Defusing a Population Explosion? Jobs and Fertility in sub-Saharan Africa

Céline Zipfel*
London School of Economics and Political Science

December 2020

For latest version please click here

Abstract

Over three-quarters of global population growth between now and 2100 is expected to occur in sub-Saharan Africa. This paper leverages micro data from 153 reproductive health surveys to shed light on a macro puzzle: fertility rates are exceptionally high in sub-Saharan African countries, conditional on GDP per capita. I establish an important empirical fact: the relationship between income and desired fertility is, on average, steeper in sub-Saharan Africa relative to other countries. I hypothesize that a key factor behind this steeper gradient is that poorer households face especially limited salaried employment opportunities in this region. A quantity-quality trade-off model of fertility choice featuring a fixed human capital requirement for entry into salaried employment predicts that a feedback loop can arise, where poorer families get stuck in a high fertility - informal occupation equilibrium in which they also under-invest in their children's education. Rich micro data assembled from reproductive health surveys, censuses and household expenditure surveys provide empirical support for the model's key assumptions and predictions. The findings suggest that differences in occupational choice sets across the income distribution represent an important driver of sub-Saharan Africa's exceptional fertility trend.

*I am indebted to Nava Ashraf, Oriana Bandiera, Gharad Bryan, Robin Burgess and Rachael Meager for their continuous guidance and support throughout this project. I owe special thanks also to Tim Besley, Maitreesh Ghatak, Jay Euijung Lee, Niclas Moneke, Torsten Figueiredo Walter and Silvana Tenreyro for very insightful discussions. Clare Balboni, Svetlana Chekmasova, Alexia Delfino, Martina Fazio, James Fenske, Tatiana Komarova, Gabriel Leite Mariante, Virginia Minni, Rachel Ngai, Tsogsag Nyamdavaa, Wolfgang Ridinger, Pauline Rossi, Caterina Soto Vieira, Matthieu Teachout, Heidi Thysen, Martina Zanella and seminar participants of the LSE Development Economics and STICERD Work-in-Progress seminars provided helpful comments. Financial support from the European Social Research Council is gratefully acknowledged. Contact: Department of Economics, LSE, UK. www.celinezipfel.com, c.p.zipfel@lse.ac.uk.

1 Introduction

Population growth lies at the heart of many debates on the tension between economic growth and sustainability. Ever since Malthus (1798), scholars across disciplines have worried about a population explosion leading to the depletion of the Earth's resources (Meade 1961, 1967; Ehrlich 1968). Set against this, other researchers regard population growth as one of the key drivers of technological change and thus economic growth (Boserup 1965; Romer 1990; Kremer 1993). Unpacking the determinants of population growth is crucial to inform debates about its consequences.

Figure 1

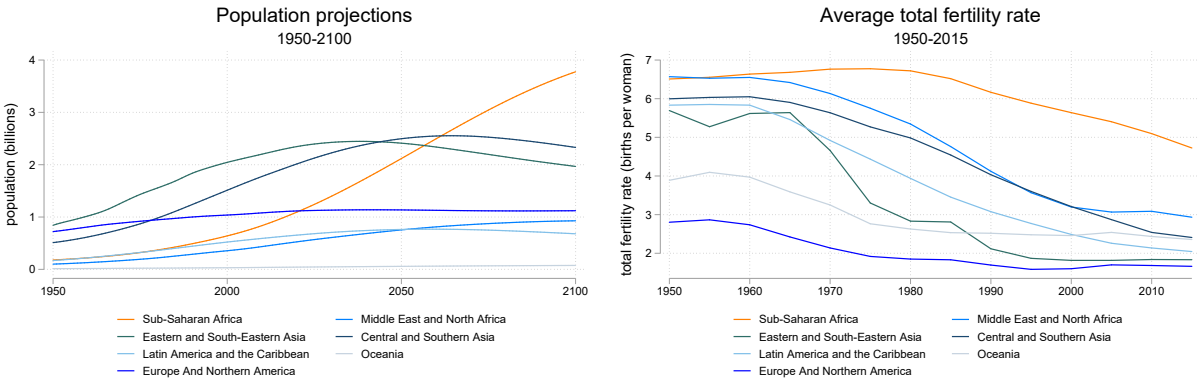

(a) Population projections

(b) Total fertility rates

This paper investigates the determinants of the highest rates of population growth in the world, currently concentrated in sub-Saharan Africa. Panel (a) of Figure 1 illustrates that over three-quarters of global population growth between now and 2100 is expected to take place in sub-Saharan Africa. The region's population is projected to double by 2050 and almost quadruple by 2100. Panel (b) shows that persistently higher fertility rates underlie this pattern, with total fertility rates averaging around 5 births per woman in 2015, compared to 1.8-2.8 in other low-income countries.[1] Importantly, this disparity cannot solely be explained by differences in national income. Conditional on GDP per capita, the total fertility rate in sub-Saharan Africa is almost two births per woman higher than anywhere else in the world (Figure 2).

Fertility is primarily determined by two factors: parents' desired number of children and the availability of contraceptives. Much of the literature to date studies the role of family planning services in reducing fertility rates.[2] However, households' demand for these services ultimately depends on their *desired* fertility, the determinants of which remain relatively under-studied in the context of developing countries.[3]

[1] Demographers recognize the region's fertility decline to be following a unique trend (e.g. Bongaarts 2017).

[2] E.g. Miller (2009), Ashraf et al. (2014), Portner et al. (2011), McKelvey et al. (2012), and Joshi and Schultz (2012).

[3] In macroeconomics, the convergence of fertility rates to about two births per woman in *developed* countries has inspired the field of unified growth theory (Galor and Weil 1996, 1999, 2000). Other important contributions include Doepke (2005) on child mortality and Doepke and Zilibotti (2005) on child labour regulations.

This paper contributes to filling this knowledge gap by leveraging micro data to shed light on the drivers of sub-Saharan Africa's exceptional fertility trend. To isolate economic determinants from cultural or historical factors fixed at the country level, I study variation in demand for children within countries.[4] I establish that the dispersion in desired fertility across the income distribution is significantly larger in sub-Saharan African countries on average. I then investigate the drivers of this differential income-desired fertility gradient across world regions through the lens of a fertility choice model.

Figure 2

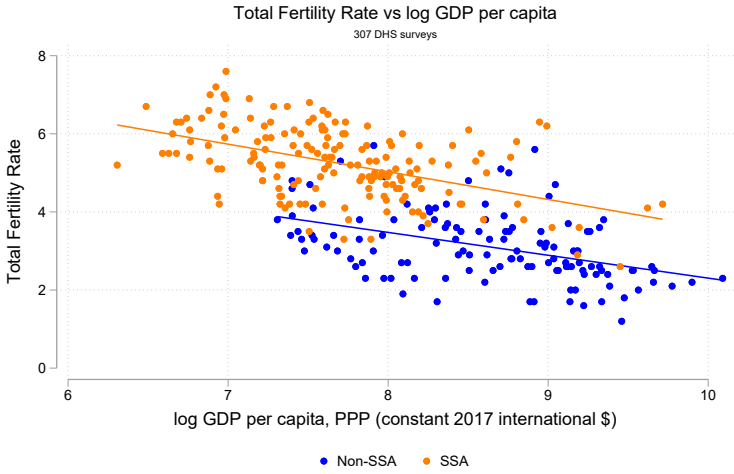

I construct gradients between income and desired fertility from the Demographic Health Surveys (DHS) micro data, using wealth as a proxy for income and the ideal number of children reported by respondents as a measure of desired fertility. I find that around one third of the difference in total fertility rates conditional on GDP per capita between sub-Saharan Africa and other regions is accounted for by disproportionately higher demand for children amongst the poorest segments of the population in sub-Saharan Africa, where the poor-rich gap in desired fertility is, on average, over twice as large as in other developing countries. Thus, to the best of my knowledge, this paper is the first to document that the income-desired fertility gradient is significantly steeper in sub-Saharan Africa.

To test that this steeper gradient does not simply reflect larger income dispersion in sub-Saharan African countries, I construct a cross-sample wealth ranking of households from the DHS surveys to compare the living standards of the different within-country wealth quintiles across countries. I can confidently reject larger income dispersion as a potential explanation: while households are on average poorer in sub-Saharan African

[4]Previous work on the determinants of desired fertility in sub-Saharan Africa has studied the role of cultural norms and practices, such as polygyny (Tertilt 2005, 2006; Rossi 2019) and traditional beliefs about maternal mortality (Ashraf et al. 2017), in fertility choices. Norms are also the topic of de Silva and Tenreyro (2017, 2020), who link the speed of the fertility decline in a number of developing countries to advertisement campaigns calling for a shift to a "small family" norm. This paper complements this literature by focusing on the economic determinants of demand for children, which may shape incentives to deviate from social norms.

countries, the poor-rich gap in living standards within countries is in fact smaller in sub-Saharan Africa relative to other regions, conditional on GDP per capita.

Instead, I propose that a key driver of the steeper gradient in sub-Saharan Africa is that the poor face especially limited job opportunities in this region, even conditional on GDP per capita. Recent, well-identified micro evidence confirms that limited occupational choice sets can generate poverty traps (Bandiera et al. 2017; Balboni et al. 2020). Accordingly, I argue that the nature of labor market opportunities available to the poorest individuals in sub-Saharan Africa is responsible for a high fertility - informal occupation feedback loop that persists via families also under-investing in their children's education. To the best of my knowledge, this study is the first to investigate the role of jobs in sub-Saharan Africa's exceptionally high fertility levels.

Starting with Becker and Lewis (1973), a long literature has put forward two key mechanisms to explain a negative relationship between income and fertility: the time cost associated with childbearing and the "quantity-quality trade-off". I argue that parents' occupational choice sets can shape fertility decisions through both of these channels. First, the absence of suitable jobs may deter poor women from reducing their fertility by keeping the opportunity cost of childbearing low (the "opportunity cost of childbearing" channel).[5] Second, parents' perception of the returns to educational investments could be shaped by their own experience of the labor market. In settings where increases in schooling are associated with better employment opportunities, parents may be more willing to substitute quantity for quality, namely reduce their fertility and increase educational investments in each child (the "parental valuation of education" channel).[6]

I incorporate both channels into a simple theoretical framework that enriches a Beckerian household optimisation model with a poverty trap framework à la Galor and Zeira (1993).[7] Parents can either work in an informal occupation - namely, subsistence self-employment - or in salaried employment. Faced with a human capital threshold for entry into salaried employment, poorer, less educated individuals are forced to choose the informal occupation, which yields lower returns to human capital. Parents then decide jointly how many children to have and how much to invest in each child's education. I show that a set of parameter values exist for which the combination of borrowing constraints and the fixed education requirement for salaried employment can generate a feedback loop where poor families can get stuck in a high fertility - informal occupation equilibrium. The latter is also characterised by low educational investments in children, thereby preventing them from entering the formal occupation themselves and contributing to the persistence of the occupational structure across generations.

[5]Two studies - set in South Asian contexts characterized by low female labor force participation - show that the arrival of new employment opportunities can lead women to delay fertility (Jensen 2012; Heath and Mobarak 2015).

[6]Two papers exploiting natural experiments in the US South in the early twentieth century - the construction of the Rosenwald schools and the eradication of the hookworm disease - provide evidence that households reduce their fertility when the returns to educational investments increase (Aaronson et al. 2014; Bleakley and Lange 2009).

[7]The household optimisation model in this paper is akin to the one in Vogl (2016), who establishes that, since the start of the demographic transition in developing countries, children from larger families are on average less educated.

Bringing the model to the data, I provide empirical support for the framework's key assumption and key prediction using 49 micro-censuses covering 29 countries from 1990 to 2013. First, I compare the relationship between salaried employment and years of schooling in sub-Saharan Africa to that in other countries, controlling for GDP per capita levels. I find that this relationship is flat for individuals with under five years of education in sub-Saharan Africa, with the probability of holding a salaried job being close to zero, whereas the slope is positive in other parts of the world. This is consistent with the assumption of a fixed human capital requirement for salaried employment, a key necessary condition for the existence of multiple equilibria in the model. From about five years of education, the average share of individuals in salaried employment increases sharply with respect to years of schooling in sub-Saharan Africa, joining the curve of the other countries around twelve years of education.

Second, my empirical analysis confirms that less educated parents are also more likely to under-invest in their children's education relative to their richer counterparts, a core prediction of the model. For mothers' education levels ranging from zero to ten years of schooling, educational investments - as measured by the fraction of their children aged 14 and above who have completed primary school - appear to be significantly lower in sub-Saharan Africa. The shape of the relationship between parental human capital and children's education is also noticeably different across the two sets of countries: at low levels of education, the slope of the relationship between mothers' years of schooling and the fraction of their children aged 14 and above who have completed primary school is less steep in sub-Saharan Africa. From around five years of schooling, the slopes are virtually the same across the two sets of countries. This is consistent with the model's prediction that a fixed human capital requirement for entry into salaried employment can generate multiple equilibria.

Taken together, these two findings provide an explanation for the steeper socioeconomic gradient in desired fertility observed in sub-Saharan Africa: the presence of a fixed human capital requirement for entry into salaried employment can generate a high fertility - informal occupation feedback loop for the poorest families.

In the final section of the paper, I test for the presence of the two channels underpinning the model's predictions: whether jobs shape parents' valuation of education, and whether the scarcity of employment opportunities lowers mothers' opportunity cost of childbearing. I complement the DHS and IPUMS data with household expenditure surveys to yield insight into the relationship between parental labor market outcomes and spending on children's education.[8] Under the assumption that men's fertility intentions are less responsive to the opportunity cost of childbearing channel than those of women, I use men and household heads as the unit of analysis when testing for the parental valuation of education channel, whereas the analysis of the opportunity cost of childbearing mechanism focuses on women.

[8] I have gathered household expenditure surveys from which I can construct relevant statistics for 12 sub-Saharan African countries. In this version of the paper I present results from the Kenya Integrated Household Budget Survey.

Indeed, demand for children is strongly correlated with the type of work individuals engage in, conditional on their level of education, urban/rural status, and religion and ethnic group fixed effects. Men report an ideal number of children almost three quarters of a child lower if they work for pay, outside of agriculture, and all year round. I also find that households where the household head has a salaried job spend more on each child's education and have fewer children. These findings suggest that parents whose employment status implies higher returns to education are more likely to substitute quantity for quality of children.

I also provide evidence in line with the hypothesis that the opportunity cost of childbearing factors into women's fertility choices in sub-Saharan Africa. Pooling together all DHS samples from sub-Saharan Africa including labor market outcomes, I find a strongly negative correlation between women's reported ideal number of children and the availability of off-farm and salaried work in their province of residence, conditional on the same set of covariates as above. The magnitude of this correlation is twice as large for salaried employment as for a broader definition of off-farm work. This points to the greater incompatibility of salaried employment with raising many children.

I exploit variation in the destination of female migrants from the same district in the Malawi 2008 census to test that this correlation is not only capturing hidden heterogeneity in social norms - which may affect both women's reported ideal number of children and the share of women working outside the home or household farm. Conditional on religion, ethnicity, age fixed effects and husband's characteristics, I find a strongly negative association between women's total number of children ever born and the rate of salaried employment in their current location, but virtually no correlation with the rate of salaried employment in their district of birth. This lends further credibility to the existence of the opportunity cost of childbearing channel in sub-Saharan Africa.

This paper establishes a new empirical fact: the exceptionally high demand for children reported by the poorest segments of the sub-Saharan African population is an important driver of its unique fertility transition, which will be responsible for a large share of projected global population growth over the 21st century. I argue that the scarcity of salaried employment opportunities available to the poorest families is an important determinant of the region's steeper socioeconomic gradient in desired fertility. Insights from a fertility choice model combined with rich micro data suggest that sub-Saharan African labor markets feature a fixed education requirement for salaried employment that contributes to the persistence of a segregated occupational structure by income.

The paper is organised as follows. Section 2 describes the data used to produce the cross-regional comparisons of socioeconomic gradients in desired fertility and participation in salaried employment, which are presented in section 3. Section 4 discusses the theoretical framework. Section 5 provides empirical support for the model's key assumption and predictions as well as the two main mechanisms underpinning the model's predictions. Section 6 concludes and discusses promising avenues for future research stemming for this study.

2 Data

2.1 Demographic Health Surveys

The DHS Program collects and disseminates nationally representative data on fertility, family planning, and a broad set of health indicators including maternal and child health outcomes. With samples for 83 LMICs in the public domain, geographical coverage of the DHS data is extensive, and the most complete for Sub-Saharan Africa.

Due to the program's focus on reproductive, maternal and child health, the target population of the DHS surveys is women of childbearing age. The most common sample universe is therefore women aged 15-49. Most survey waves also include a questionnaire for adult males aged 15 and older, which I use in my analysis of correlates of men's desired fertility in section 5.

To construct income-desired fertility gradients within countries, I pool together all the surveys that include information on household wealth and individuals' desired fertility in the log GDP per capita bin of common support [7.2;9] where sub-Saharan African samples overlap with samples from other countries. These amount to 153 surveys in total, covering 56 countries: 28 from sub-Saharan Africa and 28 from other regions. I use two key variables: the wealth index constructed by the DHS and the ideal number of children reported by respondents.

Since the middle of the 1990s, all DHS surveys collect information on asset ownership, dwelling characteristics and access to utilities at the household level. This permits the construction of a wealth index as a proxy for socioeconomic status using principal components analysis. Principal component analysis (PCA) is a data dimensionality reduction tool which involves the construction of new uncorrelated variables as linear combinations of an original set of variables that are correlated with each other. This method is commonly used in development to produce a uni-dimensional measure of socioeconomic status from different multi-categorical indicators of dwelling characteristics, access to utilities, and asset counts (Montgomery et al. 2000; Filmer and Pritchett 2001; McKenzie 2005).

These variables are typically predictive of socioeconomic status, but many of the categories included in these indicators will be correlated with each other. As an example, dwellings where the roof is made of palm leaf may also be more likely to have walls made of mud and not have piped water than dwellings where the roof is made of tiles. PCA decomposes the variation in these variables into orthogonal components, and outputs weights for each variable. These weights can then be used to construct a score which will be lower for households living in mud houses relative to those living under tiled roofs, thereby providing a socioeconomic ranking. Grouping households into quintiles ensures that these are sufficiently large bins for socioeconomic status differentiation. I use these quintiles as a proxy for relative income.[9]

[9]Since I am interested in the slope of the relationship only, the actual value of earnings in each bin is irrelevant.

Most rounds of the DHS surveys also include a module collecting information on respondents' employment and work. Female respondents are typically asked whether they are currently working, what was their primary occupation over the last 12 months, do they typically work throughout the year, seasonally, or only occasionally; is this work paid or unpaid; and who employs them for this work (i.e. whether they are self-employed, work for a family member or are self-employed). Given that DHS surveys are representative at the first administrative level (province or region) I can construct female employment shares for different occupations at the province level to study the link between the availability of suitable labor market opportunities for women and their desired fertility (see section 5.2.2).[10]

2.2 IPUMS census data

IPUMS is the world's largest collection of publicly available census microdata. The data are coded and documented consistently across countries and over time to facilitate comparative research. Samples are available for 67 LMICs. In the log GDP per capita bin that I restrict my attention to in this paper - the range of common support for sub-Saharan African and other countries - there are 50 samples that provide information on years of schooling, employment type and education of children in the household. This corresponds to about 38 million individuals. Most of the census microdata samples assembled by IPUMS are 10% simple random samples of the country's entire population.

Most census waves contain information on years of schooling, labor force participation, occupation, and a broad classification of an individual's type of employment contract. Unlike the DHS data, IPUMS provides a distinction between the "unemployed" and "economically inactive" status, which is crucial for accurately measuring female labor force participation. To construct the share of people in salaried employment, I make use of the "status in employment" variable, which records whether a person is self-employed, working for a wage or salary, or an unpaid worker.

2.3 Kenya Integrated Household Budget Survey 2005-2006

The main limitation of the DHS and IPUMS databases is the absence of information on income and consumption expenditure. Therefore, to complement these two data sources, I use the 2005-2006 Kenya Integrated Household Budget Survey. It includes information on household demographics, a fertility section addressed to mothers, sources of labor income, and detailed consumption expenditure diaries. I use the section on schooling expenditure per child in the last 12 months, including tuition fees and schooling inputs such as textbooks and uniforms, to construct a measure of average educational investments per child at the household level.

[10]The coverage of these variables is not as good for men's samples - for instance, only a quarter of these ask about the employer of male respondents. Therefore, to fully characterise the labor markets of countries in my GDP per capita bin, I make use of additional data, namely the IPUMS censuses.

3 Income-desired fertility gradients

This section establishes an important extension to the well-known fact that fertility has remained persistently higher in sub-Saharan Africa compared to other developing countries in the same national income range. Demographers have coined this fact sub-Saharan Africa's "unique fertility transition" (Bongaarts 2017). I show that this difference is largest amongst the poorest segments of the population, making the income-desired fertility gradient steeper in sub-Saharan Africa, conditional on GDP per capita. I quantify the role of this differential in explaining the total fertility gap observed between sub-Saharan African and other countries. I then provide evidence that the gradient is not steeper in sub-Saharan Africa simply as a result of a more disperse distribution of living standards within countries on average. Finally, I introduce the hypothesis that the findings presented throughout this section are consistent with a more skewed distribution of salaried work opportunities across the income spectrum in sub-Saharan Africa.

Figure 3

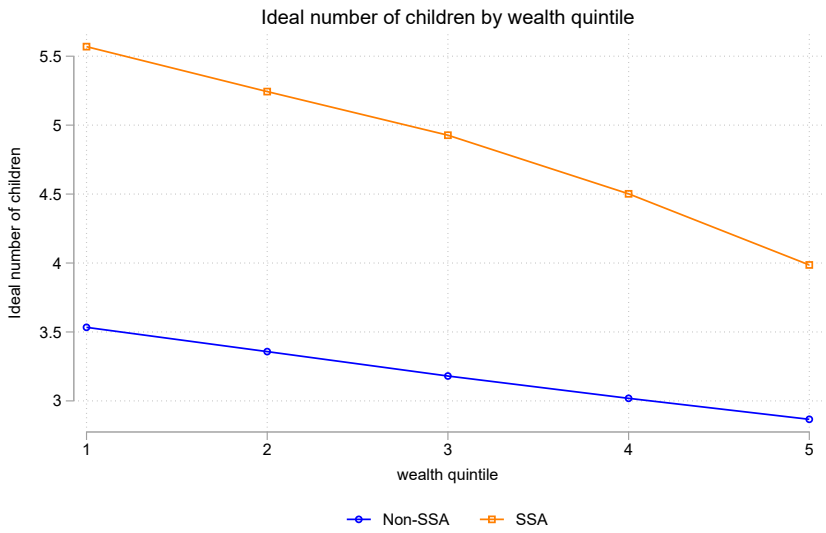

3.1 The gradient

The socioeconomic gradients in desired fertility are displayed in Figure 3. On the vertical axis is the average ideal number of children reported by women across DHS surveys. On the x-axis are wealth quintiles, i.e. proxies for income rankings. The relationship between income and desired fertility is over twice as steep in sub-Saharan Africa, where the average gap in the desired fertility reported by women in the poorest and richest quintile is 1.58 children, than in other countries, where this difference is only 0.67. In terms of ratios, the ideal number of children of women in the poorest quintile is 1.40 that of women in the richest quintile in sub-Saharan Africa versus 1.23 in other regions.

These differentials subsist when I residualise the ideal number of children variable on GDP per capita (see the left-hand side panel of Figure 6 in section 3.4). The difference in the two gradients is also virtually unchanged when I control for country fixed effects in addition to GDP per capita (see Appendix Figure A.4).

3.2 The role of the gradient in explaining sub-Saharan Africa's exceptional fertility

To quantify the role of the gradient in explaining sub-Saharan Africa's exceptionally high fertility rates, I proceed in two steps. First, I compute what the average ideal number of children in sub-Saharan African samples would be if the gradient was the same as in other regions. I find that the average gap in desired fertility between the two sets of countries would be 1.12 children if the gradients were the same, instead of the 1.65 difference observed in the data.

Figure 4

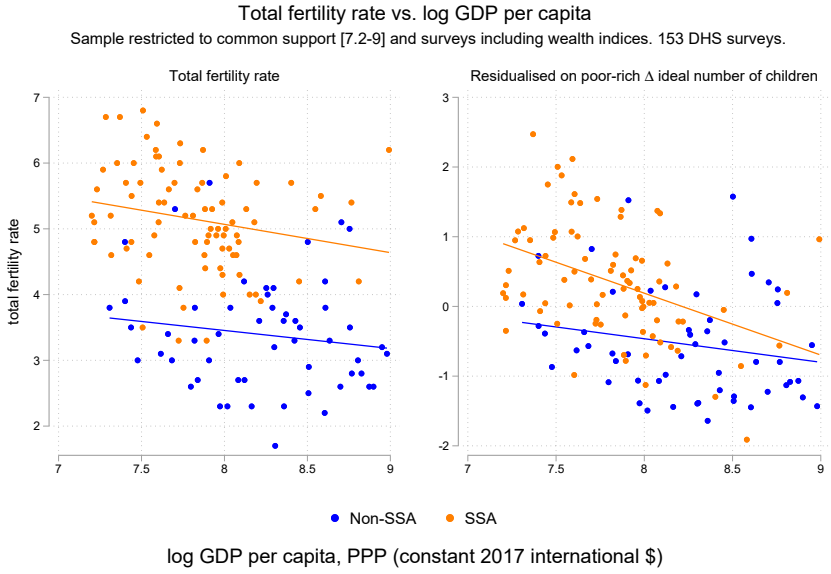

In the second step, I run a simple regression of total fertility rates on log GDP per capita and a dummy variable equal to 1 for sub-Saharan African samples. I then compare the coefficient estimate for the sub-Saharan Africa indicator when I also control for the gradient - as measured by the poor-rich difference in desired fertility. Table A.1 in the appendix displays the results: it shows that controlling for the gradient makes the coefficient estimate on the sub-Saharan Africa dummy drop from 1.689 to 1.147, which corresponds to a 32% decrease in the overall gap in total fertility captured by the dummy when the gradient is excluded from the specification. This suggests that the steeper income-desired fertility gradient accounts for 32% of the fertility gap between sub-Saharan Africa and other regions conditional on GDP per capita.

Figure 4 plots the total fertility rate for the 153 DHS surveys included in this analysis against log GDP per capita on the x-axis. As is clear in panel (a), the cloud of sub-Saharan

African observations lies above the cloud of points from other countries, with only a few exceptions. Panel (b) plots the same relationship, but where the outcome variable (total fertility rate) is residualised on the poor-rich gap in desired fertility. The two clouds of residuals are distinctly closer together.

In sum, these findings suggest that the differential gradient in sub-Saharan Africa explains about one-third of the overall disparity between the region's fertility rates and that of other countries, even conditional on national income.

3.3 Comparability of gaps in living standards

One natural interpretation of the difference in gradients observed in Figure 3 is that we may simply be looking at different gaps in living standards. If households in the poorest quintile are disproportionately poorer relative to the richest quintile in sub-Saharan Africa, the gradients are not comparing desired fertility across the same gap in household assets and living conditions in sub-Saharan Africa relative to other countries. To test this hypothesis, I construct a cross-sample wealth index using principal component analysis (PCA) within my GDP per capita bin of common support.

The notable difference between a standard PCA and the present exercise is that the goal of the latter is to rank households according to an index capturing relative inequality in living standards *across* samples from multiple countries instead of relative inequality in living standards within a given sample from a specific country.

To achieve this, I select a subset of 56 DHS surveys covering 32 countries out of the original 153 surveys in my GDP per capita bin because they all include the same set of variables related to household living standards which have been shown to be important for differentiating households by socioeconomic status while avoiding problems of clumping and truncation as much as possible (McKenzie 2005). These cover eight indicators of dwelling conditions (electricity, source of drinking water, type of toilet, wall, roof and floor material, and cooking fuel, number of household members per sleeping room) and five indicators of durable goods ownership (radio, fridge, bicycle, motorcycle, television). I then harmonize the multiple categories corresponding to each indicator across samples and perform the PCA on this harmonized sample of 650,374 households, 53% of which are from sub-Saharan African samples.

Figure 5 plots the distributions of the resulting cross-sample wealth index, residualised on log GDP per capita, for the two sets of countries. The distribution of the index for Sub-Saharan Africa is more skewed to the left, suggesting a clear negative gap in living standards overall. Indeed, I find that the residualised wealth score of sub-Saharan African households is 1.2 standard deviations lower on average. This suggests that, for the same level of GDP per capita, there are significantly higher rates of poverty in sub-Saharan African countries.

However, this difference in absolute living standards across the two regions is not more pronounced for the poorest quintiles. In other words, the fact that the poor-rich

gap in desired fertility is higher in sub-Saharan Africa than other countries, conditional on GDP per capita, is not due to a larger poor-rich differential in absolute living standards. Table A.2 in the appendix shows this in regression form, namely that the gap in living standards between the top quintile and the two poorest ones is statistically significantly *smaller* in sub-Saharan Africa. This is also illustrated by Appendix Figure A.5, where we see that the gradient between within-sample wealth quintiles and the cross-sample wealth index (on the y-axis) is indeed less steep in sub-Saharan Africa.

Figure 5

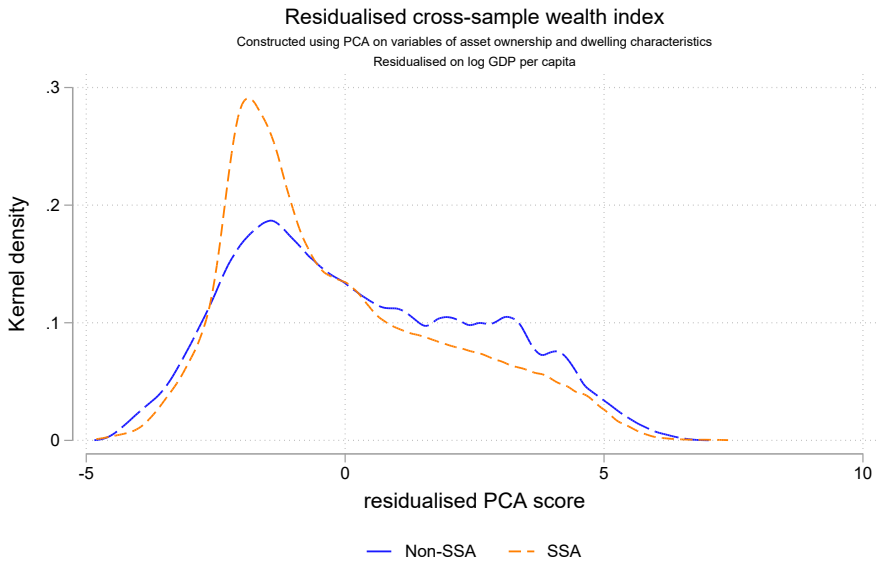

3.4 From the income-desired fertility gradient to jobs

The findings of this section raise a new question: what could explain the steeper gradient between income and desired fertility in sub-Saharan Africa? The hypothesis I propose in this paper is that access to labor market opportunities - and in particular, off-farm salaried jobs - can account for at least part of the differential in gradients across sub-Saharan African and other countries.[11]

This is motivated by the facts displayed in Figure 6. The difference in the gradients plotted in the left-hand side panel (where the y-axis variable is ideal number of children residualised on GDP per capita) is mirrored by a similar differential in the right-hand side panel, which plots the share of women in salaried employment (also residualised on GDP per capita) by wealth quintile. The salaried employment indicator, constructed from variables in the DHS, corresponds to the share of women reporting working for

[11] I rule out two other explanations for the gradient differential. The first is female education: in the final section of the paper, which tests for the opportunity cost of childbearing channel, Figure 12 shows that the gradient in years of education is remarkably similar across the two sets of countries. Appendix Figure A.7 also shows that the gradient differential survives when I control for differential exposure to family planning information by wealth quintile.

someone else and receiving payments for their work. While the difference in levels is narrower across the two sets of countries, the gradient is again clearly steeper in sub-Saharan Africa, as the gap in the share of salaried workers is about twice as small for the top quintile than for the bottom quintile.

Figure 6

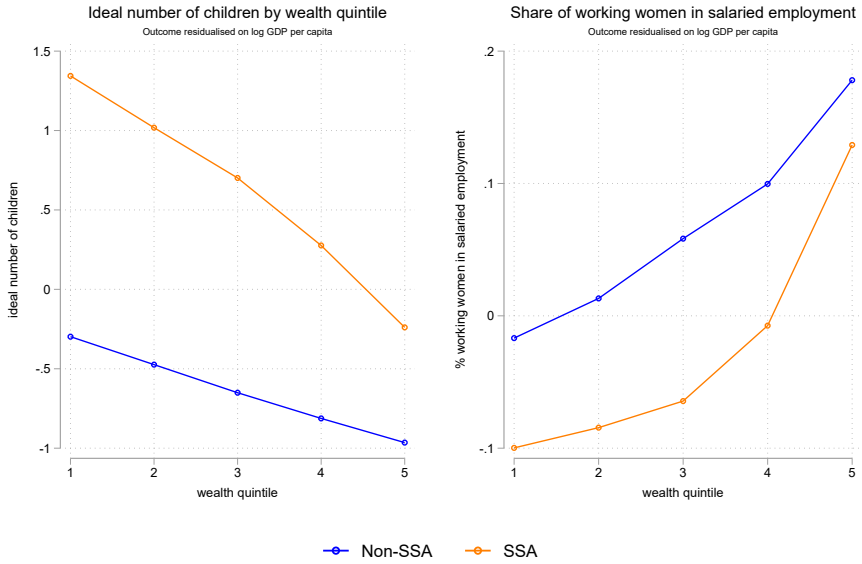

4 Theoretical framework

4.1 Background on fertility theories

A large body of theoretical work attempts to rationalise the negative relationship between fertility and income that has been documented for many developed and developing countries since the start of the demographic transition. A useful review of fertility theories is provided by Jones et al. (2011). Two key takeaways, which guide my thinking throughout this paper, can be drawn from these models.

The first is that theoretical frameworks that predict a negative income-fertility relationship must account for the opportunity cost of parental time, under the assumption that raising children reduces the time that parents can allocate to the labour market. If the substitution effect of higher parental wages offsets the income effect (children being considered a normal good), then higher-wage parents will choose a lower number of children.[12] Beyond the monetary value of parents' time, the time cost parameter in my version of the model will also capture the opportunity cost of childbearing as a function of the *type* of work parents are engaged in. For instance, women working from

[12]Note: in this paper I take household wealth as a proxy for parental **labor** income. Otherwise, under the assumption that children are a normal good, the relationship would be positive.

home, on the farm, or self-employed in off-farm activities where they can set their own hours around their childcare duties, will have a lower opportunity cost of children than women working in salaried employment.

In standard quantity-quality trade-off models, another channel that may contribute to fertility differentials across the income distribution arises if richer parents also place more value on child quality. To formally investigate the link between the steeper income-desired fertility gradient and the fact that individuals in the poorest quintiles are more likely to work in informal occupations in sub-Saharan Africa (Figure 6), I re-interpret this mechanism in terms of the distribution of occupations by level of education. If the local labor market features returns to schooling in the sense that people with even only a few years of education hold more stable and higher-return jobs than people with no education, then low-skilled parents might be more convinced of the benefits of reducing their fertility in favour of increased educational investments in children. In other words, the occupational structure can also shape the value that parents place on education.

The assumption that parents' perception of the returns to educational investments in children may be heterogeneous across levels of income or rural/urban status is not unreasonable in the context of sub-Saharan Africa. Recent evidence suggests that some parents place relatively little value on education: Ashraf et al. (2020) find that parents have lower educational attainment goals for their children than the latter in Zambia.[13]

4.2 Model

This section presents a simple framework, which enriches a standard quantity-quality trade-off model à la Becker and Lewis (1973) with a human capital threshold for participation into salaried employment.[14]

In the first period, parents choose an occupation between informal subsistence work and salaried employment. Labor earnings are an affine function of human capital: $y = a + wh$, where y denotes total labor income, h denotes human capital, w is the wage rate per unit of human capital, and a is a positive constant which guarantees everyone a minimum level of income; one can think of this last term as a unit of productive capacity that everyone is endowed with. The formal occupation yields higher returns to human capital $y = a + \bar{w}h$, but is only accessible to those with human capital above a certain threshold, $h \geq \bar{h}$. Parents below the threshold $\bar{h}$ are thus "forced" to choose the informal occupation, which yields lower earnings: $y = a + \underline{w}h$, where $\underline{w} < \bar{w}$.[15]

[13] Recent research also suggests that there is large dispersion in the quality of schooling within and across developing countries (Figueiredo Walter 2020). Exploring the role of this variation in explaining differences in income-desired fertility gradients is left for future work.

[14] The quantity-quality trade-off optimisation problem is borrowed from Vogl (2016), but makes one simplification by removing the goods cost of children. Assuming a positive goods cost of children would unnecessarily complicate the model's predictions, as its effect on fertility operates in the opposite direction as the time cost. This model thus uses the simplest optimisation problem that unambiguously predicts a negative relationship between parental income and fertility (Jones et al. 2011), since that is what I find in the data (see Figure 3). This is because the goal of this theoretical exercise is to propose an explanation for the difference in the slope of this relationship across regions.

[15] I take the wage rate $w = \{\underline{w}, \bar{w}\}$ as completely exogenous as a starting point. Note that even without assuming a

In the second period, parents take their occupation as fixed, and optimise with respect to three goods: household consumption c, how many children to have n, and how much to invest in each child's education e. In the utility function, I substitute the quality production function $h_{child} = \theta_0 + \theta_1 e$.[16] θ_0 is a human capital endowment; θ_1 is the return to educational investments; and each child costs $\tau \in (0, 1)$ units of time.[17]

Three parameters take on different values from one side of the threshold to the other, capturing the key differences across formal and informal occupations that could affect desired fertility, which were introduced in section 4.1. The opportunity cost of childbearing is "low" ($\underline{\tau} \underline{w} h$) in the informal occupation and "high" ($\bar{\tau} \bar{w} h$) in the formal occupation. The time cost parameter τ widens the wedge in the opportunity cost of childbearing between the two occupations: $\bar{\tau} > \underline{\tau}$. This captures the fact that mothers working in salaried employment cannot take their children to work with them, and typically work fixed hours set by their employer. Likewise, parental valuation of education is "low" ($\underline{\theta_1}$) in the informal occupation and "high" ($\bar{\theta_1}$) in the formal occupation. As discussed in section 4.1, the intuition underlying this assumption is that poorer parents with low levels of education $h < \bar{h}$ are less likely to think that their children could ever access the formal occupation (which yields higher returns to education), thereby biasing their perception of the returns to investing in children's education downwards.

For a given w, θ_1 and τ, households face the following optimisation problem:

$$\max_{c,n,e} U(c, n, e) = \alpha \ln(c) + (1 - \alpha)(\ln(n) + \beta \ln(\theta_0 + \theta_1 e))$$

subject to the budget constraint:

$$c + ne = \underbrace{(1 - \tau n)(a + wh)}_{parental\, labor\, income}$$

where w is parental wage. Importantly, households face borrowing constraints: parents cannot borrow against their children's future earnings to invest in their education.

Given that the vast majority of the households in the samples I assemble offer at least some education to their children, I restrict my attention to the set of interior solutions below:

$$c^* = \alpha(a + wh)$$

$$n^* = \frac{(1 - \alpha)(1 - \beta)(a + wh)}{\tau(a + wh) - \frac{\theta_0}{\theta_1}}$$

$$e^* = \frac{\beta\tau(a + wh) - \frac{\theta_0}{\theta_1}}{1 - \beta}$$

different wage rate for the two occupations, the fixed human capital requirement implies that parental labor earnings are always higher in the salaried employment occupation, by construction.

[16] This follows Becker and Tomes (1976), who show that using an affine quality production function with a positive constant ensures that the model predicts a negative income-fertility relationship.

[17] It would be more realistic, if parents typically rely on older children for help with childcare of young children, to define $\tau(n)$ as a strictly concave function. As emphasized by Moav (2005), this would simply predict an even higher poor-rich gap than the current framework by making the relative marginal cost of quantity even lower for the poor. This is not useful for establishing theoretical pathways from jobs to desired fertility, which is the goal of this model.

The condition for this set of solutions to exist is:

$$a + wh > \frac{\theta_0}{\theta_1 \beta \tau} \tag{1}$$

At every level of parental human capital h for which condition (1) is satisfied, comparative statics predict the following monotonic relationships:

$$\frac{\partial n}{\partial h} < 0, \quad \frac{\partial n}{\partial w} < 0, \quad \frac{\partial n}{\partial \tau} < 0 \quad \frac{\partial n}{\partial \theta_1} < 0, \quad \forall\ \theta_0,\ \theta_1 > 0$$

The sign of the last three derivatives captures the two channels linking jobs to fertility choices in this study. First, the fact that the number of children declines with the time cost τ and the wage rate w corresponds to the opportunity cost of childbearing. According to the model, women will be more likely to reduce their fertility if there are more opportunities to engage in salaried, off-farm work in their location. This can operate through a higher monetary return to their time (w), a higher "effective" time cost of children (τ), or both. Second, fertility declines with the returns to educational investments θ_1. According to the hypothesis outlined in section 4.1, the value parents place on education may be shaped by their experience of the local labor market: individuals working in areas where the occupational structure is largely dominated by informal work may underestimate the benefits of more education for their children.[18]

Without the indivisibility in human capital for obtaining a salaried job, the model therefore predicts the relationship between parental education and the number of children they would bring into the world to be monotonically decreasing. Adding the human capital requirement $\bar{h}$ for entry into salaried employment leads to a discontinuity in the relationship between parental human capital and fertility n. The differences in this relationship across both sides of the threshold are a result of the dichotomies ($\underline{\tau}$ vs. $\bar{\tau}$ and $\underline{\theta_1}$ vs. $\bar{\theta_1}$) I impose in this framework. This is illustrated by Figure 7a: holding everything else constant, parents at the right of the threshold who have secured a salaried job will have fewer children simply as a result of their higher opportunity cost of children $\bar{\tau}\bar{w}h$ and their higher valuation of education $\bar{\theta_1}$.

The choices predicted by the model have important implications for the intergenerational transmission of human capital, which, in this set-up, is key to households' asset accumulation and hence the likelihood that any dynasty with initial human capital stock below $\bar{h}$ will eventually access the salaried occupation. To examine the long-run dynamics, I make a small tweak in notation: let h_t denote human capital in generation t and h_{t+1} denote that of their children. Plugging in the optimal choice of educational investment per child e^* into the child's quality production function $h_{t+1} = \theta_0 + \theta_1 e_t$ yields the following transition equation for human capital:

$$h_{t+1} = \begin{cases} \frac{\beta}{1-\beta}[\underline{\theta_1}\underline{\tau}(a + \underline{w}h) - \theta_0] & \text{if } h_t < \bar{h} \\ \frac{\beta}{1-\beta}[\bar{\theta_1}\bar{\tau}(a + \bar{w}h) - \theta_0] & \text{if } h_t \geq \bar{h} \end{cases}$$

Recall that θ_1 and $\tau w h$ represent the benefit of an additional unit of educational expenditure and the opportunity cost of one more child respectively. As such, they are key

[18] One can expect this channel to be stronger in communities characterised by very low education levels, i.e. far below the human capital threshold for salaried employment, thus making the latter seem "out of reach".

Figure 7

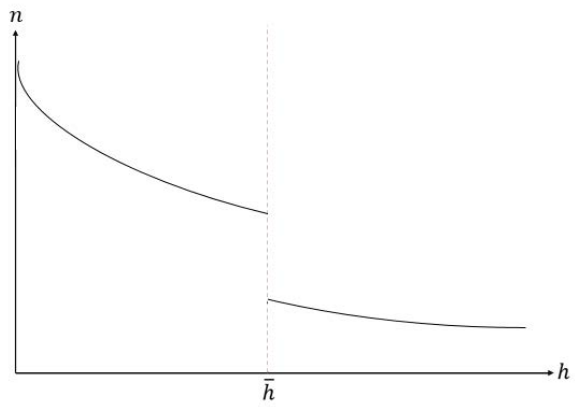

(a) Number of children vs. parental human capital

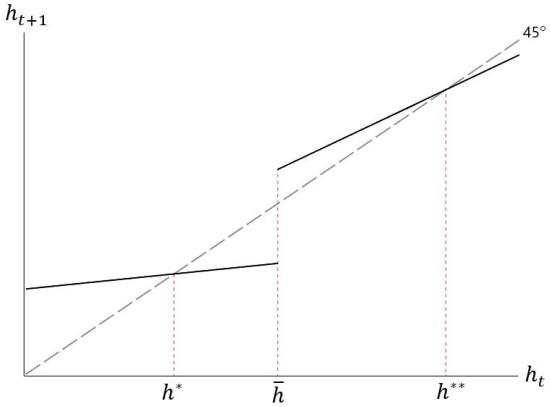

(b) Human capital transition equation

determinants of the quantity-quality trade-off. As θ_1 increases, educational investments become relatively more appealing: the shadow price of quantity increases and parents substitute quantity for quality, thereby increasing h_{t+1} for a given level of parental human capital h_t. Beyond the monetary value of parental time wh, the opposite is true for a decrease in τ: when the time cost associated with one more child drops, parents find quantity relatively cheaper, and so will have more children but these will have lower human capital h_{t+1}. When either of these parameters shifts, so do both the intercept and the slope of the transition equation.

This model embeds the possibility of two steady states: these are depicted as h^* and h^{**} in Figure 7b. h^* represents the "low" steady state, where dynasties that started with low levels of h will converge in the long run. These households work in informal occupations, earn $\underline{w}h^*$, and will persistently choose high n and low e for their children. At the "high" steady state h^{**}, where dynasties that started with $h \geq \bar{h}$ will converge in the long run, households engage in salaried employment, earn $\bar{w}h$, and persistently choose low n and high e for their children.

Conditions for the existence of multiple equilibria. As demonstrated in the theoretical literature on poverty traps (Ghatak 2015), the combination of the human capital indivisibility for entry into salaried employment with borrowing constraints is not sufficient for the existence of the two equilibria depicted in Figure 7b. These exist only for the set of parameter values that meet condition (2):

$$\bar{h} > \frac{\beta(\underline{\theta_1}\underline{\tau}a - \theta_0)}{1 - \beta(1 + \underline{\theta_1}\underline{\tau}\underline{w})} \tag{2}$$

It is worth noting that the wage rate also enters the human capital transition equation. If the average marginal return to an extra year of schooling in the informal occupation $\underline{w}$ is high enough, households who start with $h < \bar{h}$ can invest their way out of the low equilibrium h^*, leading dynasties to eventually converge to h^{**} by reducing their fertility in the process.

Possible extensions. There exist multiple extensions to this framework which could help yield further insight into the drivers of the cross-regional disparity in the income-desired fertility gradient observed in the data. The first one is to extend this framework into a general equilibrium model. In its present form, the framework yields partial equilibrium predictions, which rely on the strong assumption of exogenous labor market conditions. Extending the framework to a general equilibrium set-up would allow one to endogenize the human capital threshold for entry into salaried employment $\bar{h}$ and/or the returns to the two types of work, i.e. wage rates $\bar{w}$ and $\underline{w}$. For instance, a possible implication of the low equilibrium h^* is that, by increasing the stock of low-skilled labour in the informal occupation, low-skilled parents' fertility decisions in period 2 have a negative feedback effect on the informal labor market, by depressing $\underline{w}$.[19] This type of extension would yield predictions for the presence of multiple equilibria at the aggregate, rather than individual, level.

Another potentially useful tweak to the model would be to examine how the predictions change under homothetic preferences. As is commonly used in standard Beckerian fertility frameworks (Jones et al. 2011; Vogl 2016), the current set-up assumes a non-homothetic utility function, ensuring that the share of household expenditure allocated to education e increases with parental income. This implies that inequality in income may be partly responsible for the existence of multiple stable steady states here. As discussed in Ghatak (2015), the presence of strong income effects might lead to multiple stable steady states too and in such cases, the policy recommendations for breaking households out of the high fertility - informal occupation equilibrium may be different from when it is driven by non-convexities and borrowing constraints.

4.3 Examining the gradient differential through the lens of the model

I now use the framework presented in this section to interpret the difference in gradients observed across sub-Saharan African and other countries through the lens of the model's parameters. Figure 8 displays the same relationships as those displayed in Figure 7, but plotting two different curves and lines for the segments of the population whose human capital levels lie below the fixed education threshold $\bar{h}$. The orange curve (line) represents sub-Saharan African countries while the blue dotted one represents other developing countries in the same GDP per capita range.

Figure 8a shows one representation of the model that would predict the socioeconomic gradient differential observed in desired fertility across the two sets of countries. On the y-axis is desired fertility and on the x-axis is parental human capital.[20] The number of children born to the average household in sub-Saharan Africa stands

[19] One straightforward extension that would achieve this prediction is to assume that the production function characterising the informal subsistence activity involves labour and another fixed factor of production, such as land.

[20] In my analysis of the IPUMS micro-censuses, I use parental human capital instead of wealth quintiles as a measure of socioeconomic rankings. This is primarily because I need to test for the assumption that a fixed human capital requirement exists for entry into salaried employment, which I do in section 5.1. I show in two separate figures - Figure 12 in section 5.2.2. and Figure A.8 in Appendix C - that neither the distribution of years of education nor

Figure 8

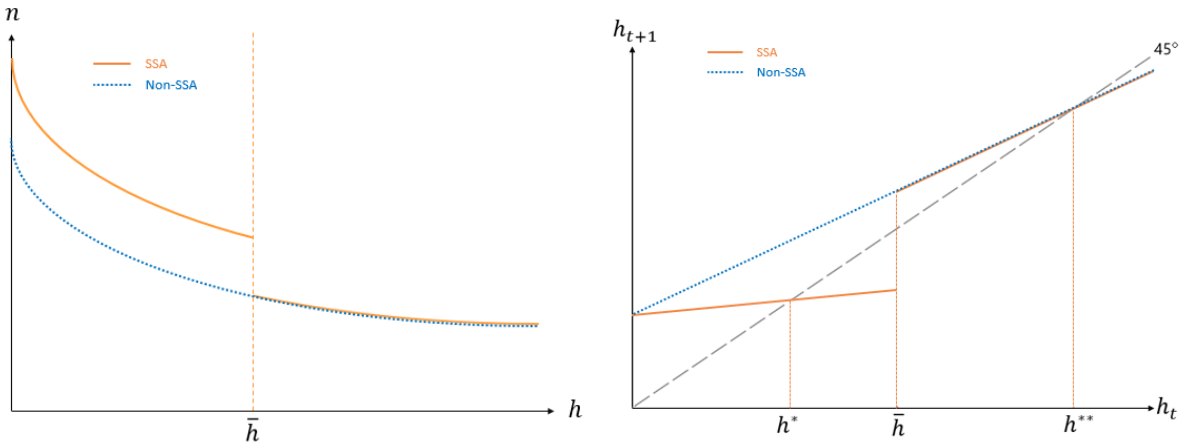

(a) Number of children vs. parental human capital

(b) Human capital transition equation

above that in other countries for levels of parental human capital below the threshold, after which the relationships are the same. This reflects two possible scenarios according to the model. The first is that the human capital threshold does not exist outside of sub-Saharan Africa.[21] In this case, the blue dotted curve depicts the predicted relationship between parental human capital and desired fertility in both sets of countries when $h \geq \bar{h}$, whereas, for $h < \bar{h}$, the sub-Saharan African curve is the orange one.

Under the second scenario, the fixed human capital requirement for entry into the formal occupation is the same in both sets of countries, but at least one of the parameters linking parents' jobs to desired fertility in the informal occupation is higher outside of sub-Saharan Africa, namely parental valuation of education $\underline{\theta_1}$ or the opportunity cost of children $\underline{\tau}$.[22] What is depicted in Figure 8a is an extreme case of this scenario, where the opportunity cost and valuation of education parameters are the same whether parents work in the formal or informal occupation (namely a given τ and θ_1) outside of sub-Saharan Africa, leading to the same predicted relationship between h and n as in the absence of a human capital threshold for entry into salaried employment.

Figure 8b shows the human capital transition paths corresponding to these two scenarios respectively. Here again, all countries are assumed to follow the same path for levels of parental human capital above $\bar{h}$ - this is depicted by the two lines (the blue dotted line and the smooth orange one) overlapping from $\bar{h}$. The human capital transition path followed by households in sub-Saharan African countries depicted here is one where dynasties starting with $h < \bar{h}$ will be stuck in equilibrium h^* in the long run. That is, parameters $\bar{h}$, $\underline{\theta_1}$ and $\underline{\tau}$ meet condition (2), and there exist two equilibria. This

the relationship between wealth and years of education is statistically significantly different across the two sets of countries. Therefore, using years of education should preserve the socioeconomic ranking.

[21] I consider that the relative scarcity of manufacturing firms, a well-documented characteristic of sub-Saharan Africa's structural transformation process (Osei and Jedwab 2013; Gollin et al. 2016) makes this a plausible scenario.

[22] One example of this could be casual wage labour in agriculture. This type of informal work takes place outside the home and may not feature the flexibility required for raising many children at the same time.

implies that there is a kink in the relationship between h_t and h_{t+1} at $\bar{h}$: intergenerational mobility in education is lower at levels of parental education below $\bar{h}$ in these countries.

Assuming, first, that the fixed human capital requirement for entry into salaried employment does not exist outside of sub-Saharan Africa, then the relationship between parents' human capital and that of their children is linear in these countries. This transition path is represented by the blue dotted line. This line also depicts the intergenerational transmission of human capital outside of sub-Saharan Africa under the second scenario outlined above, in which both sets of countries face the same human capital threshold $\bar{h}$ but $\underline{\theta_1}$ and/or $\underline{\tau}$ are higher outside of sub-Saharan Africa. What is represented here is, again, the extreme case where θ_1 and τ are the same whether parents work in the formal or informal occupation, in which case the human capital transition equation is the same on both sides of $\bar{h}$.[23]

5 From Theory to Empirics

In this section I use micro data from three different sources to validate the model presented in the previous section. In the first sub-section, I provide empirical support for both the key assumption and the key prediction of the model. Section 5.2 then provides evidence for the two main channels linking jobs to fertility in the framework: parental valuation of education and jobs as the opportunity cost of children.

5.1 Model validation

The two empirical facts presented in this section, Figures 9 and 10, were produced using IPUMS micro-census data. The figures plot cross-regional comparisons constructed from a set of 49 IPUMS samples: 23 from sub-Saharan Africa (the orange curves) and 29 from other developing countries (the blue curves) within the same log GDP per capita bin of common support [7.2;9] to which I restrict my attention throughout this paper. To account for the fact that sub-Saharan African samples are more concentrated towards the lower end of this range and that my outcomes of interest (salaried employment rates and educational investments) are bound to vary with national income, I residualise these on GDP per capita. Each curve plots the smoothed values of a local polynomial regression of the y-variable on the x-variable, along with 95% confidence bands.

Figure 9 provides empirical support for the key assumption of the model: the presence of a discontinuity in the relationship between education and the probability of holding a salaried job. Breaking down the share of individuals in salaried employment by years of schooling reveals an important fact: the relationship between education and entry into salaried employment is convex in both sets of countries, but this feature is more

[23] Even in a less extreme version of this scenario, where $\underline{\theta_1}$ and/or $\underline{\tau}$ are higher outside of sub-Saharan Africa but still lower than in the formal occupation, leading to a kink also in the blue dotted line, this line could lie above the orange one at levels of h below $\bar{h}$ but without generating multiple equilibria.

Figure 9

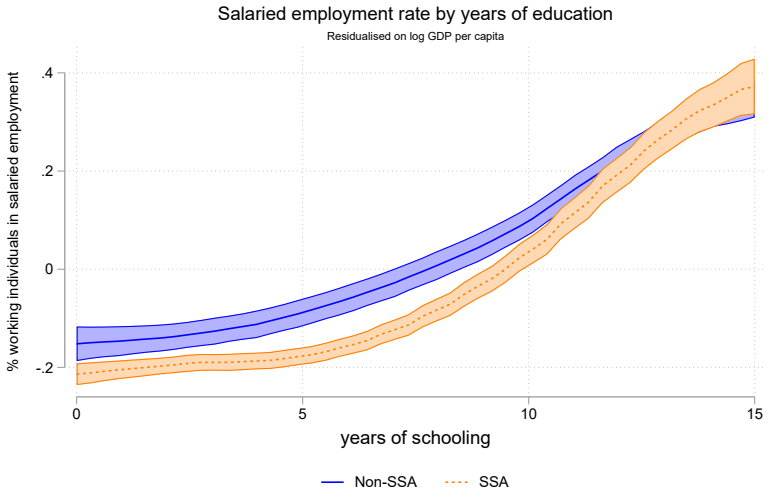

49 IPUMS samples: 23 from SSA, 26 from other developing countries.
Total observations: 37,499,290.

pronounced in sub-Saharan Africa. The share of individuals in salaried employment is virtually flat (around less than 10% of individuals on average), for the first five years of education in sub-Saharan Africa. There is a significant gap between sub-Saharan Africa and other regions in the 5-9 years of schooling range, after which the slope becomes steeper in sub-Saharan Africa, with the curves joining up around 13 years of schooling.[24]

This new finding is in line with the two theoretical explanations I propose for the steeper socioeconomic gradient in desired fertility in sub-Saharan Africa in section 4.3, illustrated by Figure 8a. Either the fixed human capital requirement for salaried employment exists in sub-Saharan Africa only (or at least this feature is more common in sub-Saharan African labor markets than that of other countries) or it exists in both sets of countries but the type of low-skilled jobs available outside of sub-Saharan Africa is such that the time cost τ and/or the value of education θ_1 is comparatively higher in those countries. The latter interpretation supposes that the outcome examined here - whether the individual is recorded as a "wage/salary worker" in the census - encompasses some wage workers in *informal* occupations, such as casual wage labourers in agriculture.[25]

Figure 9 is also consistent with the framework's underlying mechanisms linking fertility decisions to labor market opportunities. Parents in the poorest segments of the population, facing the flat part of the curve, may have a comparatively lower perception of the returns to education in sub-Saharan Africa. If women's labor market options

[24] As shown in section B of the Appendix and later in the paper in Figure 12, the distribution of adult individuals' years of schooling is not significantly different across the two sets of countries.

[25] Note that these cross-regional comparisons inevitably hide substantial heterogeneity across countries. This is illustrated in Appendix Figure A.9, which plots the relationship between years of schooling and the share of workers in salaried employment for 25 samples separately. Some countries outside of sub-Saharan Africa, like Bangladesh and Vietnam, feature a much higher share of salaried workers at very low levels of education than all sub-Saharan African countries **and** a non-linear relationship between years of schooling and the probability of working in salaried employment. Other countries, like El Salvador or Morocco, present a much more linear relationship, which is more consistent with the absence of a human capital threshold for entry into salaried employment.

Figure 10

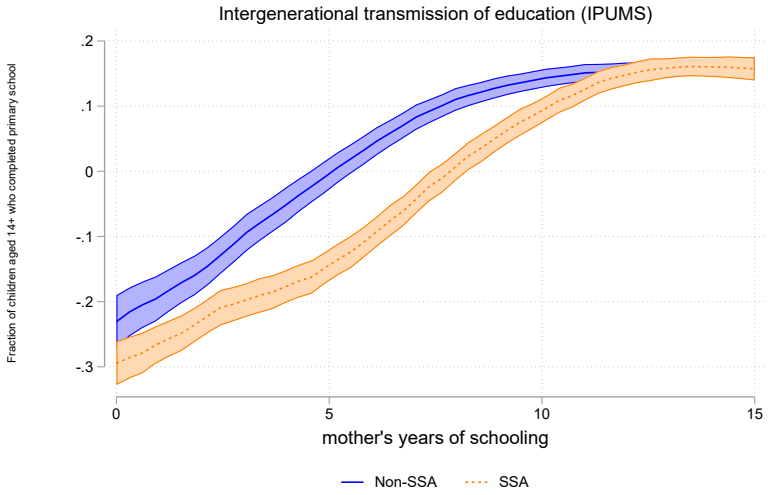

49 IPUMS samples: 23 from SSA, 26 from other developing countries.
Total observations: 25,542,840 mothers.

are less likely to involve salaried work outside the home, this would also reduce the opportunity cost of children and thereby curb their incentives to reduce their fertility.

Figure 10 shows the fraction of children aged 14 and above who have completed primary school (residualised on GDP per capita) on the y-axis against their mother's years of schooling on the x-axis. It is the empirical counterpart to Figure 8b in the previous section. There are two key takeaways to draw from this figure. First, the fact that the sub-Saharan African curve lies below that of the other low-income countries at almost every level of mother's education except for the top of the range points to a relative under-investment in children's education in sub-Saharan Africa. Second, the figure also displays a notable difference in slopes at lower levels of mother's years of education, which causes the gap between the two curves to widen until around 5 years of education. The slope of the orange curve then increases from about 5 years of education, and eventually the two curves overlap, from around the 12-year mark. This mirrors the behaviour of the curves plotted in Figure 9.

Examining these patterns through the lens of the model's parameters yields different possible interpretations of Figure 10. Assuming that the human capital requirement for salaried employment exists in both sets of countries, the flatter slope of the sub-Saharan African curve could be attributed to a lower time cost of children $\underline{\tau}$ and/or a lower perception of returns to educational investments $\underline{\theta_1}$ by parents engaged in the informal occupation in sub-Saharan Africa.[26] These could explain comparatively lower educational investments in children among the poorest segments of the population in sub-Saharan Africa.

Another possible interpretation of Figure 10, when examined jointly with Figure 9,

[26]Of course, the slope of the human capital transition equation is also partly determined by another parameter, returns to human capital w. I abstract from considerations about the wage rate for explaining these cross-regional disparities, as it would be unreasonable to assume it to be comparable across countries.

is that the fixed human capital requirement for salaried employment is simply absent in at least some countries outside of sub-Saharan Africa.[27] The fact that there is a kink in the orange curve around 5 years of education but no detectable change in the slope of the blue curve at that point provides particular support for this interpretation. Taken together, these facts point to the existence of multiple equilibria, where the poorest families are trapped in a high fertility - informal occupation equilibrium characterised by lower educational investments in children.

In sum, these two figures validate both the core assumption and the key prediction of the model outlined in section 4. Figure 9 suggests that the presence of a human capital requirement for participating in salaried work is characteristic of sub-Saharan African labor markets. Importantly, for this fact alone to offer an explanation for the steeper wealth-desired fertility gradient, it has to be interpreted as this threshold being absent in at least some countries outside of sub-Saharan Africa. Figure 10 shows that most parents make comparatively lower educational investments in children in sub-Saharan Africa, where the slope of the human capital transition path is also flatter at low levels of parental education, and then becomes similar to that of other countries at around the same level of schooling associated with a steep increase in the share of salaried workers.

These findings suggest that at least part of the explanation for the steeper socioeconomic gradient in desired fertility observed in sub-Saharan Africa resides in the existence of fixed human capital requirements for entry into salaried employment that make it particularly difficult for the poorest of the poor to transition from informal to formal occupations in these countries.

5.2 The role of jobs in shaping fertility preferences in sub-Saharan Africa

This section investigates the labor market determinants of variation in desired and realised fertility within sub-Saharan Africa. I test for the presence of the two key channels linking parents' jobs to fertility choices in the framework. In order to isolate the variation in fertility preferences that is correlated with individuals' labor market outcomes from the variation that could be attributed to cultural factors, all specifications control for ethnic group and religious group fixed effects.

Sub-section 5.2.1 presents evidence in line with the parental valuation of education channel, whereby parents' employment outcomes can affect their perception of returns to education and, through this, their willingness to substitute quantity for quality of children. Since this channel can operate through the labor market outcomes of either parent, I use men and household heads as the unit of analysis in this first sub-section.

Sub-section 5.2.2 provides empirical support for the other key auxiliary prediction of the model, namely that desired fertility is negatively correlated with the availability of suitable job opportunities outside the home because these raise the opportunity cost of

[27] As shown by Figure 8 in section 4.3, the existence of a human capital threshold only in sub-Saharan Africa can account for different intergenerational dynamics of human capital investments across the two sets of countries.

bearing children. Since women traditionally bear most of the time cost associated with having children, I focus on women as the unit of analysis to test for the presence of this second channel in the data.

Table 1: Correlates of men's desired fertility in sub-Saharan Africa.

	(1)	(2)	(3)	(4)	(5)
urban	-1.074***	-0.857***	-1.013***	-1.031***	-0.842***
	[0.145]	[0.125]	[0.143]	[0.142]	[0.126]
years of education	-0.163***	-0.150***	-0.160***	-0.161***	-0.150***
	[0.009]	[0.008]	[0.009]	[0.009]	[0.008]
works outside of farm		-0.515***			-0.422***
		[0.077]			[0.071]
works for pay			-0.372***		-0.203**
			[0.094]		[0.086]
works all year				-0.233***	-0.116**
				[0.065]	[0.053]
R^2	0.273	0.275	0.274	0.273	0.275
Ethnic group FE	X	X	X	X	X
Religion FE	X	X	X	X	X
Survey FE	X	X	X	X	X
Outcome mean	6.318	6.318	6.318	6.318	6.318
Observations	150531	150531	150531	150531	150531

Standard errors clustered by DHS survey (country x year) in brackets. Controls for age.
* p<0.1, ** p<0.05, *** p<0.01.

5.2.1 Jobs and parental valuation of education

Table 1 displays correlations between men's desired fertility and their labor market outcomes in the DHS data, pooling together all the men's samples from sub-Saharan Africa that contain information on men's employment, ethnic group and religion. The first two rows of estimates consistently show that higher levels of education and urban residence are associated with lower desired fertility in men: column (1), for instance, shows that an additional year of male education is associated with a drop in 0.163 ideal number of children.[28] Urban residence is also associated with lower desired fertility of up to one child. This large correlation partly captures the fact that more men work outside of agriculture in urban areas, since the size of the coefficient magnitude decreases when adding an indicator for off-farm work to the specification (columns (2) and (5)).

Three measures of men's labor market outcomes - those with the best coverage in the DHS men's samples - are controlled for separately in columns (2), (3) and (4). The coefficient estimates suggest that each of these indicators is negatively correlated with

[28]Since the dependent variable is an integer here, this suggests that a significant fraction of men with an additional year of education report a lower ideal number of children, relative to those with less education.

Figure 11: Fertility and educational investments, by household head's type of employment.
Kenya Integrated Household Budget Survey 2005-2006.

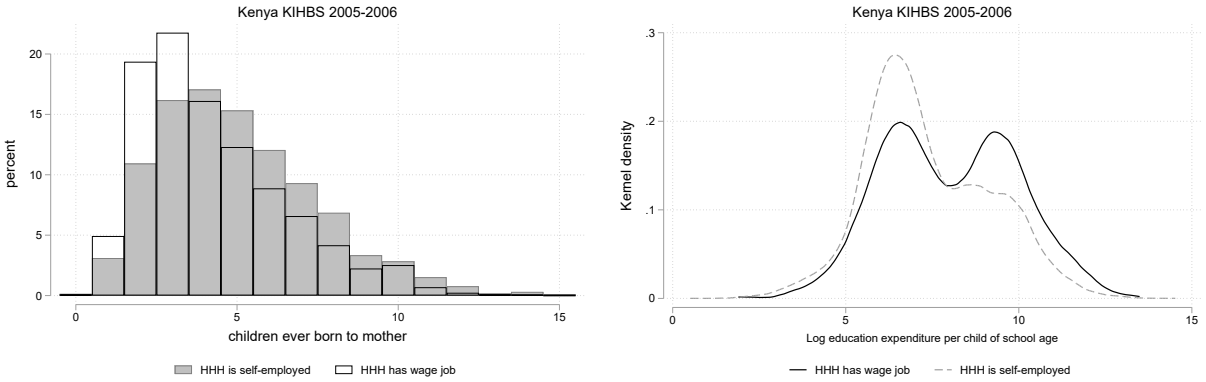

(a) Children ever born to mother

(b) Distribution of education expenditure

men's desired fertility. For instance, column (2) reports that men working outside of agriculture, conditional on the specification's set of covariates, report an ideal number of children about half a child lower on average. Since working outside of agriculture is also likely to be correlated with working for pay and seasonally, I also control for these three characteristics together. The estimates from this specification are reported in column (5). They show that men's ideal number of children is about three quarters of a child lower if they work for pay, outside of agriculture, and all year round, conditional on their level of education, whether they live in an urban or rural location, and a set of religion and ethnic group fixed effects. This corresponds to a 12% decrease relative to the sample mean of men's reported ideal number of children.

A key takeaway from Table 1 is, therefore, that men's employment characteristics are strong predictors of their fertility preferences, even conditional on their level of education. While these estimates do not reflect the presence of a causal relationship from men's jobs to fertility intentions, they are consistent with the "valuation of education" channel playing a role in the relationship between labor market outcomes and desired fertility.[29] In terms of the model's parameters, I interpret these findings as suggestive of parents' experience of the labor market shaping their perception of the returns to educational investments θ_1.

To yield insight into the relationship between parents' labor market outcomes and educational investments in children, I complement the DHS surveys with a household expenditure survey from Kenya. I find that Kenyan households where the household head has a salaried job spend more on each child's education and have fewer children. This is illustrated in Figure 11: panel (a) depicts the distribution of the total number of children ever born to the spouse of the household head, and panel (b) plots the kernel density distribution of log education expenditure per child.[30] Panel (a) shows that the

[29] Using men as the unit of analysis here makes it less likely that these correlations capture the "opportunity cost of childbearing" mechanism, which is more likely to affect women's fertility preferences.

[30] 5% of households report no education spending and are excluded from this analysis due to the log transformation.

distribution of total number of children born to couples where the head holds a salaried job is shifted to the left relative to that of households where the head is self-employed (the grey bars). Panel (b) shows that the distribution of expenditure per child is more distinctly bimodal - as well as shifted to the right - for households where the head is in salaried employment (the smooth black line).

Taken together, these facts are consistent with the idea that parents in more stable forms of employment are also more prone to substituting quantity for quality of children.[31] In the Appendix, Table A.3 presents these associations in regression form, controlling for urban/rural status and the household head's years of schooling. It shows that these correlations remain statistically significant when controlling for these covariates. The interpretation of these findings in line with the "valuation of education" channel is that household heads in salaried employment are more likely to reduce fertility in favour of larger educational investments in their children.

Figure 12: Women's years of education by wealth quintile

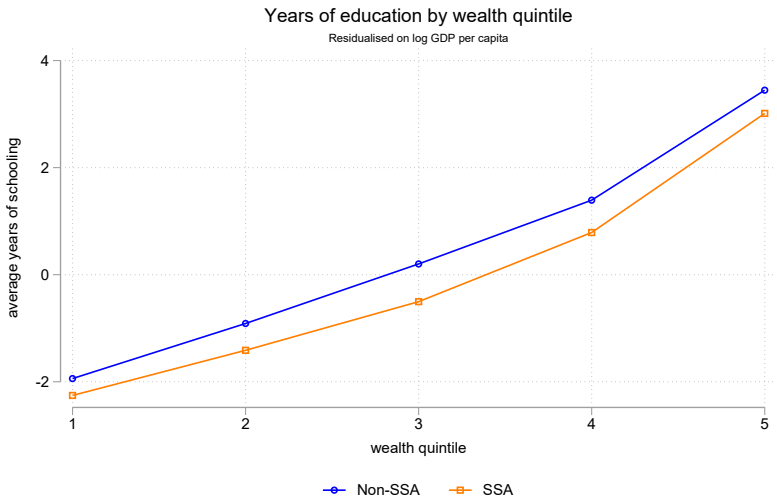

153 DHS surveys. 83 from sub-Saharan Africa, 70 from other developing countries. Total observations: 2,575,486 women aged 15-49.

5.2.2 Jobs as the opportunity cost of bearing children

The opportunity cost of childbearing argument is simple: if raising children is time-intensive, parents whose time could yield higher monetary returns in the labor market should be more inclined to reduce their fertility. Since women traditionally bear most of the time cost associated with having children, the steeper income-desired fertility gradient observed in sub-Saharan Africa could thus point to a comparatively lower monetary value of poor women's time on average in this region.

[31] The discussion of the role of jobs in fertility choices abstracts from direct income considerations. Of course, under credit constraints, higher income parents can afford more education for their children. Despite the abolition of primary school fees in Kenya in 2003, households still face other costs (see Appendix Figure A.11).

Figure 12 shows that the average years of female education are virtually indistinguishable across the two sets of countries, controlling for GDP per capita. The level differences are statistically insignificant for each quintile, and the slopes of the relationship with wealth are almost identical. This suggests that if the female opportunity cost of time channel is at play in explaining the cross-regional differences in the desired fertility gradient, it does not operate through female education.[32]

If we consider years of education to be a reliable measure of women's productive capacity, then one explanation for the fact that the gradient in years of education is similar across the two sets of countries while desired fertility is much higher among the poorest sub-Saharan African women is that the returns to education are comparatively lower for the poor in sub-Saharan Africa, due to the relative scarcity of job opportunities.[33]

Figure 13

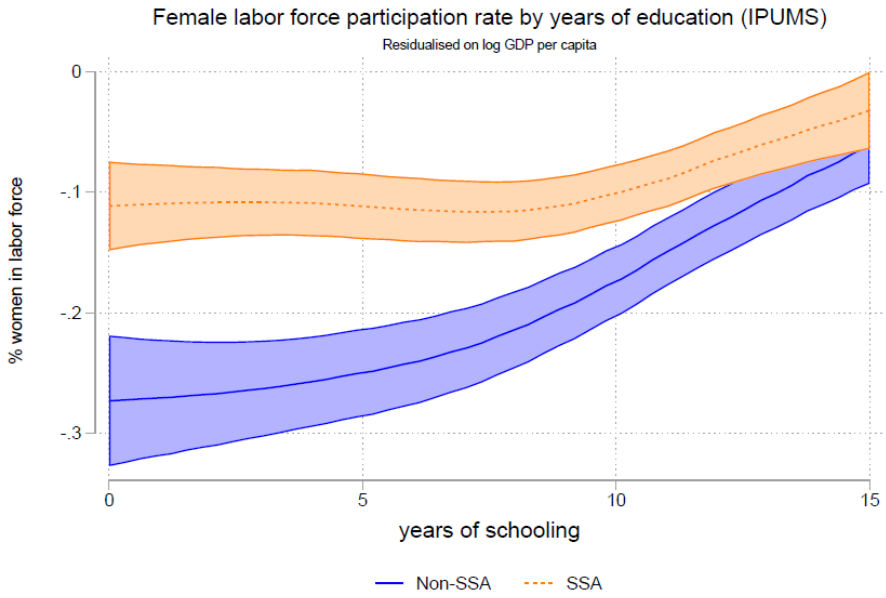

49 IPUMS samples. 23 from sub-Saharan Africa, 26 from other developing countries. Total observations: 22,909,009 women aged 25-50.

Figure 13 plots the relationship between female labor force participation and years of schooling. Two things are of note. The first is that there is a very large level difference - of almost 20 percentage points - across the two regions, controlling for GDP per capita: labor force participation is significantly higher in sub-Saharan Africa at almost every level of education (except the very top of the distribution). This stands in contrast to the documented fact that female labor force participation typically follows a U-shaped relationship with economic development (Goldin 1995). In other countries, it has indeed

[32] Heath and Jayachandran (2018) review the impacts of female education on fertility in developing countries. They find that most studies find a delaying effect of education on *early* fertility (e.g. Breierova and Duflo 2004; Monstad et al. 2008). Several channels may be responsible for the effects of education on early fertility, including increased use of contraceptives and delayed marriage (e.g. Keats 2018). The only study (set in Nigeria) to provide robust evidence of a labor market mechanism also finds that education leads to fewer *total* children born (Osili and Long 2008).

[33] Years of education only measure skills accurately if the quality of schooling is comparable across countries.

been observed that the early stages of the structural transformation process, typically characterised by a shift from informal (such as subsistence agriculture) to formal work (such as manufacturing) in the occupational structure, coincide with a drop in female labor force participation. Here, we see a much larger share of women participating in the labor market in sub-Saharan Africa relative to other countries, for the same level of GDP per capita.[34] In line with the overarching takeaway in this paper, this suggests that sub-Saharan Africa's occupational change process is, much like its fertility transition, unique.

The second finding to emerge from Figure 13 is that the relationships are vastly different: in sub-Saharan Africa, female labor force participation is flat with respect to years of schooling until about nine years of education, while in other countries, the curve is (almost) monotonically increasing, with a clear positive relationship appearing from around 4 years of schooling. One natural interpretation of the pattern observed for the countries outside of sub-Saharan Africa is that, as the labor market starts to yield returns to their human capital, women increase their participation in the labor market and reduce their fertility. In this case, women's decision on the extensive margin (whether or not to participate) is influenced by the options available to them on the intensive margin (how the labor market rewards their time if they participate). Salaried employment opportunities in manufacturing or services, for instance, may provide them with an incentive to both increase their educational attainment and reduce their fertility.[35]

Figure 13 thus points to the importance of defining the extensive and intensive margins with respect to specific features of the sub-Saharan African context. Poor women are more likely to be in the labor force in sub-Saharan Africa; but the work they do inside the labor market tends to be low-income, informal occupations. The patterns uncovered thus far suggest that these typically yield low returns to their education - at least in terms of the probability of obtaining a salaried job - all the while providing them with the flexible working hours required for raising many children. I therefore focus on variation in the *nature* of women's work in sub-Saharan Africa to test for the presence of the "opportunity cost of childbearing" channel from here onwards.

I make use of the DHS and IPUMS individual-level data to correlate the desired and realised fertility outcomes of female respondents to measures of the availability of employment opportunities in their location of residence. Tables 2 and 3 display results of these correlations, focusing on the availability of two types of jobs as explanatory variables: off-farm work and salaried employment. I define off-farm work as a broader occupational category that encompasses all activities outside of farm work, whether the individual is self-employed or a salaried worker employed by someone else.

[34] Since the distribution of female years of schooling is virtually the same in the two sets of countries (see Appendix Figure A.8), it must be that overall female labor force participation rates are also substantially higher overall in sub-Saharan Africa, conditional on GDP per capita.

[35] Jensen (2012) and Heath and Mobarak (2015) exploit random and quasi-random variation in India and Bangladesh respectively to test this hypothesis, and find that the arrival of new employment opportunities - business process outsourcing in Jensen (2012) and garment manufacturing in Heath and Mobarak (2015) - do indeed raise women's labor force participation by increasing their educational attainment and delaying the onset of their fertility.

Table 2: Correlations between women's ideal number of children and the availability of non-farm and salaried employment in sub-Saharan Africa (DHS).

	Ideal number of children					
	(1)	(2)	(3)	(4)	(5)	(6)
urban	-0.545***	-0.542***	-0.536***	-0.540***	-0.516***	-0.516***
	[0.038]	[0.038]	[0.038]	[0.038]	[0.038]	[0.039]
years of education	-0.113***	-0.111***	-0.111***	-0.112***	-0.111***	-0.111***
	[0.004]	[0.004]	[0.004]	[0.005]	[0.005]	[0.005]
% salaried in province	-1.978***	-1.900***	-2.289***			
	[0.542]	[0.536]	[0.577]			
salary worker		-0.122***	-0.376***			
		[0.026]	[0.047]			
salary worker*% salaried in province			1.676***			
			[0.453]			
% non-farm in province				-0.890***	-0.828***	-0.835***
				[0.194]	[0.197]	[0.208]
non-farm worker					-0.129***	-0.134**
					[0.021]	[0.060]
non-farm worker*% non-farm in prov.						0.013
						[0.152]
R^2	0.361	0.361	0.361	0.361	0.361	0.361
Outcome Mean	5.228	5.228	5.228	5.228	5.228	5.228
Ethnic group FE	X	X	X	X	X	X
Religion FE	X	X	X	X	X	X
Survey FE	X	X	X	X	X	X
Observations	600968	600968	600968	600968	600968	600968

Note: * p<0.1, ** p<0.05, *** p<0.01. Controls for age. Standard errors clustered by province in brackets.

Table 2 displays correlations produced from a sample of over 600,000 women in sub-Saharan Africa, constructed by pooling all the individual-level DHS data on women. The outcome variable here is women's reported ideal number of children. One striking fact is that the coefficient highlighted in column (1) is much more negative – over twice as large in magnitude – than the one in column (4). Column (1) shows that an increase of 10 percentage points in the share of women working in salaried employment in a province is associated with a drop in average ideal number of children of about 0.2. The corresponding coefficient estimate for off-farm work in column (4) is only 0.09. I interpret this correlation as in line with the paper's overarching hypothesis that the availability of salaried employment - which is typically higher-return and less compatible with raising many children - increases the opportunity cost of children.

Column (3) displays the results of a specification in which I interact individuals' employment status (whether or not the respondent is a salaried worker) with the share of salaried employment in the province. It is more intuitive to interpret the coefficient estimate on the interaction term in column (3) as women in salaried employment reporting an even lower ideal number of children when salaried jobs are particularly scarce. Then the estimate can be interpreted as follows: conditional on a woman being a salaried worker, a 10 percentage point decrease in the rate of salaried employment is associated with a 0.17 drop in her ideal number of children. This suggests that the value of having one of these salaried jobs increases the scarcer these jobs are – and this effect manifests itself with an even lower desired fertility. The fact that the interaction yields an insignif-

Table 3: Correlations of women's participation in salaried employment and fertility with the share of salaried employment in their district of residence, by migration status (Malawi IPUMS 2008).

	Non-migrant	Migrated for marriage	Last migrated before marriage	Last migrated after marriage
	(1)	(2)	(3)	(4)
		Salaried employment		
years of schooling	0.003***	0.014***	0.012***	0.019***
	[0.000]	[0.001]	[0.001]	[0.001]
% salaried employment in current location	0.472***	0.263***	0.448***	0.201***
	[0.044]	[0.052]	[0.058]	[0.040]
% salaried employment in birth district		0.162***	-0.019	0.031
		[0.062]	[0.058]	[0.049]
R^2	0.079	0.119	0.117	0.131
Outcome Mean	0.045	0.090	0.087	0.116
		Number of children ever born		
years of schooling	-0.064***	-0.089***	-0.109***	-0.095***
	[0.002]	[0.006]	[0.007]	[0.005]
% salaried employment in current location	-0.862***	-1.239***	-1.774***	-2.021***
	[0.291]	[0.210]	[0.294]	[0.295]
% salaried employment in birth district		-0.239	-0.314	-0.935***
		[0.257]	[0.371]	[0.297]
R^2	0.563	0.629	0.542	0.465
Outcome Mean	4.052	2.532	3.503	4.283
Ethnic group FE	X	X	X	X
Religion FE	X	X	X	X
Age FE	X	X	X	X
Observations	104590	10140	10031	22464

Also controls for rural/urban location and for husband's years of education and type of work (farm/off-farm and informal/formal). Standard errors clustered by sub-district in brackets. *p<0.1, **p<0.05, ***p<0.01.

icant coefficient estimate for off-farm work in column (6) is, once again, consistent with salaried employment being associated with a higher opportunity cost of children.

The correlations between the share of female salaried employment in a province and women's desired fertility reported in columns 1-3 of Table 2 could only be capturing hidden heterogeneity in social norms. For instance, areas characterised by more conservative gender norms may exhibit both a lower share of women in salaried employment and higher levels of desired fertility. To test this hypothesis, I exploit variation in the destination of female migrants from the same district in the Malawi 2008 census.

Table 3 displays correlations between the availability of salaried employment opportunities for women and (a) the probability that they are in salaried employment, and (b) the total number of children they ever gave birth to. I split the sample of married women according to four groups: those who never left their district of birth (column (1)); those who crossed district borders around the year of their marriage, which I interpret to mean that their primary motive for moving was marriage and not the search for work opportunities (column (2)); those whose last migration dates back to before their marriage (column (3)); and finally, those whose last migration to date took place after their marriage. To assuage concerns about the role that assortative matching in marriage

outcomes may play in explaining these correlations, I also control for husband's years of education and type of work (whether they are working outside the farm, and whether they are in salaried employment) in all specifications.

Two main patterns emerge from the estimates displayed in the bottom panel of Table 3. First, the different magnitudes of the coefficient estimate for the rate of salaried employment in women's current location variable across columns (1)-(4) undoubtedly reflect that selection plays a role in the decision to migrate away from one's district of birth. For example, the largest estimate, in column (4), could capture the fact that women's decision to move locations after the year of their marriage - so, presumably with their husband - is likely to be driven by unobservable characteristics of these individuals or couples that are also correlated with lower desired fertility.[36]

Nonetheless, the results in columns (2)-(4) suggest that women who move to areas with higher rates of salaried employment tend to give birth to fewer children. Column (2) reports the results of the specification I consider most likely to isolate the opportunity cost of childbearing channel from selection factors. This is because the timing of these women's last migration suggests that marriage, more than the search for appealing job opportunities, was their primary motive for moving. This estimate states that a 10 percentage point increase in the share of female salaried employment in a woman's current location is associated with a 0.12 drop in the number of children ever born, conditional on a set of husband's characteristics, ethnic group and religion fixed effects, age fixed effects, woman's education, rural/urban residence, and the share of female salaried employment in the woman's district of birth.

The second noteworthy finding emerging from Table 3 is that the coefficient estimate on the availability of jobs is much more negative (and significant) for women's current location of residence than their location of birth. These estimates are reported in the bottom two rows of Table 3. This finding lends support to the opportunity cost of childbearing hypothesis, because it suggests that fertility is not solely determined by social norms and cultural practices. If norms around fertility were driving the negative association between the availability of jobs in a province and women's desired fertility in Table 2, one would expect the rate of salaried employment in a woman's location of *birth* to be a strong predictor of her fertility choices. The underlying intuition here is that women born in districts where more conservative gender norms are more widespread are also likely to have grown up in settings with lower rates of female salaried employment and higher fertility rates. The fact that the estimates in the bottom row of Table 3 are either insignificantly different from zero (columns (2) and (3)) or at least much smaller in magnitude than the coefficient on the rate of female salaried employment in women's *current* location (column (4)) goes against this hypothesis.

Taken together, the strength of the correlations reported in Tables 2 and 3 therefore

[36] These could be couples who move in search of appealing employment prospects or of a higher quality of schooling for their children, for instance. Such couples might have higher aspirations for their own economic situation and/or their children's economic future than the average couple, leading them to consciously restrict their fertility in favour of higher educational investments.

provides empirical support for the hypothesis that salaried employment opportunities are tightly linked to women's fertility choices.

6 Discussion

This paper uses micro data on desired fertility across the income distribution within countries to shed light on the drivers of sub-Saharan Africa's exceptionally high fertility rates. I find that almost a third of the overall gap in fertility between sub-Saharan Africa and other developing countries is explained by the fact that there exists a steeper income-desired fertility gradient in the former, even conditional on levels of GDP per capita. Through the lens of a Beckerian framework supported by micro data from several sources, the paper shows that this is related to the fact that the poor face especially limited salaried employment opportunities in sub-Saharan Africa.

Investments aimed at relaxing the constraints faced by the poor in broadening their occupational choice set are important for poverty alleviation, as demonstrated by recent studies on poverty traps and occupational choice (Balboni et al. 2020; Bandiera et al. 2017). These show that providing households with a productive asset and the skills to turn this into a profitable activity enables poor families to break out of poverty. By linking jobs to fertility choices, this paper highlights another set of mechanisms via which the nature and supply of jobs can affect the intergenerational persistence of poverty. Occupations may affect parents' willingness to substitute quantity for quality as well as the opportunity cost of childbearing. I find that a key feature of jobs that may shape fertility choices is whether they involve working for someone else, as opposed to informal self-employment - the flexibility of which makes it less incompatible with having many children, even if the work takes place outside the home.

The key takeaways from this paper thus complement the policy recommendations stemming from previous work. They suggest that large investments aimed at increasing the set of salaried employment opportunities available to the poorest segments of the population are arguably just as important as improving access to contraceptives - if not more - for reducing fertility in sub-Saharan Africa.[37]

These findings open up several avenues for future research. A natural follow-up study would attempt to reconcile the fact that the income-desired fertility gradient is steeper in sub-Saharan Africa with the small but growing literature on gender gaps in fertility in sub-Saharan Africa. These have been documented both in terms of realised (Field et al. 2016) and desired fertility (Doepke and Tertilt 2018) at the country level. This suggests that accounting for differences in desired fertility across spouses may allow

[37] Assessing the general equilibrium implications of such big push expansion policies on the slope of the income-desired fertility gradient in sub-Saharan African countries is left for future work.

Figure 14

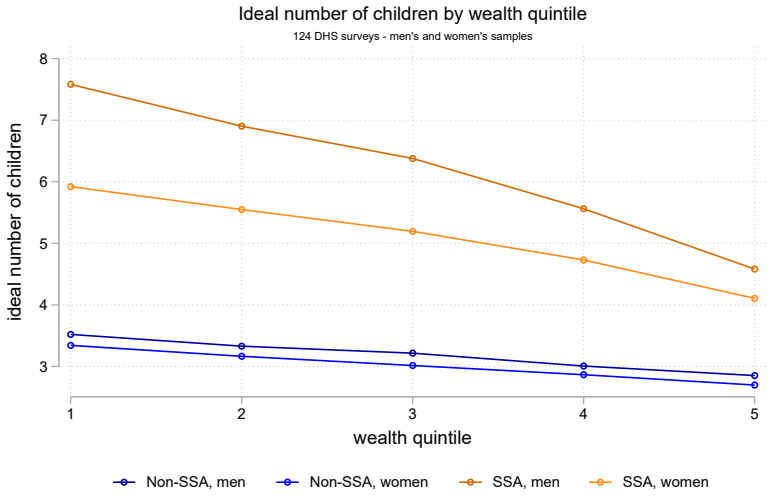

us to better understand the pace of fertility transitions across and within sub-Saharan African countries.[38]

Examining differences in desired fertility across men and women across the income distribution reveals that not only do men typically report a higher ideal number of children than their wives on average, but the slope of the relationship between men's desired fertility with household wealth is also significantly steeper than that of women's in sub-Saharan Africa, whereas the gradients are virtually identical on average across men and women in other regions. This result is displayed in Figure 14.

Unpacking the determinants of this additional cross-regional disparity in wealth- desired fertility gradients could shed further light on the macro puzzle that motivated this paper: the fertility gap between sub-Saharan Africa and other countries. Part of my future research agenda will involve relating gender gaps in desired fertility by income to the findings in this study.[39] Indeed, one possible interpretation of the facts I document in this paper is that, if jobs are especially scarce in sub-Saharan Africa, the poorest women may be relatively more dependent on men for their livelihoods than elsewhere. As a result, their ideal number of children may be shaped at least in part by the will to meet their husbands' own fertility aspirations.

Such a finding would raise two additional important questions. First, it would be fundamental to understand why men's socioeconomic gradient in desired fertility is steeper than women's in sub-Saharan Africa. Second, extending this paper's model from a unitary to collective framework with intra-spousal bargaining could allow us to assess the extent to which improving women's bargaining power might reduce the poor-rich gap in demand for children in sub-Saharan Africa.[40]

[38] Importantly, studies such as Ashraf et al. (2014) and Ashraf et al. (2017) also highlight the possible negative welfare implications of misalignment in desired fertility across spouses.

[39] Appendix section D provides some detail on this, including preliminary evidence about the gradient between wealth and the share of polygynous unions and measures of relative earnings within couples.

[40] The role of spousal bargaining power in fertility has been documented in other settings. In Malaysia, Rasul (2008)

One particular feature of sub-Saharan African marriage markets might also require further examination through the lens of socioeconomic gradients: polygyny. This practice - the cultural norm that men can have multiple wives - is widespread through much of sub-Saharan Africa and has been shown to be a key contributor to the region's persistently high fertility rates (Tertilt 2005, 2006; Rossi 2019). Investigating whether this practice is more common amongst the poorest segments of the population - for instance, if the lack of job opportunities deters deviation from this historically determined norm - could represent another fruitful avenue for future research.[41]

Finally, an important follow-up research avenue on the role of jobs in explaining sub-Saharan African's exceptional fertility trends would be to investigate the determinants of cross-country variation - and spatial variation within countries - in the relationship between socioeconomic status, desired fertility and occupations within sub-Saharan Africa.[42] It could help inform the design of policies aimed at lowering fertility to quantify the role of job opportunities in explaining fertility differentials across the income distribution on a country-by-country basis. This is left for future work.

finds that the outcome of spousal bargaining over fertility decisions depends on the distribution of bargaining power in the couple. Doepke and Kindermann (2019) account for the distribution of childcare duties between mothers and fathers. One could fit a collective model to DHS samples where the husbands of female respondents are also interviewed: these include the ideal number of children reported by both spouses.

[41] The top row of Appendix A.8 provides preliminary evidence, constructed from the latest wave of IPUMS-DHS surveys from sub-Saharan Africa, that the incidence of polygyny is indeed negatively correlated with income, and women in polygynous marriages tend to report higher desired fertility, consistent with the findings of Rossi (2019).

[42] Appendix Figure A.13 overlays the gradient in desired fertility with the gradient in the share of women working for a salary separately for 25 sub-Saharan African countries. While the slope of the two gradients is very similar in most countries (similar magnitude, opposite sign), this graph also reveals substantial heterogeneity in these gradients across countries in the region. In some countries, like Congo and Rwanda, the two gradients are virtually flat, whereas in others, like Nigeria, Namibia or Zimbabwe, they are very steep.

References

Aaronson, D., Lange, F., and Mazumder, B. (2014). Fertility Transitions along the Extensive and Intensive Margins. *American Economic Review*, 104(11):3701–24.

Ashraf, N., Bau, N., Low, C., and McGinn, K. (2020). Negotiating a Better Future: How Interpersonal Skills Facilitate Intergenerational Investment. *The Quarterly Journal of Economics*, 135(2):1095–1151.

Ashraf, N., Field, E., and Lee, J. (2014). Household Bargaining and Excess Fertility: An Experimental Study in Zambia. *American Economic Review*, 104(7):2210–37.

Ashraf, N., Field, E., Rusconi, G., Voena, A., and Ziparo, R. (2017). Traditional Beliefs and Learning about Maternal Risk in Zambia. *American Economic Review*, 107(5):511–15.

Baird, S., Chirwa, E., McIntosh, C., and Özler, B. (2010). The Short-Term Impacts of a Schooling Conditional Cash Transfer Program on the Sexual Behavior of Young Women. *Health Economics*, 19(S1):55–68.

Balboni, C., Bandiera, O., Burgess, R., Ghatak, M., and Heil, A. (2020). Why Do People Stay Poor? *mimeo*.

Bandiera, O., Burgess, R., Das, N., Gulesci, S., Rasul, I., and Sulaiman, M. (2017). Labor Markets and Poverty in Village Economies*. *The Quarterly Journal of Economics*, 132(2):811–870.

Becker, G. and Tomes, N. (1976). Child Endowments and the Quantity and Quality of Children. *Journal of Political Economy*, 84(4):S143–62.

Becker, G. S. and Lewis, H. G. (1973). On the Interaction between the Quantity and Quality of Children. *Journal of Political Economy*, 81(2, Part 2):S279–S288.

Bleakley, H. and Lange, F. (2009). Chronic Disease Burden and the Interaction of Education, Fertility, and Growth. *The Review of Economics and Statistics*, 91(1):52–65.

Bongaarts, J. (2017). Africa's Unique Fertility Transition. *Population and Development Review*, 43(S1):39–58.

Bongaarts, J. and Casterline, J. (2013). Fertility Transition: Is sub-Saharan Africa Different? *Population and Development Review*, 38(1):153–168.

Boserup, E. (1965). The Conditions of Agricultural Growth : The Economics of Agrarian Change Under Population Pressure.

Breierova, L. and Duflo, E. (2004). The Impact of Education on Fertility and Child Mortality: Do Fathers Really Matter Less Than Mothers? (10513).

de la Croix, D. and Doepke, M. (2003). Inequality and growth: Why differential fertility matters. *The American Economic Review*, 93(4):1091–1113.

de la Croix, D. and Doepke, M. (2004). Public versus private education when differential fertility matters. *Journal of Development Economics*, 73(2):607 – 629.

de Silva, T. and Tenreyro, S. (2017). Population Control Policies and Fertility Convergence. *Journal of Economic Perspectives*, 31(4):205–28.

de Silva, T. and Tenreyro, S. (2020). The Fall in Global Fertility: A Quantitative Model. *American Economic Journal: Macroeconomics*, 12(3):77–109.

Doepke, M. (2005). Child Mortality and Fertility Decline: Does the Barro-Becker Model Fit the Facts? *Journal of Population Economics*, 18(2):337–366.

Doepke, M. and Kindermann, F. (2019). Bargaining over Babies: Theory, Evidence, and Policy Implications. *American Economic Review*, 109(9):3264–3306.

Doepke, M. and Tertilt, M. (2018). Women's Empowerment, the Gender Gap in Desired Fertility, and Fertility Outcomes in Developing Countries. *AEA Papers and Proceedings*, 108:358–62.

Doepke, M. and Zilibotti, F. (2005). The Macroeconomics of Child Labor Regulation. *American Economic Review*, 95(5):1492–1524.

Duflo, E., Dupas, P., and Kremer, M. (2015). Education, HIV, and Early Fertility: Experimental Evidence from Kenya. *American Economic Review*, 105(9):2757–97.

Dupas, P., Duflo, E., and Kremer, M. (2019). The Impact of Free Secondary Education: Experimental Evidence from Ghana. *Working Paper*.

Ehrlich, P. (1968). *The Population Bomb*. Ballantine books. Ballantine Books.

Ferre, C. (2009). *Age At First Child: Does Education Delay Fertility Timing? The Case Of Kenya*. The World Bank.

Field, E., Molitor, V., Schoonbroodt, A., and Tertilt, M. (2016). Gender Gaps in Completed Fertility. *Journal of Demographic Economics*, 82(2):167–206.

Figueiredo Walter, T. (2020). Misallocation in the Public Sector? Cross-Country Evidence from Two Million Primary Schools. *Working Paper*.

Filmer, D. and Pritchett, L. H. (2001). Estimating Wealth Effects without Expenditure Data-or Tears: An Application to Educational Enrollments in States of India. *Demography*, 38(1):115–132.

Galor, O. and Weil, D. N. (1996). The Gender Gap, Fertility, and Growth. *American Economic Review*, 86(3):374–387.

Galor, O. and Weil, D. N. (1999). From Malthusian Stagnation to Modern Growth. *American Economic Review*, 89(2):150–154.

Galor, O. and Weil, D. N. (2000). Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond. *American Economic Review*, 90(4):806–828.

Galor, O. and Zeira, J. (1993). Income Distribution and Macroeconomics. *The Review of Economic Studies*, 60(1):35–52.

Ghatak, M. (2015). Theories of Poverty Traps and Anti-Poverty Policies. *The World Bank Economic Review*, 29:77–105.

Goldin, C. (1995). *The U-Shaped Female Labor Force Function in Economic Development and Economic History*, pages 61–90. University of Chicago Press.

Gollin, D., Jedwab, R., and Vollrath, D. (2016). Urbanization with and without industrialization. *Journal of Economic Growth*, 21.

Heath, R. (2017). Fertility at Work: Children and Women's Labor Market Outcomes in Urban Ghana. *Journal of Development Economics*, 126:190 – 214.

Heath, R. and Jayachandran, S. (2018). The Causes and Consequences of Increased Female Education and Labor Force Participation in Developing Countries. *Oxford Handbook on the Economics of Women, ed. Susan L. Averett, Laura M. Argys and Saul D. Hoffman. New York: Oxford University Press.*

Heath, R. and Mobarak, A. M. (2015). Manufacturing Growth and the Lives of Bangladeshi Women. *Journal of Development Economics*, 115:1 – 15.

Jensen, R. (2012). Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India. *The Quarterly Journal of Economics*, 127(2):753–792.

Jones, L., Schoonbroodt, A., and Tertilt, M. (2011). Fertility Theories: Can They Explain the Negative Fertility-Income Relationship? *Demography and the Economy*.

Joshi, S. and Schultz, T. (2012). Family Planning and Women's and Children's Health: Long-Term Consequences of an Outreach Program in Matlab, Bangladesh. *Demography*, 50.

Keats, A. (2018). Women's schooling, fertility, and child health outcomes: Evidence from Uganda's free primary education program. *Journal of Development Economics*, 135:142 – 159.

Kremer, M. (1993). Population Growth and Technological Change: One Million B.C. to 1990*. *The Quarterly Journal of Economics*, 108(3):681–716.

Lavy, V. and Zablotsky, A. (2015). Women's Schooling and Fertility Under Low Female Labor Force Participation: Evidence from Mobility Restrictions in Israel. *Journal of Public Economics*, 124(C):105–121.

Malthus, T. R. (1798). *An Essay on the Principle of Population*. Number malthus1798 in History of Economic Thought Books. McMaster University Archive for the History of Economic Thought.

McKelvey, C., Thomas, D., and Frankenberg, E. (2012). Fertility Regulation in an Economic Crisis. *Economic Development and Cultural Change*, 61(1):7–38.

McKenzie, D. J. (2005). Measuring Inequality with Asset Indicators. *Journal of Population Economics*, 18(2):229–260.

Meade, J. E. (1961). Mauritius: A Case Study in Malthusian Economics. *The Economic Journal*, 71(283):521–534.

Meade, J. E. (1967). Population Explosion, the Standard of Living and Social Conflict. *The Economic Journal*, 77(306):233–255.

Miller, G. (2009). Contraception as Development? New Evidence from Family Planning in Colombia. *The Economic Journal*, 120(545):709–736.

Moav, O. (2005). Cheap Children and the Persistence of Poverty. *The Economic Journal*, 115(500):88–110.

Monstad, K., Propper, C., and Salvanes, K. G. (2008). Education and Fertility: Evidence from a Natural Experiment. *The Scandinavian Journal of Economics*, 110(4):827–852.

Montgomery, M., Gragnolati, M., Burke, K., and Paredes, E. (2000). Measuring Living Standards with Proxy Variables. *Demography*, 37:155–174.

Osei, R. and Jedwab, R. (2013). Is structural change in sub-saharan africa different? new historical evidence from ghana.

Osili, U. and Long, B. (2008). Does female schooling reduce fertility? Evidence from Nigeria. *Journal of Development Economics*, 87(1):57–75.

Ozier, O. (2018). The Impact of Secondary Schooling in Kenya: A Regression Discontinuity Analysis. *Journal of Human Resources*, 53(1):157–188.

Portner, C. C., Beegle, K., and Christiaensen, L. (2011). Family Planning and Fertility: Estimating Program Effects Using Cross-Sectional Data. (5812).

Rasul, I. (2008). Household Bargaining over Fertility: Theory and Evidence from Malaysia. *Journal of Development Economics*, 86(2):215 – 241.

Romer, P. M. (1990). Endogenous Technological Change. *Journal of Political Economy*, 98(5):S71–S102.

Rossi, P. (2019). Strategic Choices in Polygamous Households: Theory and Evidence from Senegal. *The Review of Economic Studies*, 86(3):1332–1370.

Tertilt, M. (2005). Polygyny, Fertility, and Savings. *Journal of Political Economy*, 113(6):1341–1371.

Tertilt, M. (2006). Polygyny, Women's Rights, and Development. *Journal of the European Economic Association*, 4(2-3):523–530.

Vogl, T. S. (2016). Differential Fertility, Human Capital, and Development. *The Review of Economic Studies*, 83(1):365–401.

A Appendix to section 1.

Figure A.1

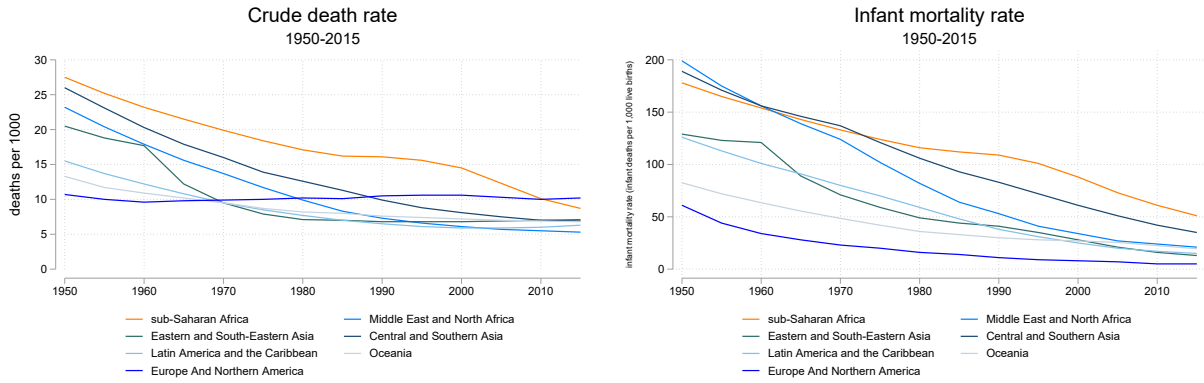

(a) Crude death rates

(b) Infant mortality rates

Source: United Nations Population Division (2019). World Population Prospects 2019.

B Appendix to section 3.

Figure A.2

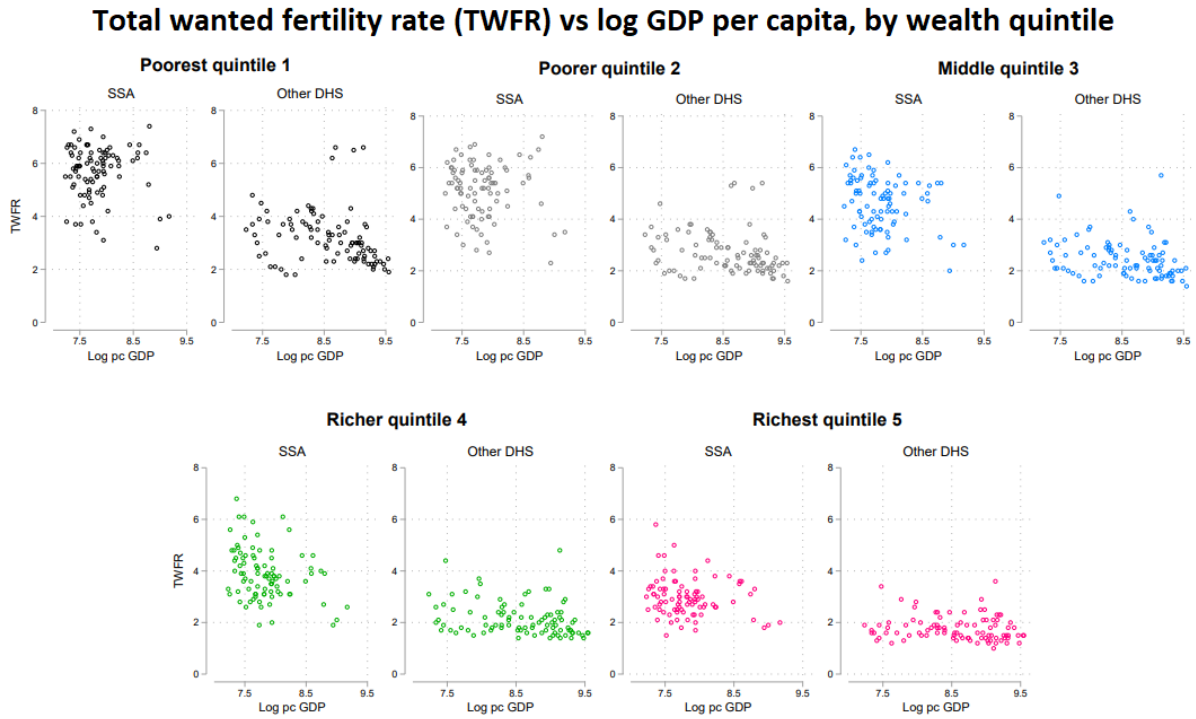

Each scatterpoint corresponds to a DHS survey. X-axis: log per capita GDP, in 2011 PPP USD (WDI). Y-axis: Total Wanted Fertility Rate (DHS).

153 Demographic and Health Surveys. Y-variable plots the average total wanted fertility rate in each sample (country x year). Left-hand-side panel includes sub-Saharan African samples only, while right-hand-side panel includes all other DHS samples. Each colour corresponds to a different quintile, in ascending order from top left to bottom right.

Figure A.3

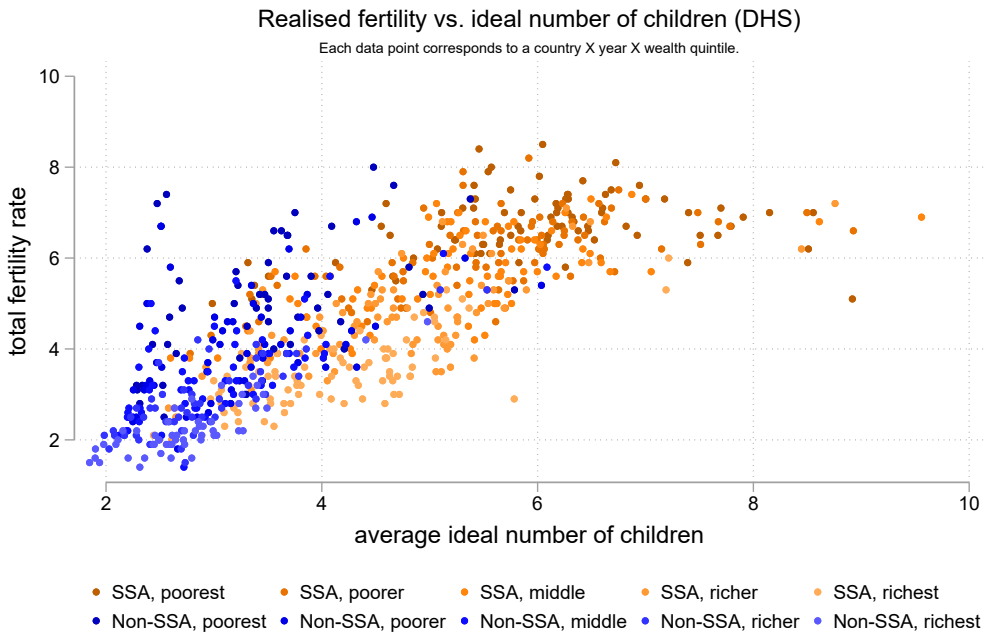

153 Demographic and Health Surveys. This graph plots the total fertility rate against the average ideal number of children reported by respondents in each country x year x wealth quintile in the log GDP per capita bin of common support [7.2-9]. It is important to note that the fertility rate is flat relative to ideal number of children at very high levels of desired fertility (7+ children), possibly pointing to biological constraints for achieving an average of more than 7 births per woman. Estimates of gradient differentials are virtually unaffected if these observations are dropped. The Pearson correlation coefficient for these two variables is 0.78.

Figure A.4

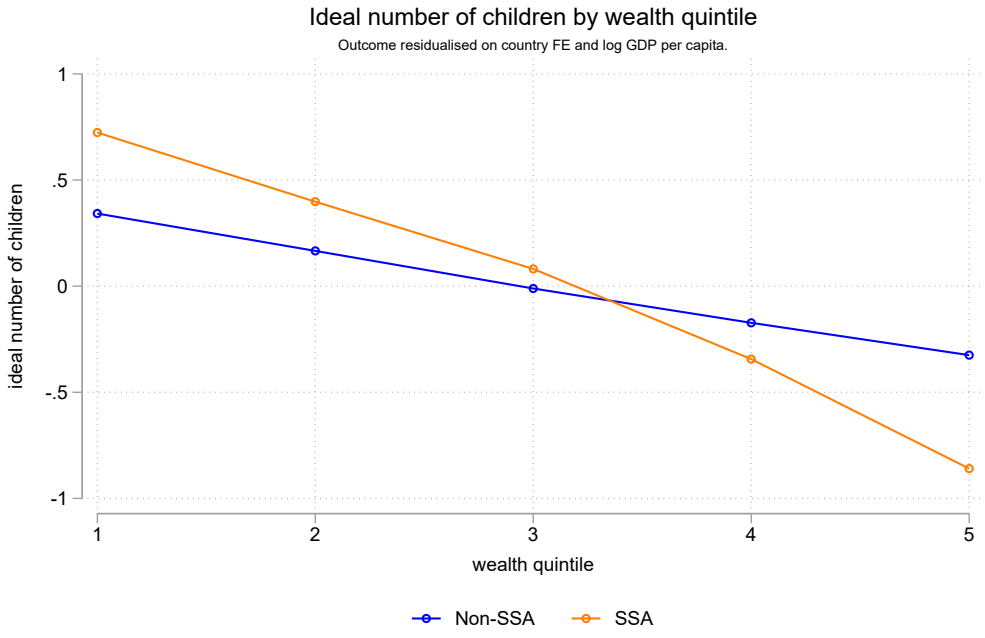

153 DHS surveys. Y-variable is ideal number of children residualised on country fixed effects and log GDP per capita.

Table A.1: Accounting for the poor-rich gap in desired fertility

	Total fertility rate		
	(1)	(2)	(3)
log GDP per capita	-0.944***	-0.150	-0.233
	[0.284]	[0.246]	[0.227]
sub-Saharan Africa		1.689***	1.147***
		[0.225]	[0.223]
poor-rich Δ desired fertility			0.569***
			[0.114]
R^2	0.148	0.496	0.584
Outcome Mean	4.399	4.399	4.399
Observations	153	153	153

Observations weighted by number of waves per country.
* p<0.1, ** p<0.05, *** p<0.01.

Figure A.5

This graph plots average values of the cross-sample wealth index - constructed using PCA on a subset of 56 DHS surveys - by within-sample wealth quintile. The main finding to emerge here is that income dispersion in living standards as measured by this international wealth index, is **not** larger in sub-Saharan African countries on average.

Table A.2: Cross-regional differences in the poor-rich gap in living standards (DHS)

	cross-sample wealth index
log GDP per capita	1.135**
	[0.430]
sub-Saharan Africa (SSA)	-1.700**
	[0.666]
first quintile	-6.156***
	[0.176]
second quintile	-4.875***
	[0.109]
third quintile	-3.350***
	[0.330]
fourth quintile	-1.812***
	[0.288]
SSA x quintile 1	2.013***
	[0.491]
SSA x quintile 2	1.110***
	[0.399]
SSA x quintile 3	0.346
	[0.394]
SSA x quintile 4	0.008
	[0.311]
R^2	0.790
Outcome Mean	0.029
Observations	603640

Note: * p<0.1, ** p<0.05, *** p<0.01.
Standard errors clustered by survey.

Figure A.6

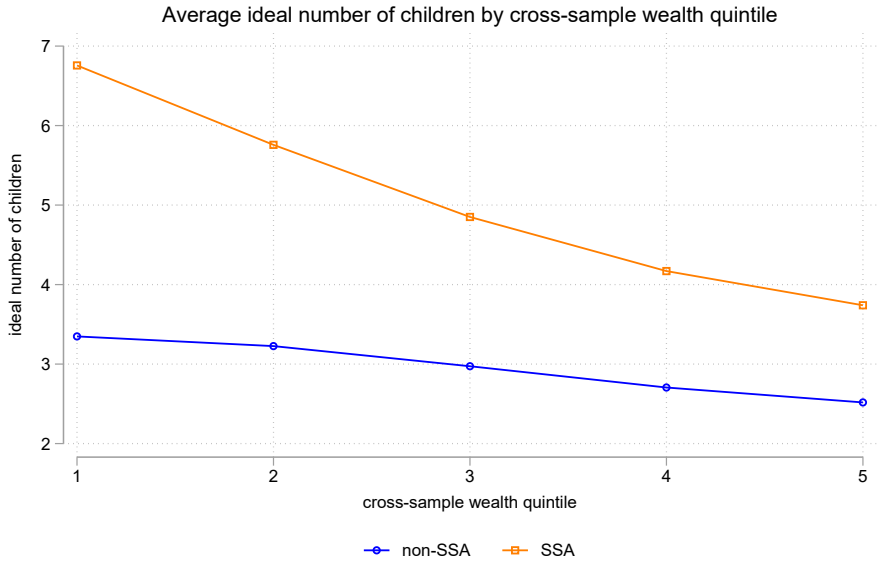

On the y-axis is women's reported ideal number of children. On the x-axis are quintiles of the cross-sample wealth index I construct using PCA on a subset of 56 DHS surveys - covering 32 countries out of the original 153 surveys in my GDP per capita bin - that all include the same set of variables related to household living standards. These are: electricity, source of drinking water, type of toilet, wall, roof and floor material, and cooking fuel, number of household members per sleeping room, and ownership of fridge, bicycle, motorcycle and television.

Figure A.7

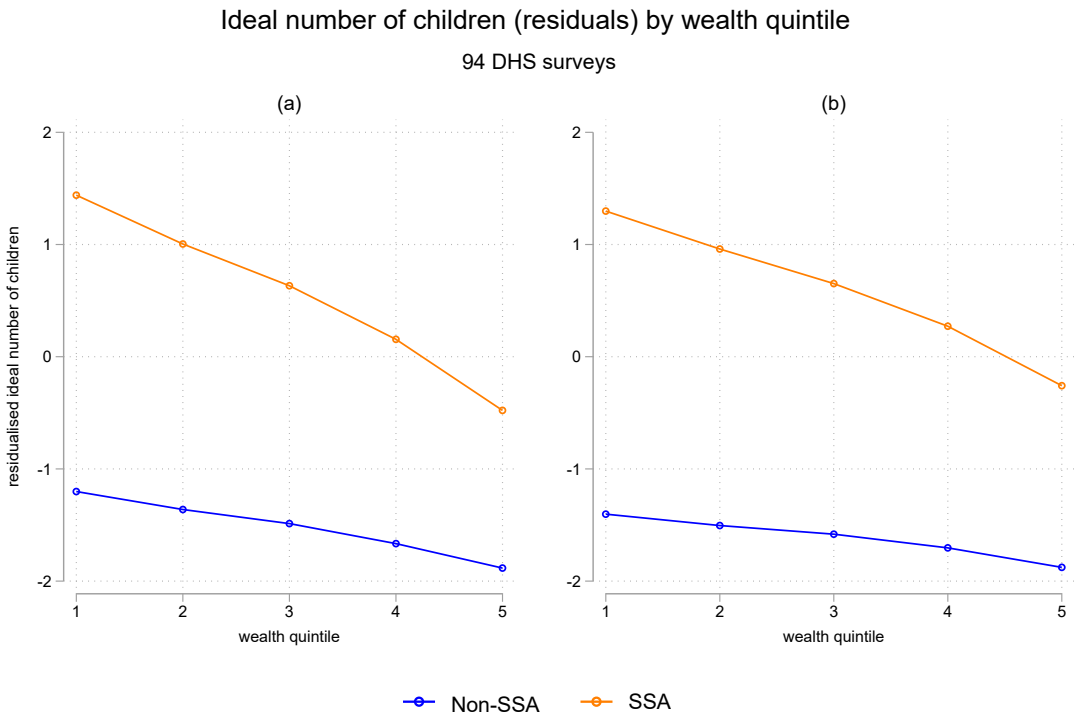

Graph (a) plots ideal number of children residualised on log GDP per capita. Graph (b) plots the same outcome residualised on log GDP per capita **and** share of women exposed to family planning campaigns (radio). This analysis uses only 94 out of 153 surveys in the GDP per capita range of common support, i.e. the ones for which information was collected on exposure to family planning information (via radio, newspaper or posters and leaflets).

Figure A.8

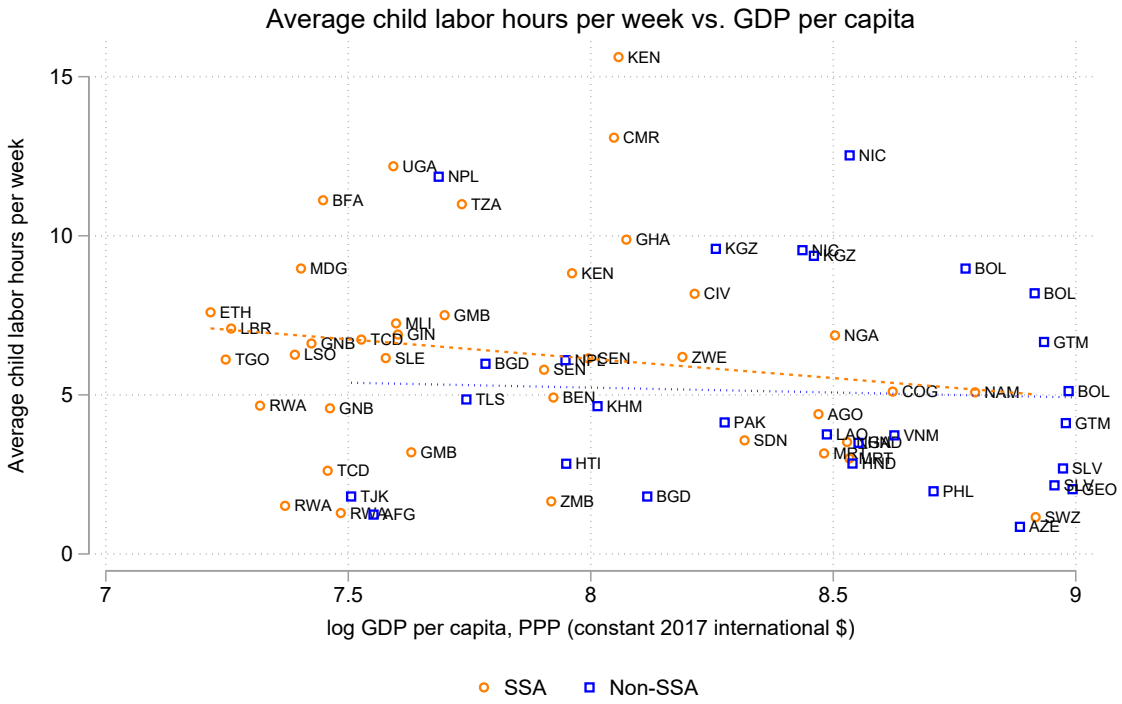

Source: World Development Indicators 2020. This variable is only available for 67 observations from 49 countries for the log GDP per capita bin of common support [7.2-9]. 30 countries are in sub-Saharan Africa and 19 in other regions. To the best of my knowledge, no publicly available data exists that would allow me to construct this statistic at the country x year x wealth quintile level at this stage.

C Appendix to section 5.

Figure A.9

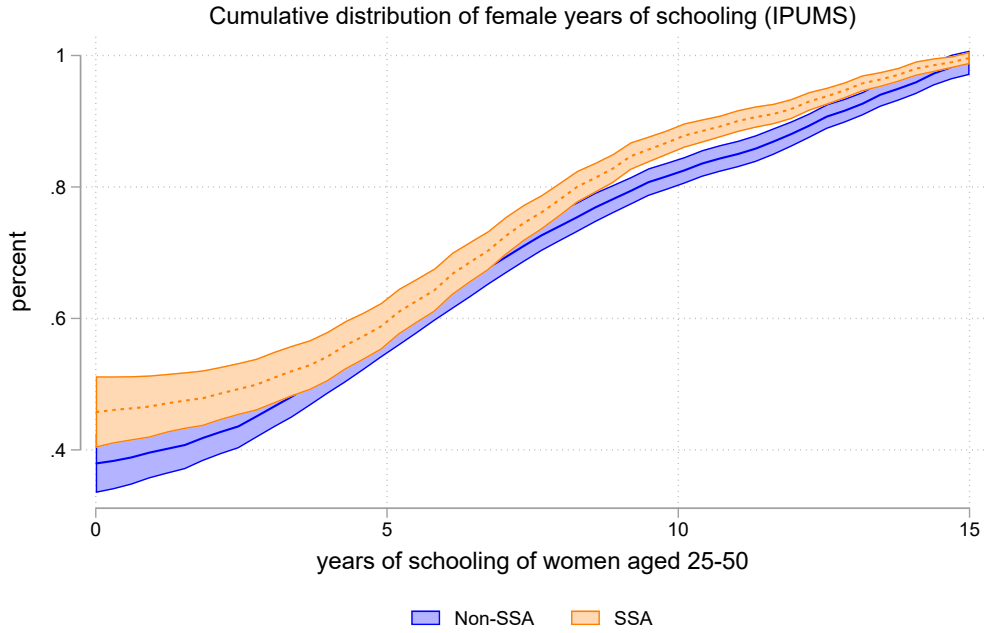

49 IPUMS samples. Using the same set of micro-censuses in the GDP per capita bin of common support (log GDP per capita=[7.2-9]) as for other cross-regional analyses in section 5.1, this graph plots the cumulative distributions of female years of schooling for sub-Saharan and non-sub-Saharan African years of schooling separately. These include 95% confidence bands. As is clearly apparent here, one cannot reject the null that these distributions are the same across the two sets of countries. This allows me to reject the hypothesis that the non-sub-Saharan African countries are at the "high" (h^{**}) equilibrium whilst the sub-Saharan African ones are in the "low" one (h^{*}).

Figure A.10

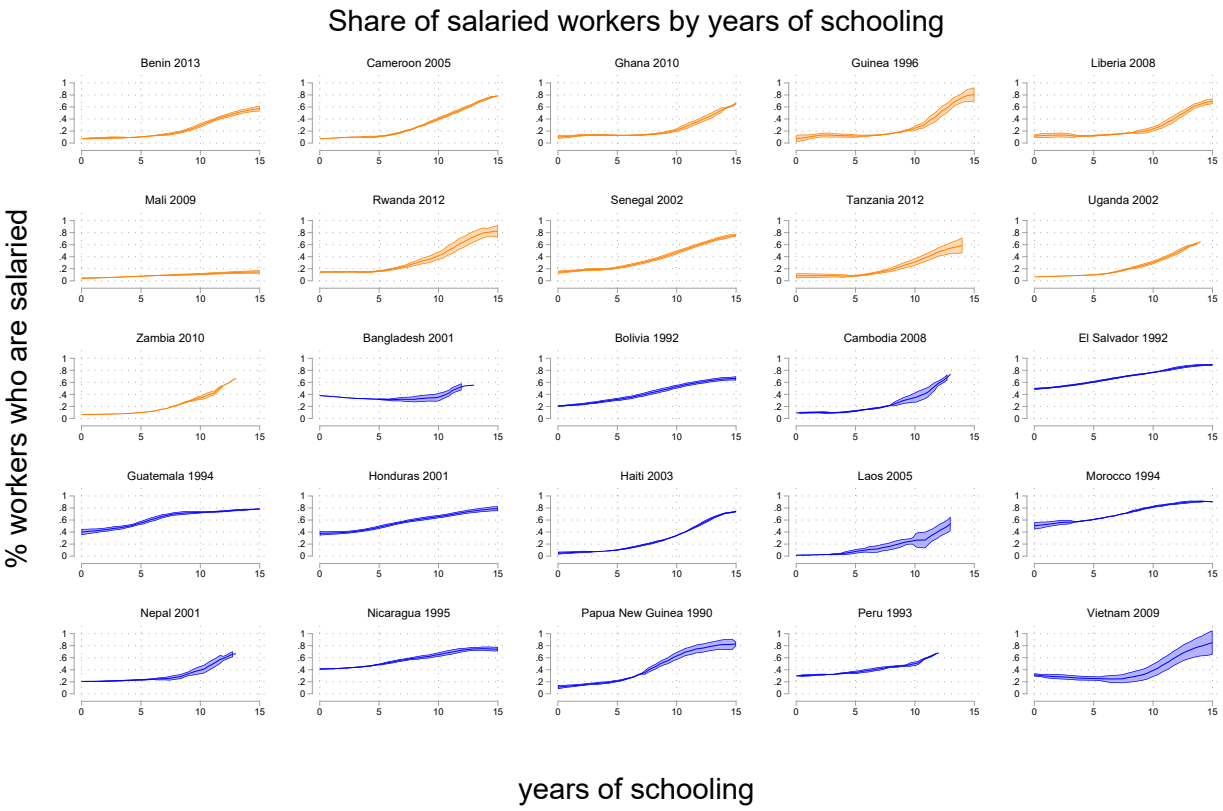

25 IPUMS samples. Local polynomial regression estimates, by sample, with 95% confidence bands. The orange plots correspond to sub-Saharan African samples. On the Y-axis is the raw measure of salaried employment rates. The estimates for the other 24 samples are similar, so I exclude them here for visualisation purposes.

Figure A.11: Spatial variation in desired fertility and years of education of women (pink) and men (green)

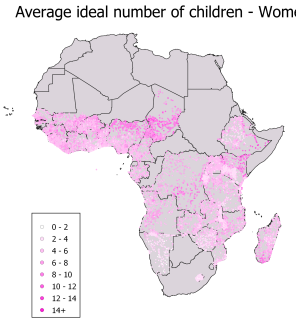

(a) Women

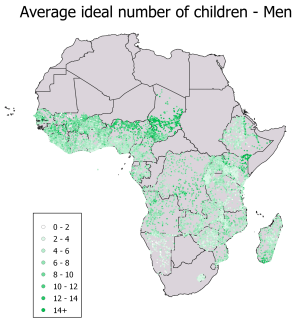

(b) Men

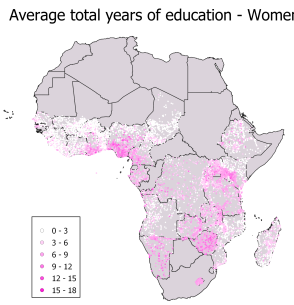

(c) Women

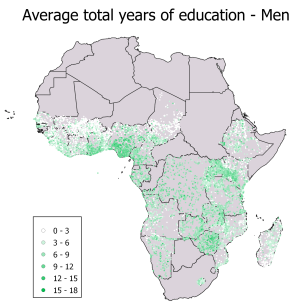

(d) Men

Figure A.12: Schooling costs per child (Inv. Arcsine)
Kenya Integrated Household Budget Survey 2005-2006

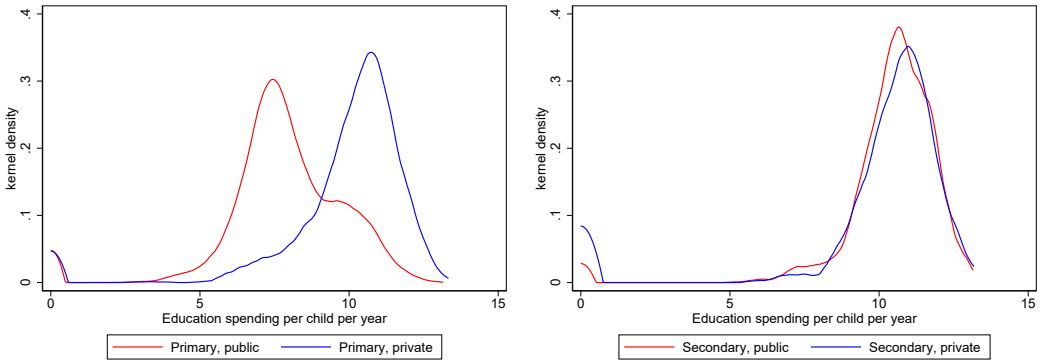

(a) Primary school **(b)** Secondary school

Figure A.13: Relationship between household head education and (a) fertility; (b) educational investments
Kenya Integrated Household Budget Survey 2005-2006

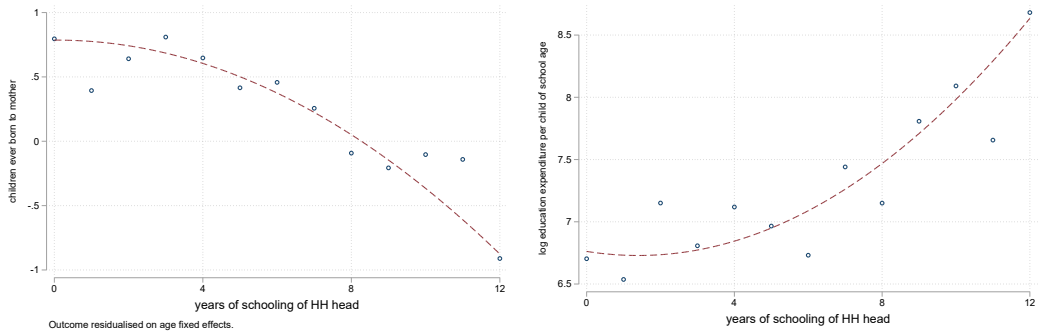

(a) Children ever born to mother **(b)** Log education expenditure per child

Table A.3: Kenya Integrated Household Budget Survey 2005-2006

	Number of children ever born to mother (1)	Log education expenditure per child (2)
urban	-1.122***	1.041***
	[0.105]	[0.109]
years of schooling of household head	-0.153***	0.144***
	[0.011]	[0.009]
household head has a salaried job	-0.203**	0.186**
	[0.087]	[0.081]
R^2	0.126	0.205
Outcome Mean	4.861	7.458
Observations	5514	5514

* p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered by enumeration area in brackets. Sample restricted to all households with at least one child of schooling age.

D Appendix to section 6.

Figure A.14

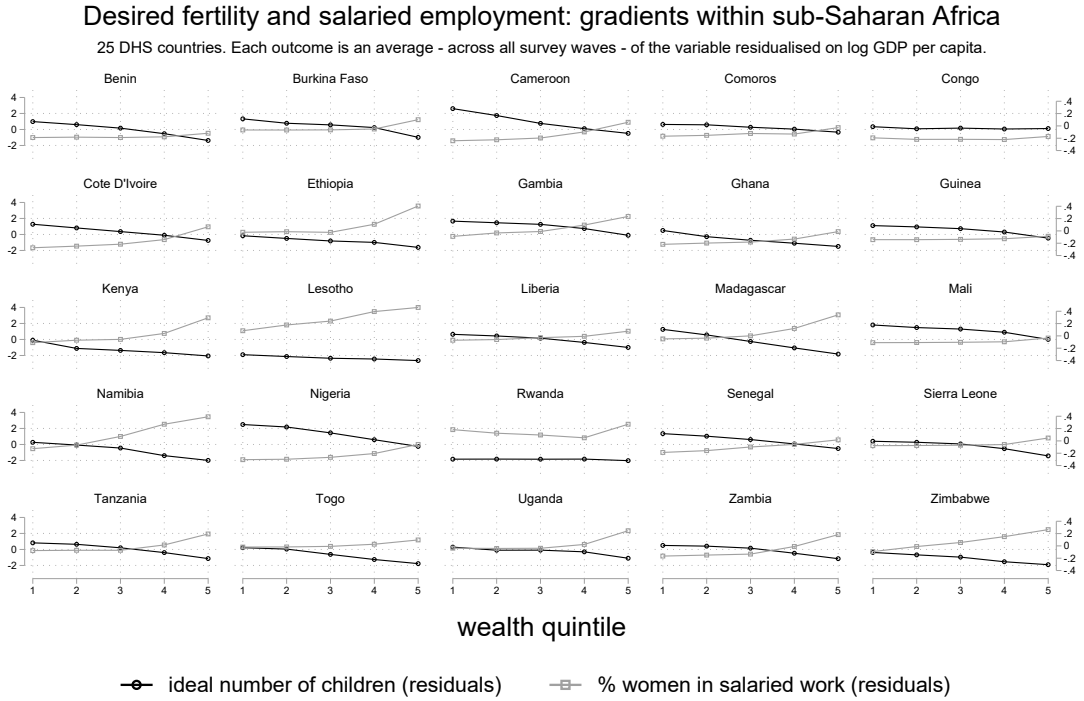

Figure A.15: Desired fertility of husbands and wives in monogamous unions

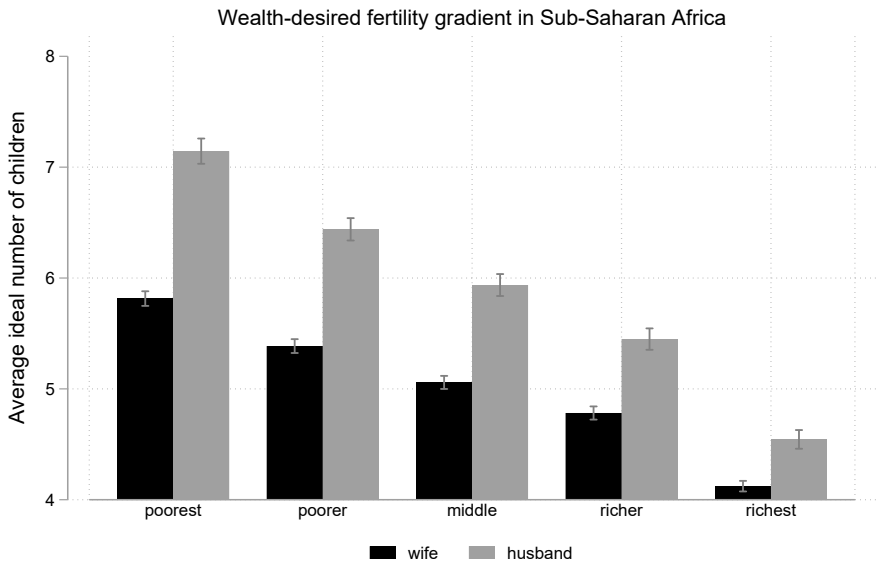

Figure A.16: Desired fertility gradient by woman's relative earnings (DHS)

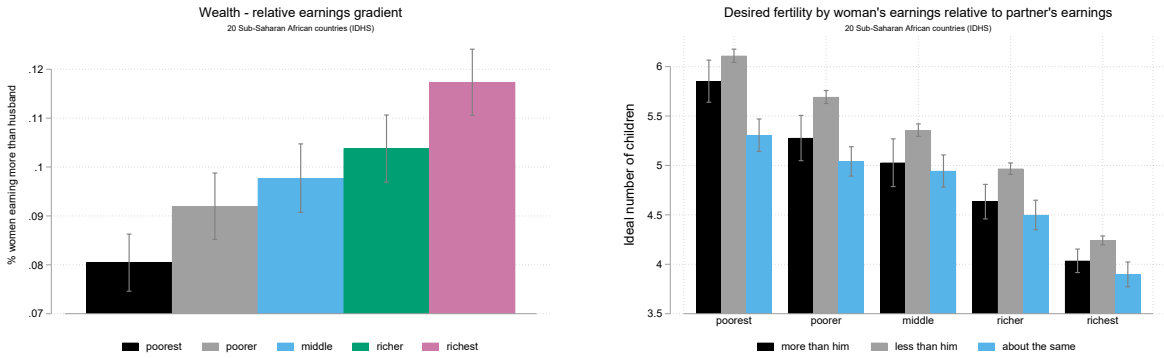

Figure A.17: Polygyny in sub-Saharan Africa (IPUMS-DHS)

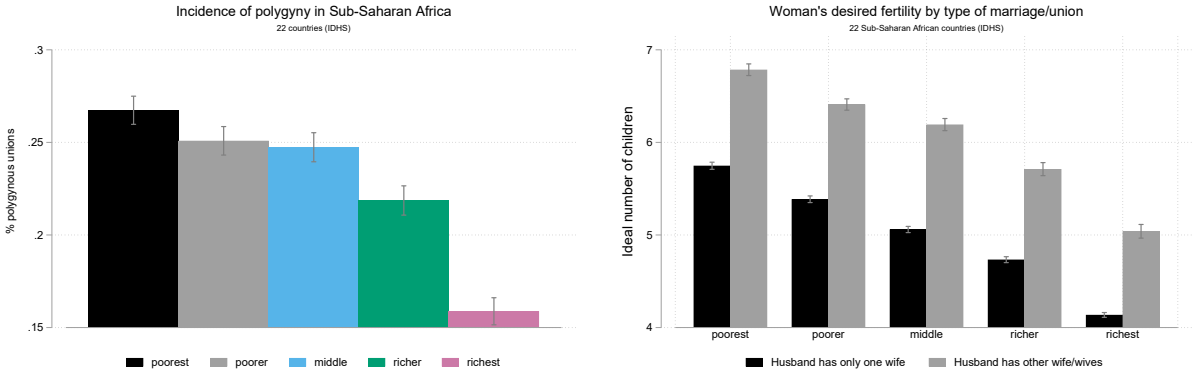

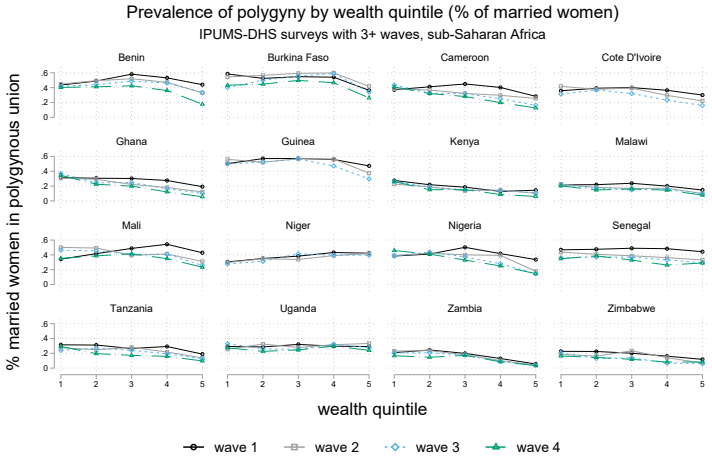

Table A.4: Husbands' and wives' desired fertility in sub-Saharan Africa (DHS). Assortative matching? Or do men internalise the opportunity cost?

	Ideal number of children			
	Wives		**Husbands**	
	(1)	(2)	(3)	(4)
wife works in non-agriculture	-0.113***	-0.143***	-0.109*	0.049
	[0.032]	[0.037]	[0.066]	[0.079]
husband works in non-agriculture	-0.145***	-0.107***	-0.333***	-0.315***
	[0.037]	[0.038]	[0.070]	[0.072]
wife works all year		0.063*		-0.301***
		[0.037]		[0.079]
husband works all year		-0.135***		-0.101
		[0.036]		[0.065]
R^2	0.344	0.345	0.254	0.255
Outcome Mean	5.316	5.316	6.454	6.454
Observations	36046	36046	36046	36046

Standard errors clustered by enumeration area in brackets. Controls for age and years of schooling of both spouses, and urban/rural status of residence.
* p<0.1, ** p<0.05, *** p<0.01.