The London School of Economics and Political Science

Essays in Development Economics

Céline Zipfel

A thesis submitted to the Department of Economics
for the degree of Doctor of Philosophy

June 2021

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 28278 words.

Statement of conjoint work

I confirm that Chapter 3 was jointly co-authored with Tavneet Suri, Daniel Keniston and Esther Duflo, and I contributed 25% of this work.

Abstract

The first chapter of this thesis leverages micro data from 179 reproductive health surveys to shed light on a macro puzzle: fertility rates are exceptionally high in sub-Saharan Africa conditional on GDP per capita. The paper first establishes an important empirical fact: the relationship between wealth and desired fertility is, on average, steeper in sub-Saharan Africa. It then explores the links to the relative scarcity of salaried employment opportunities in these countries. 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.

The second chapter exploits variation in exposure to the reform across birth cohorts and locations to evaluate the impacts of primary school fee abolition on female education, employment and fertility outcomes in Malawi, Ghana, Uganda and Tanzania. The findings suggest that fee elimination improved educational attainment among cohorts most exposed to the reform: the effects range from a 5% increase in years of schooling in Malawi to a 17% increase in Ghana. The probability of completing primary school also increased in Tanzania, Ghana and Uganda, but not Malawi, for which I find a small but negative effect. The FPE reform is also associated with an increase in employment and a reduction in fertility. In Malawi, the increase in women's labor market participation is concentrated in the education sector, with a small share of women shifting away from agriculture. In Tanzania, employment effects are driven by increased participation in the education sector and self-employed retail. In Ghana and Uganda, the share of women engaging in self-employed non-agricultural activities increases. The employment effects are largest in Ghana, where they are driven by entry into self-employed retail, manufacturing and food/accommodation services. Correlational evidence suggests that heterogeneity in policy impact may be partly

attributable to differences in schooling productivity before the reform and in how the governments accommodated the enrolment increase.

The third chapter is joint work that leverages a randomized experiment to study the diffusion of the impacts of an agronomy training program through the social and geographic networks of coffee farmers in Rwanda. We find no evidence of diffusion through geographic networks or from the treatment group to the control group. Our results suggest a reinforcement of treatment effects within the treatment group, concentrated around leaf health improvements. This is driven by farmers who attended the training sessions with people to whom they already had a social connection at baseline. The program also caused a re-sorting of social networks: we find that both treatment and control group households increased social links to individuals trained in their village.

Acknowledgements

I am enormously grateful to my supervisors Oriana Bandiera, Gharad Bryan, Robin Burgess and Rachael Meager for their invaluable guidance and support throughout my PhD. Their insights and continuous encouragement have greatly contributed to my development as a researcher. In particular, I want to thank them for encouraging me to take risks and ask big questions. I want to thank Nava Ashraf for sharing her expertise, encouraging me to pursue research on a topic she and I are both passionate about, and taking the time to help me improve my communication skills. I have been lucky to benefit from the advice and insights of many other professors at the LSE. I am grateful to Maitreesh Ghatak, Tim Besley, Silvana Tenreyro, Rachel Ngai and Rocco Machiavello for very helpful conversations and comments.

I thank Lubala Chibwe, Jane Dickson, Rhoda Frith, Annie Nicholas, Emma Taverner and Mark Wilbor for their help in a variety of requests and their support over the last six years. I gratefully acknowledge the financial support of the Economic and Social Research Council (ESRC).

The third and final chapter of this dissertation is the culmination of many years of work on a project very dear to my heart, as it played a huge role in helping me find my calling as a development economist. I am extremely grateful to Tavneet Suri for taking a chance on me by hiring me to manage the RCT fieldwork in Rwanda, and for her encouragement and mentorship ever since. I thank her, as well as Esther Duflo and Daniel Keniston, for making me co-author on this project. Last but not least, a big thank you to the Nyarubaka team in Rwanda for their hard work and dedication to stellar data collection. A special thank you to Damien Munyandekwe for all his help over the years.

Several colleagues and friends at the LSE have helped to make the years of my PhD such special ones. I am lucky to have formed friendships with Svetlana Chekmasova, Ruth Fortmann, Dita Eckardt, Wolfgang Ridinger, Menna Bishop, Amanda Dahlstrand-Rudin and Xuezhu Shi. Clare Balboni, Jay Euijung Lee, Martina Zanella, Virginia Minni and Torsten Figueiredo Walter in particular have been a wonderful source of continuous support, both intellectually and emotionally. For their support, advice and solidarity during a difficult search on the academic job market, I thank Adrien Bussy, Alexia Delfino, Nicola Limodio, Tillman Hoenig, Nicola Fontana and Heidi Thysen. I cannot thank Niclas Moneke enough for all his time and advice over the last year.

I feel extremely fortunate to have benefited from a strong network of friends, especially since becoming a parent, during my PhD. I am deeply grateful to the Kalbreier family for their emotional and practical support throughout this journey. My Hammersmith friends provided empathy and much-needed breaks during hard times. A special thank you goes to Perrine for her honesty and joie de vivre, Jule for a professional standard photo shoot, and Joanna for her reassurance and humour during that first year of parenthood.

The support and encouragement of my family has been immeasurably important in helping me pursue a career that I love. I am very grateful to my parents, Charlotte and Jean Marc, for everything they have done to help me achieve this, and for teaching me the value of hard work.

My deepest gratitude is to Joseph, my greatest supporter. Thank you for doing the lion's share of the work at home so I could focus on my research over the past year, and for always believing in me. And finally, I am so thankful for the existence of Yohan, who brought an incredible amount of joy to the final stages of this journey.

Contents

List of Figures

List of Tables

Chapter 1

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

1.1 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, 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.

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).

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

Figure 1.1: Global population projections and fertility trends

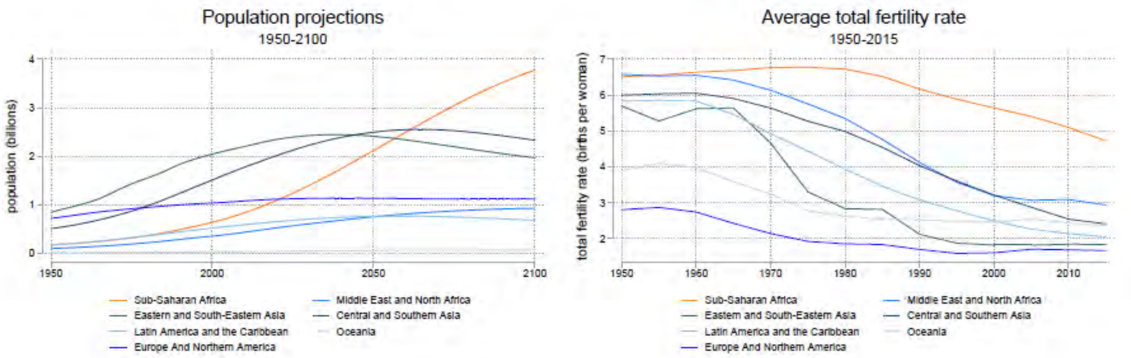

(a) Population projections **(b)** Total fertility rates

Fertility is primarily determined by two factors: parents' desired number of children and the availability of contraceptives. While desired fertility has been recognized as a key contributor to cross-country variation in fertility rates (Pritchett (1994)), much of the literature to date has focused on the impact of "supply-side" interventions on fertility decline in the developing world.[2] The determinants of demand for children in these countries thus remain relatively under-studied.[3]

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.

I exploit micro data from 179 Demographic and Health Surveys (DHS) to construct gradients between income and desired fertility, using a household-level wealth index as a proxy for income and two different measures of desired fertility: the ideal number of

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

[3]In macroeconomics, an extensive body of research investigates the role of demand-side factors in the convergence of fertility rates to about two births per woman in developed countries, which has inspired the field of unified growth theory (Galor and Weil, Galor and Weil, Galor and Weil (1996, 1999, 2000)). Other important contributions include Doepke (2005), who shows that the fall in child mortality rates cannot account for much of the fertility transition in developed countries, and Doepke and Zilibotti (2005), who develop a political economy model to explain why the fertility decline coincided with the introduction of child labor regulations in Britain.

[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, 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, 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.

children reported by female respondents and the Total Wanted Fertility Rate. I find that around 25% 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, about twice as large as in other developing countries.

Figure 1.2: Sub-Saharan Africa's exceptionally high fertility rates

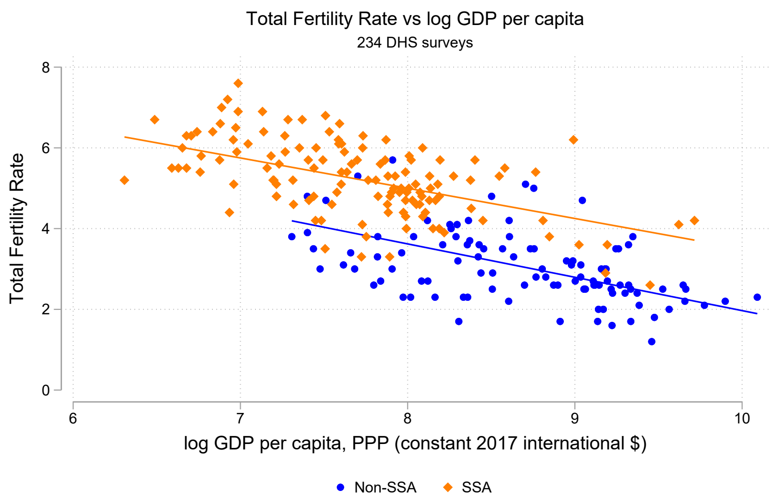

To test that this steeper gradient does not simply reflect larger income dispersion in sub-Saharan African countries, I construct a cross-sample wealth ranking of households from the DHS surveys to compare the living standards of the different within-country wealth quintiles across countries. I can confidently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

This study proposes that a key driver of the steeper wealth-desired fertility gradient in sub-Saharan Africa is the relative scarcity of stable job opportunities in this region. Two additional empirical facts motivate this hypothesis. First, the shape of the socioeconomic gradient in salaried employment mirrors the shape of the wealth-desired fertility gradient in both sets of countries. In sub-Saharan Africa, the relationship is linear for the first three quintiles and becomes more steeply positive from the third quintile, at which point the slope of the relationship with desired fertility also becomes more steeply negative. In the other countries, the slope of the gradient in both desired fertility and salaried employment is virtually linear throughout. Second, a simple mediation analysis reveals that salaried employment rates can account for about 20% of the gap in desired fertility across the two sets of countries.

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.

Bringing the model to the data, I provide empirical support for the framework's key assumption and key prediction using 64 micro-censuses covering 34 countries from 1990 to 2015. 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

[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.

find that this relationship is flat for individuals with under 4.5 years of education in sub-Saharan Africa, with the probability of holding a salaried job being close to zero, whereas the slope is (mildly) 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 with respect to years of schooling in sub-Saharan Africa, joining the curve of the other countries around 11 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. Mothers with up to 10 years of schooling have significantly less educated children - as measured by the fraction of their children aged 14 and above who have completed primary school - 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 lead to multiple equilibria in the intergenerational transmission of human capital.

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 a household expenditure survey from Kenya to yield insight into the relationship between parental labor market outcomes and spending on children′s education. Under the assumption that men's fertility intentions would be less responsive to the opportunity cost of childbearing than women's, 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.

Indeed, demand for children is strongly correlated with the type of work individuals engage in, conditional on their level of education, urban/rural status, knowledge of modern contraceptive methods, 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 the broader occupational category of off-farm work. This points to the greater incompatibility of salaried employment with raising many children. For women who have a salaried job, the negative correlation with desired fertility is even stronger in regions where such jobs are particularly scare.

To test that this correlation does not only capture 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 - I exploit variation in the destination of female migrants from the same district in the Malawi 2008 census. 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. For women who migrated for marriage - the group for whom I am least concerned about selection biasing the results - I find 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 results of this final test thus lend 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 the bulk of projected global population growth over the 21st century. I argue that the relative scarcity of salaried employment opportunities in the region is an important determinant of its steeper socioeconomic gradient in desired fertility. Insights from a fertility choice model combined with rich micro data suggest that a fixed education requirement for salaried employment may contribute to the persistence of an occupational structure dominated by informal self-employment.

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.

1.2 Data

1.2.1 Demographic and 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 low and middle-income countries 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. These amount to 179 surveys in total, covering 65 countries: 105 surveys across 34 countries in sub-Saharan Africa and 74 surveys across 31 countries in other regions. I use three key variables: the wealth index and the Total Wanted Fertility Rate provided by the DHS, and the average ideal number of children reported by female respondents in the DHS surveys, which I construct directly from the micro data at the quintile level.

DHS wealth index. 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 (PCA). 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.

Desired fertility indicators. The analysis in this paper focuses on two different measures of desired fertility that can be constructed from the DHS data: women's reported ideal number of children and the Total Wanted Fertility Rate (TWFR).

The ideal number of children variable is collected by customising a survey question according to women's fertility history. Women with no living children at the time of the interview are asked "If you could choose exactly the number of children to have in your whole life, how many would that be?", while women with at least one living child are asked "If you could go back to the time you did not have any children and could choose exactly the number of children to have in your whole life, how many would that be?"

As emphasized by Pritchett (1994), we should be cautious about interpreting variation in reported ideal number of children as capturing inherent differences in demand for children. Indeed, women with higher *realised* fertility may be more likely to report a higher ideal number of children to avoid admitting that they have children they did not want. Such ex-post rationalisation of prior births may lead us to overestimate desired fertility. If women - on average and/or in the poorest quintiles in particular - tend to report higher fertility ideals in sub-Saharan Africa than elsewhere in part because of their higher *re-*

alised fertility, then comparing this indicator across regions may lead us to overestimate the difference in desired fertility between these two sets of countries.

Another indicator of desired fertility - demographers' measure of choice - is the Total Wanted Fertility Rate (TWFR). This methodology was instilled by demographers in the 1990s to address concerns about ex-post rationalization of prior births in the ideal number of children variable (Bongaarts (1990)). The TWFR is constructed in the same way as the Total Fertility Rate (TFR), which is defined as the total number of children that would be born to the average woman if she were to live to the end of her childbearing years and give birth to children according to the age-specific fertility rates over the preceding three years preceding the survey. The TWFR, however, only includes "wanted births", by excluding past births in excess of desired family size (ICF (2012)).

Since this measure is also built on retrospective data, the TWFR is unlikely to entirely correct for ex-post rationalisation of prior births. However, measuring the fraction of unwanted births based only on questions about future desires instead of past behavior, actually produces very similar estimates of desired fertility (Bongaarts (1990); Pritchett (1994)).

Employment variables. 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.

1.2.2 IPUMS census data

IPUMS International is the world's largest collection of publicly available census micro data. The data are coded and documented consistently across countries and over time to facilitate comparative research. Samples are available for 65 LMICs, and most include information on individuals' socioeconomic characteristics such as education and employment. The majority of the census 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 also 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. In most censuses, the unit of data collection is the household, which also allows me to construct variables about children's schooling levels and document how this relates to their parents' education.

1.2.3 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.

1.3 Socioeconomic gradients in desired fertility

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 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. Section 3.2 then quantifies the role of this differential in the total fertility gap observed between sub-Saharan African and other countries. In section 3.3, I 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 may be linked to the relative scarcity of salaried work opportunities in sub-Saharan Africa. This is backed by another set of empirical findings presented in section 3.4.

Figure 1.3: Socioeconomic gradients in desired fertility

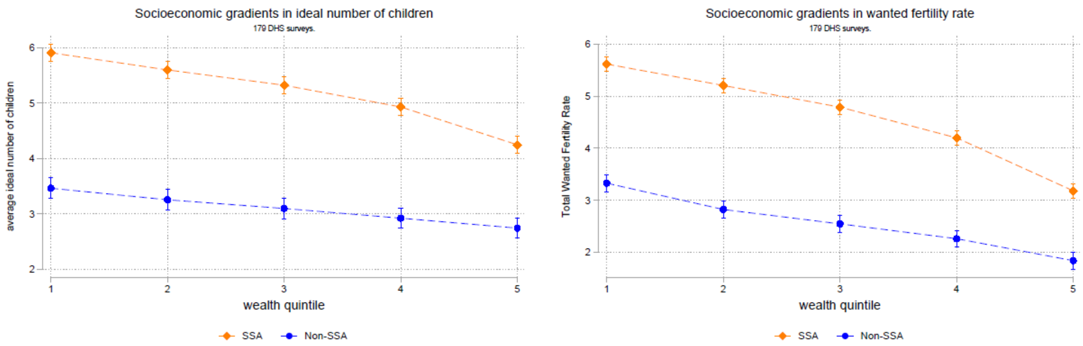

(a) Ideal number of children **(b)** Total Wanted Fertility Rate

1.3.1 A steeper gradient in sub-Saharan Africa

The socioeconomic gradients in desired fertility are displayed in Figure 3. On the horizontal axis are wealth quintiles. Figure 3a displays the gradient in the average ideal number of children reported by female respondents, whereas 3b plots the Total Wanted Fertility Rate on the y-axis. These gradients are constructed by taking the average of each indicator at the quintile-region level (i.e. for sub-Saharan Africa vs. all other DHS countries), weighed by the inverse of the number of DHS waves available per country.[8] The graphs also plot 95% confidence intervals which capture the dispersion in these estimates across surveys separately for sub-Saharan African and other countries.

A key takeaway from this figure is that the poor-rich gap in desired fertility is, on average, around twice as large in sub-Saharan Africa than in other countries. The average gap in women's reported ideal number of children between the poorest and richest quintile is 1.67 children in this region, compared to a gap of only 0.72 in other countries. In terms of wanted fertility rates, the difference between the bottom and top quintile stands at 2.45 wanted births in sub-Saharan Africa and 1.47 in other countries.[9] Thus, regardless of the indicator used to measure desired fertility, Figure 3 documents a significant gap between sub-Saharan Africa and other countries, not only on aggregate, but also in terms of the slope of these socioeconomic gradients.[10]

While this empirical fact stems from a comparison of averages at the country-year-quintile level, it is not driven by outliers. This is illustrated in Appendix Figure A.2, which plots total wanted fertility rates by log GDP per capita separately for each quintile

[8]The total number of DHS surveys available per country that contain all the relevant variables for this paper's analysis - namely, the wealth index and labor market and desired fertility measures - ranges from 1 (e.g. Afghanistan) to 7 (Senegal).

[9]In terms of ratios, the difference, while smaller, is also statistically significant at the 1% level.

[10]GDP per capita differences can only partly account for the aggregate gap in desired fertility, as Figure 6a illustrates later on in section 3.4.

of each DHS sample. The gap in wanted fertility rates across the two sets of countries is noticeably larger among the poorer quintiles than the richer quintiles, conditional on GDP per capita, with the sub-Saharan African estimates standing starkly above those of other countries for the bottom three quintiles especially.

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

To quantify the role of the gradient in explaining sub-Saharan Africa's exceptionally high fertility rates, 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.[11] I then compare the coefficient estimate for the sub-Saharan Africa indicator when I also control for the gradient - as measured by the gap in Total Wanted Fertility Rates between the poorest and richest quintiles.

Figure 1.4: Total Fertility Rates residualised on socioeconomic gradients in desired fertility

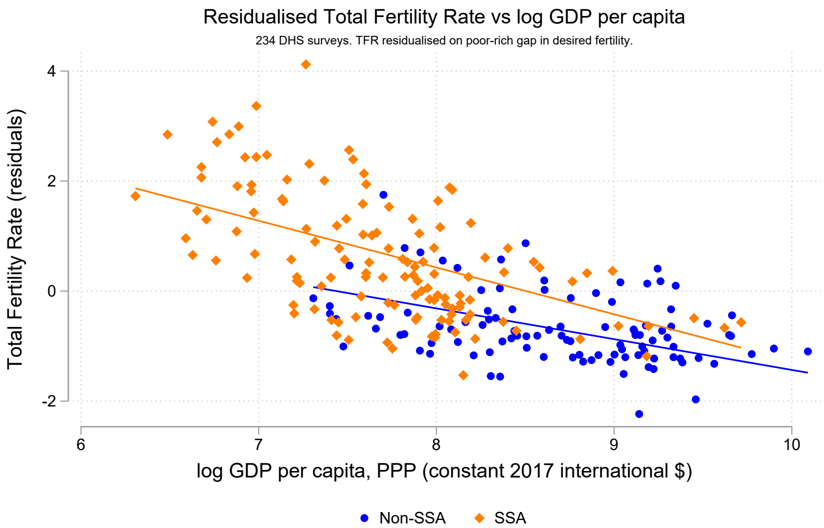

Figure 4 plots the Total Fertility Rate for the 233 DHS surveys for which the Total Wanted Fertility Rate is available at the country-year-quintile level, against log GDP per capita on the x-axis. It is the counterpart to Figure 2 in section 1, but where this time the outcome is residualised on the poor-rich gap in wanted fertility. As is visually evident, the two clouds of residuals are distinctly closer together in Figure 4 than the raw Total Fertility Rates are in Figure 2. Appendix Table A.1 displays the regression estimates from this exercise: it shows that controlling for the gradient makes the coefficient estimate on the

[11] As illustrated by Appendix Figure A.4, the ideal number of children indicator is positively correlated with total fertility rates at the country-year-quintile level. The correlation coefficient is 0.78.

sub-Saharan Africa dummy drop from 1.401 to 1.051, which corresponds to a 25% decrease in the total fertility gap captured by the dummy when the gradient is excluded from the specification. This suggests that the steeper wealth-desired fertility gradient accounts for 25% of the fertility gap between sub-Saharan Africa and other regions conditional on GDP per capita.

In sum, these findings suggest that the differential gradient in sub-Saharan Africa accounts for a significant share of the overall disparity between the region's fertility rates and that of other countries, even conditional on national income.

1.3.3 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) pooling together 56 DHS surveys.

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 179 surveys included in the analysis 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 overall living standards. Indeed, I find that the residualised wealth score of sub-Saharan African households is 1.2 standard deviations lower on average. Hence, for the same level of GDP per capita, there must be significantly higher rates of poverty in sub-Saharan African countries. However, the difference in absolute living standards across the two regions is by no means more pronounced for the poorest quintiles. This suggests that the steeper socioeconomic gradient in desired fertility in sub-Saharan Africa is not simply the result of larger income dispersion in these countries conditional on GDP per capita.

Figure 1.5: Distribution of international wealth index (residualised on GDP per capita) in sub-Saharan Africa vs. other developing countries

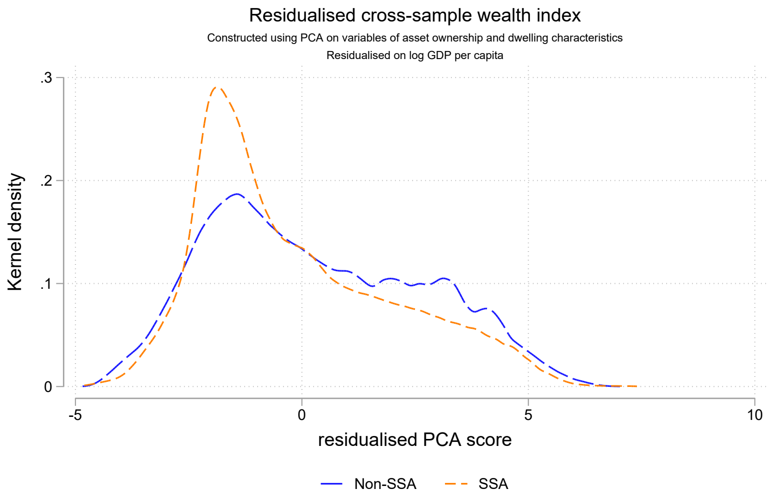

Table A.2 in the appendix presents this finding in regression form, namely that the gap in living standards between the top quintile and the two poorest ones is statistically significantly *smaller* in sub-Saharan Africa. This is also illustrated in Appendix Figure A.5, where we see that the gradient between within-sample wealth quintiles and the cross-sample wealth index (on the y-axis) is indeed less steep for the sub-Saharan African countries in this sample.

1.3.4 From the income-desired fertility gradient to jobs

The findings uncovered in this section raise a new question: what could explain the steeper socioeconomic gradient in desired fertility in sub-Saharan Africa? This paper proposes to focus on the hypothesis that access to labor market opportunities - and in particular, off-farm salaried jobs - can account for at least part of the difference in demand for children across sub-Saharan African and other countries.

Figure 1.6: Gradient in desired fertility vs. gradient in salaried employment rate

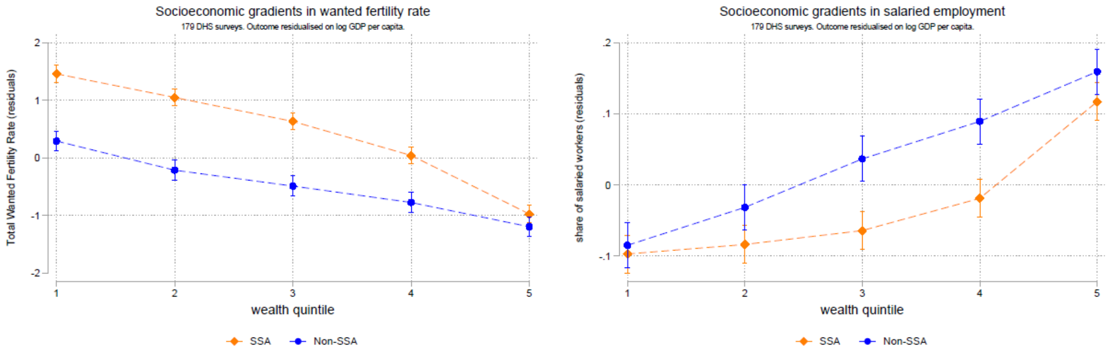

To characterise the type of employment that individuals are most likely to engage in, I construct two variables from the DHS micro data. The first is the share of working women engaging in agricultural activities. The second is the share of working women who report receiving payments from working for someone else. A simple regression of these two variables on a binary indicator for sub-Saharan Africa and log GDP per capita reveals two facts. First, the share of working women in agriculture is not significantly different in sub-Saharan African relative to other DHS countries, as shown in columns (4) and (5) of Appendix Table A.3. Second, female salaried employment rates are 9.6 percentage points lower in sub-Saharan African countries on average, conditional on GDP per capita (column (2) of Appendix Table A.3). This difference is statistically significant at the 10% level and corresponds to 35% of the average salaried employment rate in countries outside of sub-Saharan Africa. This association is also robust to controlling for differences in average female years of schooling across countries: the coefficient estimate on the sub-Saharan African dummy in column (3) of Appendix Table A.3 indicates an 8.2 percentage point difference in salaried employment rates conditional on GDP per capita and average years of schooling.

Figure 6 illustrates this: it shows that the gap in total wanted fertility rates (residualised on GDP per capita) in panel (a) is mirrored by a substantial gap in the share of women in salaried employment (also residualised on GDP per capita). Figure 6b shows that this gap is concentrated in the three middle quintiles, and is largest for quintiles 3 and 4 in particular. The slope of the socioeconomic gradient in salaried employment is also noticeably different across the two sets of countries: the slope is, on average, smaller in sub-Saharan Africa for the three poorest quintiles, then becomes steeper between quintiles 3 and 4, and even more so between quintiles 4 and 5, leading the residualised salaried employment rate of the top quintile to almost catch up to its value outside of sub-Saharan Africa. This more convex shape mirrors that of the shape of the sub-Saharan African

Table 1.1: The role of salaried and off-farm employment rates in sub-Saharan Africa's exceptionally high desired fertility.

	Total Wanted Fertility Rate					
	(1)	(2)	(3)	(4)	(5)	(6)
sub-Saharan Africa	1.443***	1.163***	1.379***	0.746***	0.622***	0.741***
	[0.222]	[0.234]	[0.247]	[0.209]	[0.218]	[0.204]
% salaried employment		-2.915***			-1.653***	
		[0.462]			[0.487]	
% non-farm employment			-1.838***			0.048
			[0.382]			[0.566]
SSA x quintile 1				0.947***	0.997***	0.953***
				[0.225]	[0.222]	[0.239]
SSA x quintile 2				1.042***	1.026***	1.051***
				[0.169]	[0.169]	[0.196]
SSA x quintile 3				0.902***	0.805***	0.913***
				[0.134]	[0.139]	[0.182]
SSA x quintile 4				0.596***	0.486***	0.603***
				[0.109]	[0.114]	[0.144]
Quintile dummies				Yes	Yes	Yes
R^2	0.478	0.593	0.577	0.676	0.705	0.676
Outcome Mean	3.677	3.677	3.677	3.677	3.677	3.677
Observations	895	895	895	895	895	895

* p<0.1, ** p<0.05, *** p<0.01. Each observation is a DHS survey x quintile. Standard errors are clustered at the country level. All regressions are weighted by the inverse of the number of DHS waves per country and control for log GDP per capita.

gradient in Total Wanted Fertility Rates, pictured in Figure 6a, which is also linear for the first three quintiles and then becomes more negative between quintiles 3 and 4 and even more so between quintiles 4 and 5.

The facts displayed in Figure 6 and Appendix Table A.3 thus motivate the paper's focus on the role that especially low rates of salaried employment opportunities in sub-Saharan African labor markets may play in the persistently high demand for children documented in this region.

To estimate how much of the gap in desired fertility between sub-Saharan African and other countries can be accounted for by differences in the nature of employment across these two sets of countries, I first regress the Total Wanted Fertility Rate on log GDP per capita and a binary indicator equal to 1 for sub-Saharan African countries at the country-year-quintile level. The estimates of this baseline specification are reported in column (1) of Table 1. Columns (2) and (3) report the results of also controlling for the salaried employment rate and the share of non-agricultural employment respectively. I then observe the effect of adding each of these controls to the specification on the sub-Saharan Africa dummy coefficient. What is evident from the estimates reported in the top row of columns (1)-(3) in Table 1 is that salaried employment rates can account for a much larger fraction of the wanted fertility gap (nearly 20%) than non-agricultural employment (less than 5%):

controlling for the former leads the sub-Saharan Africa dummy coefficient to decrease by 28 percentage points whereas controlling for the latter lowers it by only 6.4 percentage points.

Columns (4)-(6) of Table 1 display the results of a second specification which also includes dummy variables for quintiles 1-4 and an interaction term between the sub-Saharan Africa dummy and each of these four dummies. The key take-away to draw from this additional test is that salaried employment accounts not only for a sizeable fraction of the overall level difference in desired fertility between sub-Saharan African and other countries at the same level of GDP per capita, but also accounts for part of the differential slope with respect to socioeconomic ranking. Indeed, examining the last four rows of columns (4) and (5) shows that the coefficient estimate on the interaction term is most affected by the addition of the salaried employment rate covariate for quintiles 3 and 4 - precisely those for which the wage employment gap is largest across the two sets of countries, as seen in Figure 6b. Here again, I find no such difference between the estimates of column (4) and (6), which further corroborates the conclusion that the share of agricultural employment cannot account for much of the desired fertility gap between sub-Saharan African and other countries.

Taken together, the findings in Table 1 and Figure 6 suggest that the relative scarcity of salaried employment opportunities in sub-Saharan Africa may play a role in women's exceptionally high desired fertility in this region. The next section formalises possible channels linking these two facts in a theoretical framework. One determinant of both labor market and fertility outcomes will be a key factor in this relationship: education. Indeed, the types of jobs that women - and their children in the future - have access to may crucially depend on human capital investments.

1.4 Theoretical framework

1.4.1 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 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 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 wealth-desired fertility gradient and the fact that individuals in the poorer quintiles are more likely to work in informal occupations in sub-Saharan Africa (Figure 6b), 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]

1.4.2 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]

[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.

[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

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]

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:

find in the data. 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 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.

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

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:

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

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

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

The condition for this set of solutions to exist is:

$$a + wh > \frac{\theta_0}{\theta_1 \beta \tau} \tag{1.1}$$

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 in 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 τwh represent the benefit of an additional unit of educational expenditure and the opportunity cost of one more child respectively. As such, they are key 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.

[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 1.7: Predicted relationship between parental education and (a) fertility (b) educational investments in children

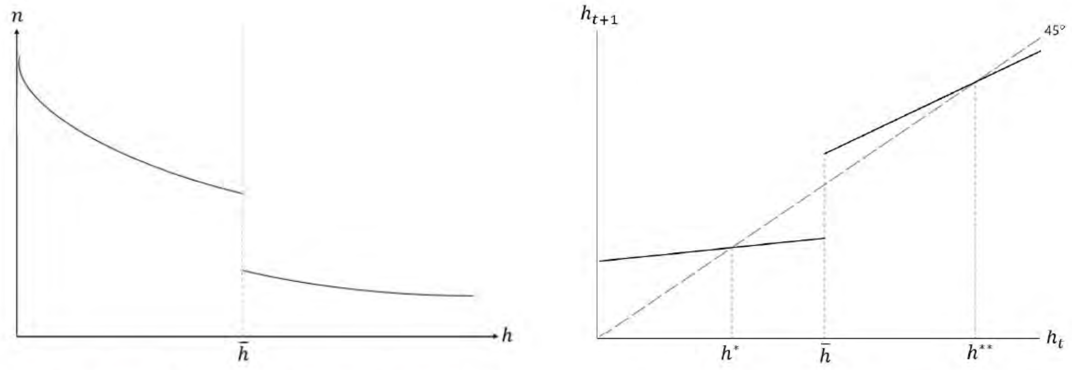

(a) Number of children vs. parental human capital

(b) Human capital 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):

$$\bar{h} > \frac{\beta(\underline{\theta_1 \tau} a - \theta_0)}{1 - \beta(1 + \underline{\theta_1 \tau w})} \tag{1.2}$$

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.

1.4.3 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.

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

[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.7 in the Appendix - that neither the distribution of years of education nor the relationship between wealth and years of education is statistically

Figure 1.8: Gradient differential and human capital transition equation

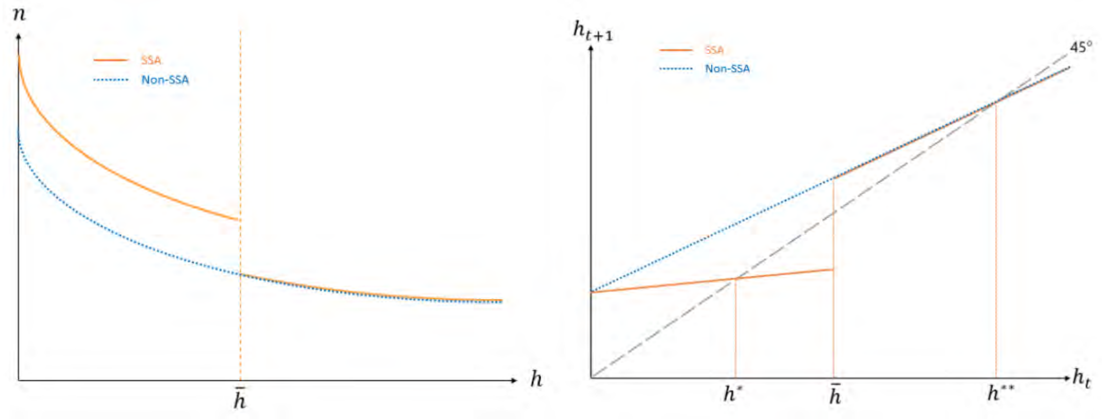

of children born to the average household in sub-Saharan Africa stands above that in other countries for levels of parental human capital below the threshold, after which the relationships are the same. This reflects two possible scenarios according to the model. The first is that the human capital threshold 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

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.

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 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 wage 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]

1.5 From Theory to Empirics

This section leverages micro data from three different sources to validate the model discussed above. Section 5.1 presents empirical support for both the key assumption and the key prediction of the model. Section 5.2 provides evidence for the two main channels linking jobs to fertility in the framework: parental valuation of education and the opportunity cost of children.

1.5.1 Model validation

The two empirical facts presented in this section, Figures 9 and 10, were produced from IPUMS micro-census data. The figures plot cross-regional comparisons constructed from a set of 64 IPUMS samples: 25 from sub-Saharan Africa (the orange curves) and 39 from other developing countries (the blue curves). To account for the fact that most sub-Saharan African countries have lower GDP per capita in this sample and that my outcomes of interest (salaried employment rates and educational investments in children) 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.

[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 1.9: Probability of holding a salaried job by years of education completed conditional on GDP per capita

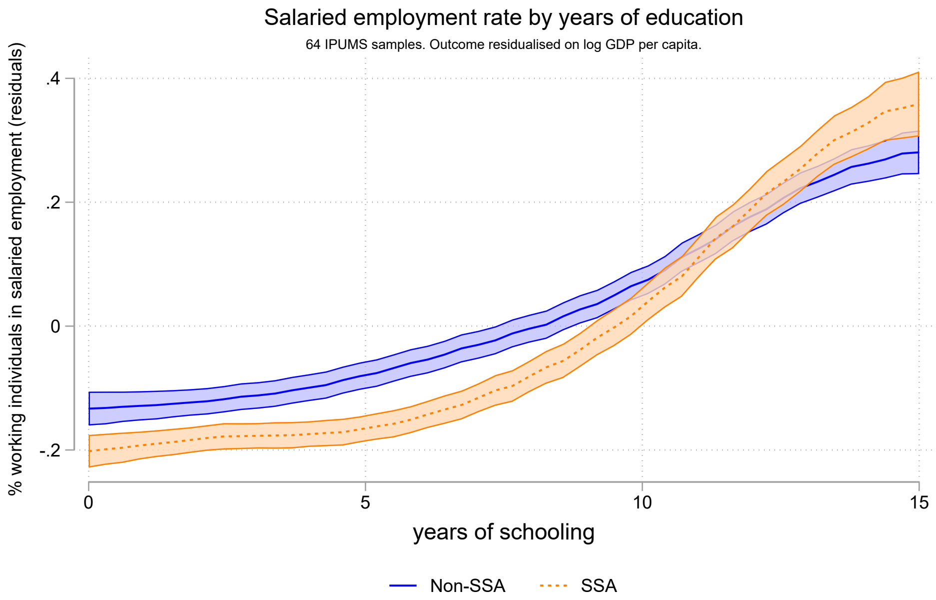

Smoothed local polynomial regression estimates with 95% confidence bands. Estimates weighted by the inverse of the number of waves per country. Residuals from regressing salaried employment indicator on log GDP per capita. Salaried employment indicator is defined for individuals who report being employed at the time of the census. 64 IPUMS samples: 25 from sub-Saharan Africa, 39 from other developing countries. Number of observations: 36,173,556 individuals aged 25-35.

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 pronounced in sub-Saharan Africa. The share of individuals in salaried employment is virtually flat (less than 10% of individuals on average) from 0 to around 4.5 years of education in sub-Saharan Africa. The figure shows a gap in wage employment rates across the two sets of countries that is largest in the 4-8 years of schooling range, after which the slope becomes steeper in sub-Saharan Africa, with the curves joining up around 11 years of schooling.

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 in 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

Figure 1.10: Educational attainment of children conditional on mother's total years of education

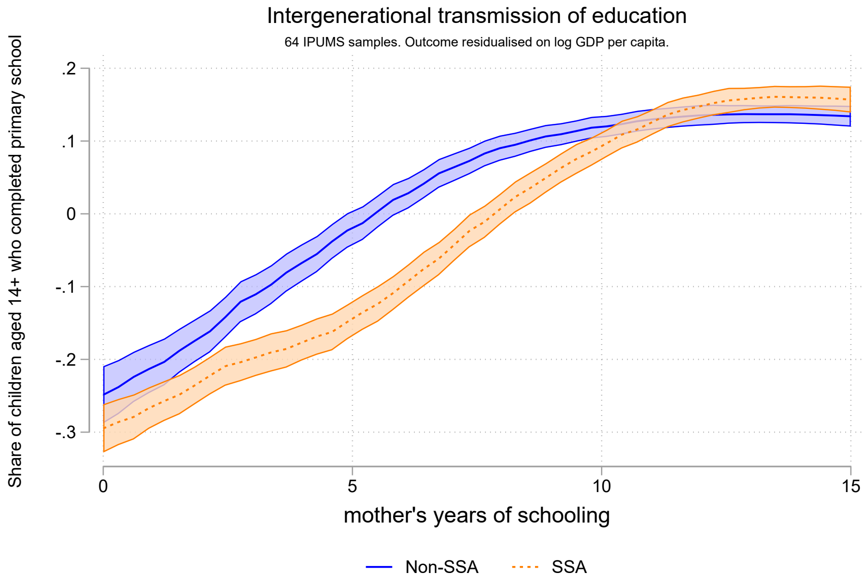

Smoothed local polynomial regression estimates with 95% confidence bands. Estimates weighted by the inverse of the number of waves per country. Residuals from regressing household share of children aged 14+ who completed primary school on log GDP per capita. 64 IPUMS samples: 25 from sub-Saharan Africa, 39 from other developing countries. Total number of observations: 14,977,065 mothers residing with at least one child aged 14 or above.

τ is comparatively higher in those countries.[24] 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 labourers in agriculture. Such contracts, which typically pay low wages, may also be less compatible with raising many children than self-employment.

Figure 9 is also consistent with the mechanisms linking fertility decisions to labor market opportunities that underlie the framework's predictions. 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 θ_1 in sub-Saharan Africa. If women's labor market options are less likely to involve salaried work outside the home, this would also reduce the opportunity cost of children (via a lower time cost τ) and thereby curb their incentives to reduce their fertility.

Figure 10 plots the fraction of children aged 14 and above who have completed primary

[24] Appendix Figure A.8, which plots the relationship between years of schooling and the share of workers in salaried employment for 25 samples separately, reveals substantial heterogeneity across countries outside of sub-Saharan Africa. For example, Bangladesh features 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, while Morocco presents a much more linear relationship, which is more consistent with the absence of a human capital threshold for entry into salaried employment.

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. Two main findings emerge 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 4.5 years of education, and eventually the two curves overlap, from around the 11-year mark. This mirrors the behaviour of the curves plotted in Figure 9 to some extent.

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.[25] These could explain comparatively low educational investments in children among the poorer segments of this region's population.

Another possible interpretation of Figure 10, when examined jointly with Figure 9, is that the fixed human capital requirement for salaried employment is simply absent in at least some countries outside of sub-Saharan Africa.[26] The fact that there is a kink in the orange curve just before 5 years of education but no detectable change in the slope of the blue curve at that point supports this interpretation. Taken together, these facts may 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, this threshold must be absent in at least some countries outside of sub-Saharan Africa. Figure 10 shows that most parents make comparatively lower

[25] The slope of the human capital transition equation is also partly determined by another parameter, returns to parents' human capital w. I deliberately ignore the role of this parameter in explaining these cross-regional disparities for lack of representative data that would allow me to construct a reliable estimate of its value comparably across countries. This may be the subject of follow-up work on this topic.

[26] As shown by Figure 8 in section 4.3, the existence of a human capital threshold in sub-Saharan African labor markets only could account for different intergenerational dynamics in human capital investments.

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 an increase in the probability of holding a salaried job.

These findings suggest that part of the explanation for the steeper wealth-desired fertility gradient observed in sub-Saharan Africa may in the existence of a fixed human capital requirement for salaried employment that makes it particularly difficult for the poorest families to move away from informal self-employment in these countries.

1.5.2 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 theoretical 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 include ethnic group and religion fixed effects. I also control for individuals' total years of education, rural/urban residence status, and knowledge of modern contraceptive methods.[27]

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 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.

[27]While the role of access to contraceptives in sub-Saharan Africa's exceptionally high fertility rates is beyond the scope of this paper, I include this variable in the set of covariates because contraceptive knowledge and availability could plausibly influence women's reported desired family size by giving them a greater sense of control over their fertility, or shifting norms about ideal family size downwards (Pritchett (1994)). The specifications control for knowledge only, however, as this measure is less subject to reserve causality concerns than current use of contraception.

Table 1.2: Correlates of men's desired fertility in sub-Saharan Africa.

	Ideal number of children			
	(1)	(2)	(3)	(4)
non-farm occupation	-0.455***			-0.375***
	[0.054]			[0.055]
works for pay		-0.331***		-0.185**
		[0.077]		[0.078]
works all year			-0.204***	-0.109**
			[0.055]	[0.054]
Country x ethnic group FE	Yes	Yes	Yes	Yes
Country x religion FE	Yes	Yes	Yes	Yes
Country x year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
R^2	0.272	0.271	0.271	0.272
Outcome Mean	6.283	6.283	6.283	6.283
Observations	147921	147921	147921	147921

All specifications control for the following individual-level characteristics: rural/urban, years of education and knowledge of modern contraceptive methods. All regressions are weighed by sampling weights (normalized by population size and number of DHS surveys per country). Standard errors clustered by province x survey year in brackets. * p<0.1, ** p<0.05, *** p<0.01.

Jobs and parental valuation of education

Table 2 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, education, desired fertility, ethnic group and religion.

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 (1), (2) and (3). The coefficient estimates suggest that each of these indicators is negatively correlated with men's desired fertility. For instance, column (1) reports that, holding everything else constant, men working outside of agriculture report an ideal number of children almost half a child lower on average. Since non-farm workers are more likely to work for pay and less likely to engage in seasonal work, I also control for these three characteristics together in specification (4). The estimates in column (4) suggest that men's ideal family size is about 0.7 children lower if they work for pay, outside of agriculture and all year round, conditional on the specification's rich set of covariates. This represents a drop of about 11% relative to the sample mean.

Figure 1.11: Distribution of (a) number of children ever born and (b) educational expenditure per child by household head's type of employment.

Kenya Integrated Household Budget Survey 2005-2006.

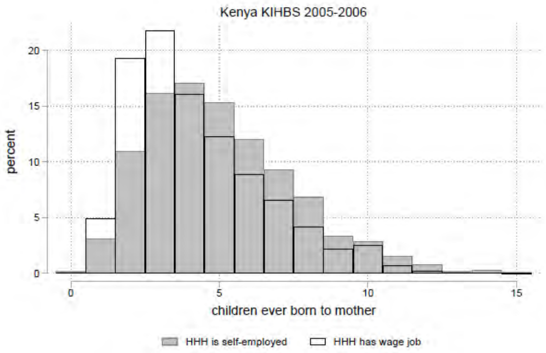

(a) Children ever born to mother

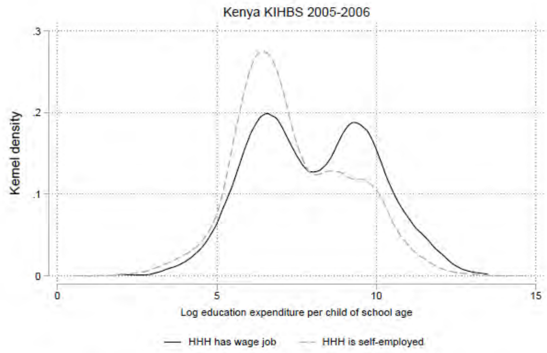

(b) Distribution of education expenditure

A key takeaway from Table 2 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. 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.[28] Figure 11a shows that the 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). Figure 11b 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.[29] In the Appendix, Table A.4 presents these associations in regression form, controlling

[28] 5% of households, who report no education spending, get dropped from the sample in this analysis as a result of the log transformation.

[29] The discussion of the role of jobs in fertility choices abstracts from income considerations. Of course,

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 households in which the head has a salaried job are more likely to reduce fertility in favour of larger educational investments in their children.

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 in this region.

Figure 1.12: Socioeconomic gradient in female years of schooling

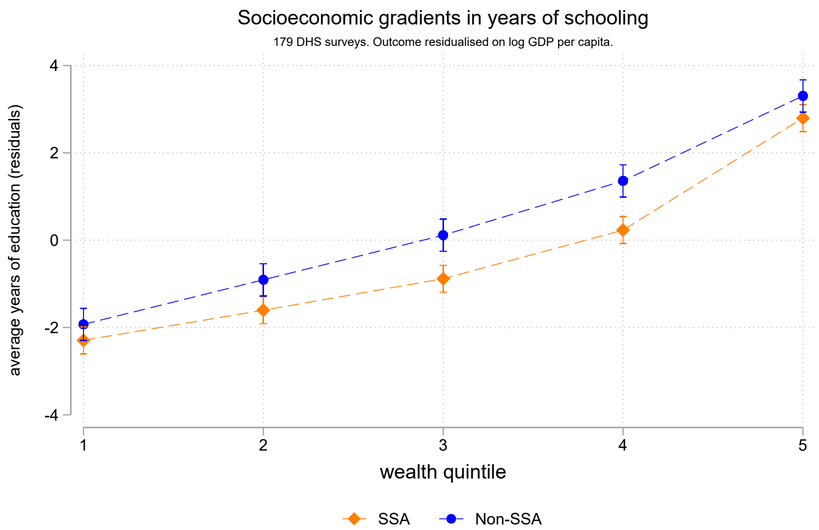

179 DHS surveys. 105 from sub-Saharan Africa, 74 from other low-income countries.
Number of observations: 3,141,292 women aged 15-49.

Figure 12 shows that average years of female education - residualised on GDP per capita - are notably closer in value across the two sets of countries than the share of women in salaried employment indicator displayed in Figure 6b. The level differences are statistically significantly different for quintiles 3 and 4 only. This suggests that if the opportunity cost

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 (Appendix Figure A.9).

of time channel plays a role in the desired fertility gradient differential across sub-Saharan African and other countries, it must operate only partly through female education.[30]

Figure 1.13: Female Labor Force Participation conditional on education

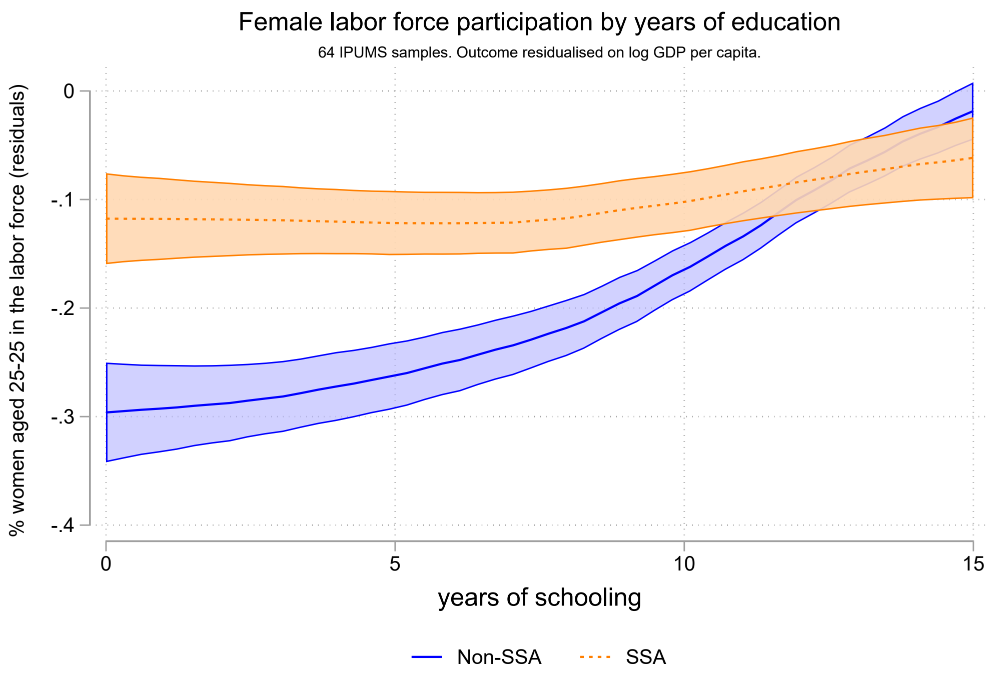

Smoothed local polynomial regression estimates with 95% confidence bands. Estimates weighted by the inverse of the number of waves per country. Residuals from regressing Female Labor Force Participation indicator on log GDP per capita. 64 IPUMS samples: 25 from sub-Saharan Africa, 39 from other developing countries. Total number of observations: 14,088,269 women aged 25-35.

Figure 13 plots the relationship between female labor force participation and years of schooling separately for sub-Saharan African and other countries. Two main findings are of note. First, there is a very large level difference across the two sets of countries, controlling for GDP per capita, of up to 20 percentage points. This gap is statistically significant for the 0-11 years of schooling range. This stands in contrast to the documented fact that female labor force participation typically follows a U-shaped relationship with economic development (Goldin (1995)). It has indeed been observed that the early stages of the structural transformation process, typically characterised by a shift from informal (e.g. subsistence agriculture) to formal employment (e.g. manufacturing) in the occupational structure, coincided with a drop in female labor force participation in many countries.

[30] Heath and Jayachandran (2018) review the impacts of female education on fertility in developing countries. They find that most studies detect 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)).

Here, we see a much larger share of women participating in the labor market in sub-Saharan Africa relative to other countries at the same level of GDP per capita.[31] In line with the overarching takeaway from this paper, this suggests that sub-Saharan Africa may be undergoing a unique occupational change process - much like its fertility transition.

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 8 years of education, after which it increases gently. In other countries, the curve is (almost) monotonically increasing, with a steep positive relationship appearing from around 3 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 to participate) is influenced by the options available to them on the intensive margin (how the labor market rewards their time if they participate). Job opportunities in manufacturing or services, for instance, may provide them with an incentive to both increase their educational attainment and reduce their fertility.[32]

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 their occupations tend to be informal self-employment activities. These typically yield low returns to their education but could also provide 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 use the DHS and IPUMS micro data to correlate the desired and realised fertility outcomes of female respondents with the availability of employment opportunities in their location of residence. Tables 3 and 4 report these correlations, focusing on the availability of two types of jobs: off-farm work - which includes self-employment activities - and salaried employment.

[31] Since the distribution of female years of schooling is quite similar in the two sets of countries (see Figure 12 and Appendix Figure A.7), it must be that female labor force participation rates are also substantially higher on aggregate in sub-Saharan Africa, conditional on GDP per capita.

[32] 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 1.3: Correlations between women's ideal number of children and the availability of off-farm and salaried employment in sub-Saharan Africa (DHS).

	Ideal number of children			
	(1)	(2)	(3)	(4)
% wage work in province	-1.706***	-1.991***		
	[0.533]	[0.574]		
works for a wage		-0.377***		
		[0.046]		
works for a wage * % wage work in prov.		1.567***		
		[0.456]		
% non-farm work in province			-0.858***	-0.785***
			[0.174]	[0.185]
non-farm work				-0.121**
				[0.056]
non-farm work * % non-farm work in prov.				-0.027
				[0.144]
Country x ethnic group FE	Yes	Yes	Yes	Yes
Country x religion FE	Yes	Yes	Yes	Yes
Country x year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
R^2	0.362	0.363	0.362	0.363
Sample Mean	5.212	5.212	5.212	5.212
Observations	600968	600968	600968	600968

All specifications control for the following individual-level characteristics: rural/urban residence, years of education and knowledge of modern contraceptive methods. Regressions are weighted by sampling weights normalized by population size and number of DHS waves per country. Standard errors clustered by province X year in brackets.
*p<0.1, **p<0.05, ***p<0.01.

Table 3 displays correlations produced from a sample of over 600,000 women in sub-Saharan Africa, constructed by pooling all the DHS women's samples in the IPUMS-DHS database at the time of writing. The outcome variable here is the ideal number of children reported by female respondents of childbearing age.[33] Specifications (1)-(4) control for the same set of covariates as those used to produce Table 2 for the men's samples. The key difference here is that the explanatory variable of interest is the share of employment opportunities available in female respondents' province, which I use as a proxy for the opportunity cost of childbearing.

The first key take-away one can draw from this table is that the coefficient on the share of salaried employment in column (1) is much more negative – over twice as large in magnitude – than the one on the share of non-farm employment in column (3). Column (1) shows that, conditional on the specification's rich set of covariates, a 10 percentage point increase in the share of women working in salaried employment in a province is associated with a drop in ideal number of children of about 0.17 on average. The corresponding

[33]As a robustness check, I run the same set of specifications used to produce Table 3 on a sample restricted to *childless* women, to address concerns that the ideal number of children variable may be biased by ex-post rationalisation of prior births (as discussed in section 2). Results are reported in Appendix Table A.6. The coefficient estimates are smaller in magnitude but their sign and statistical significance survive.

coefficient estimate for off-farm work in column (3) is only about 0.09. This finding supports the hypothesis that the availability of salaried employment - assumed to be higher-return and less compatible with raising many children than non-agricultural self-employment - increases the opportunity cost of children.

Columns (2) and (4) display the results of a specification which interacts our variable of interest (i.e. the share of salaried and non-farm employment respectively) with the respondent's own employment status - whether they are a salaried worker and whether they work outside of agriculture respectively. Here, the significantly positive coefficient on the interaction term in column (2) indicates that women in salaried jobs report a lower ideal number of children in provinces where such opportunities are particularly scarce. It shows that, conditional on a woman being a salaried worker, a 10 percentage point decrease in the share of female salaried employment in the province is associated with a 0.16 drop in her ideal number of children. In other words, the correlation between having a salaried occupation and a woman's desired fertility becomes stronger (i.e. more negative) with the relative scarcity of such jobs in her province. The fact that the interaction yields an insignificant coefficient estimate for off-farm work in column (4) is, once again, consistent with salaried employment being associated with a higher opportunity cost of children than non-agricultural work more generally.

It is important to keep in mind that these estimates represent correlations, and as such, may be subject to more than one interpretation. For instance, an alternative interpretation of the association between the prevalence of female salaried employment and women's desired fertility reported in Table 3 is that it may only be capturing heterogeneity in social norms across provinces. Areas characterised by more conservative gender norms may exhibit both a lower share of women in salaried employment and higher levels of desired fertility, for example. To test this hypothesis, I exploit variation in the destination of female migrants from the same district in the Malawi 2008 census.

Table 4 reports 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

Table 1.4: Correlations between the share of salaried employment in a woman's district of residence and her (i) participation in salaried employment and (ii) fertility, by migration status (Malawi IPUMS 2008).

	Never Migrated (1)	Migrated for marriage (2)	migrated before marriage (3)	Migrated after marriage (4)
	Salaried employment			
% 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]
Outcome Mean	0.045	0.090	0.087	0.116
R^2	0.079	0.119	0.117	0.131
	Number of children ever born			
% 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]
Outcome Mean	4.052	2.532	3.503	4.283
R^2	0.563	0.629	0.542	0.465
Ethnic group FE	X	X	X	X
Religion FE	X	X	X	X
Age FE	X	X	X	X
Observations	104590	10140	10031	22464

All specifications control for rural/urban residence, woman's years of education, husband's years of education and type of employment (farm/off-farm and salaried/self-employment). Standard errors clustered by sub-district in brackets. *p<0.1, **p<0.05, ***p<0.01.

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 results displayed in the bottom panel of Table 4. First, the different coefficient estimates reported in columns (2), (3) and (4) on the rate of female salaried employment in a migrant woman's destination suggest that women who move to areas with higher rates of salaried employment tend to give birth to fewer children. Of course, the different magnitudes of this coefficient 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 be capturing the fact that women's decision to migrate after the year of their marriage - so, presumably with their husband - is likely to be driven by unobservable characteristics of these couples that are also correlated with lower desired fertility.[34]

[34] 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.

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 4 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 the table. This finding corroborates the opportunity cost of children hypothesis, because it suggests that fertility is not solely determined by social norms. If norms around fertility were the main factor driving the negative association between the share of women engaged in salaried work and women's ideal number of children in a province reported in Table 3, 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 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. The fact that the estimates in the bottom row of Table 4 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 3 and 4 thus supports the hypothesis that salaried employment opportunities are tightly linked to women's fertility choices because of the opportunity cost of children associated with this type of work.

1.6 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 a quarter of the overall gap in fertility between sub-Saharan African and other developing countries is explained by the presence of a steeper income-desired fertility gradient in the former, even when controlling for GDP per capita. Through the lens of a

Beckerian framework supported by micro data from several sources, the paper links this to the relative scarcity of salaried employment opportunities in sub-Saharan Africa. The latter accounts for about 20% of the gap in desired fertility between these two sets of countries.

Recent evidence has shown that limited occupational choice sets can generate poverty traps (Bandiera et al. (2017); Balboni et al. (2020)). This paper complements the findings from these studies by highlighting another set of mechanisms through which the availability of different types of jobs may affect the intergenerational persistence of poverty. I argue that the distribution and nature of labor market opportunities available in sub-Saharan Africa may be responsible for a feedback loop between informal self-employment and high fertility that persists via families also under-investing in their children's education. This study also emphasizes that the fertility decisions are more strongly correlated with the prevalence of salaried work than that of self-employment in non-agricultural activities. Thus, one key policy implication from this paper is that large investments aimed at increasing the set of salaried employment opportunities available to the poorer segments of the population may be just as important as improving access to contraceptives - if not more - for reducing fertility in sub-Saharan Africa.[35]

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. Accounting for differences in desired fertility across spouses may allow us to better understand the pace of fertility transitions across and within sub-Saharan African countries.[36] 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 is displayed in Appendix Figure A.12.

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:

[35]Quantifying the impacts of expanding access to salaried jobs relative to contraceptives on fertility and the general equilibrium effects of such big push expansion policies on the income-desired fertility gradient in sub-Saharan African countries is left for future work.

[36]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.

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.[37] 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.[38]

Polygyny might also merit further examination through the lens of socioeconomic gradients. While cross-sectional evidence, as emphasized in section 3, suggests that this practice is not more common amongst the poorest segments of the population, the distribution of job opportunities may interact with social norms in shaping incentives to adhere to this practice. The top panel of Appendix A.18 illustrates that women in polygynous marriages tend to report higher desired fertility, consistent with the findings of Rossi (2019). The bottom panel of Appendix A.18 plots the evolution of polygyny rates by wealth quintile over repeated cross-sections of the DHS surveys for 16 sub-Saharan African countries. While the fraction of married women in a polygynous union appears relatively stable over time on aggregate, the figure presents suggestive evidence that polygyny rates have been on the decline in recent years for richer women in some countries (e.g. Benin, Cameroon and Ghana). One possible explanation could be that, as occupational structures evolve, the outside options available to more educated women improve, inciting them to deviate from this historically determined norm if they wish to. Future research could be directed at testing this hypothesis.

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

[37]Appendix 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.

[38]The role of spousal bargaining power in fertility has been documented in other settings. In Malaysia, Rasul (2008) 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. Future work could involve fitting an intra-spousal bargaining model of fertility to sex-disaggregated household survey data on fertility, educational investments and expenditure.

socioeconomic status, desired fertility and occupations within sub-Saharan Africa.[39] 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.

[39] Appendix Figure A.11 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 across countries. In some countries, like Congo and Rwanda, the two gradients are virtually flat, whereas in others, like Nigeria and Zimbabwe, they are very steep.

References

Daniel Aaronson, Fabian Lange, and Bhashkar Mazumder. Fertility Transitions along the Extensive and Intensive Margins. *American Economic Review*, 104(11):3701–24, November 2014.

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

Nava Ashraf, Erica Field, Giuditta Rusconi, Alessandra Voena, and Roberta Ziparo. Traditional Beliefs and Learning about Maternal Risk in Zambia. *American Economic Review*, 107(5):511–15, May 2017.

Nava Ashraf, Natalie Bau, Corinne Low, and Kathleen McGinn. Negotiating a Better Future: How Interpersonal Skills Facilitate Intergenerational Investment. *The Quarterly Journal of Economics*, 135(2):1095–1151, 01 2020.

Clare Balboni, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil. Why Do People Stay Poor? *mimeo*, 2020.

Oriana Bandiera, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman. Labor Markets and Poverty in Village Economies*. *The Quarterly Journal of Economics*, 132(2):811–870, 03 2017. ISSN 0033-5533.

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

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

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

John Bongaarts. The Measurement of Wanted Fertility. *Population and Development Review*, 16(3):487–506, 1990. ISSN 00987921, 17284457.

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

Esther Boserup. The Conditions of Agricultural Growth : The Economics of Agrarian Change Under Population Pressure. 1965.

Lucia Breierova and Esther Duflo. The Impact of Education on Fertility and Child Mortality: Do Fathers Really Matter Less Than Mothers? May 2004.

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

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

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

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

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

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

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

Erica Field, Vera Molitor, Alice Schoonbroodt, and Michèle Tertilt. Gender Gaps in Completed Fertility. *Journal of Demographic Economics*, 82(2):167–206, 2016.

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

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

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

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

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

Oded Galor and Joseph Zeira. Income Distribution and Macroeconomics. *The Review of Economic Studies*, 60(1):35–52, 01 1993.

Maitreesh Ghatak. Theories of Poverty Traps and Anti-Poverty Policies. *The World Bank Economic Review*, 29:77–105, 04 2015. ISSN 0258-6770.

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

Douglas Gollin, Remi Jedwab, and D. Vollrath. Urbanization with and without industrialization. *Journal of Economic Growth*, 21, 2016.

Rachel Heath and Seema Jayachandran. 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.*, 2018.

Rachel Heath and A. Mushfiq Mobarak. Manufacturing Growth and the Lives of Bangladeshi Women. *Journal of Development Economics*, 115:1 – 15, 2015.

ICF. The DHS Program STATcompiler, 2012.

Robert Jensen. 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, 03 2012.

Larry Jones, Alice Schoonbroodt, and Michèle Tertilt. Fertility Theories: Can They Explain the Negative Fertility-Income Relationship? *Demography and the Economy*, 2011.

Shareen Joshi and T. Schultz. Family Planning and Women's and Children's Health: Long-Term Consequences of an Outreach Program in Matlab, Bangladesh. *Demography*, 50, 12 2012.

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

Michael Kremer. Population Growth and Technological Change: One Million B.C. to 1990*. *The Quarterly Journal of Economics*, 108(3):681–716, 08 1993.

Thomas Robert Malthus. *An Essay on the Principle of Population*. Number malthus1798 in History of Economic Thought Books. McMaster University Archive for the History of Economic Thought, 1798.

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

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

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

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

Grant Miller. Contraception as Development? New Evidence from Family Planning in Colombia. *The Economic Journal*, 120(545):709–736, 08 2009.

Omer Moav. Cheap Children and the Persistence of Poverty. *The Economic Journal*, 115 (500):88–110, 12 2005.

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

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

R. Osei and R. Jedwab. Is Structural Change in Sub-Saharan Africa Different? New Historical Evidence from Ghana. *mimeo*, 2013.

Una Osili and Bridget Long. Does Female Schooling Reduce Fertility? Evidence from Nigeria. *Journal of Development Economics*, 87(1):57–75, 2008.

Claus C Portner, Kathleen Beegle, and Luc Christiaensen. Family Planning and Fertility: Estimating Program Effects Using Cross-Sectional Data. 2011.

Lant Pritchett. Desired Fertility and the Impact of Population Policies. Policy Research Working Paper Series 1273, The World Bank, 1994.

Imran Rasul. Household Bargaining over Fertility: Theory and Evidence from Malaysia. *Journal of Development Economics*, 86(2):215 – 241, 2008. ISSN 0304-3878.

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

Pauline Rossi. Strategic Choices in Polygamous Households: Theory and Evidence from Senegal. *The Review of Economic Studies*, 86(3):1332–1370, 09 2019.

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

Michèle Tertilt. Polygyny, Women's Rights, and Development. *Journal of the European Economic Association*, 4(2-3):523–530, 05 2006.

Tom S. Vogl. Differential Fertility, Human Capital, and Development. *The Review of Economic Studies*, 83(1):365–401, 07 2016. ISSN 0034-6527.

IPUMS Demographic and Health Surveys. Version 8. Elizabeth Heger Boyle, Miriam King, and Matthew Sobek. Minneapolis, MN: IPUMS and ICF. 2020. https://doi.org/10.18128/D080.V8.

Demographic and Health Surveys (various). Funded by USAID. Rockville, Maryland: ICF. 1990-2017.

Integrated Public Use Microdata Series, International: Version 7.3. Minnesota Population Center. Minneapolis, MN: IPUMS. 2020. https://doi.org/10.18128/D020.V7.2.

Kenya National Bureau of Statistics. Integrated Household Budget Survey (IHBS) v01. 2005-2006.

Appendix to Chapter 1.

Figure A.1: Mortality trends over time by region

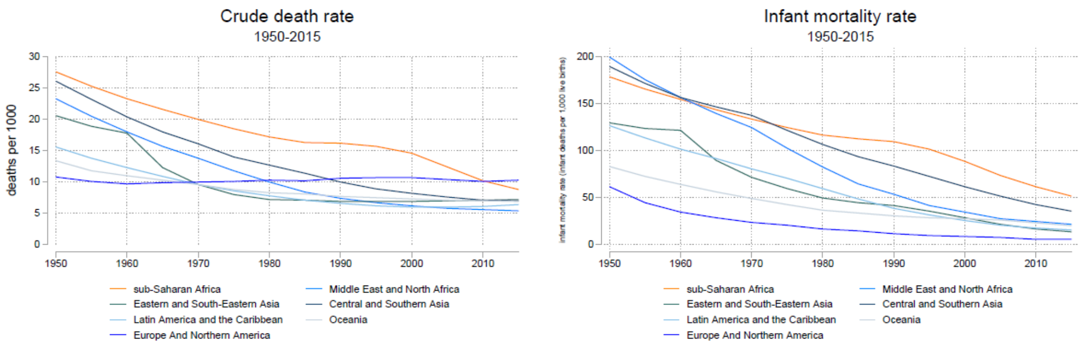

Figure A.2: Total Wanted Fertility Rate by country x year x wealth quintile.
179 DHS samples.

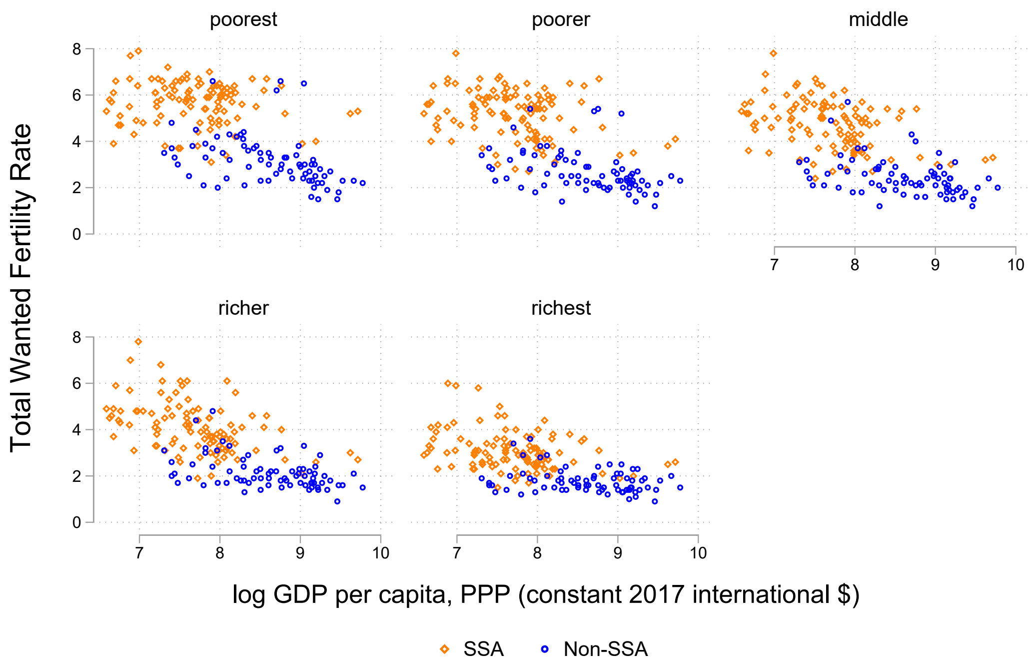

Figure A.3: Share of working women in salaried jobs by country x year x wealth quintile. 179 DHS samples.

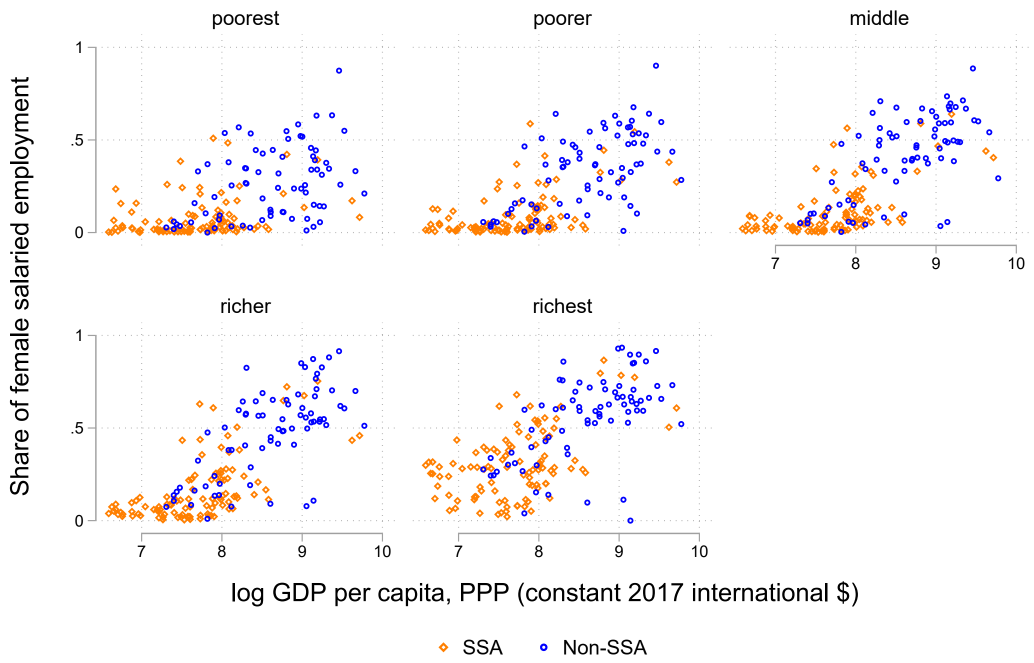

Figure A.4: Correlating the ideal number of children variable with the TFR

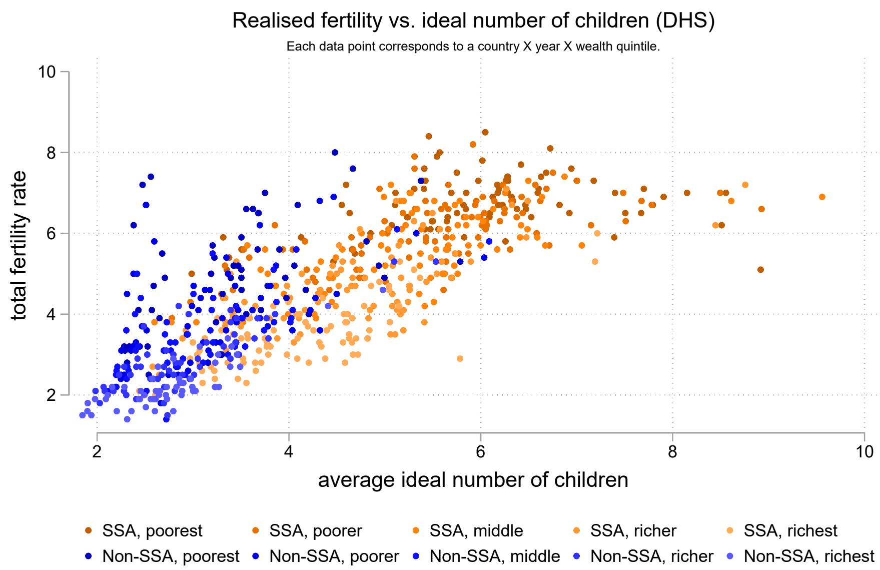

179 Demographic and Health Surveys. This graph plots the total fertility rate against the average ideal number of children reported by respondents in each country x year x wealth quintile. It is important to note that the fertility rate is flat relative to ideal number of children at very high levels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

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

	Total Fertility Rate (TFR)		
	(1)	(2)	(3)
log GDP per capita	-1.243***	-0.780***	-0.758***
	[0.124]	[0.134]	[0.115]
sub-Saharan Africa		1.401***	1.051***
		[0.233]	[0.211]
poor-rich Δ TWFR			0.360***
			[0.093]
R^2	0.506	0.674	0.731
Outcome Mean	4.278	4.278	4.278
Observations	233	233	233

* p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered at the country level in brackets. Poor-rich Δ TWFR = gap in Total Wanted Fertility Rate across top and bottom wealth quintiles.

Figure A.5: Socioeconomic gradients in harmonized international wealth index

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

Standard errors clustered by country-year in brackets. * p<0.1, ** p<0.05, *** p<0.01.

Figure A.6: Gradients in desired fertility residualised on exposure to family planning

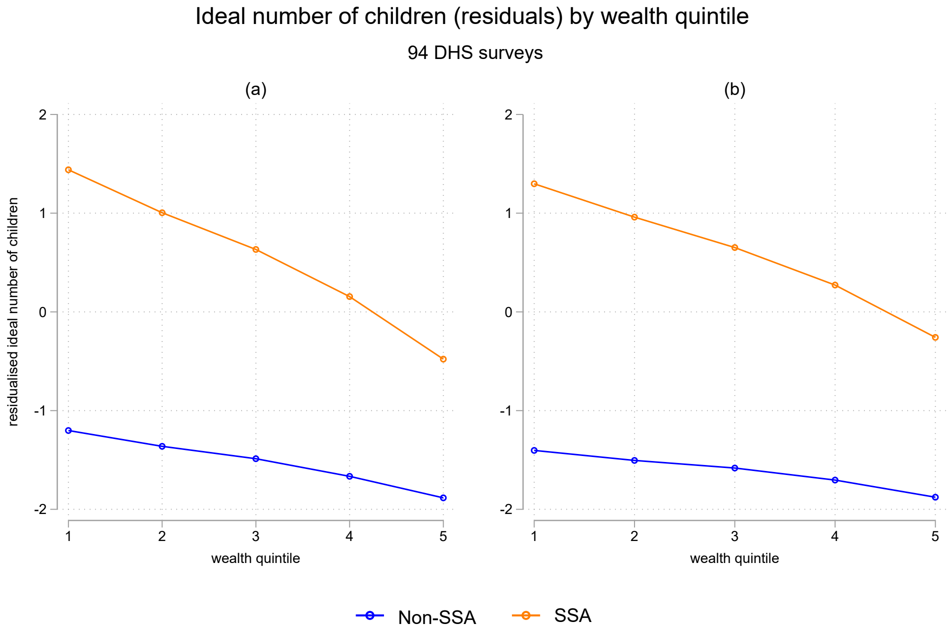

(a) plots ideal number of children residualised on log GDP per capita. Graph (b) plots the same outcome residualised on log GDP per capita **and** share of women exposed to family planning campaigns (radio). This analysis uses only the 94 DHS surveys for which information was collected on exposure to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

Table A.3: Differences in women's education and type of employment across sub-Saharan African and other countries

	Years of schooling	Share of salaried workers	Share of salaried workers	Share of agricultural workers	Share of agricultural workers
	(1)	(2)	(3)	(4)	(5)
sub-Saharan Africa	-1.106	-0.096*	-0.082*	0.035	0.030
	[0.699]	[0.048]	[0.047]	[0.052]	[0.055]
log GDP per capita	2.112***	0.181***	0.154***	-0.194***	-0.185***
	[0.413]	[0.031]	[0.037]	[0.033]	[0.032]
years of schooling			0.013		-0.004
			[0.010]		[0.007]
R^2	0.464	0.560	0.574	0.495	0.497
Outcome Mean	5.963	0.272	0.272	0.349	0.349
Observations	179	179	179	179	179

Years of schooling: average years of schooling across all female respondents aged 15-49 in DHS survey. Share of salaried employment: share of working women in sample who are working for someone else for pay. Share of agricultural employment: share of working women in sample engaging in agricultural work. All regressions are weighted by the inverse of the number of DHS waves available per country. Standard errors clutered by country in brackets. * p<0.1, ** p<0.05, *** p<0.01. 179 DHS surveys.

Figure A.7: Distribution of female years of education

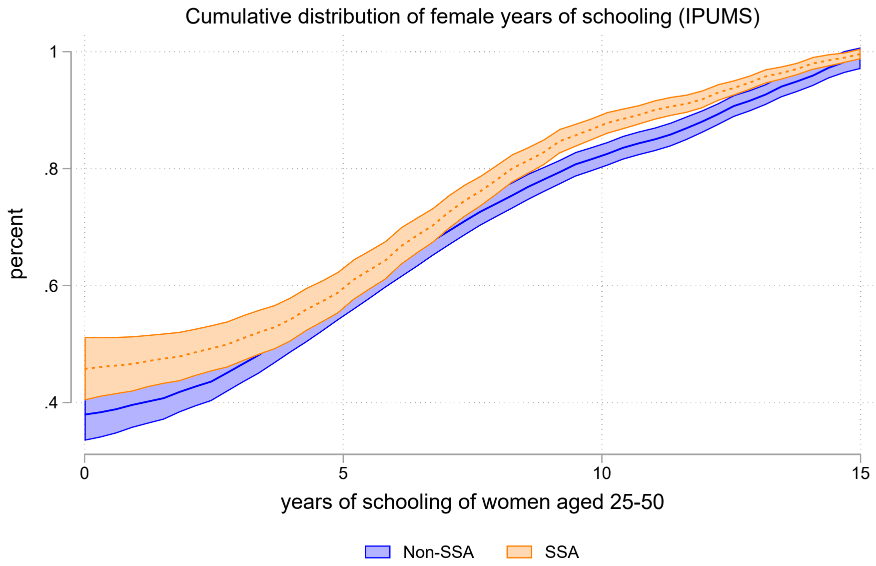

64 IPUMS samples. This graph plots the cumulative distributions of female years of schooling for sub-Saharan and non-sub-Saharan African years of schooling separately. These include 95% confidence bands. As is clearly apparent here, one cannot reject the null that these distributions are the same across 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.8: Salaried employment rates conditional on schooling by country

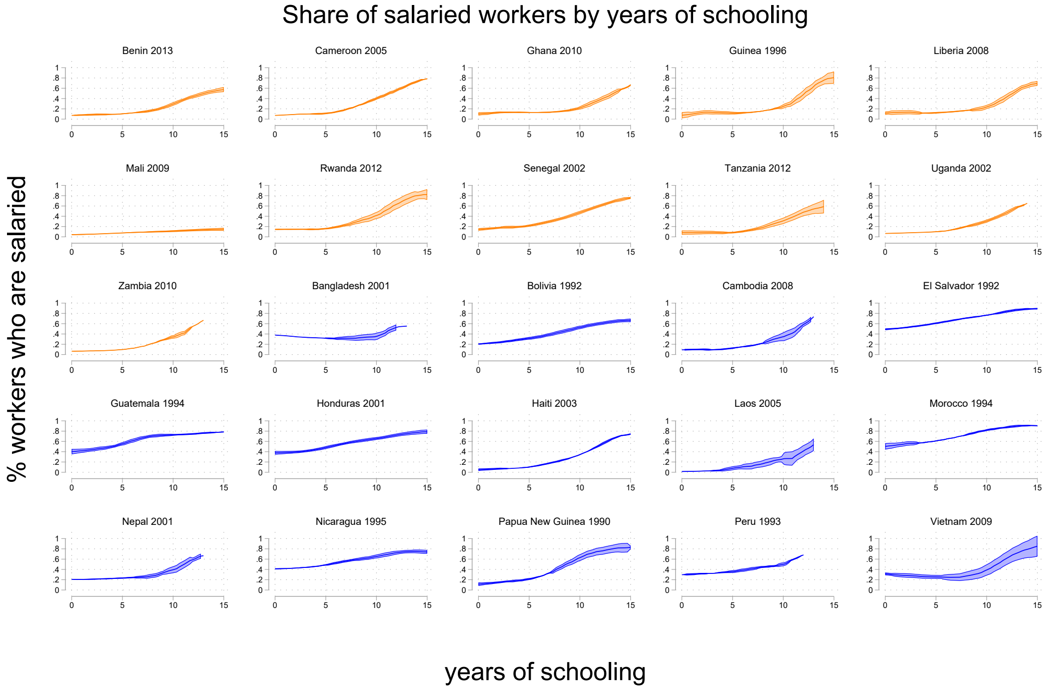

25 IPUMS samples. Local polynomial regression estimates with 95% confidence bands. The orange curves plot the results for sub-Saharan African samples. On the Y-axis is the raw measure of salaried employment rates. Estimates for the other samples are similar so I exclude them here for visualisation purposes.

Figure A.9: Schooling costs per child (Inv. Arcsine)

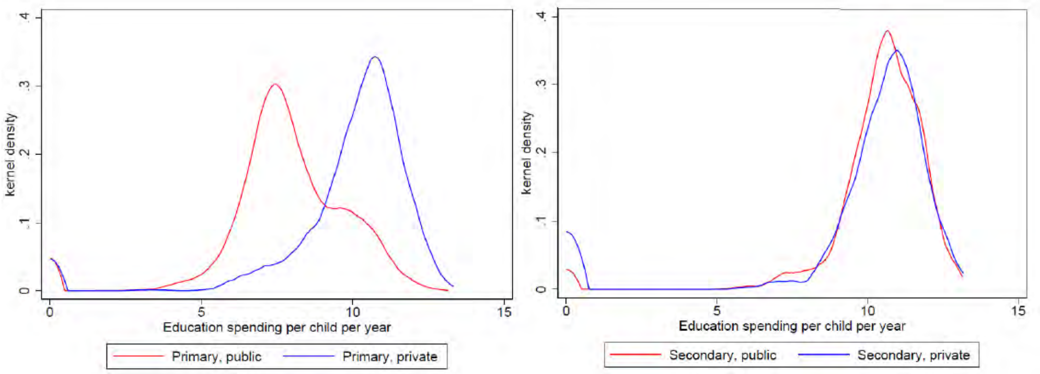

(a) Primary school **(b)** Secondary school

Kenya Integrated Household Budget Survey 2005-2006.

Figure A.10: Relationship between household head education and (a) fertility (b) educational expenditure per child in Kenya 2005-2006

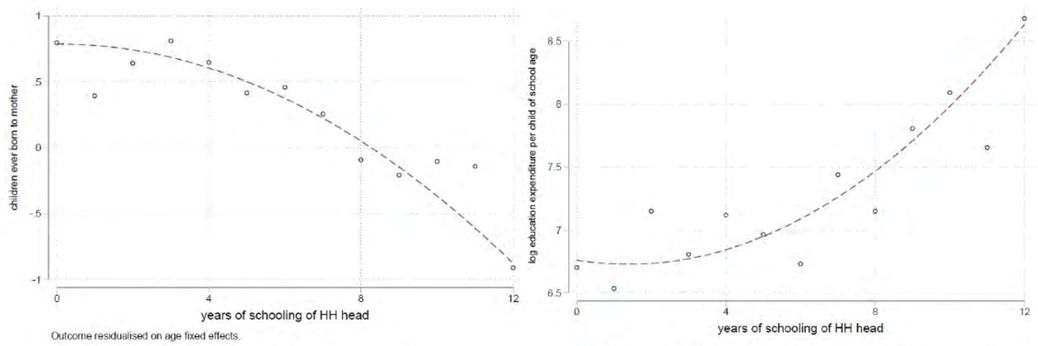

(a) Children ever born to mother **(b)** Log education expenditure per child

Kenya Integrated Household Budget Survey 2005-2006.

Table A.4: Estimates of the relationship between household head education and (a) fertility (b) educational expenditure per child in Kenya 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. Kenya Integrated Household Budget Survey 2005-2006.

Figure A.11: Gradient in desired fertility and salaried employment rate for 25 sub-Saharan African countries

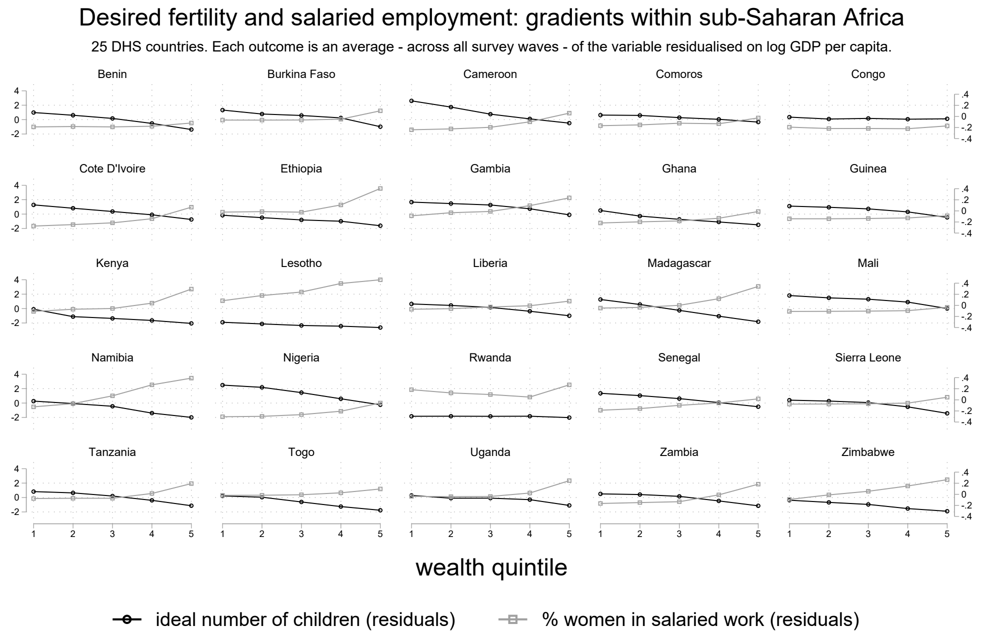

Figure A.12: Gender gaps in desired fertility

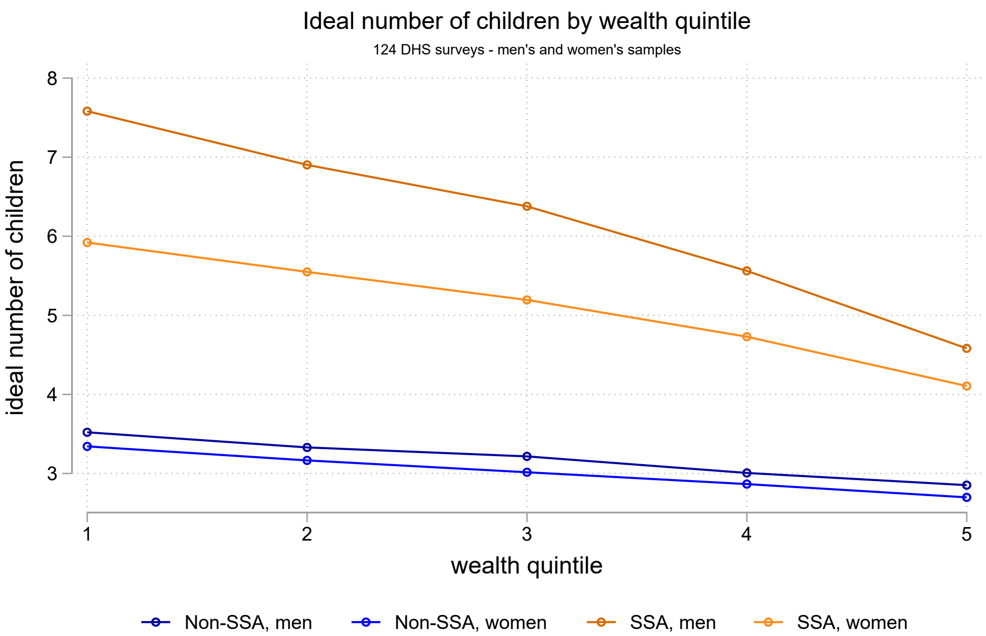

Figure A.13: Women's relative earnings and desired fertility

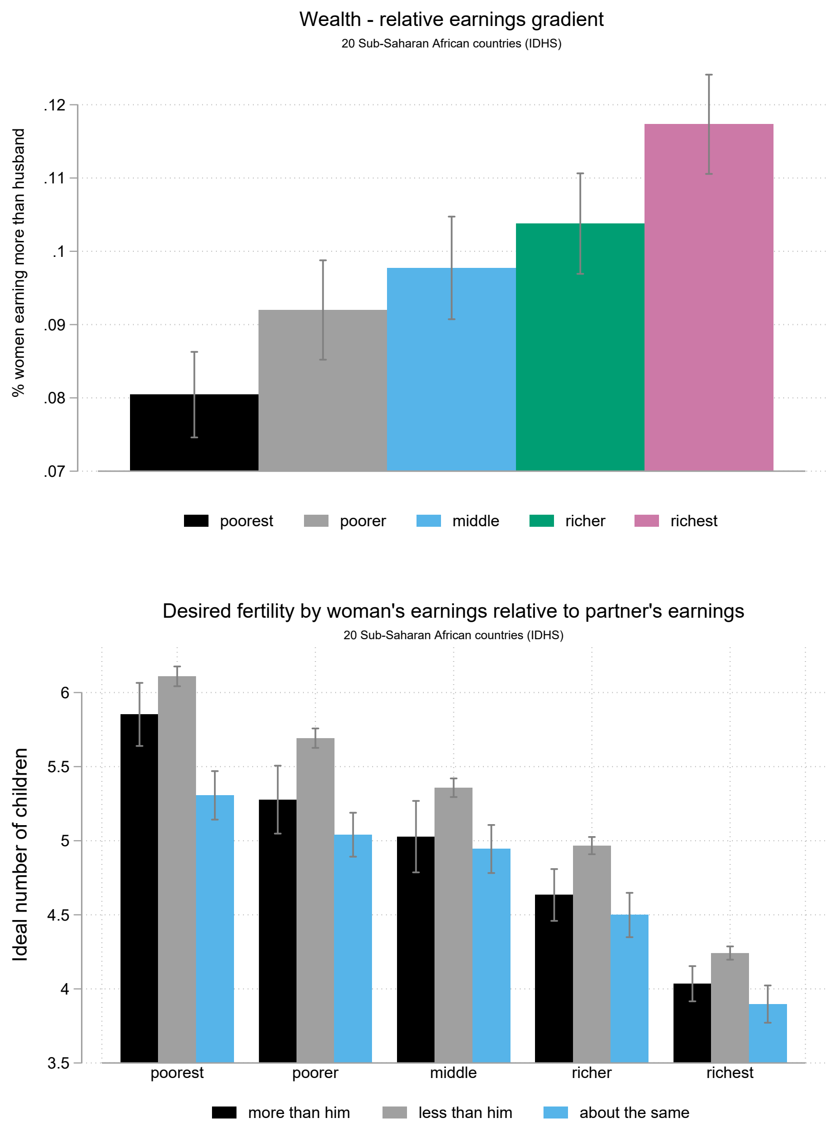

Figure A.14: Polygyny in sub-Saharan Africa (IPUMS-DHS)

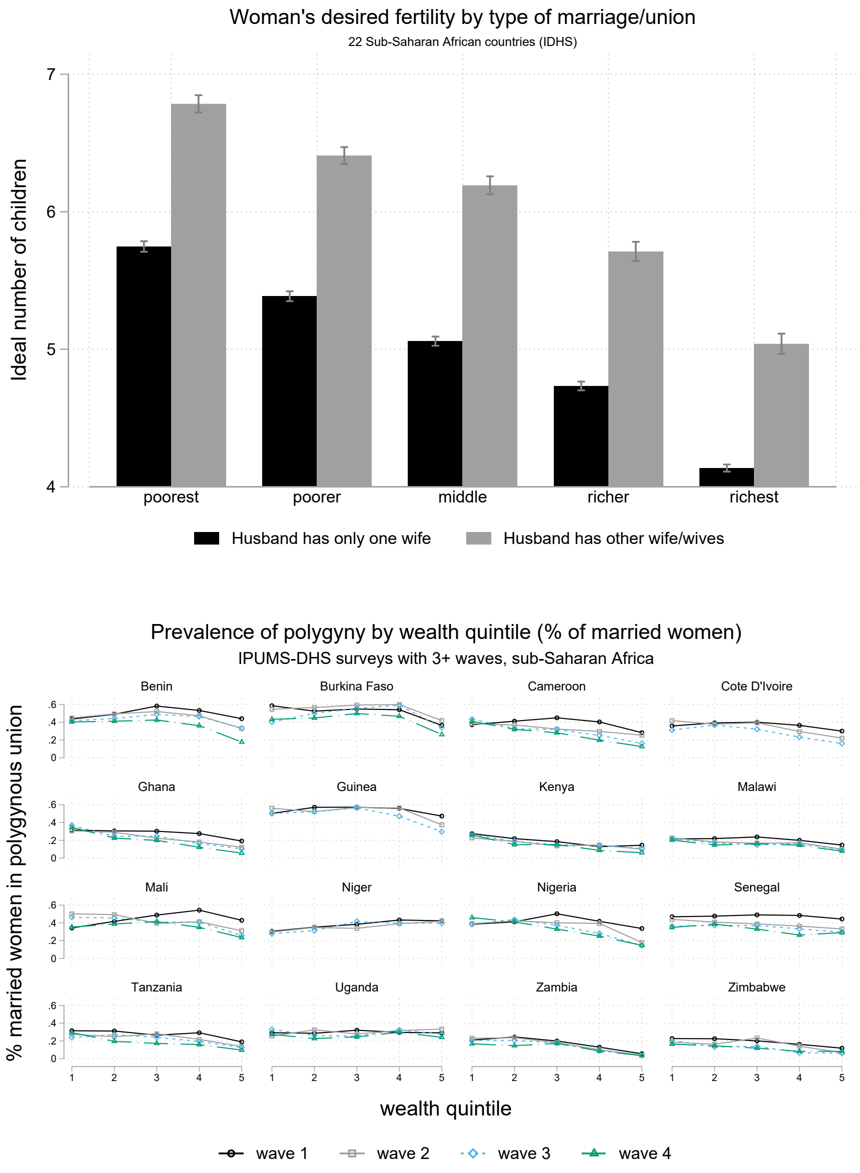

Table A.5: Husbands' and wives' desired fertility in sub-Saharan Africa (DHS)

	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.

Table A.6: Correlations between childless women's ideal number of children and the availability of off-farm and salaried employment in SSA (DHS).

	Ideal number of children			
	(1)	(2)	(3)	(4)
% wage work in province	-0.847*	-1.184**		
	[0.504]	[0.568]		
works for a wage		-0.204***		
		[0.050]		
works for a wage * % wage work in prov.		1.454***		
		[0.489]		
% non-farm work in province			-0.698***	-0.724***
			[0.185]	[0.188]
non-farm occupation				-0.019
				[0.054]
non-farm occupation * % non-farm work in prov.				0.066
				[0.114]
Country x ethnic group FE	Yes	Yes	Yes	Yes
Country x religion FE	Yes	Yes	Yes	Yes
Country x year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
R^2	0.311	0.311	0.311	0.311
Outcome Mean	4.327	4.327	4.327	4.327
Observations	158431	158431	158431	158431

All specifications control for the following individual-level characteristics: rural/urban residence, years of education and knowledge of modern contraceptive methods. Regressions are weighted by sampling weights normalized by population size and number of DHS waves per country. Standard errors clustered by province-year in brackets. *p<0.1, **p<0.05, ***p<0.01.

Chapter 2

The Impacts of Free Primary Education on Female Employment in four sub-Saharan African countries

2.1 Introduction

Access to education is considered both a basic human right and a critical factor in the eradication of poverty. The United Nations' fourth Sustainable Development Goal is to "Ensure inclusive and equitable quality education and promote lifelong learning opportunities for all" by 2030. Indeed, in 2019, over 58 million children - and 31.5 million girls - of primary school age were still out of school (UIS (2021)). This paper studies the impacts of expanding access to primary education in sub-Saharan Africa, focusing on female employment. A large body of evidence has shown that female education and labor market participation increase women's bargaining power and children's health and delay fertility in developing countries (Heath and Jayachandran (2018)). However, we still know relatively little about the direct labor market impacts of expanding access to education on women's workforce participation and occupations in sub-Saharan Africa.

I evaluate the effects of primary school fee elimination - a reform known as Free Primary Education (FPE) - in Ghana, Malawi, Tanzania and Zambia on women's education levels and occupations. After decolonization, parents were required to pay tuition fees to enrol their children in primary school in most of sub-Saharan Africa. These have since been abolished in many countries, in an effort towards universal primary education. Since the occupational structure in most sub-Saharan African countries is still largely dominated by subsistence agriculture, whether increasing access to primary education can raise productivity in agriculture and/or allow young people to shift into other higher-productivity sectors should be of primary interest to policymakers. Due to their comparative advantage in "brain-based" work, women may also be in the best position to reap the benefits from these reforms and, in doing so, accelerate the sectoral shift away from "brawn-based" industries (Ngai and Petrongolo (2017)).

This paper uses micro-census data harmonized by IPUMS-International to evaluate the impacts of FPE on women's education, employment and fertility. For identification, I exploit variation in year and location of birth to construct a differences-in-differences approach à la Duflo (2001) and Larreguy and Marshall (2017). Comparing the outcomes of eligible cohorts with those of individuals who were too old to attend primary school in the year that the policy was implemented provides variation pre- and post-reform. I combine this difference with variation in pre-reform primary school completion rates across locations of birth, thus providing geographic variation in the intensity of the reform.

I find that the elimination of primary school fees led to an increase in women's years of education completed and literacy rates in all four countries in this sample. These effects are associated with a rise in labor force participation and employment. The education effects are most modest in Malawi, where the reform even led to a small drop in the probability of completing primary school for the group of women exposed to the reform.

Analysing effects on women's *type* of employment reveals substantial heterogeneity across countries. The reform is associated with an increase in the share of women working outside of agriculture in all countries except Tanzania, where it is unaffected. In Ghana, Tanzania and Uganda, the employment boost is driven entirely by women increasing their participation in self-employed activities, while in Malawi, it is driven by an increase in the share of salaried workers, although this effect is of small magnitude. Turning to women's sector of occupation, I find that the employment boost associated with the reform is concentrated in the retail, manufacturing and food/accommodation sectors in Ghana, and in the retail and education sectors in Tanzania and Malawi. (No data on sector of occupation was available for Uganda.)

In the final section of the paper, I provide suggestive evidence that the differences observed in the size and direction of these effects across the four countries in this sample may be partly attributable to differences in schooling quality. Indeed, because of the potential deterioration of learning outcomes that may result from large enrolment surges, the direction of the impacts of FPE on the distribution of skill levels and occupations in these countries is ambiguous a priori. School fee elimination, by creating a significant strain on already limited resources, could improve equality of opportunity at the expense of adverse quality effects on educational standards.[1] To investigate this question, I examine two proxy measures for schooling productivity. The first is the mean primary pupil-teacher ratio (provided by the UNESCO Institute for Statistics) of each country, for which data is available before and after the reform. Second, I examine the relationship between years of schooling and a literacy assessment measure collected as part of the Demographic and Health Surveys (DHS).

The findings from this analysis, albeit of a correlational nature, suggest that the large inflow of students in Malawi and Uganda may have lowered learning productivity at the primary level. In particular, the data support the interpretation that the reform did not boost skill accumulation in Malawi as much as in the other three countries. Tanzania and Ghana, on the other hand, may have managed to accommodate the enrolment surge induced by FPE better. These countries also experienced larger increases in educational attainment and employment. In Ghana, the country for which I find the largest employment effects, the pupil-teacher ratio was substantially lower (at around 30) pre-reform than in the other countries (all in the 40-70 range), and literacy rates conditional on years of education appear to have improved after the reform.

This study contributes to an emerging literature on the impact of tuition fee elimination on female labor market participation in developing countries. Since most FPE reforms were implemented in the 1990s and 2000s in sub-Saharan Africa, researchers have only recently started to study their labor market impacts. In this respect, the findings in this study resonate most with those of Chicoine (2020), who finds that schooling increased, fertility declined, and women's labor market participation expanded as a result of the removal of primary school fees in Ethiopia. On the relationship between female education and fertility, Heath and Jayachandran (2018) provide a helpful review. They show that most well-identified studies on this question detect 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)). Interestingly, 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)). In the micro-development literature, many papers have studied the returns to education in developing countries (see Patrinos and Psacharopoulos (2020) for a review). A seminal contribution in terms of causal evidence is Duflo (2001), who finds evidence of positive returns to men's schooling

[1] Afrobarometer surveys of individuals administered a few years after the reforms took place show that over half of the respondents would prefer to pay user fees to raise educational standards in these four countries (see Appendix Figure A.1).

in Indonesia. This study adds to this literature by focusing on women's occupations and relating the heterogeneity in program impact to differences in schooling productivity across countries.

The remainder of this paper is structured as follows: section 2 describes the data. Section 3 presents the identification strategy. Section 4 discusses the results. Section 5 concludes.

2.2 Data

In all countries examined in this paper, the objective of the FPE policy was simple: abolish all fees that might present a barrier to primary education for children from poor households. According to a comparative study on the implementation of the reform in several countries, prior to FPE, the fees were used to fund all school-level inputs except for teacher salaries (World Bank (2009)).

2.2.1 Sample Countries

Of the 24 sub-Saharan African countries for which IPUMS micro-census data is publicly available, 20 countries have officially documented FPE reforms. IPUMS datasets containing information on education, employment status, occupation, fertility and location of birth were available for 14 of these countries. These are listed in Table 1. I analyze the effect of the reform in the four countries for which (i) data was collected at least 10 years after the reform, to allow identification of effects on both education and labor market outcomes, (ii) the variables allowed me to construct the same differences-in-differences approach exploiting variation in exposure to the reform across birth locations and cohorts and (iii) variables relating to educational attainment, employment and fertility were available. In countries where more than one primary school fee abolition policy was implemented since independence, I analyze the latest reform. For each of the four countries, I restrict my sample to all women who were in age of having completed primary school at the time of the data collection, so aged 15 and above.

2.2.2 IPUMS data

Education. In all four samples, I use as my main education outcomes two variables harmonized by IPUMS - years of schooling completed and literacy status - and a binary indicator for whether the respondent completed primary school, which I construct from the harmonized "educational attainment" variable.

Table 2.1: Latest sub-Saharan African censuses harmonized by IPUMS

Country	Year of latest census (IPUMS)	Year of FPE	Earliest birth cohort exposed	Notes
Benin	2013	2006		FPE reform too recent
Botswana	2011	1980		No literacy variable
Burkina Faso	2006	2007		FPE reform too recent
Cameroon	2005	2000		FPE reform too recent
Ethiopia	2007	1994		No information on sector of employment, location of birth or location of previous residence
Ghana	**2010**	**1996**	1990	Birth location variable: region (admin 1)
Guinea	2014			
Kenya	2009	2003		FPE reform too recent
Lesotho	2006	2000-2006		FPE reform too recent
Liberia	2008	2007		FPE reform too recent
Malawi	**2008**	**1994**	1988	
Mali	2009			
Mozambique	2007	2004		FPE reform too recent
Rwanda	2012	2003		FPE reform too recent
Senegal	2013			
Sierra Leone	2004	2004		FPE reform too recent
South Africa	2011			
South Sudan	2008	2011		FPE reform too recent
Sudan	2008	2011		FPE reform too recent
Tanzania	**2012**	**2001**	1994	Birth location variable: region (admin 1)
Togo	2010	2008		FPE reform too recent
Uganda	**2014**	**1997**	1991	Use previous district of residence as proxy for district of birth; no sector of employment
Zambia	2010	2002		FPE reform too recent
Zimbabwe	2012	1980		No location of birth or literacy variable

Employment. The IPUMS censuses include an "employment status" variable which classifies individuals as employed, unemployed, or inactive. To measure labor force participation, I use this classification to construct a binary outcome variable indicating whether a female is recorded as economically active. Using the same variable, I also construct a binary indicator for whether the respondent is employed. The "class of worker" variable classifies individuals into three types of employment: self-employed, wage/salary worker, or unpaid worker. I use this to construct binary indicators for self- and salaried employment. These are equal to 0 for the other categories of employment type, as well as for the unemployed and economically inactive individuals.

Another key outcome measures whether the respondent is working outside of agriculture, which I construct using the "primary occupation" variable. This variable, harmonized according to the International Standard Classification of Occupations (ISCO), is available for Malawi, Tanzania and Ghana. For Uganda, I use the source variable provided by the country's statistical institute, which cannot be harmonized according to the ISCO classification but does allow me to identify the agricultural professions. For the sector of occupation, I construct binary variables for each of the 13 categories in the "INDGEN" variable provided by IPUMS which harmonizes the source variables into groupings that roughly conform to the International Standard Industrial Classification (ISIC). Unfortunately, no information on the sector of occupation is available for Uganda.

Fertility. Finally, I use the total number of children ever born alive reported by each respondent as my measure of fertility.

Treatment definition. To build my identification strategy, I use the district (for Malawi and Uganda) or region (for Ghana and Tanzania) of birth variable as a proxy for the district or region where the individual went to primary school. Since the Uganda sample does not include a location of birth variable, I impute it as the current district of residence of individuals who have never having moved away from their district (inferred from the "years residing in locality" variable) and, for those who report having migrated, as the district where they lived before their last move. This includes some measurement error, as some people may have moved districts more than once since they were born, but excluding the individuals who have moved districts from the analysis leaves the estimates virtually unchanged for Uganda.

The samples provided by IPUMS are representative at the second level of administration (district). This feature allows me to construct a measure of "intensity" of the reform, defined as the primary school completion rate of the older (unexposed) cohorts born in each district or region. Intuitively, the lower this is, the more the district's population might benefit from school fee elimination. This treatment intensity variable is similar to the one constructed in the evaluation of the Nigerian UPE reform by Larreguy and Marshall (2017).

Birth cohorts are constructed using the age of surveyed individuals and the year of the census wave. Importantly, to smooth out some of the measurement error due to age misreporting, I group cohorts into two-year intervals. Indeed, the distribution of ages recorded in each census suggests that there was significant reporting bias: for example, even ages are much more likely frequently reported than odd ages.

2.3 Identification Strategy

The identification strategy exploits variation in the intensity of the reform across locations of birth and in individuals' exposure to the reform according to their age in the year of its introduction. The average treatment effects of FPE on educational attainment, labor market and fertility outcomes are therefore estimated using the following differences-in-differences specification:

$$Y_{idt} = \alpha_t + \delta_d + \beta\, Treatment_{id} \times Post_{it} + \epsilon_{idt} \tag{2.1}$$

where Y_{idt} is the outcome variable for individual i from birth cohort t born in location (district or region) d, α_t is a cohort of birth fixed effect and δ_d is a location of birth fixed effect. As discussed in the previous section, birth cohorts comprise a two-year interval, to smooth out some of the measurement error introduced by age misreporting.

$Treatment_{id}$ measures individual i's "intensity of exposure$''$ to the reform, constructed using pre-reform primary school completion rates in individual i's location of birth: $Treatment_{id} = 1 - \overline{Primary_d}$ where $\overline{Primary_d}$ is the mean primary school completion rate of the older cohorts who were educated in location d in the decade preceding the reform. The main effects ($Post_{it}$ and $Treatment_{id}$) are not visible in equation (1) because they are controlled for by the cohort of birth and location of birth fixed effects.

Defining cohort exposure to FPE. $Post_{it}$ is a binary variable indicating whether individual i from birth cohort t was eligible to benefit from the FPE reform - i.e. whether they were in age of attending primary school (or attending it in the future, if too young) at the time of its implementation. The individuals who were older than the official age for *finishing* primary school (14 in the three countries examined in this study) in the year that the reform was implemented constitute the "pre-reform" group. Those who were younger - i.e. 14 years old or less when the reform was launched - constitute the "post-reform" group. Officially the age of primary school entry is 6 in Ghana, Malawi and Uganda and 7 in Tanzania. I include the "partially treated" individuals - those aged 7-14 in the year of the reform - in the "exposed" group, because many of them likely benefited from the elimination of fees. First, as a result of delayed entry or grade repetition - both very common in these countries prior to the reform (World Bank (2009)) - many children in this group may in fact have been in the early stages of their primary school education when FPE was launched. Second, other individuals might have simply extended the time spent in school thanks to the reform, if tuition fees previously led some financially-constrained parents to pull their children out of school before they could reach the end of the cycle.

The main identifying assumption that has to hold in this framework to ensure that the differences-in-differences strategy will yield estimates of the causal effects of the reforms is that of 'parallel trends'. That is, it must be plausible that the difference in education levels between high- and low-intensity areas would have remained the same after the year of the FPE reform if it had not been implemented. If this assumption holds, any change in the relative trends in education levels between the 'treatment' and 'control' locations after FPE can be credibly attributed to the causal impact of the reform.

It is therefore necessary to examine the trends of the control and treatment groups before the launch of FPE to be able to draw accurate conclusions about any changes in the relative trends in education levels of treatment and control locations post-reform. To test the validity of this identification assumption, I therefore estimate:

$$Y_{idt} = \alpha_t + \delta_d + \sum_t \beta_t Treatment_{id} \times Birthcohort_{it} + \epsilon_{idt} \tag{2.2}$$

where $Birthcohort_{it}$ is a dummy equal to 1 if individual i was born in cohort t. Evidence that the parallel slopes assumption holds will be reflected in β_t estimates that are insignificantly different from zero for all cohorts preceding the oldest cohort of individuals exposed to the reform (i.e. the 'first' of the 'post-reform' cohorts). The results of this test are discussed in section 4.1.

Another possible threat to this identification strategy would be the simultaneity of the FPE reform with other policies that could have affected the treatment and control areas differentially and around the same time. However, to the best of my knowledge, the introduction of FPE did not coincide with other policy changes in any of these countries.

Finally, it is important to note that I restrict my analysis to a reduced-form approach. First, this study seeks to estimate the average effects of the FPE reforms on female education and employment, whereas using these policies as an instrument for education would

Figure 2.1: Treatment x birth cohort estimates from specification (2) where Y = years of education completed

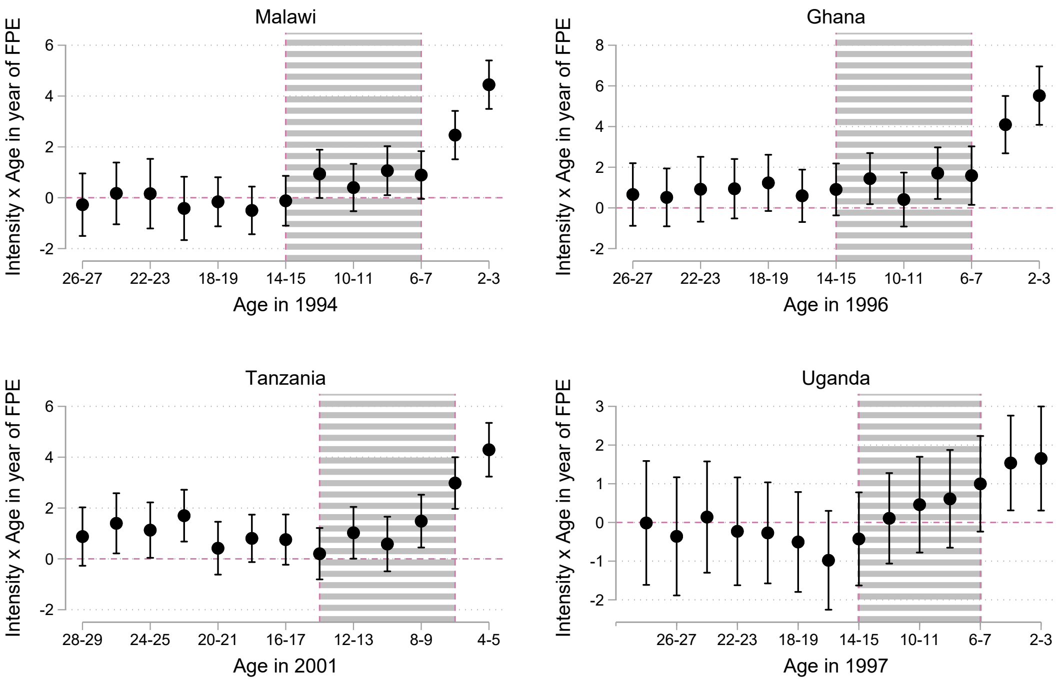

$Treatment_{id} \times 1[Birthcohort_{it}]$ coefficient estimates from specification (2) along with their 95% confidence intervals. The omitted category is the interaction of $Treatment_{id}$ with the oldest cohort. Standard errors are clustered at the location of birth X birth cohort level. Location of birth variable defined as (i) region of birth for Ghana and Tanzania; (ii) district of birth for Malawi and (iii) district of residence before last migration for individuals who previously resided in a different district and current district for those who never moved districts.

imply estimating the causal effect of education on occupational choice, which is beyond the scope of this paper. It is also arguable that the reform would not make a plausibly excludable instrument for education. For instance, another possible channel of impact of the reform on children's occupations than education is an income effect. Indeed, families who would have sent their children to primary school even in absence of the reform could have reallocated these savings towards other productive investments.

2.4 Results

2.4.1 Testing for pre-reform parallel trends

Figure 1 plots the estimates of the β_t coefficients from specification (2), where the outcome variable is years of schooling. Cohorts to the left of the shaded area are classified as "pre-reform" as they were too old to be eligible to benefit from FPE at the time of its implementation. Cohorts within the shaded area and to the right of the shaded area are classified as "post-reform". Cohorts within the shaded area constitute the "partially treated" group and those to the right of the shaded area make up the "fully treated" group, since they were eligible to benefit from the reform throughout the entire length of their primary school cycle.

Overall, these estimates are fairly consistent with the identifying assumption that female years of schooling were following parallel trends across locations with different primary school completion rates prior to the launch of the reform. In Malawi and Ghana (the two upper panels), all estimates of the coefficients on the interaction between the $Treatment_{id}$ and birth cohort variables lie in a 95% confidence interval around zero. We see some evidence of differential time effects for the partially treated cohorts, and the estimates increase sharply for the ″fully treated cohorts″. This is in line with the assumption that the reform led the higher-intensity locations to catch up to their lower-intensity counterparts in terms of educational attainment.

The Tanzania results paint a similar picture, except that some of the estimates for cohorts aged 22-29 in 2001 also lie significantly above zero, suggesting that the higher-intensity locations were starting to show signs of catching up to the lower-intensity regions in terms of female educational attainment about a decade ahead of the reform. However, the estimates for cohorts aged 14-21 in 2001 are all insignificantly different from zero, and from there onward we see the difference-in-difference estimate increase with the length of exposure to the reform.

Finally, the Uganda estimates are insignificantly different from zero for all pre-reform cohorts *and* for the partially treated cohorts, and the β_t coefficient estimates are statistically significantly positive only for the fully treated cohorts. While the estimates for pre-reform cohorts are all insignificantly different from zero, they point to a widening gap in female years of schooling between higher- and lower-intensity districts in the years leading up to the reform. This could raise concerns that the differential trend break that follows could simply be the result of mean reversion (namely, that the higher-intensity districts would have caught up to the lower-intensity districts even without the reform). Nonetheless, the upward jump that can be observed for the fully treated cohorts should provide some reassurance that this alternative interpretation is less likely than that of the parallel slopes assumption.

2.4.2 Average treatment effects

Education effects. The main effects of the FPE reform on female educational attainment are reported in Table 2. The estimates show that years of schooling and literacy rates

Table 2.2: Education effects

	Malawi	Uganda	Ghana	Tanzania
Years of schooling				
Treatment x Post	0.217* [0.114]	0.752*** [0.139]	0.976*** [0.320]	0.722*** [0.247]
Sample Mean, Post=0	4.181	4.756	5.675	5.644
Completed primary				
Treatment x Post	-0.030*** [0.008]	0.105*** [0.020]	0.091*** [0.027]	0.220*** [0.026]
Sample Mean, Post=0	0.239	0.414	0.521	0.702
Literate				
Treatment x Post	0.068*** [0.022]	0.145*** [0.022]	0.180*** [0.039]	0.256*** [0.026]
Sample Mean, Post=0	0.612	0.531	0.627	0.753
Cohort of birth FE	Yes	Yes	Yes	Yes
Location of birth FE	Yes	Yes	Yes	Yes
Observations	258,626	617,551	541,524	951,568

Standard errors are clustered at the location of birth x cohort of birth level.
*** p<0.01; ** p<0.05; * p<0.10.

increased significantly with the FPE reform. The estimated effects on years of schooling range from a 0.22 year increase in Malawi to a 0.98 increase in Ghana. These correspond to 5% and 17% of the mean of unexposed cohorts respectively. Literacy effects range from a 7 percentage point increase in Malawi to a 26 percentage point increase in Tanzania. The final outcome analysed here is the probability of finishing primary school, for which I find positive effects in Uganda, Ghana and Tanzania, but a 3 percentage point *drop* (significant at the 1% level) for Malawi. In the other countries, the estimated effect ranges from a 9 percentage point increase in Ghana to a 22 percentage point increase in Tanzania.

Taken together, these results thus reveal substantial heterogeneity in the magnitude of the educational attainment impacts of the reform, with Malawi recording the most modest while the other three countries exhibit much larger effects. The negative effect on primary school completion detected in the Malawi sample could point to congestion effects having led some individuals - who would have completed primary school in absence of the reform - to drop out before the end of the primary school cycle.

Labor market participation and employment. Since this paper seeks to evaluate the reduced form impacts of the FPE reform on female labor force participation and employment, the next part of the analysis follows the same steps as those used to assess the effects of the policy on educational attainment. First, Figure 2 plots the estimates of the

Figure 2.2: Treatment X birth cohort estimates from specification (2)
where Y = Labor Force Participation

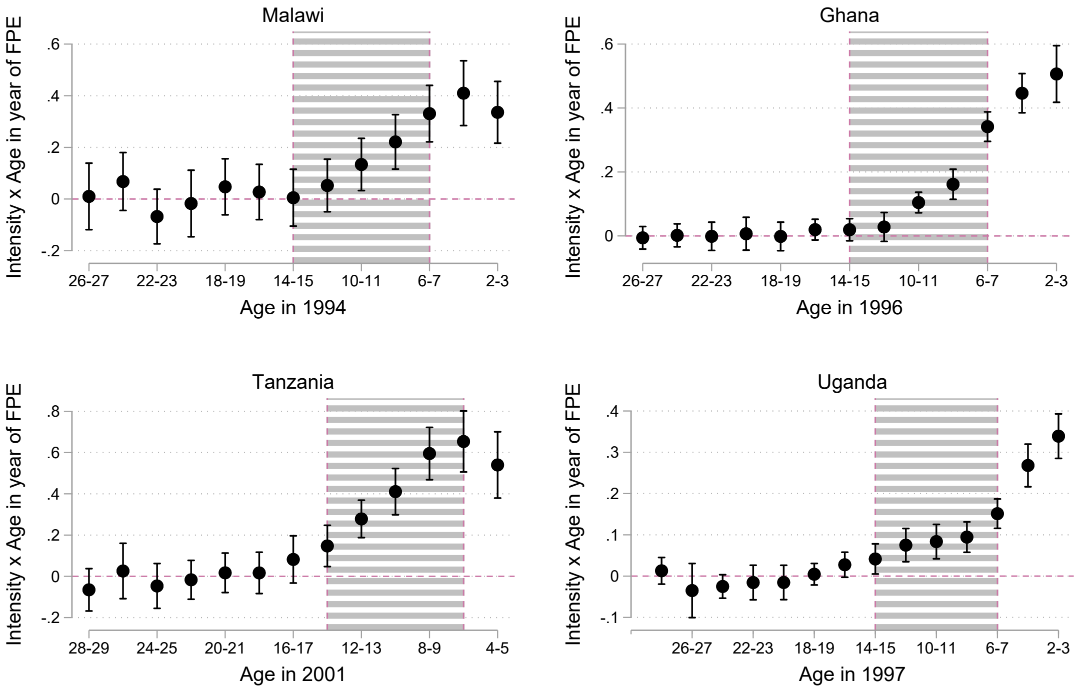

$Treatment_{id} \times 1[Birthcohort_{it}]$ coefficient estimates from specification (2) along with their 95% confidence intervals. The omitted category is the interaction of $Treatment_{id}$ with the oldest cohort. Standard errors are clustered at the location of birth X birth cohort level. Location of birth variable defined as (i) region of birth for Ghana and Tanzania; (ii) district of birth for Malawi and (iii) district of residence before last migration for individuals who previously resided in a different district and current district for those who never moved districts.

β_t coefficients from specification (2), but this time using female labor force participation as the outcome. The Malawi and Ghana estimates support the hypothesis that the reform led to an increase in female labor force participation: in both samples, the coefficient estimates are insignificantly different from zero for all cohorts who were too old to benefit from FPE at the time of its implementation, while women who were aged 10-11 or younger and grew up in higher-intensity locations are more likely to participate in the labor market. This pattern is more pronounced for the youngest cohorts, where the divergence is most evident in Ghana for individuals aged 6-7 and younger in 1996. This first set of results thus suggests that the FPE policies were associated with a boost in female labor force participation in these two countries.

The coefficient estimates for Tanzania and Uganda, displayed in the two lower panels of Figure 2, are noticeably less convincing. While all estimates are significantly positive for the partially and fully treated cohorts in both countries - with a significant jump in their magnitude for the fully treated cohorts observed in Uganda - the fact that the pre-reform estimates are (almost) monotonically increasing over time, starting from before the reform, raises the concern that female employment would have undergone differential trends even in absence of the reform in these two countries. To isolate the effect of the reform from differential trends in labor demand factors, I control for birth location-specific linear trends when running specification (1) on the Tanzania and Uganda samples (region and district trends respectively).

The first five rows of Table 3 display the results of running specification (1) on different employment outcomes. As these are the results of a linear probability model specification, I do not discuss the actual value of the estimates themselves, and instead focus on the relative magnitudes of the effects observed across the four countries.

The top two rows present "extensive margin" effects: the changes in labor market participation (i.e. the probability of being recorded as "economically active" in the census) and employment status (whether the individual is recorded as employed) associated with the FPE policy. These effects are statistically significantly positive for all four countries, with the smallest effects observed in Malawi and the largest in Ghana. The relative magnitude of these estimates across countries is thus consistent with that of the educational attainment effects reported in Table 2.

The third, fourth and fifth row of Table 3 report the changes in the share of women in different *types* of employment associated with FPE: whether they work outside of agriculture, whether they have a salaried job and whether they are self-employed. Each of these three indicators is equal to 1 for women whose primary occupation is outside of agriculture, who work for a wage or for themselves respectively, and 0 for all other women (including those who are unemployed or outside the labor force). This allows me to evaluate the effects of FPE on the overall distribution of adult women's economic activities.

Two key facts emerge from this section of the table. In the third row, we see that the reform was associated with an increase in the share of women reporting working outside of agriculture in Malawi, Uganda and Ghana, but not in Tanzania, where the estimated effect is not statistically significantly different from zero. The fourth and fifth rows show that in Uganda, Ghana and Tanzania, the boost in labor force participation and employment is driven entirely by increases in self-employment, whereas in Malawi it is driven by an increase in the share of women in salaried jobs.[2]

[2]Interestingly, the joint findings on non-agricultural work and self-employment for Tanzania suggest that the FPE-induced employment boost was driven by an increase in self-employment of similar proportion across the agricultural and non-agricultural sector.

Table 2.3: Labor Force Participation, Employment and Fertility

	Malawi	Uganda	Ghana	Tanzania
Economically Active				
Treatment x Post	0.037**	0.042***	0.177***	0.139***
	[0.017]	[0.011]	[0.038]	[0.050]
Sample Mean, Post=0	0.793	0.894	0.884	0.773
Employed				
Treatment x Post	0.041**	0.066***	0.202***	0.172***
	[0.016]	[0.015]	[0.040]	[0.049]
Sample Mean, Post=0	0.611	0.884	0.848	0.756
Non-agricultural occupation				
Treatment x Post	0.026**	0.057***	0.056***	-0.032
	[0.011]	[0.016]	[0.020]	[0.039]
Sample Mean, Post=0	0.594	0.254	0.704	0.494
Wage/salary worker				
Treatment x Post	0.048***	0.010	0.020	0.040
	[0.012]	[0.008]	[0.015]	[0.046]
Sample Mean, Post=0	0.091	0.322	0.105	0.066
Self-employed				
Treatment x Post	-0.009	0.046***	0.218***	0.129**
	[0.009]	[0.011]	[0.038]	[0.050]
Sample Mean, Post=0	0.470	0.506	0.664	0.662
Number of children ever born				
Treatment x Post	-0.604***	-0.877***	-1.131***	-0.469**
	[0.070]	[0.088]	[0.141]	[0.209]
Sample Mean, Post=0	4.725	5.869	3.313	4.095
Year of birth FE	Yes	Yes	Yes	Yes
Location of birth FE	Yes	Yes	Yes	Yes
Location of birth trend	No	Yes	No	Yes
Observations	258,626	617,551	541,524	951,568

Changing the definition of labor force participants to include students (which IPUMS records as "inactive") yields very similar estimates. The Uganda esti--mates are robust to excluding the movers who were over primary school age when they last moved districts. Standard errors are clustered at the location of birth x cohort of birth level. *** p<0.01; ** p<0.05; * p<0.10.

Finally, the last row of Table 3 shows that the reform is associated with a significant decrease in women's fertility, as measured by the number of live children they gave birth to. These estimates range from a drop in about 0.5 births in Malawi and Tanzania to over 1 birth in Ghana.

In sum, the findings laid out in Table 3 suggest that FPE led a significant share of women to reduce their fertility in favour of increased participation in labor market activities. One possible mechanism behind these results is that FPE raised the opportunity cost of women's time in these countries, leading them to substitute away from childbearing.[3] This explanation would imply that there exist positive returns to primary education in the labor market.

That the employment boost associated with FPE is accompanied by a drop in fertility in all four countries is perhaps not surprising, as a number of studies have demonstrated that increasing female schooling leads to reductions in fertility (see Heath and Jayachandran (2018) for a review). One may worry that, in fact, the employment effects documented in Table 3 are only a consequence of this drop in fertility, if having fewer children increases women's ability to participate in the labor market. While this study does not provide evidence for the order of the links in this causal chain, the fact that we find substantial fertility effects across four countries with a variety of employment effects, is somewhat reassuring. Indeed, this drop in fertility is associated with a boost in self-employment - which is arguably more compatible with raising many children than salaried employment - in Uganda, Ghana and Tanzania, whereas it is associated with an increase in the share of salaried workers in Malawi. This combination of findings is, on the whole, consistent with the hypothesis that an increase in schooling enlarges women's opportunity set in the labor market.

In the final step of this reduced-form analysis, I break down the employment results from Table 3 by sector of occupation. As mentioned in section 3, this variable is constructed by IPUMS using the ISIC classification of industries to harmonize sector codes across samples. It is available for Malawi, Ghana and Tanzania, but not Uganda. Figure 3 highlights the source of the occupational shifts observed for individuals most exposed to the FPE reform in each country.

The first panel shows a small shift away from agriculture towards education and retail in Malawi. The largest positive effect is also the most precisely estimated: women young enough to have benefited from FPE are 2.4 percentage points more likely to work in the

[3] What we cannot infer from these results is whether the reduction in the number of children born is driven by women simply delaying their fertility as a result of staying in school longer or by a drop in desired fertility, which would result in lower total fertility.

education sector after the reform. This is consistent with the positive salaried employment effect detected for Malawi: it suggests that much of the employment boost associated with the FPE reform is driven by an increase in the share of teachers. This is perhaps not surprising, as the FPE policy itself boosted demand for additional teachers (World Bank (2009)). Thus, an important implication of this analysis is that FPE was not associated with substantial occupational change in Malawi, aside from an increase in the occupational share of the education sector to accommodate the enrolment increase induced by the reform.

Figure 2.3: DD coefficient estimates from specification (2) where Y is a binary indicator for each sector of occupation

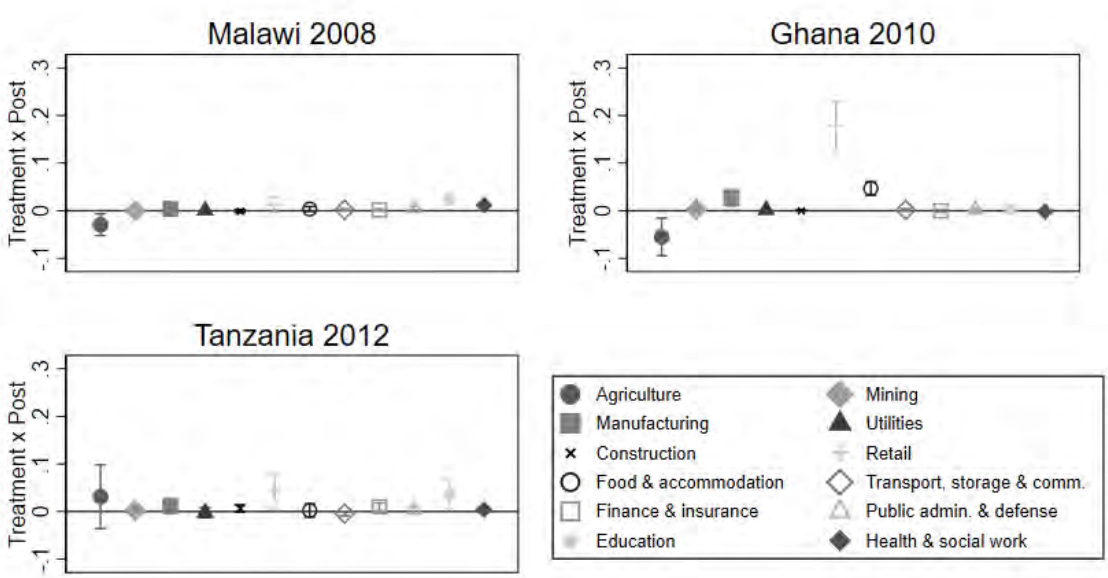

In Ghana, lower participation in agriculture is associated with a substantial shift into retail activities as well increased participation in manufacturing and food and accommodation services. The magnitude of the retail effect is particularly large, at almost 18 percentage points, up from a sample mean of 25% for the unexposed cohorts. These results, when examined jointly with the substantial increase in the share of women engaging in self-employed activities reported in Table 3 and the educational attainment effects of Table 2, form an internally coherent story. They suggest that expanding access to primary education led a substantial share of women to acquire basic literacy and numeracy skills, which may have enabled them to start engaging in self-employed activities in the retail and food service industry.

The bottom left panel plots the coefficient estimates for Tanzania. Here, women most exposed to the reform seem to increase their participation in agriculture - although the standard errors are large for this estimate so it is not statistically significant - and in the education and retail sectors. Here again, the boost to female self-employment reported in Table 3 suggests that the FPE policy may have equipped more women with the basic skills necessary for engaging in self-employed retail activities. The results on agriculture are somewhat less intuitive. However, they are in line with the findings of a recent study of the 1962 Universal Primary Education policy in Tanzania which shows that increased access to primary education induced a statistically significant increase in the probability of working in agriculture (Delesalle (2019)). This suggests that there may be positive returns to basic education in the agricultural sector in Tanzania.

2.4.3 Impact heterogeneity and schooling quality

The analysis presented in section 4.2 uncovers substantial heterogeneity in the magnitude and nature of the impact of FPE on female education and employment outcomes. This section relates these differences to variation in schooling quality indicators over time across the four countries.

Figure 4 plots the evolution of primary pupil-teacher ratios over time at the national level for each country. The first striking difference that emerges from this set of figures is the enormous gap between Malawi and the other countries. Before the reform, the average pupil-teacher ratio at the primary level stood at almost 70:1 in Malawi, compared to just below 40:1 in Uganda and Tanzania and around 30:1 in Ghana. This finding is suggestive of lower education quality in Malawi relative to the other countries prior to the reform. The pupil-teacher ratio is much lower in Ghana, the country for which I find the largest impact of the reform on years of schooling and employment (see Tables 2 and 3 respectively).

Interestingly, there is no significant jump in the pupil-teacher ratio in either Ghana or Malawi around the time of the reform - in fact, in Malawi this indicator was on a downward trend between 1991 and 1996.[4] Indeed, governments invested in the recruitment and deployment of teachers in both countries to sustain the large, FPE-induced enrolment surge - although many of these new teachers had not been properly trained by the time they were deployed (World Bank (2009)).

Figure 4 suggests that Uganda and Tanzania experienced similar patterns in the evolution of pupil-teacher ratios at the primary level. This indicator was on a relatively stable trend, averaging around 40:1, in the years leading up to FPE, at which point it jumped

[4]Note that this does not mean that the reform did not lead to any increases in the pupil-teacher ratio. As documented by Figueiredo Walter (2020), there exists large dispersion in class size and pupil-teacher ratios at the primary level within developing countries, and policy evidence suggests that such disparities were exacerbated by FPE (World Bank (2009)). Since the enrolment surge was largest in high-intensity areas, it is very possible that it is precisely in those regions that the pupil-teacher ratio increased, despite nationwide deployments of additional teachers.

by almost 50%, and then gradually declined from about 4 years after the reform. This points to possible congestion effects immediately after the reform in these countries, which may have dampened over time as the governments increased capacity and allocated more resources to the primary education sector in response to the enrolment surge.

Figure 2.4: Pupil-Teacher Ratio before and after FPE

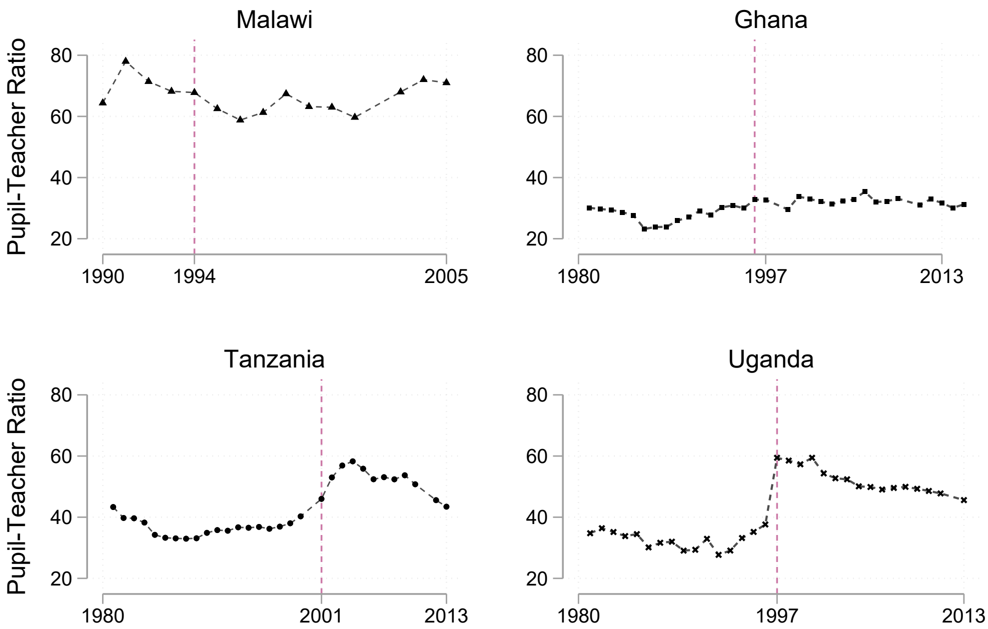

The data used to produce the top left figure comes from the Malawi Ministry of Education. The UNESCO Institute for Statistics provided the data for Ghana, Tanzania and Uganda.

How do the trends observed in Figure 4 translate to learning outcomes? The IPUMS data do not contain information that would allow me to assess the impacts of school fee abolition on skill levels. The literacy measure is self-reported and as such does not provide an objective indicator of literacy skills. Therefore, for this final piece of analysis I make use of the Demographic and Health Surveys (DHS). Recent DHS waves include a literacy variable that was collected by asking respondents with less than secondary school education to read aloud a sentence on a card. I use this variable to construct an indicator equal to 1 for respondents who were classified as “able to read [the] whole sentence” and 0 otherwise.

Figure 5 plots the average of this measure by years of education completed, separately for women exposed and not exposed to the reform. Here, I focus on two groups: women aged 0-7 at the time of the reform, who were therefore eligible to fully benefit from it, and women aged 15-25, who were too old to enrol in primary school when the reform was launched. In Malawi and Uganda, these graphs, albeit only indicative, suggest that the productivity of primary schooling may have decreased after the reform: individuals who

Figure 2.5: Literacy rates conditional on years of education before and after FPE

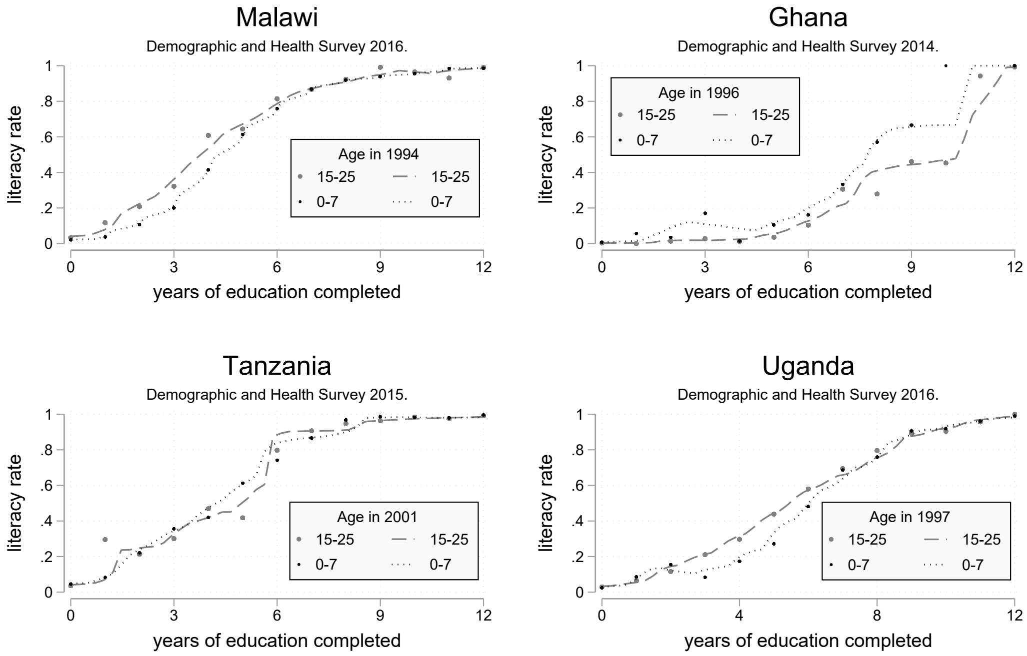

completed 1-5 years of schooling in Malawi and 2-7 years of schooling in Uganda were on average less likely to know how to read a whole sentence after FPE than before the reform. Simple regression estimates (not reported for conciseness) indicate that the literacy rate of the exposed cohorts, conditional on years of schooling, was 4 and 2 percentage point lower than that of the older cohorts in Uganda and Malawi respectively. In Ghana, a gap is observed in the other direction: individuals exposed to the 1996 policy are 11 percentage point more likely, conditional on total years of schooling completed, to be able to read a whole sentence compared to the older cohorts. Finally, in Tanzania, the evidence does not point to a change in the literacy-years of schooling relationship in either direction.

The changes documented in Figure 5, while by no means conclusive due to their correlational nature, provide some insight into the differences in program impact observed across the four countries. First, they point to a decrease in the learning productivity of the primary schooling system induced by the large inflow of students in Malawi and Uganda.[5] Second, they suggest that Tanzania and Ghana may have managed to accommodate the

[5] Another possibility is that the students who were induced to enrol in primary school as a result of the reform were of lower ability than those that would have enrolled even if fees had not been eliminated. This trait could be correlated with their families' lower willingness to pay for education prior to the reform. I cannot address this selection concern but I consider this a much less plausible explanation for the gaps observed in Uganda and Tanzania than the schooling productivity mechanism. The adverse quality effects of the reform have been widely documented (e.g. World Bank (2009), Riddell (2003)) and this channel is not consistent with the fact that Uganda and Malawi had the lowest rates of primary school completion prior to the reform (0.41 and 0.24 respectively) among the four countries in this sample, so we would not expect the set of individuals affected by FPE to be of substantially worse ability than those in Tanzania and Ghana.

enrolment surge induced by the FPE policy better than Malawi and Uganda, since the relationship between years of schooling and literacy appears to have been unaffected in Tanzania (at least from comparing women aged 0-7 to those aged 15-25 in 2001) and to have actually improved in Ghana seven years after the reform. This would be consistent with the conclusion from Table 2 that the FPE reform had relatively large impacts on educational attainment and is also associated with a substantial increase in female employment (Table 3) in these two countries - and in Ghana especially.

Malawi is the country for which section 4.2 documents the most modest education and employment effects. The findings presented in this section lend support to the causal interpretation of the reduced-form effects of the FPE policy on employment in Table 3 and provide a possible explanation to the negative effect on primary school completion reported in Table 2. The latter suggests that a disproportionate share of the students who enrolled in primary school as a result of the reform eventually dropped out before the end of the cycle, possibly because of low learning productivity. On the whole, the results presented in sections 4.1 and 4.2 provide suggestive evidence that the increase in primary school enrolment induced by the reform in Malawi did not boost skill accumulation as much as in the other countries.

2.5 Conclusions

This paper finds that the elimination of primary school fees led to an increase in female schooling and employment in four sub-Saharan African countries. This is also associated with a drop in fertility. On the whole, these results suggest that expanding access to education enables women from the poorest families to engage in economic activities that require basic skills, such as self-employment opportunities in retail and services. This also raises the value of women's time, inciting them to reduce the number of children they give birth to. A key takeaway is thus that fee elimination can help remove some of the barriers to transitioning from unpaid family work (or being out of the workforce) to paid work outside of non-agriculture.

The findings also reveal substantial heterogeneity in the magnitude and direction of the effects of the reform on educational attainment and employment outcomes across the four countries. The country for which the effects are most modest (both in terms of educational attainment and employment outcomes) is Malawi, where the rate of primary school completion was lowest and the pupil-teacher ratio was substantially higher than in the other countries prior to the reform. In other words, the country furthest away from achieving universal primary education prior to FPE also appears to have achieved the smallest gains from school fee elimination. Suggestive evidence points to low levels of learning productivity ex-ante, that might have deteriorated further with the large enrolment surge induced by the reform, as the possible culprit for these modest effects. In contrast, the country found to have experienced the largest impacts of primary school fee abolition, Ghana, was starting from a much higher primary school completion rate ex-ante (52% of women vs. 24% in Malawi) and seems to have minimised adverse quality effects from the enrolment surge, at least as suggested by the evolution of pupil-teacher ratios and the literacy-years of schooling relationship.

The main limitation of this study is that it follows a purely reduced form approach. The channels of impact I put forth in the interpretation of the results thus remain suggestive, and alternative mechanisms cannot be ruled out. First, the reform could also have affected occupational outcomes through an income effect, if its impact on employment is driven by families who would have been willing to pay tuition fees for sending their children to school and were thus able to reallocate the fees towards other productive investments that could influence their children's occupations. Second, assuming the effect of FPE on employment materialised solely through increased schooling, we should keep in mind that education can affect labor market outcomes through other channels than cognitive skill accumulation. These include signaling of ability (Spence (1973)), the non-cognitive effects of schooling (Duckworth and Gross (2014)), or improved marriage prospects (Ashraf et al. (2020)). Delving into the underlying factors responsible for the differences in occupational choice effects associated with FPE across these four countries may constitute a fruitful avenue for future research.

Finally, it is important to note that, even in Ghana, the education boost is not accompanied by any detectable change in the share of salaried workers. Indeed, the substantial increase in female employment associated with the reform in this country is largely concentrated in self-employed retail activities. While it is beyond the scope of this paper to determine whether this type of job creation can generate labor market amplification effects in the longer run, this finding raises a new set of questions. Do there exist salaried jobs that women could aspire to but for which the skill requirements are beyond those acquired at the primary level? Or, does the absence of a dynamic modern sector result in a lack of demand for salaried workers in Ghana? Little is known about whether the type of structural change undergone by many sub-Saharan African countries includes the sectoral shift from "brawn-based" industries to services that has been shown to boost female employment in other settings (Heath and Jayachandran (2018); Ngai and Petrongolo (2017)).[6] The absence of such demand-side factors may be limiting the labor market impacts of the FPE

[6]See Osei and Jedwab (2013) for a case study of Ghana's structural transformation process. They

reform in the four countries studied in this paper. The nature of the relationship between structural change and female education in sub-Saharan African economies remains poorly understood, thus opening doors for future research.

conclude that "structural change remains different in Ghana — it has occurred without a green revolution, an industrial revolution, or a service revolution of the types seen in Asia."

References

Nava Ashraf, Natalie Bau, Nathan Nunn, and Alessandra Voena. Bride price and female education. *Journal of Political Economy*, 128(2):591–641, 2020.

Lucia Breierova and Esther Duflo. The Impact of Education on Fertility and Child Mortality: Do Fathers Really Matter Less Than Mothers? May 2004.

Luke Chicoine. Free Primary Education, Fertility, and Women's Access to the Labor Market: Evidence from Ethiopia. *The World Bank Economic Review*, 35(2):480–498, 01 2020.

Esther Delesalle. The Effect of the Universal Primary Education Program on Labor Market Outcomes: Evidence from Tanzania. Technical Report NBER 25789, 2019.

Angela Duckworth and James J. Gross. Self-Control and Grit: Related but Separable Determinants of Success. *Current Directions in Psychological Science*, 23(5):319–325, 2014.

Esther Duflo. Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *American Economic Review*, 91 (4):795–813, September 2001.

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

Rachel Heath and Seema Jayachandran. 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.*, 2018.

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

Horacio Larreguy and John Marshall. The Effect of Education on Civic and Political Engagement in Nonconsolidated Democracies: Evidence from Nigeria. *The Review of Economics and Statistics*, 99(3):387–401, 07 2017.

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

L. Rachel Ngai and Barbara Petrongolo. Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4):1–44, October 2017.

R. Osei and R. Jedwab. Is Structural Change in Sub-Saharan Africa Different? New Historical Evidence from Ghana. *mimeo*, 2013.

Una Osili and Bridget Long. Does Female Schooling Reduce Fertility? Evidence from Nigeria. *Journal of Development Economics*, 87(1):57–75, 2008.

Harry A. Patrinos and George Psacharopoulos. Chapter 4 - returns to education in developing countries. In Steve Bradley and Colin Green, editors, *The Economics of Education (Second Edition)*, pages 53–64. Academic Press, second edition edition, 2020.

Abby Riddell. The Introduction of Free Primary Education in Sub-Saharan Africa. *Paper commissioned for the EFA Global Monitoring Report 2003-2004, The Leap to Equality*, 2003.

Michael Spence. Job Market Signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 08 1973.

UIS. UNESCO Institute for Statistics: http://data.uis.unesco.org. Accessed February 2021., 2021. URL `http://data.uis.unesco.org`.

World Bank. Abolishing School Fees in Africa : Lessons from Ethiopia, Ghana, Kenya, Malawi, and Mozambique. 2009.

Integrated Public Use Microdata Series, International: Version 7.3. Minnesota Population Center. Minneapolis, MN: IPUMS. 2020. https://doi.org/10.18128/D020.V7.2.

IPUMS Demographic and Health Surveys. Version 8. Elizabeth Heger Boyle, Miriam King, and Matthew Sobek. Minneapolis, MN: IPUMS and ICF. 2020. https://doi.org/10.18128/D080.V8.

Appendix to Chapter 2.

Figure A.1: Share of Afrobarometer 2002-2003 survey respondents willing to pay tuition fees for quality schooling

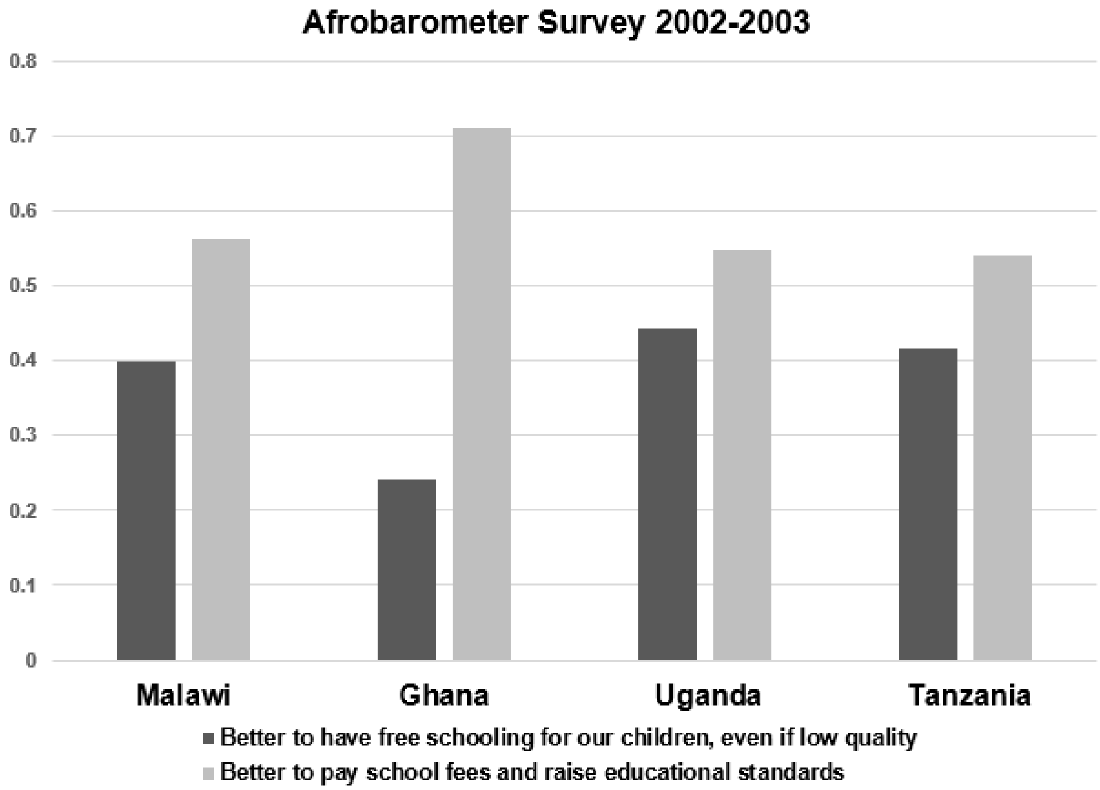

This graph plots the share of survey respondents who agree with A and B respectively in response to the question: *Which of the following statements is closest to your view? A: It is better to have free schooling for our children, even if the quality of education is low. B: It is better to raise educational standards, even if we have to pay school fees.*

Chapter 3

Chat Over Coffee? Technology Diffusion Through Social and Geographic Networks in Rwanda

3.1 Introduction

Do agricultural technologies diffuse through information networks? This question has been the subject of much research in development economics, motivated by the growing gap in agricultural productivity between developed and developing countries – and in particular sub-Saharan Africa – over the past half-century (e.g. World Bank (2007)). Given the critical importance of agriculture to low-income economies, the successful adoption of yield-enhancing technologies has the potential to substantially reduce global poverty. In this paper, we ask whether social learning can help alleviate information frictions that inhibit the diffusion of best agronomic practices.

Our study employs an RCT to study this question, in partnership with a leading NGO in Rwanda that provides agronomy training to coffee farmers to help them improve their yields. Our experiment applies an oversubscription design to a sample of approximately 1,600 farmers who applied to the training program, across all 27 villages in a given administrative unit (a sector).[1] Half of the applicants were randomly selected for the trainings, while the other half received no training and serves as our control group. Selected farmers received monthly instruction modules for the first year, followed by five refresher modules over the next year. These modules covered nutrition, pest and disease management, weed management, mulching, rejuvenation and pruning, shade, soil and water conservation, and

[1] Sector is the third level of administration in Rwanda, so corresponds to a sub-district.

record keeping. Farmers were grouped for trainings and picked a lead farmer whose plot was used for demonstrations of the agricultural practices. We also randomly varied treatment concentration at the village level: in a third of villages, 25% of farmers who signed up were treated, in another third, 50% and in the final third, 75%.

We collected data on a wide range of indicators to measure the impacts of the trainings on the farmers' knowledge and adoption of the agricultural practices, as well as their yields. In total, we collected ten rounds of survey data, in addition to a census. We also collected GPS data on these plots and distributed scales and calendars to farmers so they could weigh and record their harvests. To allay concerns surrounding error in self-reported adoption, we also conducted audits of five trees on each coffee plot. To measure the diffusion of the effects to farmers' social connections, we conducted 3 rounds of social networks surveying, listing the friends and agricultural contacts of household heads and spouses. We use this data to construct complete social network maps of all the coffee-growing households in the sector, amounting to over 5,000 households across 29 villages. This also allows us to track any changes in the map of social ties over the duration of the program. Finally, we use the GPS data to construct two measures of neighbours: people farmers live close to and people who have coffee plots next to their coffee plots.

We harness this data to build a range of outcomes. From the household surveys, we construct an index of knowledge and an index of (self-reported) adoption of improved agricultural practices. We also measure use of chemical fertilizer, compost, labor inputs, and coffee harvests (monthly and seasonal). Here, we complemented the surveys with yield-recording calendars. From the audit data, we construct an index of adoption of improved agricultural practices, which encompasses mulching, weeding, pruning and pest management. We also conduct an index of leaf health, which includes yellowing, curling and rusting of leaves, and an index of all outcomes.

We find that the training program had significant effects on knowledge and self-reported adoption of the agronomic practices: the treatment led to a 0.32 standard deviation increase in the knowledge index and a 0.14 standard deviation increase in the adoption index. The effects we detect in the tree audits data are of more modest magnitude. The "main audits" index is 0.043 standard deviations higher in the treatment group than in the control group, suggesting that the program had some effect on how well and how much the farmers applied this subset of practices (weeding, mulching and smoothing/banding the tree to control white borer). We also find suggestive evidence that the trees of treatment farmers are better nourished, as the audits data reveal a decrease in the prevalence of leaf defects. Importantly, the increase in adoption that we observe in the treatment group had an impact on yields, which are about 3% higher in the treatment group than in the control group. In terms of input use, the program led to a 1.9 percentage point increase in the share of farmers using chemical fertilizer (NPK), corresponding to a 24% increase relative to the control group. The treatment also induced farmers to allocate more labor inputs to their coffee plots: we find a 5.5% increase in the number of household labor days per tree and a 22% increase in the number of paid labor days per tree.

Our analysis of the diffusion of treatment effects through farmers' social and geographic networks focuses on two types of networks: baseline friends and neighbors. The identification strategy exploits the exogenous variation caused by the random allocation of the program in the number of treatment friends (treatment neighbors), conditional on the total number of friends (neighbors) in the sample. We do not find any evidence of spillovers to the control group, as adoption rates of control group farmers are unaffected by their number of treatment friends. Our results also suggest that there is no diffusion of the program's effects through geographic networks (household and plot neighbors).

We estimate some diffusion of effects through the treatment group, observed primarily on leaf health outcomes. Conditional on their total number of friends in the sample, having an additional friend in the treatment group reduces the leaf defects index by 0.036 standard deviations - a sign of better tree nutrition. To investigate the mechanisms driving this effect, we test whether diffusion effects are stronger for friends who attended trainings together. Our results confirm this hypothesis. Conditional on their total number of friends in the sample and the number of farmers in their training group, we find that the effect of having one more pre-existing friend in their training group is associated with a reduction in treatment farmers' leaf defects index by 0.057 standard deviations. This estimate is significant at the 1% level. This suggests that information related to tree health acquired in the trainings disseminated among treatment farmers who attended the trainings together, leading to a reinforcement of treatment effects.

Another explanation for why the diffusion of effects through farmers' baseline networks appears relatively weak on all but the leaf health channel is that friendship networks may have changed as a result of the program. We test this hypothesis in our final piece of analysis, by examining the impact of the training groups on farmers' reported social network links to other farmers with whom they discuss coffee. We find clear evidence that the training program causes a re-sorting of social networks and an increase in links with others who have learned from the trainings. These effects may be partly due to time spent together at trainings, but control group households also increase links to trained individuals in their villages.

The paper is structured as follows. Section 2 describes the setting and the program. Section 3 presents the experimental design and describes the data collection. Section 4 discusses our results. Section 5 concludes.

3.2 Background and Program Description

3.2.1 Context: Rural Rwanda

Coffee is a major export crop for Rwanda and contributes about US$62 million in export earnings per year, almost 5 percent of the country's goods and services exports (NISR (2019)). In the decade leading up to the year of our experiment (2000-2009), Rwanda's productivity was estimated to be 6447 hectograms of green coffee harvest per hectare of harvested area, ranking it as 38th out of 75 coffee-producing countries in the world (FAO 2019).

Production is dominated by 500,000 smallholder coffee producers (OCIR-Café (2008)). In Rwanda - the most densely populated country in sub-Saharan Africa - land parcels are extremely small and most farm households live off subsistence agriculture, and therefore give priority to food crops. Higher coffee yields could make coffee a more attractive crop choice and thus help small-scale farmers shift out of subsistence farming into more profitable activities.

3.2.2 TechnoServe's Agronomy Training Program

The context for our study is a large agronomy training program run for coffee farmers in rural Rwanda, aimed at improving the health of small scale farmers' coffee trees and ultimately their coffee yields. TechnoServe (henceforth TNS), an agri-business NGO, conducted agricultural business and technology trainings in several coffee growing regions in Rwanda and other East African countries between 2010 and 2015. This study focuses on the agronomy program in one sector (an administrative unit) in Southern Rwanda, called Nyarubaka. Rwanda has ideal growing conditions for coffee and it is the main export crop, but the country's agronomic practices and processing are in dire need of improvement. Between February 2010 and October 2011, TechnoServe ran an agronomy training program for coffee farmers in Nyarubaka to improve coffee yields in the area.

The training covered several best practices in coffee growing: tree rejuvenation and pruning; fertilizer use; pest, disease and weed management; mulching; soil and water conservation; optimal shade; and record keeping. Appendix 1 provides a more detailed description of the eight main sets of practices that were covered under this training program and what the expected impacts of each were. The training sessions took place once a month for eleven months in the first year of the program (starting in February 2010). TNS conducted an additional five review training sessions after these sessions were over, with the last review session being held in October 2011. The trainings were conducted with groups of approximately thirty farmers and took place on the plot of a designated "focal farmer". The focal farmers were chosen partly because of the accessibility of their coffee plot but were also meant to be respected members of the local community, had an enthusiasm for learning and were ideally not a committee member of the coffee cooperative.

The training itself was conducted by a TNS Farmer Trainer, each of whom supported approximately 10 of these focal farmer groups. These farmer trainers received monthly training from the TNS agronomist for each agronomy module, together with lesson plans and activities. They delivered the training to each group in the field of the focal farmer, with each module being conducted on a plot of approximately forty trees which then would become the demonstration plot, and with all practical work done by the farmers in the training group. A total of four farmer trainers covered the TNS Nyarubaka program.

Nyarubaka sector comprises 29 villages, all of which were included in the study. Once TNS decided to train in Nyarubaka, they advertized the program with fliers. Advertisements for the agronomy training were translated into the national language and placed at the village level in administration offices, schools and churches. The farmer trainers were then assigned to visit the villages over a week to register the interested farmers, visiting each village at least twice. During registration, the farmer trainers collected the farmers' unique Rwandan identity number (along with some limited additional information) to ensure uniqueness of the registered farmers. Farmers who did not have coffee farms were not allowed to register for the program and only one person per household was allowed to register. In total, 1594 farmers registered as being interested in the program. Although the program was advertized in all the 29 villages in Nyarubaka, only farmers from 27 of the 29 villages registered to join the program. So, the treatment farmers are only in 27 of the 29 villages, though farmers in all 29 villages were surveyed, as we describe below.

3.3 Experimental Design and Data

3.3.1 Experimental Design

To understand the diffusion of coffee practices through social networks, we worked with TNS to design a field experiment around their Nyarubaka training program. The 1594 farmers who registered for the program were randomized into a treatment and a control group. Of the 1594, 855 farmers were assigned to the treatment group to receive the agronomy training. The remaining 739 farmers were assigned to the control group and received no training at any point. In addition to this basic farmer level randomization, we randomized the 27 villages that had farmers that registered to be part of the training into three different levels of treatment concentration: in a third of villages, 25% of the registered farmers were treated, in a second third of villages, 50% of the registered farmers were treated, and in the final third of villages, 75% of the registered farmers were treated. This explains why the number of treated and control farmers is not exactly the same.

Farmers were assigned to training groups in their village or in the village nearest to their location if the number of local farmers assigned to treatment was less than the minimum size for a training group. In larger villages where the total number of farmers assigned to treatment exceeded the optimal group size for training purposes according to TNS, treatment farmers were split into two or three groups for training, based mainly on geographical convenience. For practical purposes, this split was not randomized. Once assigned to a training group, farmers were expected to remain in the same group throughout the duration of the program.

During the first year of trainings, farmers in the training villages attended an average of 8 out of the 11 meetings (an attendance rate of about 73 percent). Attendance tended to be higher in the villages with higher proportions of farmers who were offered training: households in high-density training villages attended an average of 8.2 out of 11 meetings, whereas households in low density training villages attended around 7.2. This suggests that farmers were more motivated to attend trainings when more of their neighbors were also invited to attend.

3.3.2 Data

To fully capture the diffusion of coffee practices from the training program, we designed extensive data collection activities over the course of almost three years. In total, we collected ten rounds of survey data, in addition to a census. As we describe in detail, different modules of questions were asked in different rounds of the survey and some rounds of the survey included not just treatment and control households, but also additional coffee farmers in Nyarubaka sector. In addition to survey data, as we describe below, we audited coffee fields, collected GPS data on these fields, and distributed scales and calendars to some farmers so they could weigh and record their harvests. We describe each of these data collection activities in detail.

Survey Data

From December 2009 through October 2012, we conducted ten rounds of surveys. These surveys focused on the 1,594 farmers who are part of the field experiment: the RCT-sample farmers. However, given the social networks focus of the study, we also wanted to map the full social networks of the RCT-sample farmers. Therefore, alongside the baseline survey in December 2009, we also conducted a full census of the 5,198 farming households in all 29 villages of the Nyarubaka sector, including many who had not volunteered for the study. Out of these, we focused on the 57% who had grown/harvested coffee in the year prior to the census, given the training program was targeted on coffee alone and it takes five years for coffee trees to grow once they have been planted. This meant that there were an additional 1,327 coffee farmers in Nyarubaka who did not register for the agronomy training program. Throughout, we refer to these farmers as the non-RCT-sample farmers, implying they grow coffee and live in the same sector as the RCT-sample farmers but are neither treatment nor control farmers for the agronomy program.

The data collection was split into modules that covered different aspects of the household's behavior. The modules covered household demographics, detailed plot level data for coffee as well as all other crops (including harvests, sales, labor and other inputs), coffee plot performance, coffee farming activities and practices, a consumption module, household finance and social networks for the household head and the spouse. In the social networks module, we asked both the household head and spouse who their friends were (with no limit on the number that could be listed). In addition, we asked which of these friends grew coffee and which they spoke to about coffee. Throughout the paper, we define friends as being "coffee friends", the friends that respondents in the sample report talking to about coffee.

Not every round of data collection covered the same modules and not every module in a given round covered all farmers. We collected fewer rounds of data for the farmers in the non-RCT-sample. Appendix 3 shows the schedule of which modules were collected in which rounds of data collection separately for the farmers in the RCT-sample and for those in the non-RCT-sample, as well as the timing of each survey wave. The first nine rounds of surveys took place every 2-3 months over the course of the program (recall that the training was run monthly between February 2010 and October 2011), and the tenth and final round took place in September-October 2012.

Audit Data

In addition to the survey data, we also audited the coffee plots. We were concerned that asking farmers several times about their adoption patterns may result in them erroneously reporting positive adoption simply because they were asked about it a few times. For the audits, the TNS agronomist trained our field staff to recognize a set of agronomy practices. The field staff were then given instructions on how to inspect coffee plots and trees.[2]

Some variables were collected at the household level, some at the plot level (for each coffee plot) and others at the tree level (for five trees on each coffee plot). Data on two practices were collected at the household level that were also observed by the field staff: whether the household kept record books and whether the household has a compost heap. The field staff collected data on three practices at the plot level: whether the farmer had used any methods to control for soil erosion (such as using stabilizing grasses, water traps, etc.), whether there were any shade tress on the plot, and whether the farmer had grown other crops among the coffee. Finally, data on twelve practices were collected for each of five trees that were observed on each coffee plot: whether the tree had any antestia (an insect), how many leaves were yellowing, how many leaves were curling, use of mulch, evidence of weeding, evidence of rejuvenation, evidence of pruning, evidence of integrated pest management, whether the tree had any berry borers (an insect), evidence of damage from white borers, evidence of scales or mealy bugs or mould, and signs of leaf rust.

Weigh Scale Data

Starting in March 2011, we distributed weigh scales to all the farmers in the RCT-sample for them to keep accurate counts of their coffee harvests. The bulk of the coffee harvest usually arrives in May and June. Starting in March 2011 and through June 2012, every month we distributed a yield calendar to the farmers in the RCT-sample for them to record daily harvests for that month. An example of a yield calendar is shown in Appendix 4. The farmers were given instructions on how to use the weigh scale and how to record their coffee harvests on the calendars. At the end of every month, we collected up the calendars and distributed new ones for the following month.

3.4 Results

In this section, we describe the results. First, we show randomization balance along a number of baseline characteristics. We then describe the direct treatment effects of the

[2] More details on this are provided in the Appendix.

agronomy program for the farmers in the RCT-sample. Next, we show results on diffusion of the practices among both the treatment and control farmers. Finally, we examine whether diffusion is stronger among treatment friends who also attend trainings together.

3.4.1 Baseline Checks

We use the baseline data for the farmers in the RCT-sample to check for randomization balance. For a number of different baseline outcomes, we report results from the following specification in Table 1:

$$y_{ij} = \alpha + \beta Treat_{ij} + \gamma_j + \epsilon_{ij} \tag{3.1}$$

where y_{ij} is the outcome for household i in village j, $Treat_{ij}$ is a dummy variable for whether the household was allocated to the agronomy training program; γ_j are a set of village fixed effects.

As can be seen from Table 1, we find balance between treatment and control groups across a wide variety of baseline outcomes. The p-value on the joint F-test for these outcomes is 0.2638.

3.4.2 Agronomy Training Program Treatment Effects

Before we can study diffusion, we first need to confirm that the agronomy training program itself had effects, both on the knowledge and adoption of practices as well as on real outcomes. Therefore, we first report on the treatment effects of the program, obtained by running specification (3.1) on different outcomes. For outcomes related to the practices, we create the following indices:[3]

1. **Knowledge**: this index is the mean of fifteen standardized measures of what the farmer knows. It includes whether the farmer knows each of the ten methods used to control insects, pests and other diseases and how they should be used and whether the farmer knows each of five different fertilizers that should be used on coffee.

2. **Adoption**: this index is the mean of nine standardized measures of which practices the farmer adopts. Importantly, these are collected using survey questions, and therefore measure *self-reported* adoption, as opposed to *observed* adoption. It includes whether the farmer adopted each of eight methods used to control insects, pests and other diseases, and whether the farmers kept a compost heap.

[3] All these outcomes are standardized with respect to the mean and variance of the outcome in the control group.

3. **All audits**: this index is the mean of nine standardized measures of what the farmer adopts as per the observed tree audits. This index includes two measures of whether the farmer uses integrated pest management (whether old and dry berries are removed, whether the bark is smoothed or banded to control white borer), whether the farmer used mulch, whether the tree was weeded, whether there are signs of rejuvenation, and four measures of pruning (removal of dead branches, removal of branches touching the ground, removal of crossing branches, removal of unwanted suckers).

4. **Main audits**: this index is the mean of three standardized measures of what the farmer adopts as per the observed tree audits. We limit this to whether the tree has been mulched, whether it has been weeded and whether the bark has been smoothed or banded. We used these three measures for the main audits, at the recommendation of the agronomist at TNS who reported these were the most important practices (or practices with the largest effects on tree health).

5. **Leaf defects**: this index is the mean of three standardized measures of leaf health from the tree audits. It includes: whether there are signs of the leaves yellowing, whether the leaves are curling and whether there are signs of the leaves rusting. A decrease in any of these measures would indicate an improvement in tree nutrition.

Table 2 reports treatment effects of the training program on measures of knowledge and adoption of the agronomic practices, and on yields. The first two columns show the impact of the program on our indices of knowledge (column 1) and adoption (column 2) of the practices. The treatment effects on these outcomes are large and statistically significant: the treatment leads to a 0.32 standard deviation increase in the knowledge index and a 0.14 standard deviation increase in the adoption index. This suggests that the treated farmers gained additional knowledge about coffee agronomy thanks to the program, and also report putting this new knowledge into practice.

The tree audits data are an important complement to the self-reports data, as they allow us to test whether the treatment group's higher reported adoption translates into noticeably healthier-looking trees. Treatment effects on these outcomes are reported in columns 3-5 of Table 2. Column 3 shows that the program did not have any effect on the first index of tree audit outcomes (described in the previous sub-section). However, the "main audits" index is 0.043 standard deviations higher in the treatment group than in the control group, suggesting that the program had some effect on how well and how much the farmers applied this subset of practices (weeding, mulching and smoothing/banding the tree to control white borer).

The program may also have decreased the incidence of leaf defects, as the index of probabilities that yellow, curling or rusting leaves are observed upon inspection is 0.029 standard deviations lower in the treatment than in the control group. While this estimate is not statistically significant (the p-value is 0.175), this could indicate that treatment farmers' trees are better nourished than those of the control group as a result of the program. Although this interpretation should be treated with caution given the imprecision of the estimate, it would certainly be consistent with the result in column 6, which shows the productivity impacts of the agronomy trainings: the treatment group's yields are 2.9% higher in the treatment group than in the control - an estimate significant at the 5% level.

In Table 3, we turn to the effects of the program on farmers' use of inputs. Columns 1 and 2 report effects on the use of two types of fertilizer, compost and chemical fertilizer (NPK).[4] While the coefficient estimates are positive for both outcomes, only the effect on NPK is statistically significant. The probability of a farmer reporting that they have applied NPK on at least one of their plots is 1.9 percentage points higher among the treatment farmers, which, given the low use of chemical fertilizer in the RCT-sample prior to the intervention, represents a 24% increase relative to the control group. Compost use is only modestly more frequent in the treatment group, by 1.4 percentage points - an increase of just over 2% relative to the control group - and this estimate is not statistically significant. This suggests that the treatment did not shift compost use on the average household in this sample - where the fraction of farmers who reported using compost was already 72.5% at baseline (Table 1).

The treatment also induced treated farmers to spend more time working on their coffee plots, as the number of household labor days per tree is 0.012 higher among treatment farms, a 5.5% increase relative to the control group. Treatment farmers also report a higher number of paid labor days per tree: treated farmers hired an additional 0.002 days of paid labor per tree, corresponding to a 22% increase relative to the control group.

3.4.3 Diffusion of Treatment Effects

We now turn to the diffusion of treatment effects through treated farmers' social and geographic networks. Our analysis focuses on two types of networks: baseline friends and neighbors. We focus on household neighbors in this section, because the diffusion through plot neighbors is very similar (i.e. there is no diffusion).[5] The identification strategy

[4]These are the two fertilizers that the agronomy program focuses on - other types of fertilizer not being recommended for most common soil types in Rwanda.

[5]These results are in the Appendix.

exploits the exogenous variation in the number of treatment friends, conditional on the total number of sample friends, provided by the random allocation of the program.

Diffusion among Treatment and Control Groups

The top two panels of Table 4 (panels A and B) report the results from the following specification, which we run separately on the treatment and control groups:

$$y_{ij} = \alpha + \beta NumTreatFriends_{ij} + \delta NumFriends_{ij} + \gamma_j + \epsilon_{ij} \quad (3.2)$$

where y_{ij} is the outcome for household i in village j, $NumTreatFriends_{ij}$ is the number of treatment friends of household i, and $NumFriends_{ij}$ is their total number of sample friends.

The bottom panels of Table 4 (panels C and D) report the results from the specification:

$$y_{ij} = \alpha + \beta NumTreatNeighbors_{ij} + \delta NumNeighbors_{ij} + \gamma_j + \epsilon_{ij} \quad (3.3)$$

where $NumTreatNeighbors_{ij}$ is the number of neighbors of household i who were allocated to the treatment group and $NumNeighbors_{ij}$ is household i's total number of sample neighbors.

The estimates in columns 1 and 2 report the diffusion effect of treatment friends (in panels A and B) or treatment neighbors (in panels C and D) on the knowledge and self-reported adoption of RCT-sample farmers. The coefficients in column 1 are estimates of the average effect of having an additional friend or neighbor in the treatment group on the knowledge index of treatment (in panels A and C) or control (in panels B and D) farmers. Taken together, the four estimates in column 1 point to an absence of diffusion of knowledge about the practices on either the treatment or the control group, as these estimates are all close to zero and statistically insignificant. The results in column 2 are similar for all sections of the table. The estimate in panel C, for example, indicates that having one more neighbor who is in the treatment group increases treatment farmers' adoption index by 0.021 standard deviations.

It is important to remember here that the knowledge index measures only the awareness that the agronomic practices should be applied. Likewise, the self-reported adoption index only measures whether the farmer reports having applied them - not how well or how often. Therefore, the fact that no diffusion is observed on farmers' self-reported knowledge and adoption of the practices overall does not mean that treated farmers do not discuss the content of the trainings with other farmers in the sample. This result simply indicates that, if information is shared, it is more likely to be about *how* (or how often) practices should be applied in order for them to be most effective. The coefficient in the top row of column 1 indicates that, holding constant one's total number of friends, having an additional friend in the treatment decreases (albeit insignificantly) the average treatment farmer's knowledge index by 0.004 standard deviations.

We now turn to the indices constructed using tree audits data, for which estimates of diffusion effects are reported in columns 3-5. Taken together, the results in the bottom three panels (B, C, and D) of Table 4 for columns 3-5 do not provide evidence of any diffusion effects on the control group - from treated friends or neighbours - or on the treatment group from treated neighbors. On the other hand, panel A points to some diffusion of effects through the treatment group. This diffusion is observed primarily on the leaf defects outcome: the first row of column 5 indicates that having an additional friend in the treatment group reduces the leaf defects index by 0.036 standard deviations for an average farmer's trees in the treatment group, conditional on their total number of friends in the RCT-sample. In other words, having more friends in the treatment group seems to improve the health of treatment farmers' leaves - a sign of better tree nutrition. We interpret this as diffusion of information related to tree health taught in the trainings among treatment farmers, leading to reinforcement of the treatment effect.

The significantly negative coefficients in columns 3 and 4 are surprising. In column 3, for example, the coefficient reported in panel A suggests that each additional treatment friend is associated with a decrease of 0.019 standard deviations in treatment farmers' index of tree audit outcomes (comprised of nine standardized measures of tree outcomes) on average. In panel B, the estimate reported in column 3 indicates that, for an average farmer in the control group, having one more friend in the treatment group decreases this index by 0.021 standard deviations. Column 3 of panel D states that having one more neighbor who is in the treatment group decrease control farmers' adoption index by 0.020 standard deviations, conditional on their total number of sample neighbors.

Diffusion on the Treatment Group by Program Participation

To investigate mechanisms behind the diffusion of leaf health improvements among treatment farmers, we test whether diffusion effects are stronger for friends who attended trainings together. To this end, we run one more specification on the treatment group:

$$\begin{aligned} y_{ij} = \alpha + \beta_1 TFriendsSameGrp_{ij} + \beta_2 TFriendsOtherGrp_{ij} \\ + \delta NumFriends_{ij} + \theta GroupSize_{ij} + \gamma_j + \epsilon_{ij} \end{aligned} \tag{3.4}$$

where y_{ij} is the outcome for household i in village j, $TFriendsSameGrp_{ij}$ is the number of treatment friends of household i who were part of the same training group as household i whereas $TFriendsOtherGrp_{ij}$ denotes their number of treatment friends who were part of a different training group. $NumFriends_{ij}$ is household i's total number of sample friends, and $GroupSize_{ij}$ is the number of treatment farmers in household i's training group.

The results of specification (3.4) are reported in Table 5. The first two columns report effects on our two indices of tree audit outcomes, All Audits and Main Audits, on which we found no evidence of any diffusion in Table 4. Here again, the coefficient estimate on the number of treatment friends (both those who the household trained with and those who were part of a different training group) is negative and insignificant (except for row 2 of column 1, which is significant at the 5% level).

Columns 3 and 4, however, provide evidence consistent with a channel of reinforcement of treatment effects through friends who attend the same trainings. The top row of column 3 shows that, conditional on their total number of friends in the RCT-sample and the number of farmers in their training group, the effect of having one more treatment friend in their training group is associated with a reduction in treatment farmers' leaf defects index by 0.057 standard deviations. This estimate is significant at the 1% level. In the top row of column 4, we also see a positive diffusion effect on yields - albeit statistically insignificant - with each additional treatment friend that they trained with increasing treatment farmers' yields by 0.9% on average. In contrast, the effect on these two outcomes from an additional treatment friend who attends a different training group, is a decrease in 0.009 standard deviations in the leaf defects index and a 1% decrease in yields respectively, and both are insignificant.

Finally, columns 5 and 6 report estimates on the effects of training with other treatment friends on fertilizer use. These do not lend support to a reinforcement of treatment effects on compost or NPK use driven by farmers attending the same trainings as their friends, as neither of the coefficient estimates in the top two rows are statistically significantly different from zero.

Taken together, the results in Tables 5 suggest that any diffusion of treatment effects through the treatment group is concentrated around leaf health improvements, and this is driven by treatment friends who attend the same trainings. The presence of leaf defects is an important indicator of tree nutrition. The result on yields suggests that this may result in productivity improvements as well, although this estimate is too imprecise to lend confidence to this interpretation.

Since the effects on fertilizer use are similar across friends within and outside farmers' training groups, we cannot interpret these results as simply suggesting that farmers who attend trainings together reinforce each other's learning about the importance of using fertilizer. However, they could be helping each other understand and retain the most effective ways of applying the practices (e.g. with the right quantities of compost). Another explanation for why diffusion of effects through farmers' baseline networks appears relatively weak on all but the leaf health channel is that friendship networks changed as a result of the program. We investigate this hypothesis in the next section.

3.4.4 Program Effects on Social Networks

The group training structure is a central component of the Technoserve program. Group trainings reduce the cost of implementation, and allow farmers to interact during the demonstrations, thus learning from each other and creating relationships. In this section we examine the impact of the training groups on farmers' reported social network links to other farmers with whom they discuss coffee. To carry out this analysis we first transform the data to the household-by-training group level. Each household corresponds to 39 observations of the dependent variable: 38 containing the number of friends the respondent names in each training group, plus 1 "pseudo-group" for friends not treated (whether assigned to control, or outside the sample). To account for fixed differences in these training groups both within and across villages, all regressions include fixed effects at the village-by-training group level as well as the household level. This estimation strategy allows us to focus on the interactions between group-level characteristics and household treatment status. The group level-characteristics we examine are whether the group consists of trained households (i.e., all groups except the untrained pseudo-group), whether the group is organized in the same village as the respondent, and whether the group is the same that was attended by the respondent's household. Note that each of these categories is nested: all groups in the village were trained, and (almost) all groups attended by the respondents were in the same village where the respondent lived.

Table 6 reports the results. The first column shows the effect of the treatment on the number of friends reported in any training group, including those not personally attended by the household and perhaps in other villages. Rows 1 confirms that there was no significant difference between treatment and control households' numbers of treated or untreated friends at baseline, and row 4 shows that treated households had no change in their number of untreated friends after the program. Row 5 shows a marginally significant increase of 0.129 friends in any training groups for control households after the program, and row 8 shows that this effect was not significantly larger for treated households. These changes appear to be a re-allocation of social network links rather than a net increase in friends: the reported number of links with untreated households decreases (row 10) by almost exactly the same amount as treated household friendships increase (row 4).

Column 2 refines the analysis to focus on the change in treated friends who were trained in the respondent household's own village. Row 2 shows that, at baseline, there was no difference between treatment and control households on this outcome. Row 6 shows that control households increase their average number of friends in treatment groups in their own village by 0.064 above and beyond the increase of 0.127 in friends from treatment groups in any village. The sum of these coefficients is significant at the .02% level. Finally, row 8 shows a highly significant additional increase in the number of friends from within-village treated groups if the respondent household is itself treated. Indeed, the total increase in within-village treated friends is more than twice as large as the increase in outside-village treated friends for treated households.

Finally, column 3 displays results on the number of friends from each treated household's own training group. This evidence is qualitatively different from that presented in the first 2 columns for several reasons. First, as rows 2-3 make clear, assignment to specific training groups within a village was not randomized. Thus at baseline treated households had more friends in their own training groups (row 3) and fewer in other training groups within the same village (row 2). This is likely because households were sorted into training groups based on proximity. Second, we can only examine the effects on the number of friends in the household's own training group within the sample of trained households. This is because the control households never had the opportunity to select into a training group. For this reason there is no "In same training grp.×Post" coefficient displayed in column 3. These caveats aside, the coefficients in rows 8-9 of column 3 tell a clear story of the program's impact on the social network links of treated households. Relative to the controls, trained households are marginally less likely to form friendships with treated individuals in different training groups in their own villages (-0.02 in row 8, though this is not significant). However, compared to the same trained households prior to the program, they significantly increase the number of social network contacts with households who attended the same training group (row 9). The coefficient of 0.278 is by far the largest and seems to be entirely driving the differential effect of the program on the social networks of its beneficiaries. However, because this outcome is not defined for the control group, this parameter is identified only by changes over time within the treatment group.

Regardless of the specific outcome, it is clear that the training program causes a resorting of social network and an increase in links with others who have learned from the trainings. These effects may be partly due to time spent together at trainings, but as we see in rows 4 and 5, control group households also increase links to trained individuals in their villages.

3.5 Conclusions

This paper asks whether knowledge about best agronomic practices acquired through training programs can spread through social networks. We find limited evidence of this in our RCT in Rwanda. The results suggest that there were no spillovers of the program on the adoption rates of control group farmers, nor do we find any evidence of dissemination of the effects through the farmers' geographic networks (neighbors).

Any diffusion of treatment effects through the treatment group is concentrated around leaf health improvements. We find that this effect is driven by farmers with existing social connections at baseline, with whom they attended the trainings. This is detected in the form of reductions in leaf defects in the tree audits, which means the trees are receiving better nutrition. Since we do not find evidence of diffusion in the form of increased fertilizer use, this suggests that farmers training together with friends help each other learn *how best* to implement agronomic practices surrounding tree nutrition, leading to a reinforcement of treatment effects.

Our findings contribute to the growing body of evidence on both the role of information frictions and the complexities of social learning in agriculture. Recent studies suggest that farmers can lack information on how to best apply even familiar technologies, a phenomenon that has been attributed to poor access to extension in remote areas as well as inattention to parts of farming production functions (Cole and Fernando (2012); Hanna et al. (2014)). From this stems a key policy need to understand the determinants of social learning. If technology use spreads quickly through farmers' networks, then extension programs can hope to alleviate information frictions at scale by targeting a few well-connected farmers who can be expected to spread the information to others. However, if there is little diffusion, then programs may be better off training many farmers in a concentrated geographic area to maximize impact.

An emerging trend in the literature is that randomized studies tend to find less diffusion of agricultural technologies than non-randomized studies. In this respect, our results resonate with those of other RCTs that find limited evidence of social learning in the context of extension services (e.g. Kondylis et al. (2017)) or about new seed varities (e.g. Tjernström (2017)). Part of the explanation for the overall absence of diffusion we find may reflect the fact that the identity of recipients can be of critical importance to knowledge dissemination (BenYishay et al. (2015); Beaman et al. (2020)). Recent evidence also points to the potential benefits of harnessing ICTs to provide technological solutions to social learning frictions (Cole and Fernando (2012); Aker and Ksoll (2016); Tjernström et al. (2020)).

Finally, our paper is, to the best of our knowledge, one of the first to study how social ties evolve in response to a randomly allocated extension program. Other studies of endogenous network changes include Banerjee et al. (2018), who find that microfinance *decreases* the number of links of individuals exposed to microfinance. In our setting, we find that the training program is associated with an increase in links among treatment farmers who attended trainings together as well between control group households and trained individuals in their villages. The fact that extension services can change a community's social structure is important for the effective design of programs hoping to leverage pre-existing networks as a key vector of technology diffusion in agriculture.

Perhaps the reason that much of the RCT-based research on diffusion through social networks has found limited results is that baseline social networks are a noisy and incomplete measure of post-treatment social connections. As we demonstrate in our analysis of dynamic social networks, the treatment itself creates connections between individuals. If it is these new connections that matter most for information dissemination, then an analysis using baseline links may conclude there is no social learning while in fact social learning is occurring but via a different set of links. This perspective suggests that future interventions might be designed to leverage this mechanism by creating a new social network most conducive to learning and rapid adoption.

References

Jenny C. Aker and Christopher Ksoll. Can Mobile Phones Improve Agricultural Outcomes? Evidence from a Randomized Experiment in Niger. *Food Policy*, 60:44 – 51, 2016.

Abhijit V. Banerjee, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. Changes in Social Network Structure in Response to Exposure to Formal Credit Markets. *mimeo*, 2018.

Lori Beaman, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak. Can Network Theory-based Targeting Increase Technology Adoption? *mimeo*, 2020.

Ariel BenYishay, Maria Jones, Florence Kondylis, and Ahmed Mushfiq Mobarak. Are Gender Differences in Performance Innate or Socially Mediated? *mimeo*, 2015.

Shawn Cole and A. Fernando. The Value of Advice: Evidence from Mobile Phone-Based Agricultural Extension. *SSRN Electronic Journal*, 11 2012.

Rema Hanna, Sendhil Mullainathan, and Joshua Schwartzstein. Learning Through Noticing: Theory and Evidence from a Field Experiment. *The Quarterly Journal of Economics*, 129(3):1311–1353, 06 2014. ISSN 0033-5533.

Florence Kondylis, Valerie Mueller, and Jessica Zhu. Seeing is Believing? Evidence from an Extension Network experiment. *Journal of Development Economics*, 125:1 – 20, 2017. ISSN 0304-3878.

NISR. National Institute of Statistics of Rwanda. 2019.

OCIR-Café. Rwanda Industry Review. 2008.

Emilia Tjernström. Learning from Others in Heterogeneous Environments. 2017.

Emilia Tjernström, Travis J. Lybbert, Rachel Frattarola Hernández, and Juan Sebastian Correa. Learning by (Virtually) Doing: Experimentation and Belief Updating in Smallholder Agriculture. *mimeo*, 2020.

World Bank. World Development Report 2008: Agriculture for Development. *World Bank, Washington DC.*, 2007.

Table 3.1: Baseline Checks

	Mean	Coefficient	Std Error	P-value
Yield, Coffee Kg Per Tree	1.262	.08	.079	.316
Total Trees	264.801	15.472	14.719	.293
Fraction Unproductive Trees	.312	-.012	.019	.548
Zero Productive Trees	.116	-.016	.018	.384
Cut Stems	.101	-.009	.016	.589
Removed Dead Branches	.747	-.043	.023	.067
Removed Suckers	.903	-.018	.016	.272
Removed Weeds	.99	-.005	.005	.312
Household Used Lime	.024	.01	.009	.292
Household Used NPK	.185	-.001	.02	.978
Applied Pesticides	.761	-.022	.022	.33
Applied Mulch	.867	-.022	.018	.221
Applied Compost	.725	.017	.024	.471
Book Keeping Done	.025	-.007	.008	.367
Head, Years of Schooling	3.534	-.068	.158	.669
Female Headed Household	.33	.025	.025	.319
Household Size	5.018	-.039	.113	.731
Average Schooling of Household	3.257	.101	.099	.305

*

p<0.1, ** p<0.05, *** p<0.01. Robust standard errors in brackets.
All specifications control for village fixed effects.

Table 3.2: Treatment Effects on Knowledge, Adoption and Yields

	(1) Knowledge	(2) Adoption	(3) All Audits	(4) Main Audits	(5) Leaf Defects	(6) Inv. Arcsine Yields
Treatment	0.322***	0.141***	0.005	0.043**	-0.029	0.029**
	[0.019]	[0.017]	[0.013]	[0.020]	[0.021]	[0.013]
R-squared	.11	.05	.03	.03	.08	.45
Village FE	Yes	Yes	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4666	4666	47878	47878	47878	6143

Note: * p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered by household in brackets.
All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix 3).
In columns (1)-(2), we use self-reports from the household.
In columns (3)-(5), we use tree level audit data. In column (6), we use self-reported harvests per tree, cross-checked against yield calendar data.

Table 3.3: Treatment Effects on Input Use

	Fertilizer Inputs		Labor Inputs	
	(1) Compost	(2) NPK	(3) Inv. Arcsine HH Labor Days	(4) Inv. Arcsine Paid Labor Days
Treatment	0.014	0.019**	0.012*	0.002**
	[0.018]	[0.009]	[0.007]	[0.001]
Control Mean	0.614	0.079	0.215	0.009
R-squared	.05	.06	.19	.04
Village FE	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes
Observations	6225	6225	6145	6145

Note: * p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered by household in brackets.
All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix 2).
In all columns, we use self-reports from the household.
In columns (3)-(4), all the labor variables are measured per tree.

Table 3.4: Diffusion of Treatment Effects on Treatment and Control Group

	(1) Knowledge	(2) Adoption	(3) All Audits	(4) Main Audits	(5) Leaf Defects	(6) Inv. Arc. Yields
Treatment Group: Friends						
Treated Friends	-0.005	-0.012	-0.019**	-0.025*	-0.036**	0.001
	[0.018]	[0.014]	[0.009]	[0.015]	[0.015]	[0.009]
Sample Friends	0.006	0.021*	0.013*	0.018	0.010	0.004
	[0.013]	[0.011]	[0.007]	[0.011]	[0.010]	[0.006]
Observations	2471	2471	25978	25978	25978	3261
Control Group: Friends						
Treated Friends	-0.001	-0.001	-0.021**	-0.033**	-0.004	-0.001
	[0.007]	[0.011]	[0.009]	[0.015]	[0.015]	[0.010]
Sample Friends	0.004	0.019***	0.017***	0.022**	0.010	0.014**
	[0.004]	[0.006]	[0.006]	[0.010]	[0.009]	[0.006]
Observations	2165	2165	21700	21700	21700	2843
Treatment Group: Neighbors						
Treated Neighbors	0.024	0.021	0.004	0.013	-0.009	-0.003
	[0.019]	[0.015]	[0.009]	[0.015]	[0.017]	[0.009]
Sample Neighbors	-0.028**	-0.013	-0.008	-0.019*	0.001	0.007
	[0.014]	[0.011]	[0.007]	[0.010]	[0.011]	[0.006]
Observations	2460	2460	25783	25783	25783	3247
Control Group: Neighbors						
Treated Neighbors	0.009	0.009	-0.020**	-0.009	0.002	-0.013
	[0.007]	[0.012]	[0.010]	[0.015]	[0.015]	[0.010]
Sample Neighbors	-0.007	-0.003	0.008	0.003	0.004	0.006
	[0.004]	[0.007]	[0.007]	[0.010]	[0.010]	[0.006]
Observations	2142	2142	21510	21510	21510	2814

*p<0.1, **p<0.05, ***p<0.01. Standard errors clustered by household in brackets. Specifications control for round and village FE.

Table 3.5: Diffusion of Treatment Effects by Program Participation

	(1) All Audits	(2) Main Audits	(3) Leaf Defects	(4) Inv. Arc. Yields	(5) Used Compost	(6) Used NPK
Treatment Group						
T Friends, Trained Together	-0.015	-0.021	-0.057***	0.009	0.009	-0.010
	[0.012]	[0.017]	[0.018]	[0.010]	[0.013]	[0.009]
T Friends, Not Trained Together	-0.023**	-0.028	-0.009	-0.010	0.013	-0.004
	[0.011]	[0.018]	[0.018]	[0.011]	[0.014]	[0.011]
Sample Friends	0.013*	0.018	0.010	0.004	0.011	0.021***
	[0.007]	[0.011]	[0.010]	[0.006]	[0.008]	[0.006]
Training Group Size	0.002	0.012	0.004	0.004	-0.004	-0.003
	[0.005]	[0.008]	[0.009]	[0.004]	[0.007]	[0.004]
Outcome Mean	-0.002	0.027	-0.047	0.533	0.639	0.096
Observations	25978	25978	25978	3261	3300	3300

*p<0.1, **p<0.05, ***p<0.01. Standard errors clustered by household in brackets.
All specifications control for round and village FE.

Table 3.6: Program Impact on Social Networks

	(1) Friends in Group	(2) Friends in Group	(3) Friends in Group
In training grp. × Treatment	-0.124	-0.125	-0.127
	[0.089]	[0.087]	[0.088]
In own-village training grp. × Treatment		0.046	-0.454***
		[0.042]	[0.048]
In same training grp. × Treatment			0.955***
			[0.104]
Treatment × Post	-0.006	-0.006	-0.006
	[0.081]	[0.081]	[0.081]
In training grp. × Post	0.129*	0.127*	0.127*
	[0.074]	[0.074]	[0.074]
In own-village training grp. × Post		0.064*	0.064*
		[0.035]	[0.035]
In training grp. × Treatment × Post	0.014	0.007	0.007
	[0.081]	[0.081]	[0.081]
In own-village training grp. × Treatment × Post		0.107***	-0.020
		[0.038]	[0.037]
In same training grp. × Treatment × Post			0.278***
			[0.057]
Post	-0.129*	-0.129*	-0.129*
	[0.074]	[0.074]	[0.074]
Control Mean	0.108	0.108	0.108
R-squared	.57	.58	.6
Village × Group FE	Yes	Yes	Yes
HHID FE	Yes	Yes	Yes
Observations	123084	123084	123084

* p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered by village × group in brackets. All columns use data from the baseline and round 9 social network surveys. (see Appendix 2). The dataset has been reshaped to the household × training group level. The outcome variable in all columns is each household's number of friends in each training group. An additional pseudo-group has been created for each household's untrained friends.

Appendix to Chapter 3.

The TechnoServe Agronomy Best Practices

The agronomy training program covered the following eight basic modules:

- Rejuvenation and pruning to produce new and productive wood. A multi-stem uncapped system was promoted.

- Nutrition: a balanced nutritional program based on organic and inorganic additives, with the exact requirements determined by soil analysis. It included homemade compost, and may include liming, inorganic fertilizers and foliar feeds depending on the requirements of the area.

- Integrated pest management: multiple techniques to manage pests and diseases, such as correct nutrition, tree management, biological control, traps etc. Selective pesticides used as a last resort, but safe use of pesticides promoted.

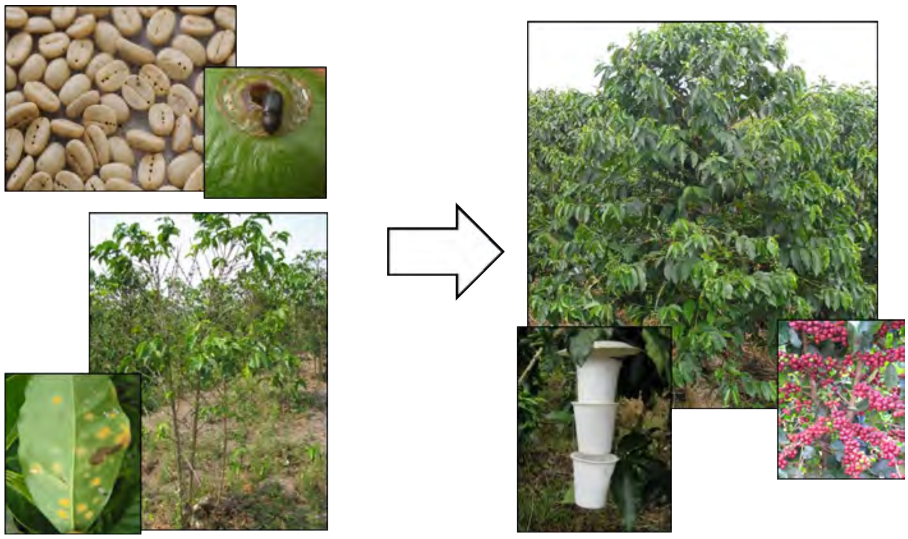

- Weed control: management of weeds through mulching, hand weeding and/or cover crops.

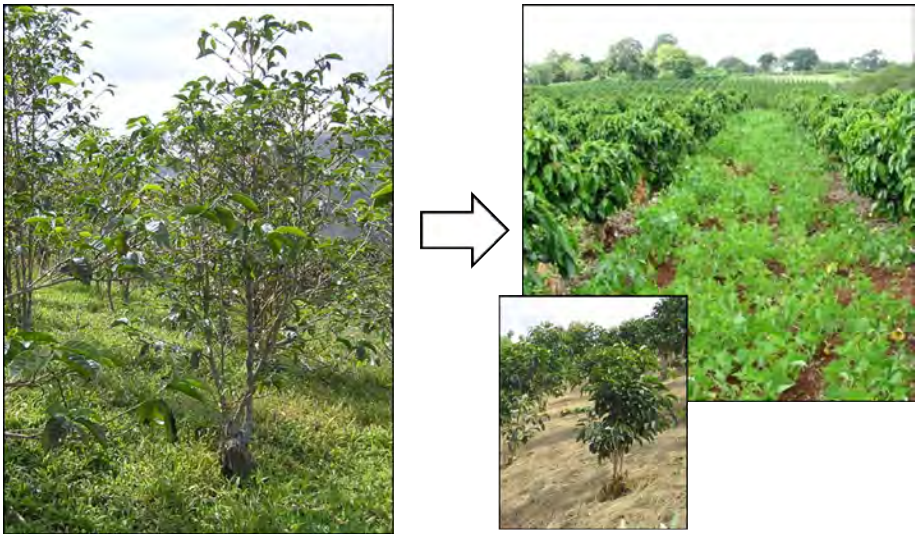

- Mulching: techniques to conserve moisture, add organic matter, and control soil erosion.

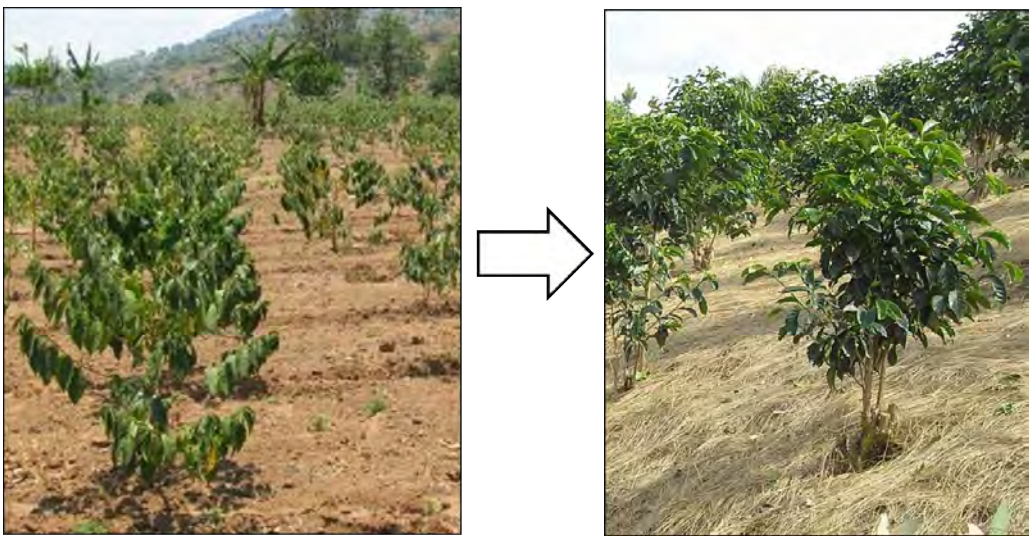

- Soil and water conservation: use of a number of techniques such as mulching, terracing and water traps to control soil erosion and maintain soil fertility. Encourage the management of water resources through conservation zones.

- Shade: use of the correct level (20-40%) of shade to reduce tree stress, conserve moisture, increase organic matter and increase biodiversity.

- Record keeping: maintenance of records of inputs, outputs, profit and loss in a record book.

The schedule of the modules covered in the training was as follows:

- February 2010: Record Keeping
- March 2010: Integrated Pest Management (IPM)
- April 2010: Coffee Nutrition
- May 2010: Coffee Harvesting
- June 2010: Weed Control
- July 2010: Mulching
- August 2010: Pruning and Rejuvenation
- September 2010: Safe use of pesticides
- October 2010: Composting
- November 2010: Erosion Control
- December 2010: Coffee Shade Management
- January 2011 to October 2011: Review

Audits Data

Our main adoption measures were constructed from plot and tree inspections data, collected using plot and tree audits. Field staff visited each coffee plot of all the coffee farmers in Nyarubaka and inspected five trees, looking for signs of adoption of the agronomic practices covered in the trainings. The enumerators were given specific instructions on how to pick the five trees on each plot. They were instructed to start at the corner of the coffee plot closest to the farmer's house and walk in the direction of the opposite corner. They were then to inspect the second tree into the field, walk to the middle of the field and inspect the tree in the middle. Starting from the middle, the field staff was to walk towards the other two corners of the field and inspect the second tree in each direction. The field staff was then to walk back to the middle of the field and continue on the original path and inspect the second to last tree in the field. For each tree, the field staff would also note the GPS coordinates of the tree.

Data Modules Collected by Sample Type

A: RCT-Sample Modules

SECTION	1	2	3	4	5	6	7	8	9	10
Cover Page	X	X	X	X	X	X	X	X	X	X
Consent	X	X	X	X	X	X	X	X	X	X
HH Roster	X	X		X	X	X	X	X	X	X
Plot Roster		X		X	X	X	X	X	X	X
Household-Level Sections										
HH Member Demographic Characteristics	X									
HH Characteristics		X				X			X	
HH Characteristics (Extended)										X
Group Memberships		X						X		
Crop Inventory	X	X		X			X	X		
Plot Questionnaire										
Long Season	X			X				X		
Short Season	X	X					X			
Other	X			X			X	X		
Labor Activities										
Long Season	X			X				X		
Short Season	X	X					X			
Crop Harvests and Sales										
Long Season	X			X				X		
Short Season	X	X					X			
Coffee Activites										
Coffee Plot Details	X	X	X							
General Household Coffee										
A. Coffee Plot Performance/Future	X									
B. Training	X									
C. Cooperative Membership	X									
D. Coffee Farming Activities/Practices	X									
Coffee Delivery	X									
Coffee Module										
Labor Activities for Coffee	X	X	X	X	X	X	X	X	X	X
Coffee Harvests		X	X	X	X	X	X	X	X	X
Coffee Sales	X	X	X	X	X	X	X	X	X	X
Coffee Inputs		X	X	X	X	X	X	X	X	X
Best Practices Schedule/Training Attendance								X	X	X
Consumption Module	X			X				X		
Household Finance										
Decisionmaking - Use of Money	X			X			X			
Bank Holdings/Savings/Debts/Credits	X			X			X			
Remittances	X									
Bank Account Location			X							
Gifts				X			X			
Non-Agricultural Income and Credit		X		X	X	X	X	X	X	
Social Networks	X	X				X			X	X
Best Practices Module: Audits										
Coffee Plot Measurements and ID		X					X		X	X
Best Practices Sheet		X					X		X	X
Tree and Plot Inspection		X					X		X	X
Feedback on TNS Training/Improvement of Knowledge of BP										X
Barriers to Adoption of Best Practices										X

B: Non-RCT-Sample Modules

SECTION	2	3	6	7	9	10
Cover Page	X	X	X	X	X	X
Consent	X	X	X	X	X	X
HH Roster				X	X	X
Plot Roster				X	X	X
Household-Level Sections						
HH Demographics (Basic)	X					
HH Characteristics (Extended)						X
Group Memberships	X					
Coffee Activites						
General Household Coffee						
A. Coffee Plot Performance/Future	X					
B. Training	X					
C. Coffee Farming Activities/Practices	X					
Coffee Module						
Labor Activities for Coffee			X	X	X	X
Coffee Harvests			X	X	X	X
Coffee Sales	X		X	X	X	X
Coffee Inputs			X	X	X	X
Best Practices Schedule/Training Attendance					X	X
Consumption Module						
Household Finance						
Decisionmaking - Use of Money						
Bank Holdings/Savings/Debts/Credits						
Remittances						
Bank Account Location		X				
Gifts						
Non-Agricultural Income and Credit						
Social Networks	X		X		X	X
Best Practices Module: Audits						
Coffee Plot Measurements and ID		X		X	X	X
Best Practices Sheet		X		X	X	X
Tree and Plot Inspection		X		X	X	X
Feedback on TNS Training/Improvement of Knowledge of BP						X
Barriers to Adoption of Best Practices						X

Timing of survey waves:

The timing of the ten survey rounds of surveys was as follows:

1. December 2009 to January 2010
2. April 2010 to May 2010
3. July 2010
4. September 2010 to October 2010
5. November 2010
6. January 2011 to February 2011
7. June 2011 to July 2011
8. October 2011
9. January 2012 to February 2012
10. September 2012 to October 2012

Yield Calendars

HHID: VILLAGE: NAME:
LOCATION:

HOW MANY KILOGRAMS DID YOU HARVEST TODAY?

1. Hang the scale to a fix and stable place ;
2. Hang the bag (with the cherries in it) to the scale ;
3. Read the number on the scale.

	MAY (05)	DAY	Write here the coffee harvest that you have just weighed	Write here the coffee harvest that you are going to sell and indicate the type of coffee (cherries, wet or dry parch)
1	1-MAY-2012	TUESDAY	Kg	Kg
2	2-MAY-2012	WEDNESDAY	Kg	Kg
3	3-MAY-2012	THURSDAY	Kg	Kg
4	4-MAY-2012	FRIDAY	Kg	Kg
5	5-MAY-2012	SATURDAY	Kg	Kg
6	6-MAY-2012	SUNDAY	Kg	Kg
7	7-MAY-2012	MONDAY	Kg	Kg
8	8-MAY-2012	TUESDAY	Kg	Kg
9	9-MAY-2012	WEDNESDAY	Kg	Kg
10	10-MAY-2012	THURSDAY	Kg	Kg
11	11-MAY-2012	FRIDAY	Kg	Kg
12	12-MAY-2012	SATURDAY	Kg	Kg
13	13-MAY-2012	SUNDAY	Kg	Kg
14	14-MAY-2012	MONDAY	Kg	Kg
15	15-MAY-2012	TUESDAY	Kg	Kg
16	16-MAY-2012	WEDNESDAY	Kg	Kg
17	17-MAY-2012	THURSDAY	Kg	Kg
18	18-MAY-2012	FRIDAY	Kg	Kg
19	19-MAY-2012	SATURDAY	Kg	Kg
20	20-MAY-2012	SUNDAY	Kg	Kg
21	21-MAY-2012	MONDAY	Kg	Kg
22	22-MAY-2012	TUESDAY	Kg	Kg
23	23-MAY-2012	WEDNESDAY	Kg	Kg
24	24-MAY-2012	THURSDAY	Kg	Kg
25	25-MAY-2012	FRIDAY	Kg	Kg
26	26-MAY-2012	SATURDAY	Kg	Kg
27	27-MAY-2012	SUNDAY	Kg	Kg
28	28-MAY-2012	MONDAY	Kg	Kg
29	29-MAY-2012	TUESDAY	Kg	Kg
30	30-MAY-2012	WEDNESDAY	Kg	Kg
31	31-MAY-2012	THURSDAY	Kg	Kg