and perhaps surprisingly a greater proportion of single-mother households, are all associated with higher observed primary completion rates, and causal place effects. Similarly, a lack of sanitation, a higher proportion employed in agriculture, and a higher proportion of Muslims, are all negatively associated with both observed primary completion rates and the causal effect of place. The proportion of individuals born abroad has little impact. In blue, causal place effects share the same association, but to a lesser degree. Percentages below characteristic names indicate the proportion of the overall variation that operates through the causal effect channel. For example, a one standard deviation higher urbanization rate is correlated with an 11.3 percentage point higher observed primary completion rate and a 5.7 percentage point higher causal effect of place on primary completion. The online appendix reports results from a multi-variate version of these regressions.

For most characteristics, at least half of the correlation between them and primary completion rates can therefore be attributed to selection or the differing characteristics of individuals across space. This implies that such intuitively known correlations like "education rates are higher in urban areas" are picking up significant effects that operate through the selection channel rather than purely correlations with the underlying distribution of local opportunity.

Figure V Correlates with observed primary completion rates and causal place effects

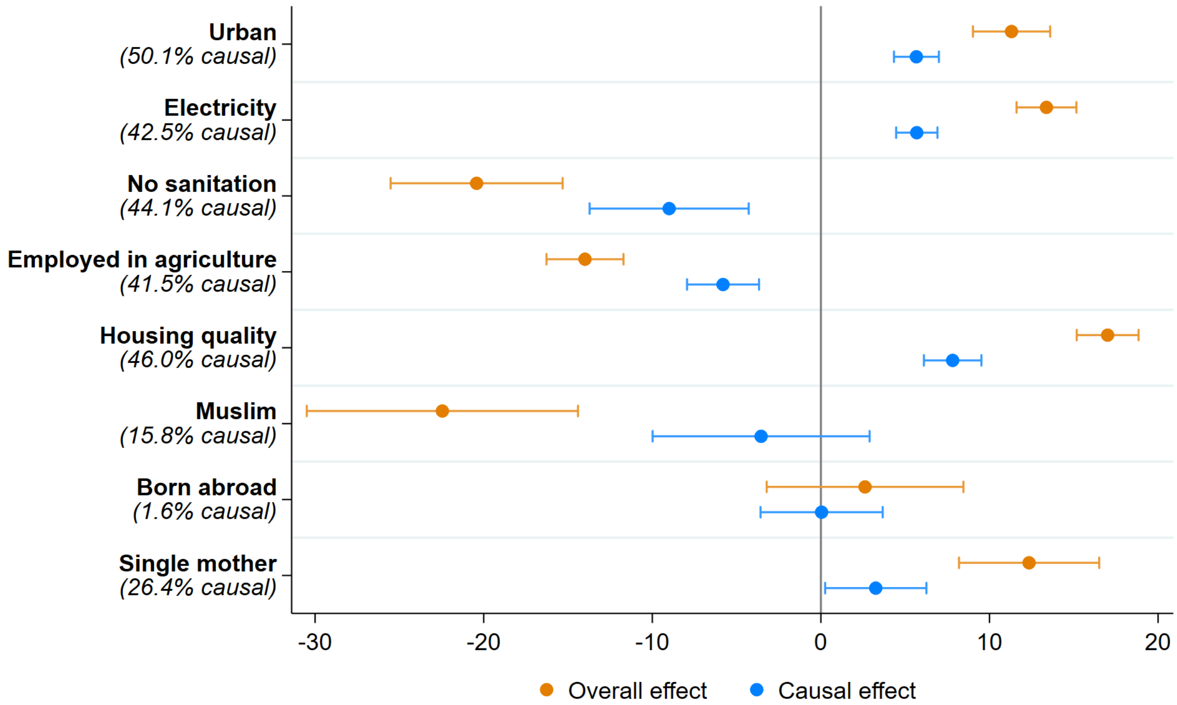

Note: This figure shows the results from running regressions of observed permanent resident primary completion rates and estimated causal place effects against place characteristics including census fixed effects. Each dot (with corresponding confidence intervals) represents a coefficient from an individual regression with the explanatory variable indicated on the y-axis. In orange, the dependent variable is observed permanent residents' primary completion, and in blue it is estimated causal place effects multiplied by thirteen. In italics, beneath characteristic labels, the proportion of the total correlation that operates through the causal place effects channel is noted. Standard errors are clustered at the locality level. Regressions are weighted by locality population.

I can also consider heterogeneity in characteristics across the three countries I study. In the online appendix, I present a figure analogous to figure VIII allowing each coefficient to vary by country. Although the broad qualitative conclusions from the pooled analysis hold, important differences arise. I find that, whereas the causal impact is very close to the overall impact for all variables in Mali, in Benin it appears that the observed correlations operate significantly less through the place effects channel. In Mali previous research on determinants of local primary completion operate mainly through the causal effects channel, whereas in Benin much of the relationship may instead be driven by selection effects.

Of the place-level characteristics considered in figure VIII, the proportion of Muslims in a location has the strongest relationship with observed primary completion rates, but one of the weakest with causal place effects. In Benin, Cameroon, and Mali, as well as many other countries across the world, religion, and education intersect [Auriol et al., 2020, Panin et al., 2021, Stoop et al., 2019]. The reasons behind known large differences in primary completion rates between those following different religions are often attributed to identity and legitimacy Manglos-Weber [2017], culture and persistence from colonization Platas [2018], the role of religious schools Tsimpo and Wodon [2014], Boyle [2014], Izama [2014], or religious violence Bertoni et al. [2019]. Armed with locality-specific causal place effects in figure VIII I go further than previous work and find that the spatial relationship between observed local primary schooling outcomes and the predominant local religion operates mainly through the selection effects channel. Table II in the appendix shows that the relationship between religion and causal place effects is robust, and considers heterogeneity by country.[20] In Benin, the association is much weaker and mainly driven by place effects. On the other hand, in Cameroon, the association is stronger and it appears that place effects have a small, but not insignificant, role to play.

In sum, these results show that the previously observed relationship between development/ opportunity and religion is more nuanced than broad cross-country averages allow. Although the predominant religion in a locality does impact the opportunities children growing up experience, observational correlations across space will significantly overestimate the role of place, and averages across countries will obscure significant cross-country heterogeneity.

5 Heterogeneity by gender and the impact of place on gender inequality of opportunity

In my sample, 39.82% of males complete primary school but only 25.68% of females do. Whereas there have been aggregate increases in education completion, the gender gap has remained stubbornly high, in Benin (1992-2013) increasing from 17 to 20 percentage points, in Cameroon (1976-2005) decreasing from 21 to 12 percentage points and in Mali (1998-2009) increasing from 6 to 8 percentage points. There is a large and well-developed literature studying the importance of female education in Sub-Saharan Africa, and the reasons behind this large gender education gap. Among the most common reasons cited are cultural norms, early marriage, the disproportionate burden of household work, teacher attitudes, teenage pregnancy, and religion [Ombati and Ombati, 2012, Baten et al., 2020, Tuwor and Sossou, 2008, Anderson, 2005, Eloundou-Enyegue, 2004, Johannes and Noula, 2011, Aikman and Unterhalter, 2005, Cooray and Potrafke, 2011]. However, little attention has been paid to the potential role of place.

[20] Information on religion is only recorded in 2005 in Cameroon and 2009 in Mali and all years in Benin, but I will focus on 2013. Therefore, for this analysis, I restrict my sample to include only the most recent census waves for which I have data. In Benin and Cameroon 25.8% and 13.9% follow Islam respectively and 46.8% and 77.7% are Christians. In Mali, however, there is much less variation with 94.9% Muslim.

In this section, I ask: does the locality in which a child grows up play as large a role in determining male as female outcomes? In addition, I consider which characteristics of locations are associated with relatively greater opportunity for females vis-à-vis males, and therefore contribute to narrowing the gender gap in primary education completion. To answer these questions I replicate the analysis in section 2 estimating overall and locality-specific causal place effects for each gender separately.

Figure VI replicates the analysis in figure III for each reported gender individually. I find that females are more responsive to locality than males, moving to a one standard deviation better location at birth causes females to increase the probability of completing primary school by almost 8pp, whereas for males the completion rate only increases by about 4pp, a statistically significant difference. This means that if a child moves at birth, a female child would converge completely to the outcomes of permanent residents in the destination location whereas a male child would only converge 50%.

Figure VI Plotting $\hat{\beta}^g$ from estimating equation 2 splitting the sample for each gender

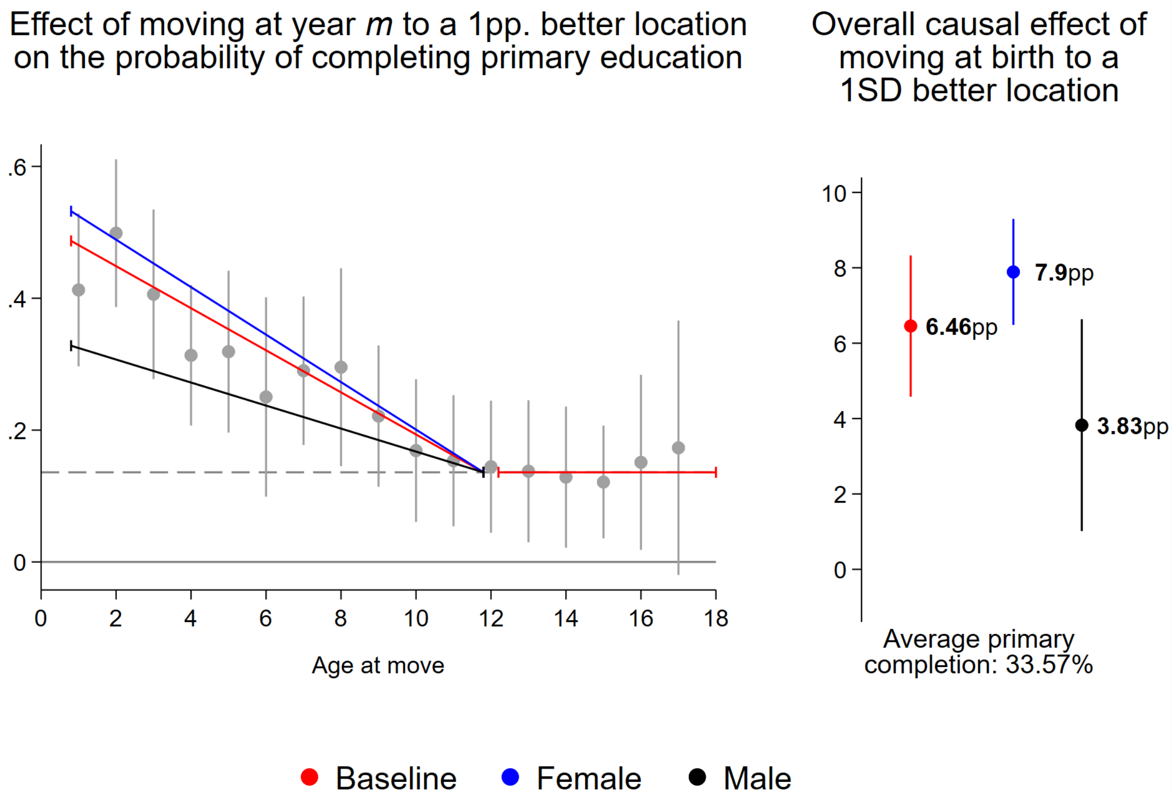

Note: The left-hand side panel of this figure shows the estimated β_m^g coefficients from the regression: $y_i = \rho_{os} + \alpha_m + \beta_m^g \cdot \Delta\bar{y}_{ods}^p + \varepsilon_i$ estimated on a sample of one time movers between the ages of 14 and 18 for each available census in Benin, Cameroon, and Mali, constricting the sample to gender g. Corresponding 95% confidence intervals are indicated by gray bars where standard errors have been clustered at the birth locality and current locality level. Superimposed on this in red are two lines of best fit with a break at age 12 and in blue and black the 1 to 12 lines of best fit for females and males respectively. The horizontal gray dashed line indicates the post 12 mean of $\hat{\beta}_m$. In the right-hand side panel, the slope of the 1 to 12 line of best fit has been translated into the overall impact of moving at birth to a one standard deviation better location where over the whole sample a one standard deviation improvement in location quality is equal to a 20pp increase in primary completion rates.

The finding that place matters more for female than male outcomes, opens up the possibility that place may have a role to play in determining the gender primary completion gap. To study this I replicate the analysis in subsection 3 on each gender-specific sub-sample. This procedure recovers two place-effect estimates for each locality-year cell, one for males and one for females. Each estimate is relative to the mean gender-specific causal place effects within that census. That is, a high female-specific causal place effect estimate in Benin in 2013 means that growing up in this location increases the probability that females complete primary school relative to females growing up in other localities in Benin in 2013. This does not allow a direct comparison between the level of female and male place effects as gender-specific means are not identified, however, I can consider whether characteristics of locations affect relative gender effects differently.

First, I consider whether places that are relatively good for males are also relatively good for females by estimating the following relationship: $\mu_{lt}^{f} = \beta \cdot \mu_{lt}^{m} + \gamma_{census} + \varepsilon_{lt}$, where γ_{census} are census fixed effects. Which recovers $\hat{\beta}$ =0.36 (0.06). This significant and positive coefficient indicates that places that have relatively more opportunity for males also hold relatively more opportunity for females. However, the coefficient is significantly less than one, so although there is a strong correlation, it's far from a perfect relationship, leaving scope for locality characteristics to be differently correlated with gender-specific causal place effects.

In figure VII, I consider whether the locality characteristics explored in section 4 differently impact male and female relative place effects. I run regressions of the following form: $13 \cdot (\hat{\mu}_{lt}^{m} - \hat{\mu}_{lt}^{f}) = \beta \cdot X_{lt} + \gamma_{census} + \varepsilon_{lt}$. A β coefficient greater than 0 indicates that a given characteristic is more strongly correlated with a high male relative to female causal place effect. For example, figure VII shows that the coefficient on urban is around two, and statistically different from zero. This means that a locality with one standard deviation higher urbanization rate is associated with opportunity for males 2 percentage points higher than that for females. Therefore, locations with higher urbanization rates are correlated with greater within-location gender inequality of opportunity. Figure VII shows a similar story for electricity availability and housing quality, but the opposite story for agricultural employment, and the proportion of the population who are Muslim. That is, if a larger proportion of the population is employed in agriculture, this is associated with better relative causal place effects for females vis-à-vis males and thus indicates that agricultural employment is associated with lower gender differences in opportunity across space.

Variables most commonly associated with development such as the urbanization rate, electricity provision, and housing quality, are found to be associated with worse gender equality of opportunity. This cautions against a uni-dimensional view of development and suggests that care needs to be taken for policy aimed at increasing development not also increasing gender inequality. This result is reminiscent of the U-shaped female labor supply development curve [Goldin, 1994, Ngai et al., 2022] whereby as a country moves away from agriculture female labor supply first decreases before increasing with the rise in services.

Figure VII Correlates with $\mu_{lt}^m - \mu_{lt}^f$

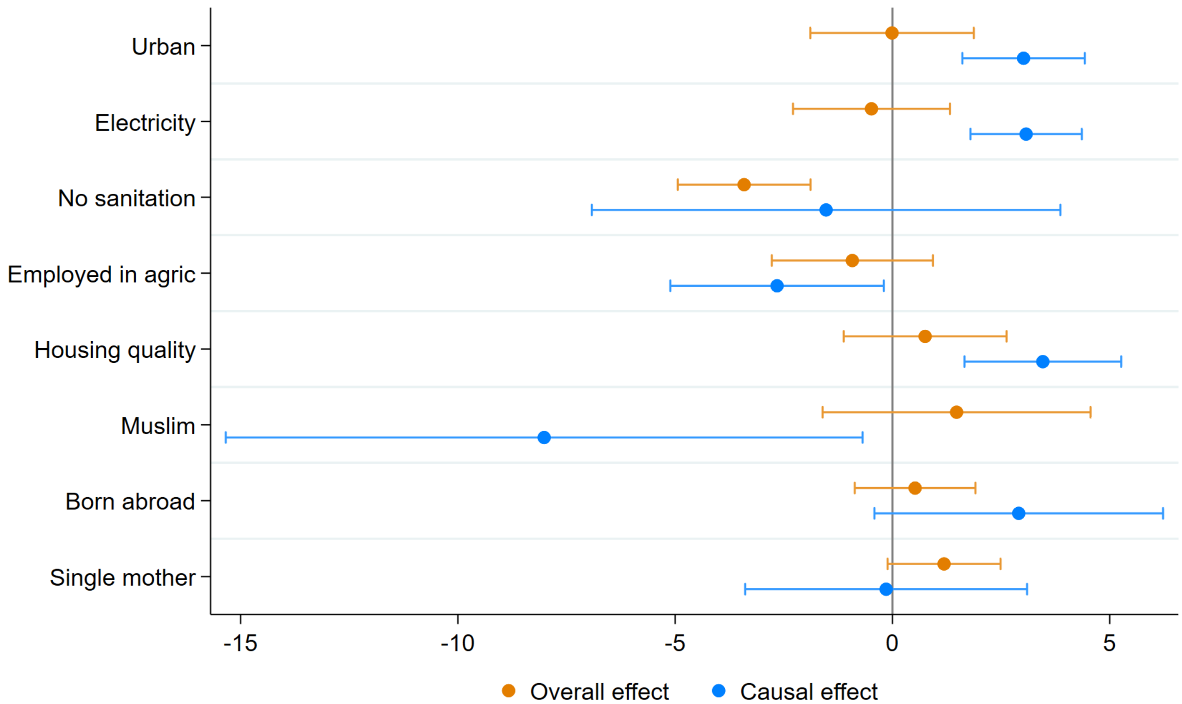

Note: This figure shows the results from running regressions of the difference between relative male and female causal place effects against various place-level characteristics. Regressions are weighted by population, include country-year fixed effects, and standard errors are clustered at the locality level.

Characteristics may have differing impacts on the inequality of opportunity by gender across each of the countries studied. In the online appendix, I present an analogous figure to VII where I allow each coefficient to vary by country. These results show broad similarities across countries but also highlight some areas of difference. Urbanization and electricity provision have a similar effect, as to increase gender inequality, in Benin, Cameroon, and Mali. Sanitation appears to have little impact in all three countries. Agricultural employment and housing quality display some differences. For both, Cameroon displays a weaker relationship than either Benin or Mali — however, power concerns preclude any strong conclusions.

Finally, in section 4 I found that although the predominant local religion is highly predictive of observed primary completion rates, it has little correlation with average causal place effects. Figure VII shows that this average null effect hides considerable gender heterogeneity. Perhaps surprisingly, the negative coefficient suggests that a greater Muslim population is relatively better for female vis-à-vis male opportunity in a region. In table III in the appendix I consider robustness of this result and heterogeneity by country. I find that this effect is driven by Benin, with Cameroon showing little impact and the estimate for Mali being imprecise no doubt due to the lack of religious variation in Mali.

In this section, I have shown that place matters more for female than male outcomes and that although place effects are correlated across genders, some locality characteristics are associated with exacerbating or reducing the gender gap in opportunity. Spatial inequality has a role to play in reducing gender inequality, however as variables commonly associated with development are also associated with increasing gender inequality, future research is needed to identify how place-based interventions may best be harnessed by policymakers to reduce gender imbalances whilst promoting development.

6 Conclusion

In this paper, I have quantified and explored the impact of spatial inequality of opportunity relating to primary school completion in Benin, Cameroon, and Mali using a movers design. I find that moving to a better location at birth has a significant impact on a child's likelihood of completing primary school. Specifically, moving to a one standard deviation better location at birth increases the probability of completing primary school by 15 percentage points. Gender is an essential factor in this analysis, with place having a more significant effect on girls than boys. Furthermore, local characteristics such as urbanization, access to electricity, and employment in agriculture are associated with place effects.

Religion also plays a role in the educational outcomes of children in these countries. Although it is weakly correlated with average causal place effects, it is strongly correlated with the gender gap in the effect of place on educational attainment. These findings highlight the importance of considering both gender and local characteristics in addressing spatial inequality of opportunity and improving access to education. This work illustrates the need for, and potential scope of, place-based policies in reducing inequality of opportunity. In future work, it will be crucial to understand the underlying causes of this variability, and arm policymakers with the tools to mediate spatial inequality of opportunity most effectively.

APPENDIX

A Additional results

Figure VIII Correlates with observed primary completion rates and causal place effects: Multivariate regression

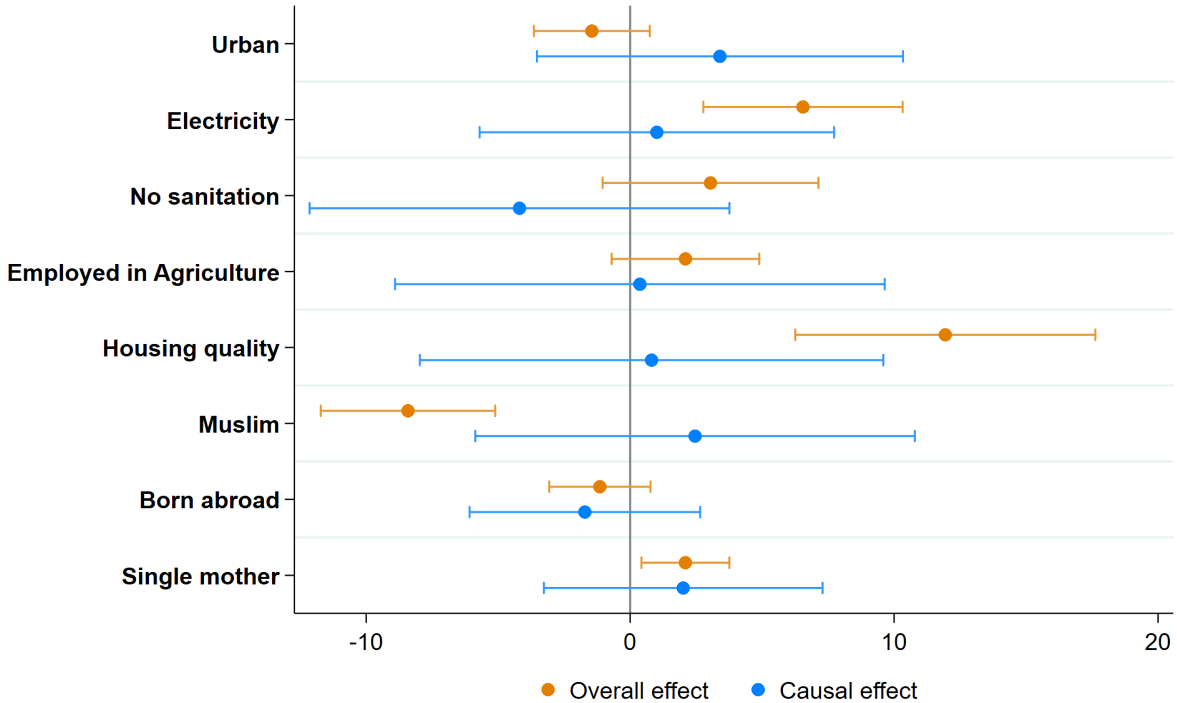

Note: This figure shows the results from running regressions of observed permanent resident primary completion rates and estimated causal place effects against place characteristics including census fixed effects. Each dot (with corresponding confidence intervals) represents a coefficient from a single multivariate regression of all characteristics. In orange, the dependent variable is observed permanent residents' primary completion, and in blue it is estimated causal place effects multiplied by thirteen. Regressions are weighted by locality population.

Table II Impact of the proportion following Islam on observed primary completion rates and causal place effects (CPE).

	(1) Primary comp.	(2) CPE	(3) Primary comp.	(4) CPE	(5) Primary comp.	(6) CPE
Prop Muslim	-0.519***	-0.123	-0.562***	-0.207***		
	(0.0906)	(0.0772)	(0.0924)	(0.0658)		
Prop Muslim x Benin					-0.257***	-0.252*
					(0.0664)	(0.134)
Prop Muslim x Cameroon					-0.795***	-0.172***
					(0.172)	(0.0499)
Drop Mali			X	X	X	X
R^2	0.641	0.0356	0.514	0.126	0.584	0.129
N	146	146	102	102	102	102

Note: This table shows the results from regressing either observed primary completion rates or estimated causal place effects against the proportion who indicate that they are Muslim in a locality. Regressions include country fixed effects, are weighted by locality population, and standard errors are robust. In Columns one, three, and five the dependent variable is the observed primary completion rate, in columns two, four, and six it is the estimated causal place effect. In columns, three to six Mali has been dropped from the sample. Due to data constraints I only use the 2005 census in Cameroon, the 2009 census in Mali, and the 2013 census in Benin when estimating these regressions.

Table III Impact of the proportion following Islam on observed gender differences in primary completion rates and causal place effects (CPE).

	(1) Male - Female CPE	(2) diff_prim	(3) Male - Female CPE	(4) diff_prim	(5) Male - Female CPE	(6) diff_prim
Prop Muslim	-0.192**	0.0354	-0.185**	0.0344		
	(0.0888)	(0.0374)	(0.0899)	(0.0383)		
Benin × Prop Muslim					-0.372*	-0.165***
					(0.191)	(0.0223)
Cameroon × Prop Muslim					-0.0517	0.176***
					(0.0578)	(0.0563)
Mali × Prop Muslim					-0.495	0.0785
					(0.475)	(0.0662)
Drop Mali			X	X		
R^2	0.0315	0.270	0.0444	0.212	0.0504	0.506
N	137	137	96	96	137	137

Note: This table shows the results from regressing either the gender difference in primary completion rates or estimated causal place effects against the proportion who indicate that they are Muslim in a locality. Regressions include country fixed effects, are weighted by locality population, and standard errors are robust. In Columns one, three, and five the dependent variable is the observed gender difference in primary completion rates, in columns two, four, and six it is the estimated gender difference in causal place effects. In columns, three to six Mali has been dropped from the sample. Due to data constraints I only use the 2005 census in Cameroon, the 2009 census in Mali, and the 2013 census in Benin when estimating these regressions.

B Validating Causal Place effects

The identifying assumption underpinning the above analysis, that selection effects into place do not vary with the age of moving, is potentially strong. For example, this prohibits those with better unobservables from systematically moving to better areas earlier. However, it doesn't preclude selection overall. I allow $\delta > 0$, that is for those with better characteristics to move to better places, just that this does not vary systematically with the child's age at move. This is also a common identifying assumption used in the literature [Chetty et al., 2016, Chetty and Hendren, 2018a,b, Chyn, 2018, Damm and Dustmann, 2014, Deutscher, 2020, Alesina et al., 2021, Laliberté, 2021].

I interrogate this assumption in four ways. Firstly, I consider whether it holds for characteristics of individuals we can observe. Secondly, I relax the assumption by including household fixed effects absorbing any time-invariant selection and only using cross-sibling variation in the age at move. Thirdly, I consider a sample of increasingly likely to be exogenous moves by looking at increasingly unusual migration periods. Lastly, I consider an over-identification outcome-based test where I show that individuals respond to their own-groups outcomes in line with the place effects, rather than selection effects, story. In all four approaches, I find evidence to support the identifying assumption. This means that any omitted variable bias or selection effects would have to operate within families proportionally to exposure times, persist in the presence of moves induces by displacement shocks, and precisely replicate permanent residents outcomes by gender in proportion to exposure times in addition to not being observed in observable dimensions such as mothers education. I find it unlikely that any qualitatively or quantitatively important threats to identification operate in this manner.

B.1 Is there evidence of varying selection effects on observables

For characteristics of individuals that are observed, I can directly consider the identifying assumption. Of course, not finding selection on some dimension doesn't preclude selection on another potentially unobserved dimension. Due to the cross-sectional nature of my data, I do not observe time-varying characteristics of individuals, however, as I can construct households I can consider selection on time-invariant characteristics of the household such as parents' education level. To investigate selection I first partition the distribution of observed differences in location quality due to moves, that is $\Delta \bar{y}^p_{ods}$ into below-average moves and above-average moves. I then compare the difference in the average quality of individuals between those who made *good* vs *bad* moves at each age at move. If I find that overall those who made better moves had better observables this is evidence of selection. If I additionally find evidence of systematic variation in this relationship by age at move - this would be evidence against the identifying assumption of constant selection effects.

Figure IX shows the results from performing this kind of analysis where we consider the individual mothers' primary completion rate as our observable measure of individual quality. Consistent with figures in the main text I find evidence of significant positive selection effects overall. In addition, and also consistent with evidence from post-12 moves I find no slope in selection effects over child's age at move. That is, I do not find evidence that individuals with more educated mothers systematically move to better areas earlier or later. On this dimension at least, there is no evidence of selection effects varying by age at move.

Figure IX Children with more educated mothers don't move to better areas earlier

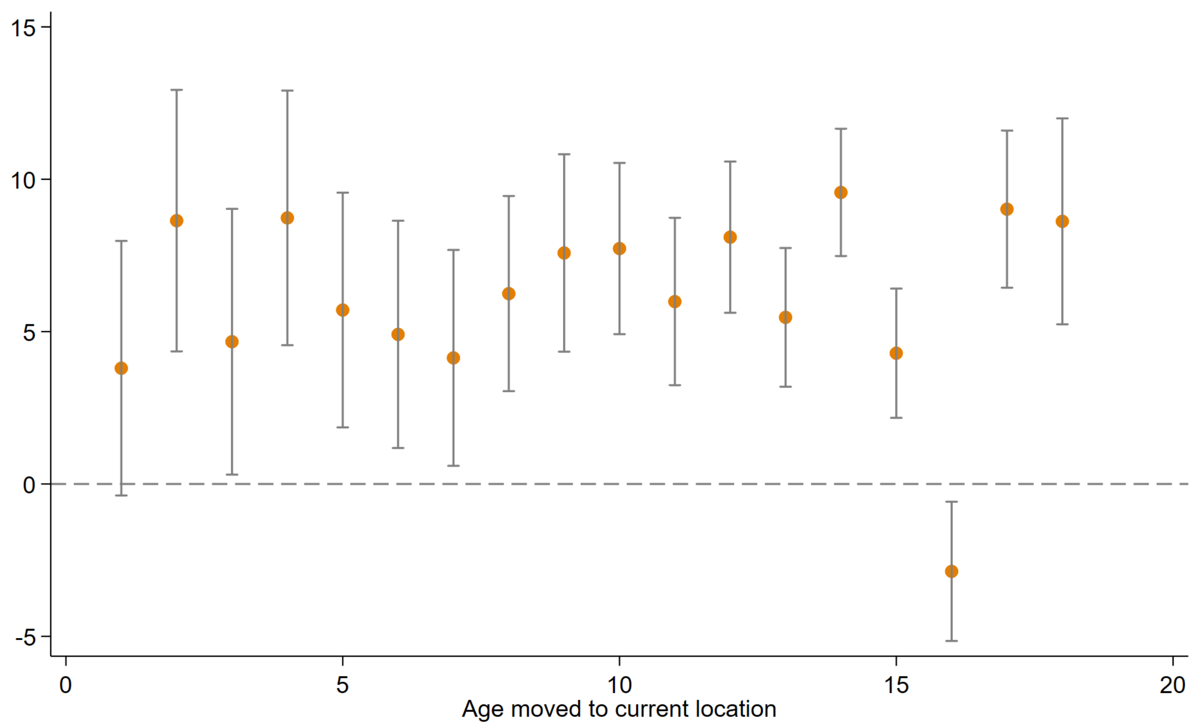

Note: This figure plots the average difference in mover children mothers' primary completion rates between those who move to better vs worse locations by the age of the child when the move occurs. Vertical lines indicate 95% confidence intervals.

B.2 Focusing on within household variation

By including household fixed effects I can relax the assumption of constant selection by allowing selection effects to vary across, but not within, households and only consider variation in age at move between children within the same household. This specification estimated the following equation: $y_i = \rho_{os} + \alpha_m + \beta_m \cdot \Delta \bar{y}^p_{ods} + \psi_h + \varepsilon_i$. Where h denotes i's household. This specification imposes considerable data requirements, firstly I can only consider multi-child households[21] and secondly ψ_h is a high dimensional object[22] to be estimated. To reduce the data burden I estimate the linear parametric model, results reported in table IV.

Table IV Parametric estimates of place effects including HH fixed effects

	Year born - birth location and age at move FE	Year born, year born by $\bar{y}_{os}$ and age move all by sample FE	Year born, year born by $\bar{y}_{os}$ and age move FE	Year born and age moved FE
	-0.0181**	-0.0180***	-0.0201**	-0.0349***
$\hat{\beta}_1$	(0.00834)	(0.00618)	(0.00785)	(0.0114)
	0.00245	-0.0141	-0.0141	0.00878
$\hat{\beta}_2$	(0.0131)	(0.0145)	(0.0130)	(0.0197)
R^2	0.779	0.745	0.741	0.727

Note: This table shows the robustness of the parametric convergence rate estimates over various fixed effects specifications where I always control for household fixed effects and therefore only use between sibling variation.

Although recovering a smaller coefficient than when estimating on the full sample without household fixed effects, the difference is not statistically significant and may be a function of attenuation bias from the demanding specification (incidental parameter problem). This points to the conclusion that across household selection is unlikely to be driving a meaningful portion of causal place effects I find. The results are also robust to alternative specifications as shown in table IV. The possibility remains that time-varying selection effects are confounding my results, to address this I consider increasingly likely to be exogenous moves in the next section.

B.3 Time varying selection effects - increasingly likely to be exogenous moves

Time vary shocks could still invalidate my identifying assumption. A household could receive a windfall and move to a better location simultaneously increasing investment in their child's education. To investigate whether this kind of force invalidates my identifying assumption I consider plausibly exogenous moves in two ways.

I use the data-driven strategy developed by Chetty and Hendren [2018a,b] to identify location-year cells that exhibited likely to be exogenous movement events. The data-driven method of Chetty and Hendren [2018a,b] first constructs a *shock* variable k_{lt} for each location l and year t which is the size of that years out movers relative to that regions average number

[21] This cuts the sample with estimating variation from 79,778 to 27,807.

[22] In my main regressions there are 12,175 households.

of out movers: $k_{lt} = \frac{Out_{lt}}{\overline{Out_l}}$. Where Out_{lt} is the number of out-movers from region l at time t and $\overline{Out_l} = \sum_t Out_{lt}$ is the locality average over time t. High values of k_{lt} represent unusually large outflows in a given year from that given area. This method hypothesizes that higher values of k_{lt} represent unusual episodes and that these moves are more likely to be exogenous. Note that we don't require some cut off after which we determine that all such moves are undoubtedly exogenous. Instead, we only suppose that the proportion of moves due to exogenous reasons is increasing with k_{lt}. If place effect estimates remain stable as the sample is restricted to larger and larger values of k_{lt} , that is more and more likely to be exogenous moves, this is evidence in favor of the identifying assumption.

Figure X reports $\hat{\beta}_1^{k_r}$ found by estimating equation equation 4 in the main text on a sample of location years for which $k_{lt} > k_r$ where k_r is the r'th percentile of k_{lt}. In both figures, I find little systematic evidence of our estimated exposure effects being driven by time-varying selection. That is, there is no significant evidence of the effect eroding as we restrict variation to moves more likely to be exogenous. As before I set the selection effects component to be the same and only plot changes in the slope to aid comparability.

B.4 Over-identification outcome-based test

Finally, I use heterogeneity in permanent residents outcomes across reported gender to consider the validity of the identifying assumption. The logic here is that in the causal place effect model movers' outcomes should converge to permanent residents outcomes for their own subgroup in this case gender. On the other hand, a selection effects model would predict convergence to general location quality. The second statement requires either that households do not observe gender-specific location quality above overall location quality or that if they do observe it they don't select on it. The former seems a reasonable informational assumption, especially when you consider the high correlation between gender-specific location qualities, and that I don't find that individuals move disproportionately to better areas as shown in the main text.

To formally test the hypothesis proposed above, I estimate the main parametric specification firstly including the own-gender measure of location quality - that is the primary completion rate of permanent residents who have the same gender as i, secondly with the other-gender measure of quality and finally with both. It's worth noting two things, first male and female primary completion rates across localities are highly correlative with correlation 0.82. Second, by reducing the number of permanent residents in each location we introduce potential measurement error in the main explanatory variable which could result in attenuation bias. Table V shows the results.

Column one of table V re-estimates my main equation on the reduced sample for which gender is reported to provide a relevant comparison unit. Column two reports results us-

Figure X Considering increasingly likely to be exogenous moves

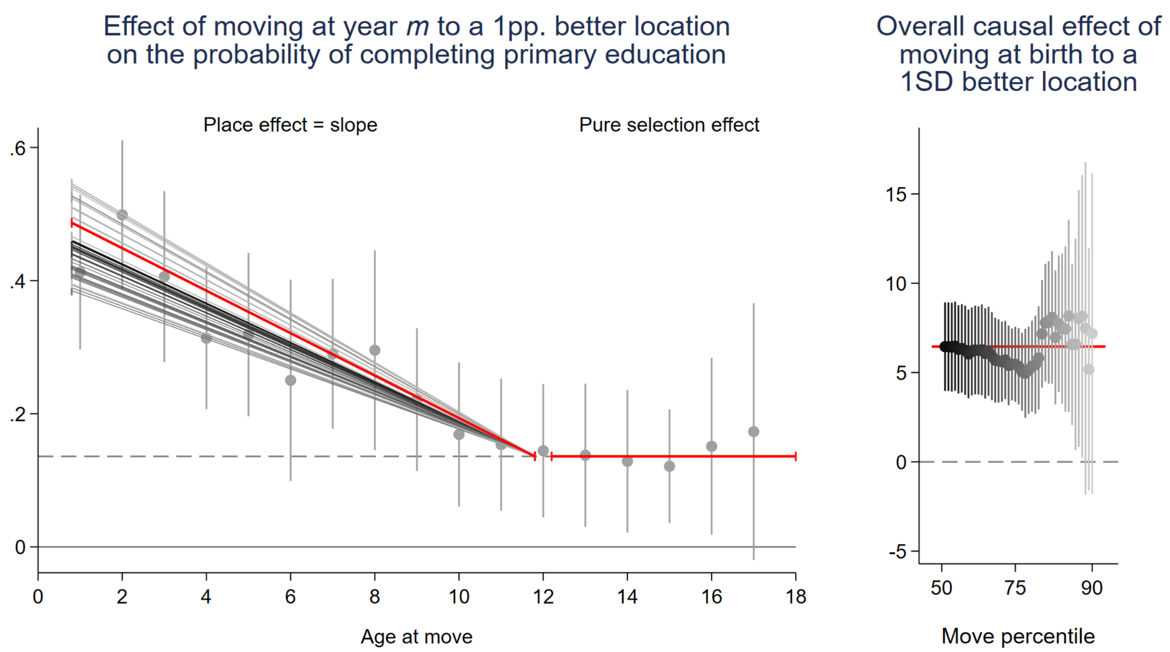

Note: This figure shows in the left hand side panel the estimated β_m coefficients from the regression: $y_i = \rho_{os} + \alpha_m + \beta_m \cdot \Delta\bar{y}^p_{ods} + \varepsilon_i$ estimated on a sample of one time movers between the ages of 14 and 18 in Benin, Cameroon and Mali. Additionally, corresponding 95% confidence intervals are indicated by gray bars where standard errors have been clustered at the birth locality and current locality level. Superimposed on this in red are two lines of best fit with a break at age 12. The horizontal gray dashed line indicated the post 12 mean of $\hat{\beta}_m$. Results from the linearised specification are also plotted in increasingly light gray where the sample is increasingly restricted to evermore unusually large movement events. In the right-hand side panel, the slope of the 1 to 12 line of best fit has been translated into the overall impact of moving at birth to a one standard deviation better location where over the whole sample a one standard deviation improvement in location quality is equal to a 20pp increase in primary completion rates. The overall effect is represented by the red vertical line and the corresponding points and standard errors in gray calculated over increasing likely to be exogenous move events are plotted for comparison.

ing the permanent resident outcomes of those individuals who have the same gender as the moving child. Column three reports results using permanent residents outcomes for those who have the opposite gender to the moving child, it's unsurprising here that the coefficient changes little from column one to column two given the high correlation between the two measures. Column four shows the results from including both measures in the same regression. Here, if both coefficients were similar and attenuated that would be evidence for the selection effects model. However, that's not what I find, instead, the same-gender coefficient remains positive and highly significant whereas the other-gender coefficient falls to a precisely estimated 0. Therefore, instead, I find evidence in favor of the causal place effects model and thus the main identifying assumption.

Table V Outcome based placebo test

	(1)	(2)	(3)	(4)
Base	0.0209***			
	(0.0029)			
Same reported gender specific		0.0176***		0.0123***
		(0.0025)		(0.0041)
Other reported gender specific			0.0164***	0.00664
			(0.0025)	(0.0041)
N	65560	65560	65560	65560
R^2	0.358	0.356	0.358	0.358

Note: This table shows the results from the outcome-based over-identification test. Column one replicates the main analysis on the slightly smaller sample where gender information is available. Column two uses a gender-specific measure of location quality, that is the primary completion rate of permanent residents of the same gender born in the same year. Column three uses the opposite gender-specific measure of location quality. In column four both measures are included in one regression. Standard errors are clustered at the origin and destination locality level. Age at move and origin by census year fixed effects are included in all regressions.

In sum, across all four approaches taken, I find consistent and strong evidence in favor of the key identifying assumption. Any omitted variable bias or selection effects would have to operate within families proportionally to exposure times, persist in the presence of moves induces by displacement shocks, and precisely replicate permanent residents outcomes by gender in proportion to exposure times, and not show up in observable such as mothers education. I find it unlikely that any qualitatively or quantitatively important threats to identification operate in this manner.

B.5 Locality-year specific identifying assumption

When estimating locality-year specific fixed effects I require a stronger identifying assumption: that selection effects are constant over age at move for each location l in each period t, rather than just on average. Formly I require that: $\mathbb{E}[e_{il} \cdot \varepsilon_i | \alpha_{odt}] = 0 \ \ \forall l \in \mathbb{L}$.

The results supporting the validity of the on-average identifying assumption summarized in section 2.2 also provide evidence for this location-specific assumption. In addition, I perform a placebo test estimating equation 3 on a sample of those who move between 14 and 18, *after* primary education completion is determined. In this case, place effects μ_l should just be noise and unrelated to the actual estimated place effects. A scatter diagram in the online appendix summarizes the relationship between actual and placebo place effects, statistics are weighted by locality population. Actual and placebo estimates exhibit little relationship with a regression slope coefficient indistinguishable from zero (0.07(se=0.05)) and a very small R-squared (0.01). In addition, less than five percent of the individual placebo coefficients are significant at the 5% level.

C Signal to noise ratio of estimated place effects

I estimate equation 3 separately for each country on a sample of (o, d, t) cells with more than 20 observations.[23] I can estimate the signal to noise ratio by decomposing estimated place effects into actual place effects and orthogonal measurement error, $\hat{\mu}_{lt} = \mu_{lt} + \eta_{lt}$, and estimating the true variation, $\hat{\sigma}^2_{\mu_{lt}} = \sigma^2_{\hat{\mu}_{lt}} - \sigma^2_{\eta_{lt}}$. A natural estimator of $\sigma^2_{\eta_{lt}}$ is given by the average squared standard errors s^2_{lt} recovered when estimating equation 3: $\sigma^2_{\eta_{lt}} = E[s^2_{lt}]$. Using this decomposition I estimate the proportion of the raw variance which is due to signal not noise as $\frac{\hat{\sigma}^2_{\mu_{lt}}}{\sigma^2_{\hat{\mu}_{lt}}} = \frac{\sigma^2_{\hat{\mu}_{lt}} - \frac{1}{N}\sum s^2_{lt}}{\sigma^2_{\hat{\mu}_{lt}}} \times 100 = 41.4$, that is 41.4% of the raw variance is due to signal not noise.[24] The standard deviation of place effects is estimated to be 15 indicating substantial spatial inequality of opportunity. Moving to a location with a one standard deviation higher place effect increases children's probability of completing primary education by 14 percentage points. Benin has larger spatial inequality of opportunity with a standard deviation of 23.7 compared to 17.5 in Mali and 2.9 in Cameroon.

References

Daniel Aaronson. Using sibling data to estimate the impact of neighborhoods on children's educational outcomes. *Journal of Human Resources*, pages 915–946, 1998.

Sheila Aikman and Elaine Unterhalter. *Beyond access: Transforming policy and practice for gender equality in education*. Oxfam, 2005.

Alberto Alesina, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou. Religion and educational mobility in africa. Technical report, National Bureau of Economic Research, 2020.

Alberto Alesina, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou. Intergenerational mobility in africa. *Econometrica*, 89(1):1–35, 2021.

Sonya Anderson. *Girls' education and the "pedagogy of difference": A study of teacher attitudes and gender-equitable pedagogy in Benin*. Harvard University, 2005.

[23] 20 observations in the 10% samples implies an expected 200 movers. Restricting the estimating sample to include only cells with more movers would decrease measurement error but increase potential selection. Results are not sensitive to varying the cut-off.

[24] Compared to 16% for the analogous commuter zone estimates in Chetty and Hendren [2018b].

Emmanuelle Auriol, Julie Lassebie, Amma Panin, Eva Raiber, and Paul Seabright. God insures those who pay? formal insurance and religious offerings in ghana. *The Quarterly Journal of Economics*, 135(4):1799–1848, 2020.

World Bank. Millennium development goals. 2015.

Mahdi Barouni and Stijn Broecke. The returns to education in africa: some new estimates. *The Journal of Development Studies*, 50(12):1593–1613, 2014.

Joerg Baten, Michiel de Haas, Elisabeth Kempter, and Felix Meier zu Selhausen. Educational gender inequality in sub-saharan africa: A long-term perspective. *African Economic History Network Working Paper Series*, 54, 2020.

Eleonora Bertoni, Michele Di Maio, Vasco Molini, and Roberto Nistico. Education is forbidden: The effect of the boko haram conflict on education in north-east nigeria. *Journal of Development Economics*, 141:102249, 2019.

Helen N Boyle. Between secular public schools and qur'anic private schools: The growing educational presence of malian medersas. *The Review of Faith & International Affairs*, 12 (2):16–26, 2014.

Bart J Bronnenberg, Jean-Pierre H Dubé, and Matthew Gentzkow. The evolution of brand preferences: Evidence from consumer migration. *American Economic Review*, 102(6): 2472–2508, 2012.

Raj Chetty and Nathaniel Hendren. The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The Quarterly Journal of Economics*, 133(3): 1107–1162, 2018a.

Raj Chetty and Nathaniel Hendren. The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *The Quarterly Journal of Economics*, 133(3):1163–1228, 2018b.

Raj Chetty, Nathaniel Hendren, and Lawrence F Katz. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902, 2016.

Eric Chyn. Moved to opportunity: The long-run effects of public housing demolition on children. *American Economic Review*, 108(10):3028–56, 2018.

Arusha Cooray and Niklas Potrafke. Gender inequality in education: Political institutions or culture and religion? *European Journal of Political Economy*, 27(2):268–280, 2011.

Flavio Cunha and James J Heckman. Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *Journal of human resources*, 43(4):738–782, 2008.

Anna Piil Damm and Christian Dustmann. Does growing up in a high crime neighborhood affect youth criminal behavior? *American Economic Review*, 104(6):1806–32, 2014.

Nathan Deutscher. Place, peers, and the teenage years: Long-run neighborhood effects in australia. *American Economic Journal: Applied Economics*, 12(2):220–49, 2020.

Parfait M Eloundou-Enyegue. Pregnancy-related dropouts and gender inequality in education: A life-table approach and application to cameroon. *Demography*, 41(3):509–528, 2004.

Robert C Feenstra, Robert Inklaar, and Marcel P Timmer. The next generation of the penn world table. *American economic review*, 105(10):3150–82, 2015.

Amy Finkelstein, Matthew Gentzkow, and Heidi Williams. Sources of geographic variation in health care: Evidence from patient migration. *The quarterly journal of economics*, 131 (4):1681–1726, 2016.

Claudia Goldin. The u-shaped female labor force function in economic development and economic history, 1994.

Joan Hamory, Marieke Kleemans, Nicholas Y Li, and Edward Miguel. Reevaluating agricultural productivity gaps with longitudinal microdata. *Journal of the European Economic Association*, 19(3):1522–1555, 2021.

IPUMS. Integrated public use microdata series, international: Version 7.2 [dataset]. 2020.

Melina Platas Izama. Muslim education in africa: Trends and attitudes toward faith-based schools. *The Review of Faith & International Affairs*, 12(2):38–50, 2014.

Tabi Atemnkeng Johannes and Armand Gilbert Noula. Gender and increased access to schooling in cameroon: A marginal benefit incidence analysis. *Journal of International Women's Studies*, 12(1):94–106, 2011.

David Lagakos. Urban-rural gaps in the developing world: Does internal migration offer opportunities? *Journal of Economic perspectives*, 34(3):174–92, 2020.

Jean-William P Laliberté. Long-term contextual effects in education: Schools and neighborhoods. *American Economic Journal: Economic Policy*, 2021.

Kirsten Majgaard and Alain Mingat. *Education in sub-Saharan Africa: A comparative analysis*. World Bank Publications, 2012.

Nicolette D Manglos-Weber. Identity, inequality, and legitimacy: religious differences in primary school completion in sub-saharan africa. *Journal for the Scientific Study of Religion*, 56(2):302–322, 2017.

Leandro Medina, Mr Andrew W Jonelis, and Mehmet Cangul. *The informal economy in Sub-Saharan Africa: Size and determinants*. International Monetary Fund, 2017.

David Molitor. The evolution of physician practice styles: evidence from cardiologist migration. *American Economic Journal: Economic Policy*, 10(1):326–56, 2018.

Rachel Ngai, Claudia Olivetti, and Barbara Petrongolo. Structural transformation over 150 years of women's and men's work. *Unpublished Working Paper*, 2022.

Victor Ombati and Mokua Ombati. Gender inequality in education in sub-saharan africa. *JWEE*, (3-4):114–136, 2012.

Amma Panin et al. Four reasons why analysis of economic policy and religion go hand-in-hand in sub-saharan africa. Technical report, PEGNet-Poverty Reduction, Equity and Growth Network, Kiel Institute for the . . . , 2021.

Melina R Platas. Culture and the persistence of educational inequality: Lessons from the muslim–christian education gap in africa. *New York University Abu Dhabi Working Paper*, 2018.

Leandro Prados de la Escosura. Augmented human development in the age of globalization. *The Economic History Review*, 2021.

George Psacharopoulos and Harry Anthony Patrinos. Returns to investment in education: a further update. *Education economics*, 12(2):111–134, 2004.

George Psacharopoulos and Harry Anthony Patrinos. Returns to investment in education: a decennial review of the global literature. *Education Economics*, 26(5):445–458, 2018.

Jeffrey D Sachs. From millennium development goals to sustainable development goals. *The lancet*, 379(9832):2206–2211, 2012.

T Paul Schultz. Evidence of returns to schooling in africa from household surveys: Monitoring and restructuring the market for education. *Journal of African Economies*, 13(suppl_2): ii95–ii148, 2004.

M Najeeb Shafiq, Robert K Toutkoushian, and Alexandria Valerio. Who benefits from higher education in low-and middle-income countries? *The Journal of Development Studies*, 55 (11):2403–2423, 2019.

Muna Shifa and Murray Leibbrandt. Spatial inequality in sub-saharan africa. *African Geographical Review*, pages 1–17, 2022.

Nik Stoop, Marijke Verpoorten, and Koen Deconinck. Voodoo, vaccines, and bed nets. *Economic Development and Cultural Change*, 67(3):493–535, 2019.

Clarence Tsimpo and Quentin Wodon. Assessing the role of faith-inspired primary and secondary schools in africa: Evidence from multi-purpose surveys. *The Review of Faith & International Affairs*, 12(2):5–15, 2014.

Theresa Tuwor and Marie-Antoinette Sossou. Gender discrimination and education in west africa: strategies for maintaining girls in school. *International Journal of Inclusive Education*, 12(4):363–379, 2008.

Felix Weinhardt. Social housing, neighborhood quality and student performance. *Journal of Urban Economics*, 82:12–31, 2014.