Essays in Development Economics

by

Rachel Sayers

Department of Economics
Duke University

Date: _____

Approved:

_____
Robert Garlick, Advisor

_____
Erica Field

_____
Duncan Thomas

_____
Xiao Yu Wang

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2020

ABSTRACT

Essays in Development Economics

by

Rachel Sayers

Department of Economics
Duke University

Date: _____

Approved:

_____
Robert Garlick, Advisor

_____
Erica Field

_____
Duncan Thomas

_____
Xiao Yu Wang

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Economics
in the Graduate School of
Duke University

2020

Abstract

This dissertation considers the role gender plays in labor markets, household decision-making, and health in sub-Saharan Africa.

The first chapter considers the impact of fast Internet access on employment outcomes and household dynamics. I find the introduction of fast Internet to sub-Saharan Africa significantly increased employment for males, but had little impact on female employment. In addition, it significantly increased perceived acceptability, among both genders, of domestic violence against women.

The second chapter considers the differential impact, by gender, of an experimental labor market intervention in South Africa, which measured skills of workseekers and provided a mechanism for workseekers to communicate their results to potential employers. I find that men experienced a larger effect of the intervention on employment outcomes than did women. This difference is largely explained by pre-existing differences between genders, rather than differential responses to treatment.

The third chapter considers the factors that contribute to female genital cutting (FGC) in Mali and tests various hypotheses to explain the persistence of the tradition. I find that maternal preference is pivotal in the decision to cut daughters. I also find that the marriage market hypothesis and the identity hypothesis of FGC decision-making are alone insufficient to explain persistence.

Acknowledgements

I am extremely grateful to Robert Garlick, Erica Field, Duncan Thomas, and Xiao Yu Wang for their help, advice, support, and encouragement at all stages of my doctoral studies. All four provided detailed feedback and clever suggestions on my research, much advice and support on the job market, and generalized encouragement along the way. I would like to thank Rob, in particular, for agreeing to be my advisor and for devoting so much time and energy to mentoring me. His knowledge and expertise not only greatly improved my own technical skills, but also helped shape my broader research agenda. More importantly, his patience, empathy, and generosity with time and advice ensured I was able to reach my goals, despite numerous hurdles along the way.

My progress at Duke would also not have been possible without the support of the greater Duke Economics community. I am thankful for all of the suggestions I received and presentation practice I got under my belt at the weekly International Population, Health and Development Seminar. I am also incredibly grateful to Kate Vyborny and Majorie McElroy for meeting with me independently to discuss my research and offer their advice. I would also like to thank Amy Clayton, Jennifer Becker, Chelsea Fillippa, and Christyn Klinck for helping me with everything else at Duke, from matriculation to graduation and everything in between.

I am also deeply indebted to my family and friends who provided me with support, encouragement, and companionship at all stages throughout my Ph.D. experience. My parents, Pauline and Colin Sayers, have always supported me in my scholarly ventures and have always helped me whenever I falter. My mother, in particular, has been my closest friend throughout my life, and she has been my major support system throughout this entire process. She has proofread papers at every stage, spent

weeks in Durham helping me during stressful times, and taken countless phone calls to provide me with the reassurance I needed. I would also like to thank my sister, Ruth Sayers, and brother-in-law, Chris Glapa, who have given me common sense, perspective, and a sense of humor during my studies. In addition, my grandparents, Bob and Brenda Turner, have provided unwavering support through all these years, even from afar.

My many friends have supported me through all the struggles, minor and major, of this program and celebrated my every achievement, small and large. My best friend, Donna Mangum, has supported me whenever I asked, and even when I didn't, and she has become a family member in the process. I would never have made it to Duke if it weren't for my friends from Loyola, Sydney Barbier, Megan Kattengell, Daniel Quick, Phuong Tran, Dwayne Fontenette, and Andrew Pinheiro. All six were by my side as I developed a passion for economics and later encouraged me to pursue a Ph.D. My Duke friends, Laurel Wheeler, Kate Maxwell Koegel, Myu Kulathungam, Chelsea Garber, and Josh Jacobs have also been a key part of my success. Laurel and Kate, in particular, saw me through many problem sets in our first year and then through my greatest hardship a few years later. Myu gave me the encouragement, sense of humor, intellectual debate, and kindness I needed to keep me going through the roughest patch.

Last but not least, I would like to thank my most constant and closest source of love and support over the past 4 years: my foxhound, Sadie. She woke me up every morning, took me out for fresh air and exercise, squeaked at me when I was down, and lay with me when I needed to rest. Her affection, silliness, and sheer stubbornness have kept me motivated, energized, and dedicated enough to finish this dissertation and my degree.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This dissertation considers the role gender plays in labor market outcomes, household decision-making, and health in sub-Saharan Africa.

Chapter 2 considers how the introduction of fast Internet access impacted employment, household decision-making roles, and approval of domestic violence for men and women. Information communication technologies (ICT) have been shown to have a large impact on labor markets, and such impacts may be particularly relevant in developing countries, where labor market frictions and information asymmetries are larger. There is also reason to believe that the impact of these technologies may differ by gender, due to gender differences in education, skills, and household structure. This topic is of high economic importance, as differential impacts of ICT on labor market outcomes by gender may result in changes in the distribution of bargaining power within a household, leading to changes in female autonomy, domestic violence and fertility.

Chapter 3 considers how a skill certification program designed to communicate skill levels both to workseekers and prospective employers impacted labor market outcomes for men and women. One reason why skill certification may affect men and women differently is due to gender differences in baseline characteristics that are related to the skills being certified or the productivity of workseekers. Another reason why skill certification may have impacts that differ by gender is due to baseline gender differences in labor market experience, education and referral networks, which are other ways for workseekers to reveal their ability to firms. Understanding whether

and how skill certification affects men and women differently is of high importance, as any differential effects by gender may lead to changes in the gender employment and wage gap.

Chapter 4 considers how parents decide whether to expose their daughters to female genital cutting (FGC). A crucial component in understanding how parents make this decision is to investigate whether it is maternal or paternal opinion on FGC that is more predictive of a daughter's cutting status. Another important component is investigating whether parents make homogeneous choices regarding FGC for all their daughters or whether there is variation within daughters of the same parents. An additional essential component is determining whether undergoing FGC as a child is associated with greater marriage market outcomes as an adult. Each of these components is important for understanding why parents choose for their daughters to undergo FGC, a household decision-making process that affects the health and well-being of future generations.

Chapter 5 concludes by reviewing the findings of this dissertation.

Chapter 2

Internet Access, Employment, and Household Bargaining

2.1 Introduction

Information communication technologies (ICT) have been shown to have a large impact on labor markets. In particular, the introduction of Internet access and other ICTs is associated with increases in labor force participation, employment rates and wages (Champion et al. 2012, Atasoy 2013, Forman et al. 2012, Efobi et al. 2018), as well as increases in firm-level productivity (Crespi et al. 2007, Commander et al. 2011). Research has also demonstrated a positive causal impact of Internet access on employment rates and employer-employee matching quality (Hjort and Poulsen 2019, Bhuller et al. 2019). These impacts may be particularly relevant in developing countries, where labor market frictions and information asymmetries are larger. For instance, ICT has had a large impact both on increasing agricultural output and increasing market integration in developing countries by reducing information asymmetries and search costs (Casaburi et al. 2014, Aker 2010, Jensen 2007).

There is reason to believe that the impact of these technologies may differ by gender, due to gender differences in education, skills, and household structure. In the U.S., the associated employment and wage effects of Internet access have been shown to vary by gender, marital status, number of children, and education (Dettling 2015,

[0]I would like to thank Jonas Hjort and Jonas Poulsen for answering questions about their paper from which I have based much of this chapter's empirical methodology.

Champion et al. 2012, Forman et al. 2012). Although the differential employment effects by gender have not been causally explored at the individual level,[1] this topic is of high economic importance due to differences in how men and women spend their income. When women control household income, they are more likely to spend that income on investments in the future and investments in child health and education than if men were to control that income (Thomas 1990, Lundberg et al. 1997, Duflo 2000, Rubalcav et al. 2009, Majlesi 2013). In addition, as women control relatively more household income, their bargaining power increases, leading to changes in female autonomy, domestic violence, and fertility, important factors in overall female welfare (Rahman and Rao 2005, Anderson and Eswaran 2009, Heath 2014, Haushofer et al. 2019). While most literature finds that increases in female income and employment increase female autonomy and decrease exposure to domestic violence, Anderson and Eswaran (2009) find this only holds for work outside of the husband's farm, while Heath (2014) finds the effects may be reversed entirely for women with low starting levels of bargaining power. Furthermore, direct provision of information and exposure to new ideas via the Internet may change individual attitudes and social norms and ultimately affect decision-making roles in the household, as demonstrated by the impact of cable television and soap operas on increasing women's autonomy, decreasing fertility, and decreasing acceptability of domestic violence (Jensen and Oster 2009, La Ferrara et al. 2012).

In this paper, I consider the differential impact of Internet access on employment by studying the introduction of fast Internet to Africa in the late 2000s and early 2010s. During this time, submarine cables were laid that connected many coastal African countries' existing Internet networks to those in Europe, India and

[1] Although Dettling (2015) uses an instrumental variable approach, she uses state-level employment rates. Champion et al. (2012) and Forman et al. (2012) employ correlational analysis.

the United Arab Emirates. These connections increased the speed and traffic capacities of African terrestrial networks. I utilize two sources of variation in Internet access to estimate the causal impact of fast Internet: timing differences between countries in when they first connected to a submarine cable and geographic differences within countries in distance to terrestrial cable networks. By combining these two sources of variation, I construct a difference-in-difference estimator to determine the impact of high-speed Internet on employment, household decision-making roles, and opinions on domestic violence, allowing for heterogeneous effects by gender. My identification is an adapted version of that used by Hjort and Poulsen (2019), altered to allow for heterogeneous treatment effects by gender.

I find that fast Internet access has substantially different labor market effects for men and women. While fast Internet access increases current male employment by 4.9 percentage points on a 73.5% mean, the impact on female employment is a statistically insignificant 2.1 percentage point increase on a 61.8% mean. The results are similar when employment is measured as having worked in the past 12 months. Fast Internet access shifts men into self-employment and work for family, and out of work for someone else; it also shifts women into year-round work. The impact of fast Internet on employment varies by education for both genders and decreases with household size and number of children in the household for women. Single women and women in female-headed households experience greater employment gains from fast Internet access. These results suggest that, although Internet access has substantively different labor market effects by gender, some of this difference is explained by differences in education and household composition by gender. Furthermore, the impact of fast Internet on employment is positive for those individuals expected to be primary earners within their household. This suggests that, although labor market effects are relatively large, these effects may only be strong enough to move one

additional household member into employment.

I find evidence that the introduction of fast Internet shifts household bargaining. Fast Internet access increases the number of women who report being the sole decision-maker for large household purchases and visits to family. This is consistent with a theory of household decision-making where decision power over the daily operation of the household is delegated to women as male attachment to the workforce, and thus time spent outside the household, increases, even though relative male bargaining power increases on the whole. I also find that the introduction of fast Internet increases the acceptability of domestic violence against women among both men and women. Fast Internet increases the number of men who report domestic violence as acceptable if a wife refuses sex with her husband or neglects her children, while it increases the number of women who report domestic violence as acceptable if a wife argues with her husband or goes out without telling her husband. This suggests that, even though male bargaining power is increasing relative to that of women, men still engage in domestic violence. This finding is consistent with a model of domestic violence where men may not receive direct utility from domestic violence but use it as a means to obtain control over household resources, and where the utility of using such violence increases as total household resources and, as such, the amount of resources extractable through domestic violence increase.

This paper contributes to the growing literature concerning the impact of ICT on employment and is, to the best of my knowledge, the first paper to consider the causal impact of fast Internet access in Africa on employment by gender. Hjort and Poulsen (2019) calculate the the impact of fast Internet on employment in this setting and find a positive relationship, but they do not separate the effects by gender. While Efobi et al. (2018) note a positive relationship between female employment and Internet access in this setting, this relationship is not causal. In addition, this is the first

paper to look at the impact of fast Internet access on household bargaining measures, including domestic violence. These findings are important, as understanding how ICT affects male and female employment differently is crucial to understanding how any associated economic gains are distributed within the household, a pattern that potentially impacts female autonomy, female welfare and levels of intergenerational investment. Furthermore, direction information provision and exposure to new ideas via the Internet may have direct implications for female autonomy and female welfare by changing attitudes and social norms. The closest prior literature on Internet access and household bargaining is correlational evidence on the impact of Internet penetration and social media use on divorce rates at the state level in the U.S. (Kendall 2011, Valenzuela 2014) and causal evidence on the impact of cable television and soap operas on female autonomy and acceptance of domestic violence in India and Brazil (Jensen and Oster 2009, La Ferrara et al. 2012).

In Section 2.2, I explain how fast Internet was introduced via submarine cables to the African continent and how these cables connected to existing terrestrial Internet networks. I also review the existing literature concerning the impact of fast Internet access on labor market outcomes and the existing literature relevant to the impact of fast Internet access on household bargaining. In Section 2.3, I describe the geographical data used to measure connectedness to the terrestrial networks, the household data I use to identify outcome variables, and the various sample restrictions. In Section 2.4, I detail the empirical method used to estimate treatment effects of fast Internet on various outcomes, along with the identifying assumptions underlying this method. In Section 2.5, I discuss the results of my analysis, including robustness checks. In Section 2.6, I conclude.

2.2 Context

2.2.1 Internet Infrastructure in Africa

Prior to and for much of the 2000s, the majority of African countries relied on Internet with low bandwidth and low capacity. Although countries had terrestrial cable networks, referred to as national "backbones", in place, these networks relied on satellites to transmit data over the Internet. These satellites had high costs per megabyte of capacity and were characterized by slow speeds (Madamombe 2007, Schumann and Kende 2013).

Beginning in the 2000s, construction began on submarine cables that would connect African backbone networks to other countries' Internet infrastructures. Submarine cables are a bundle of several strands of fiber optic cable wrapped in a protective covering and laid on the seafloor (Forden 2015). The cables run from other continents and connect to national backbone networks in Africa at designated landing points on the coast. Figure 2.1 shows the relevant submarine cables, along with the terrestrial cable networks to which the submarine cables connected (Mahlknecht 2019); the connection date of each country is determined by the country's first connection to a submarine cable.[2] It is important to note that each coastal country has its own landing point, which connects to its own national backbone, thus reducing concerns of international spillovers in my analysis.[3,4]

[2]This definition of timing is only relevant for Benin and Tanzania, for which the data's post-connection observations occurred before all submarine cables were laid. The results are robust to excluding both of these countries.

[3]Coastal countries may work with neighboring landlocked countries to connect national backbones and thereby provide the benefits of submarine cables to the landlocked country. However, determining the exact date that any spillovers to landlocked countries occurred is often difficult, so landlocked countries are excluded from this analysis.

[4]The exceptions to this are Benin and Togo, which were connected to a submarine cable through Nigeria in 2010, before they were connected directly via a coastal landing point. This is not a concern for Togo, as its coastal connection to a submarine cable was established prior to when

Figure 2.1: Submarine Cables and Terrestrial Backbone Networks

Note: The submarine cables included in this figure are The East African Marine System, TEAMS (June 2009, connected Kenya to UAE); SEACOM (July 2009; connected Kenya and Tanzania to Europe and India); Eastern Africa Submarine Cable System, EASSy (July 2010, connected Kenya and Tanzania to Europe); Main One (July 2010; connected Benin, Ghana, Nigeria and Togo to Europe); Globacom-1, GLO-1 (April 2011, connected Ghana and Nigeria to Europe); Lower Indian Ocean Network II, LION II (April 2012, connected Kenya to Europe); West Africa Cable System, WACS (May 2012, connected Democratic Republic of Congo, Ghana, Namibia, Nigeria, and Togo to Europe); and Africa Coast to Europe, ACE (December 2012, connected Benin, Democratic Republic of Congo, and Nigeria to Europe). The Lower Indian Ocean Network, LION, was also connected in 2009, but it did not impact any of the countries in my sample, so it is omitted from the map.

Once connected to submarine cables, African backbone networks benefit from much faster speed and higher capacities. From the backbone network, Internet then travels via the "last mile" of infrastructure to the end user. This can be in the form of additional cables or, as is more common in Africa, wireless transmission (Schumann and Kende 2013). As data on the "last mile" is difficult to collect, I cannot observe this infrastructure and instead set a radius around the backbone to account for the "last mile" of infrastructure.

The increase in speed and capacity attributable to the submarine cables is due in large part to the superior speed and capacity of fiber optic cables to the satellites previously used in much of Africa. However, part of this increase is also due to the large amount of African Internet traffic that is sourced and hosted internationally. For the 3 countries in my sample, Ghana, Kenya, and Nigeria, whose top 15 websites were archived during their second survey wave, all of the websites on these lists are hosted on servers in North America or Europe (Alexa 2019, Wayback Machine 2019, Hosting Checker 2019).[5]. Of the 7 countries in my sample whose top websites were archived at any point, only two countries had any of their top 15 websites hosted within Africa, and all of these African-hosted websites belonged either to local universities or local

the data's post-connection observations were collected for that country; however, Benin was not independently connected to a submarine cable until after the data's post-connection observations were collected. Nevertheless, Benin's connection via Nigeria is well-documented with a well-defined treatment date, so it is included in the sample. All results are robust to excluding Benin.

[5]To construct lists of the top 15 websites in each country, I used Wayback Machine (2019) to view archived versions of the top sites from Alexa (2019). The list of top websites is based on page views over the 3-months prior to the date archived. If information was available on top sites during the second survey wave of a country, I used the archived version nearest the end of the survey wave. If information on the top sites during the second wave was not available, I instead chose the first archived version after the second survey wave; for many countries, this occurred years after the second survey wave concluded. Once the top 15 sites for each country were obtained, I then used Hosting Checker (2019) to determine the server location of each site. Although the server location is that of the website on the date Hosting Checker (2019) was accessed (Oct. 14, 2019), the server is likely in the same country as it was years ago. Information was available on top sites for every sampled country except Togo. This information is available in detail in Tables A4, A5, A6, A7, A8, A9 and A10.

banks. Furthermore, due to a lack of direct high-capacity transmission lines between many African countries, 75% of Internet traffic moving between two African countries in 2008 traveled through Europe or North America in order to reach its destination (New Media and Development Communication 2008).

As such, the increase in speed provided by fiber optic submarine cables was an increase in speed to most websites used by and most data sent and received by Africans. It is also important to note that the increase in speed and capacity provided by the cables to the national backbones does not diminish with cable length by an appreciable amount (Hjort and Poulsen 2019).

2.2.2 ICT and Labor Market Outcomes

Pre-existing literature provides compelling evidence that fast-speed Internet access has significant labor market impacts in developed countries. In the United States, access to high-speed broadband Internet is associated with an increase in employment, total hours worked, and amount of work done from home (Champion et al. 2012, Atasoy 2013). In addition, the introduction of broadband to the U.S. was associated with wage growth and an increase in scale among existing firms (Forman et al. 2012, Atasoy 2013). In the United Kingdom, ICT is associated with firm-level productivity growth (Crespi et al. 2007). In Norway, fast Internet had dramatic effects on labor-matching, as broadband improved firms' hiring process, increased the match quality of new employees, increased the propensity with which job seekers matched with firms far from home, and boosted starting wages (Bhuller et al. 2019). In a general survey of OECD countries, an increase in broadband infrastructure raised annual per capita growth (Czernich et al. 2011).

Research on Internet access and labor markets in developed countries also indi-

cates that the effects of high-speed Internet vary by gender and other observables. High-speed home Internet usage increases labor force participation for married women in the U.S., but not for single women or men; the largest effects are for college-educated, married women with children (Dettling 2015). Dettling suggests that these effects might be explained by time saving in home production attributable to Internet use. Another study found that the correlation between Internet access and employment in the U.S. was largest for women in rural areas, less-educated individuals, and black individuals (Champion et al. 2012). A different paper finds that Internet access in the U.S. is only associated with wage growth in counties that were already educated, wealthy, and had a high concentration of IT-focused industry (Forman et al. 2012).

Research in developing countries has demonstrated that ICT also has important labor market implications in these settings. In coastal African countries, access to high-speed Internet increases employment substantially (Hjort and Poulsen 2019), and improving ICT in Sub-Saharan Africa is associated with an increase in female labor force participation (Efobi et al. 2018). At the firm level, there is a positive association between ICT capital and productivity in Brazil and India (Commander et al. 2011). In addition, in a survey of developing countries, increased Internet penetration increased exports to developed countries (Clarke and Wallsten 2006).

The literature also indicates that ICT has large labor market impacts for agricultural workers and fishers in developing countries. Distributing agricultural advice to smallholder farms by SMS in a randomized control trial increases crop yields substantially, while allowing farmers to communicate fertilizer delivery delays to the company by mobile phone substantially reduced these delays (Casaburi et al. 2014). Casaburi et al. also find that the increase in yields was concentrated among farmers with no agricultural training, demonstrating the importance of skill and education

when measuring the effect of ICT on labor market outcomes. In Niger, the adoption of mobile phones reduces search costs for farmers when searching for markets to sell their products, in turn reducing price dispersion across markets (Aker 2010). Similarly, the adoption of mobile phones among fishermen and wholesalers in India was associated with a reduction in price dispersion and better integration of markets, which increased both producer and consumer welfare (Jensen 2007).

2.2.3 ICT, Labor Markets, and Household Bargaining

To my knowledge, little literature exists concerning the impact of Internet access on household bargaining, either in developed or developing countries. The most relevant research on this topic considers the impact on Internet access on divorce rates and marriage quality. Although Internet penetration decreases search costs for alternate romantic partners, Internet penetration itself is not associated with higher divorce rates in the U.S. (Kendall 2011). In the U.S., Facebook penetration is associated with increased divorce rates, and reported use of social networking sites is negatively correlated with marriage quality (Valenzuela 2014).

Prior literature on the effects of other technologies which disseminate information, such as television, can help inform how Internet access might affect household bargaining and dynamics. In rural India, the introduction of cable television was associated with a decrease in the reported acceptability of domestic violence toward women, an increase in women's autonomy, a decrease in fertility, and a decrease in reported preference for sons over daughters (Jensen and Oster 2009). In Brazil, the expansion of soap opera programming lowered fertility, with the strongest effect on women in lower socioeconomic standing (La Ferrara et al. 2012). These results indicate that exposure to different ideas and new information through the introduc-

tion of a new technology can have significant impacts on household bargaining and household dynamics.

In addition to this literature, broader research on the impacts of income changes within the household on bargaining power suggest that Internet access may also impact household bargaining by changing employment for household members. In India, higher female wages improve women's mobility and authority, while higher male wages decrease these (Rahman and Rao 2005), and, in Bangladesh, female employment outside their husbands' farms increases women's autonomy (Anderson and Eswaran 2009). Labor demand shocks that reduce women's employment in Mexico more than that of men reduce women's relative decision-making power within households and overall household investment in children's health (Majlesi 2013). In Kenya, unconditional cash transfers to women decrease domestic violence (Haushofer et al. 2019); however, in Bangladesh, increased female employment is associated with higher rates of domestic violence for women who had low education or were young at the time of their first marriage (Heath 2014). This research indicates that differential employment effects by gender may have important implications for household bargaining and dynamics.

2.3 Data

To study the impact of fast Internet access in sub-Saharan Africa, I utilize two sources of data in this paper. The first data source is geographical data on the terrestrial Internet networks compiled by Hjort and Poulsen from AfTerFibre and www.africabandwidthmaps.com data sources, and the second is individual-level data on employment outcomes, decision-making power, and opinions on domestic violence from the Demographic and Health Survey (DHS).

The geographical data on terrestrial Internet networks is a map of all the relevant national backbone networks in place prior to the introduction of any submarine cables to my sample.[6] This map is used to determine if an individual has access to the backbone network. Specifically, an individual is deemed "connected" to the backbone if they are within 500 meters of the nearest backbone cable. The choice of 500m is based upon prior literature and the belief that this distance is a conservative proxy for the reach of "last mile" (Hjort and Poulsen 2019). In Section 2.5, sensitivity analysis of this distance is performed.

To measure outcomes, I use two DHS waves from each country included in my sample: the last wave conducted prior to the arrival of submarine cables and the first wave conducted after this arrival. The DHS is a standardized survey conducted in many countries that contains information on demographics, employment, and a number of health outcomes. The DHS is designed to be repeated approximately every 5 years per country in order to construct a repeated cross-section for the countries surveyed. The DHS is designed to be a nationally representative survey of households and of women and men of reproductive age.[7] To represent this population, the DHS uses a two-stage sampling technique. In the first stage, the DHS uses stratification to sample clusters, or small enumeration areas, with probability proportional to size; a complete household listing is then conducted in each sampled cluster. In the second stage, the DHS randomly samples households from the household listings of each sampled cluster. For each selected household, a household questionnaire, in-

[6]Hjort and Poulsen construct this map by using AfTerFibre's 2013 map of backbone networks as a base. They then remove backbones built after the construction of the first submarine cable, by using www.africabandwidthmaps.com maps on new backbones in that timeframe.

[7]Only de facto residents - residents who stayed in the household the previous night - are included in the sample, as these are the only individuals to whom the DHS male and female questionnaires are administered. This is determined during the household roster when the respondent is asked, for each household member, "Did (name) stay here last night?" This sample may exclude individuals who usually live in the household but were away for some reason, and it may include some visitors to the household. Even so, only 1.89% of men and 2.57% of women in my sample are visitors.

cluding a household roster to identify eligible women and men of reproductive age, is conducted. Since a key aim of the DHS is to collect information on marriage, fertility, reproductive health, and child health, the DHS deliberately interviews approximately twice as many women as men. In each selected household, the women's questionnaire is administered to all women aged 15-49. In a random subsample of selected households, the men's questionnaire is also administered to all men aged 15-49.[8,9] The male and female questionnaires can be linked both to each other and to the household questionnaire, household roster, and geographical information of the household's cluster.

I use DHS measures of current employment[10] and employment over the past 12 months to measure employment rates, along with measures of employment seasonality, earnings type and employer type. All of these measures are self-reported. I also utilize measures of the sector of employment, which are defined by the DHS as follows: skilled manual, unskilled manual, professional, clerical, sales, self-employed

[8]Table A12 details the size of the male subsample relative to the female subsample in each country and wave within the sample. The choice of which households are interviewed for the male subsample is random: if half the women-eligible households are also men-eligible, every second household surveyed is considered eligible for the male subsample. In this sense, the selection of the male subsample is as random as the female subsample. It is also important to note that a household is designated as eligible for the male or female survey regardless of the gender of the household members. That is, just because a men-eligible household must also be women-eligible does not mean that men are only interviewed in households with women present. For instance, 11.24% of the male subsample are the only members of their household, 33.82% of surveyed men over the age of 25 are unmarried, and 18.37% of surveyed men over the age of 25 have never been married and have never cohabitated with a romantic partner. Men are less likely than women to respond to the DHS survey, although the difference in response rate varies by country, wave and whether an observation lives in an urban or rural area. Table A13 in the Appendix details the response rates of each DHS survey by gender. The major concern regarding the lower response rates for men is that men who are away from the house for work may be disproportionately undersampled. If this is the case, any employment results will be underestimated, and the magnitude by which they are underestimated will be larger for men than for women due to the greater rates of employment away from the house among men.

[9]Some observations over the age of 49 appear in the data. These are removed, as 15-49 is the dominant age restriction for women in the countries sampled.

[10]Current employment is defined as having worked in the past 7 days.

agricultural, employed agricultural, domestic, and service. Since the distinction between self-employed versus employed agriculture is not made in 5 of the 16 waves in my sample, I calculate my own measure of self-employed and employed agriculture using data on employer type.[11] All estimates involving this distinction use these calculated measure of agricultural employment. Agricultural work in the DHS includes farmers of both crops and livestock, fishermen, foresters, and hunters.

I also use measures of household decision-making power that are derived from DHS questions on decision-making patterns. For a number of household decisions, the DHS asks respondents who usually makes various household decisions. For women, the relevant household decisions are spending her earnings, spending her husband's earnings, her health care, large household purchases, and visits to family; for men, they are spending his earnings, his health care, and large household purchases.

In addition to this measure of household decision-making, I also use data on the acceptability of domestic violence against women. The DHS asks both men and women whether a beating is justified in the following situations: a wife goes out without telling her husband, a wife neglects children, a wife argues with her husband, a wife refuses sex with her husband, or a wife burns the food. Although self-reported, these measures provide some information on attitudes toward domestic violence in my sample. Although the DHS collects data on actual experiences with domestic violence, I choose to use data on the acceptability of domestic violence for two reasons. First, DHS only administers the domestic violence portion of its women's questionnaire to one woman per household, resulting in a small sample

[11] For my calculated measures, an individual is assigned to self-employed agriculture if the DHS assigns them to either self-employed or employed agricultural work and their employer is reported as "self". Alternatively, an individual is assigned to employed agriculture if the DHS assigns them to either self-employed or employed agricultural work and their employer is reported as "family" or "someone else" or their employer is missing.

size.[12] Second, under-reporting is extremely high for these measures,[13] a problem that is commonplace for sensitive topics in household surveys like the DHS, where the surveyor is not known to the respondent, the respondent might not understand the extent of confidentiality for the survey, and the household is not always free of other members at the time the survey is administered. This combination of small sample size and under-reporting makes it extremely difficult to detect changes in incidence of domestic violence due to fast Internet access, even if such changes do occur.

The sample includes 8 coastal countries that were connected to submarine cables from Europe, India, and United Arab Emirates from 2009-2012. Each country's national backbone was connected to a submarine cable at a documented point in time. The countries are chosen from the following criteria: the country must be connected to a submarine cable during 2009-2012, the country must be observed by the DHS prior to connection, and the country must be observed by the DHS after connection. The 8 countries are Benin, Democratic Republic of Congo, Ghana, Kenya, Namibia, Nigeria, Tanzania and Togo. Table 2.1 lists the years of DHS data used for each country, along with the date each country was connected to a submarine cable.[14]

[12]Only 22,371 women (24.7%) in my sample answer questions in the domestic violence module.

[13]Among the women who answer the domestic violence module, 20.4% report ever having been insulted by their current or last partner; 9.5% report ever having been threatened with harm by their current or last partner; 10.9% report ever having been pushed, shook or had something thrown at them by their current or last partner; 8.1% report ever having been punched with a fist or hit with something harmful by their current or last partner; and 6.9% report ever being forced into unwanted sex by their current or last partner. These rates are incredibly low, even compared to more developed countries with greater gender equality. For comparison, in the United States, 47% of women have ever experienced psychological abuse from a partner, 32% have experienced physical abuse from a partner, and 16% have experienced sexual abuse from a partner (The National Center for Victims of Crime 2018). These numbers are not directly comparable, as DHS asks women about abuse from their current partner, it is evident there is a significant amount of under-reporting occurring in my sample.

[14]Results are robust to excluding Tanzania and Togo, the two countries whose baseline wave occurs

Table 2.1: Countries and years of DHS data

Country	First Wave	Year Connected	Second Wave
Benin	2001	2010	2012
Democratic Republic of Congo	2007	2012	2013
Ghana	2008	2010	2014
Kenya	2008	2009	2014
Namibia	2006	2012	2013
Nigeria	2008	2010	2013
Tanzania	1999	2009	2010
Togo	1998	2010	2013

The sample is further restricted to observations within 10 kilometers of the backbone network. This is to ensure treatment and control groups are similar enough for meaningful comparison. The choice of 10km is based on prior literature (Hjort and Poulsen 2019). Section 2.5 discusses the sensitivity of the results to this restriction.

2.4 Method

In order to identify the causal impact of Internet access, I combine variation in the timing of connection to submarine cables across countries with geospatial variation within-country in access to terrestrial backbone cables to create a difference-in-difference estimator that allows for heterogeneous treatment effects by gender. This follows the method adopted by Hjort and Poulsen, but with changes to treatment variables and fixed effects allowing for gender-specific effects.

In order to determine if an observation has access to the backbone network, I calculate the distance between each observation and the nearest terrestrial cable using nearest neighbor matching in QGIS. Those who are located within 500m of a cable are considered "connected".[15] Observations further than 10km from a cable are ex-

in the 1990s.

[15]The DHS displaces the geographic location of its clusters in a random direction and random dis-

cluded from the sample, in order to allow for comparability between connected and unconnected observations. It is important to note that the terrestrial network used to calculate distances is that which existed before any submarine cables were introduced to the sample. The key identifying assumption to this difference-in-difference approach is that observations very close to the terrestrial network and slightly further away were on parallel trends with respect to outcome measures, and that the treatment and control groups did not experience systematically different shocks after the cables arrived.

Once I have determined which observations are "connected", I can then assign observations to the treatment group. An observation is considered "treated" if they are "connected" (i.e., within 500m of a terrestrial network) and appear after submarine cables have arrived. Gender-specific treatment groups are then created in order to identify the impact of fast Internet access by gender. It is important to note that the relevant treatment is access to fast-speed Internet, and the estimated effect includes changes both at the intensive margin - faster speeds for pre-existing users - and at the extensive margin - greater take-up of Internet usage in areas that were

tance, up to 2km for urban clusters and 5km for rural clusters, while ensuring the cluster remains within the correct administrative boundary; in addition, a randomly-selected 1% of rural clusters are displaced up to 10km. This displacement will incorrectly classify some control observations as treated and some treated observations as control. Since classification error in binary errors is mean-reverting (Bound et al. 2001), this will attenuate any treatment effect. To see this, consider a true model of $y^* = \beta^* x^* + \epsilon$, where $x^* \in (0, 1)$. The independent variable, x^*, is observed as x with some error μ, such that $x = x^* + \mu$. Here, the measurement error, μ, is a classification error, such that both $Pr(x = 1|x^* = 0) \geq 0$ and $Pr(x = 0|x^* = 1) \geq 0$. Thus, if $x^* = 1$, then $x - x^* \leq 0$; similarly, if $x^* = 0$, then $x - x^* \geq 0$. As such, it must be the case that variance $\sigma_{x^*,\mu} < 0$. Since displacement is conducted at the cluster level, to all clusters, in a random direction and random distance, we know that conditional on the truth, this classification error is independent of any outcome variables. Under the non-differential measurement error assumption, we thus know that $\beta = \beta^*[1 - Pr(x^* = 1|x = 0) - Pr(x^* = 0|x = 1)]$. As such, β will be biased toward 0. β will have the correct sign as long as $Pr(x^* = 1|x = 0) + Pr(x^* = 0|x = 1) < 1$. In order to ensure that this condition holds, I use the DHS clusters as given and recreate the displacement technique, excluding the restriction of remaining in the same administrative boundary. I can then estimate the probabilities of misclassification for my sample using repeated simulation. I estimate $Pr(x^* = 1|x = 0)$ as .108 and $Pr(x^* = 0|x = 1)$ as .491. As such, although β will be attenuated, it will carry the correct sign.

already connected to the national backbone. The treatment effect does not include expansions to the national backbone prompted by the arrival of submarine cables.

I regress treatment-by-gender on the relevant outcome measures, with a number of fixed effects. The empirical specification for the regressions is as follows:

$$y_{ij(i)c(i)t} = \beta SubmarineCables_{c(i)t} * Connected_i * Gender_i$$
$$+\delta_{j(i)} * Connected_i * Gender_i + \gamma_{c(i)t} * Gender_i + \epsilon_{ij(i)c(i)t}$$

where $y_{ij(i)c(i)t}$ is an outcome for individual i in 10km-by-10km grid-cell $j(i)$, country $c(i)$, and time period t. $SubmarineCables$ is equal to 1 if country $c(i)$ has been connected to at least one submarine cable at time t and 0 otherwise. $Connected$ is equal to 1 if individual i is within 500m of the terrestrial backbone and 0 otherwise. $Gender_i$ is an indicator variable indicating the gender of individual i. $\delta_{j(i)}$ is a 10km-by-10km grid-cell fixed effect. $\gamma_{c(i)t}$ is a country-by-year fixed effect. β is the coefficient of interest.

By interacting $SubmarineCables_{c(i)t}$ with $Connected_i$ and $Gender_i$, I identify a gender-specific treatment effect. By interacting the grid-cell fixed effect, $\delta_{j(i)}$, with $Connected_i$ and $Gender_i$, I am controlling for any time-invariant differences in outcome that may be connected to location, access to Internet, and the interaction of these with gender. By interacting the country-by-year fixed effects, $\gamma_{c(i)t}$ with gender, I am controlling for any within-country-location-invariant differences in outcome that may be correlated to getting fast Internet, and I am allowing this to be gender-specific. Standard errors are clustered at grid-cell, j_i.

Balance on total and sector-specific employment rates between unconnected and connected observations prior to the arrival of any submarine cables are presented in

Table 2.2. Column 1 shows the unconditional mean of the relevant variable among unconnected male observations before connection, while column 2 shows the difference for males in connected areas, conditional on grid-cell and country-by-year fixed effects. Columns 3 and 4 follow the same pattern, but for females.

The sample is largely balanced on employment rates. Men and women are slightly less likely to be employed in any type of agriculture in connected areas, and women are more likely to be employed in domestic and service work in connected areas. These differences are small and, for the most part, the sample appears balanced in terms of employment rates. It is important to note that complete balance on these baseline characteristics is not required in order to obtain unbiased estimates, as the analysis uses an approach that exploits the difference between connected and unconnected areas in the difference of outcomes over time.

Balance on employer type, seasonality of employment and earnings type is shown in Table 2.3. Again, column 1 shows the unconditional mean of each variable for men in unconnected areas prior to connection to any submarine cables, and column 2 shows the difference for men in connected areas, conditional on grid-cell and country-by-year fixed effects; columns 3 and 4 follow the same pattern, but for women. The sample is largely balanced over these employment measures, but there are some differences. Men are more likely to work year-round and less likely to work seasonally in connected areas, while women are less likely to be self-employed and less likely to be unpaid in connected areas. Again, complete balance is not required to obtain unbiased estimates with the empirical strategy employed in this paper.[16]

[16]Balance on demographics is shown in Table A14 in the appendix. Columns 1 and 3 show the unconditional mean for men and women, respectively, in unconnected areas prior to any connection to submarine cables; columns 2 and 4 show the difference for men and women, respectively, in connected areas, conditional on grid-cell and country-by-year fixed effects.

Table 2.2: Employment Rate Balance Table: Unconnected vs Connected at Baseline

Employment Rate	Male		Female	
	Unconnected Mean (st dev)	Difference (st err)	Unconnected Mean (st dev)	Difference (st err)
Total	0.751	-0.017	0.609	-0.002
	(0.432)	(0.017)	(0.488)	(0.012)
Manual, Skilled	0.170	-0.002	0.073	0.008
	(0.376)	(0.014)	(0.261)	(0.006)
Manual, Unskilled	0.062	-0.004	0.020	-0.003
	(0.241)	(0.013)	(0.141)	(0.005)
Professional	0.098	0.008	0.071	0.001
	(0.297)	(0.011)	(0.256)	(0.007)
Clerical	0.023	0.006	0.020	0.006*
	(0.149)	(0.006)	(0.142)	(0.004)
Sales	0.098	0.008	0.071	0.001
	(0.297)	(0.011)	(0.256)	(0.007)
Agricultural, Self-Employed	0.136	-0.021**	0.080	-0.014**
	(0.343)	(0.012)	(0.271)	(0.007)
Agricultural, Employed	0.098	-0.019***	0.048	-0.007***
	(0.297)	(0.009)	(0.214)	(0.003)
Domestic	0.005	0.002	0.009	0.009***
	(0.069)	(0.002)	(0.093)	(0.003)
Services	0.061	0.008	0.062	0.015***
	(0.239)	(0.008)	(0.241)	(0.007)

Table 2.3: Employer Type, Employment Seasonality, and Earnings Type Balance Table: Unconnected vs Connected at Baseline

	Men		Women	
	Unconnected Mean (st dev)	Difference (st err)	Unconnected Mean (st dev)	Difference (st err)
Works for Self	0.335	-0.007	0.398	-0.023**
	(0.472)	(0.017)	(0.489)	(0.012)
Works for Family	0.099	-0.015**	0.066	0.002
	(0.299)	(0.008)	(0.248)	(0.005)
Works for Someone Else	0.310	0.005	0.145	0.019**
	(0.463)	(0.018)	(0.352)	(0.010)
Works Year-Round	0.688	0.050***	0.738	-0.011
	(0.463)	(0.018)	(0.440)	(0.012)
Works Seasonally	0.244	-0.042***	0.192	0.004
	(0.430)	(0.017)	(0.394)	(0.010)
Works Occasionally	0.067	-0.008	0.071	0.007
	(0.251)	(0.011)	(0.256)	(0.007)
Unpaid	0.107	-0.004	0.077	-0.021***
	(0.309)	(0.007)	(0.267)	(0.006)
Earns Cash Only	0.505	0.005	0.403	0.024*
	(0.500)	(0.017)	(0.491)	(0.015)
Earns Cash and In-Kind	0.079	-0.014	0.052	0.005
	(0.269)	(0.010)	(0.223)	(0.005)
Earns In-Kind	0.019	-0.001	0.011	-0.002
	(0.137)	(0.004)	(0.102)	(0.002)

Table 2.4: Treatment Effect on Internet Speed

	Internet speed in kbps (asinh)	Internet speed in kbps (asinh)
Treatment	0.368**	0.514***
	(0.140)	(0.142)
Observations	1226.000	848.000
Clusters	66	51
Country x time FE	Yes	Yes
City FE	Yes	Yes
Includes Large Cities	Yes	No

2.5 Results

Prior to presenting this paper's findings, I first demonstrate that treatment, as defined in Section 2.4, represents an increase in speed for the countries in my sample. I use data on Internet speed provided by Hjort and Poulsen (2009), originally compiled from Akamai data, and I replicate their findings on Internet speed for the relevant countries. The results are shown in Table 2.4.[17] These findings demonstrate that treatment represents a substantial increase in Internet speeds for observations in my sample. As Hjort and Poulsen note, these estimates likely underestimate the increase in speeds caused by connection to submarine cables, as the data almost exclusively measures Internet speed between an African user and African server. Given that most traffic accessed by African Internet users connects to servers on other continents, the effective increase in speed for most uses is likely to be much higher.

To conceptualize what this increase in speed meant in terms of what content Africans could access via the Internet, I first present descriptive statistics of Internet speed, using Hjort and Poulsen's compiled Akamai data, for treated cities pre- and post-submarine cables.[18] These means are shown in Table 2.5. I then show the

[17] Only cities with more than 10 unique IPs are included, as originally specified by Hjort and Poulsen.

[18] This is not the same as my treatment group, as Akamai data comes from different locations than the DHS data.

Table 2.5: Mean Internet Speed for Treated Cities Pre- and Post-Treatment

	Pre-Treatment Mean	Post-Treatment Mean
With Large Cities	273.225	952.990
Without Large Cities	285.767	1059.147

Table 2.6: Minimum Internet Speed for Various Activities

Activity	Required KBPS
VoIP Calls	500
General Browsing	1000
Skype Video Call	1000
Streaming SD Video	3000

minimum broadband speeds required for various Internet uses, calculated as the minimum download speed needed for "adequate" performance of each application, in Table 2.6 (Federal Communications Commission 2019). Although these tables provide only rough estimates of what Africans could access via the Internet before and after connection to the submarine cables, it is clear that the increase in speed brought by these cables allowed Africans to access more content at faster rates.

2.5.1 Fast Internet Access and Employment

I find that the introduction of fast Internet access increases current male employment, with an insignificant effect on current female employment (Table 2.7).[19] The

[19]Current employment is ascertained differently for men and women. Men are asked, "Have you done any work in the last 7 days?" Men are defined as currently employed if they respond "Yes" to this question. Women are first asked, "Aside from your own housework, have you done any work in the last 7 days?" If the woman responds, "No", she is then asked, "As you know, some women take up jobs for which they are paid in cash or kind. Others sell things, have a small business or work on the family farm or in the family business. In the last 7 days, have you done any of these things or any other work?" Women are defined as currently employed if they respond "Yes" to either question. Given the difference in question by gender, there may be some concern that men are reporting housework in employment. I do not see this as a major limitation, as

Table 2.7: Treatment Effect on Employment Status

	Currently Working	Worked in Past 12 months
Male*Treatment	0.049**	0.048**
	(0.021)	(0.022)
Female*Treatment	0.021	0.028*
	(0.015)	(0.016)
Gender Difference	0.028	0.020
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117927	118472
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

treatment effect for men is substantial: a 4.9 percentage point increase off a 73.5% mean employment rate.[20] When I measure the impact of treatment on employment over the past twelve months, a similar effect size is found for men, while a smaller but significant effect size is found for women: 2.8 percentage points on a 65.3% mean.[21]

Access to fast Internet also has a substantial effect on whom men are working for: themselves, a family member, or someone else (Table 2.8).[22] Treatment increases male self-employment by 4.1 percentage points off a 23.4% mean, increases employment for family members by 4.6 percentage points off a 6.5% mean, and decreases employment

it seems unlikely that fast Internet access would change men's propensity to report housework as employment or affect men's involvement in housework enough to cause the treatment effect observed. In addition, the gains in employment are also observed in non-domestic sectors when men report their occupation later in the survey, suggesting that an increase in housework is not driving the treatment effect for men.

[20]Sharpened q-values to account for multiple testing are presented for all results in Appendix A.4.

[21]To measure employment in the past year, both men and women are asked "Have you done any work in the last 12 months?" This question is asked after both genders have been asked about current employment and women have been primed to include informal work in their response to current employment.

[22]Both men and women are asked, "Do you do this work for a member of your family, for someone else, or are you self-employed?"

Table 2.8: Treatment Effect on Employer Type

	Works for Family	Works for Someone Else	Works for Self
Male*Treatment	0.046***	-0.054**	0.041**
	(0.010)	(0.025)	(0.019)
Female*Treatment	0.005	-0.002	0.018
	(0.008)	(0.013)	(0.016)
Gender Difference	0.041	-0.052	0.023
Mean: Men	0.065	0.227	0.234
Mean: Women	0.060	0.150	0.407
Observations	101488	101488	101488
Clusters	1429	1429	1429
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

for someone else by 5.4 percentage points off a 22.7% mean. Treatment does not significantly alter for whom women are working.

In addition to employment rates and employer type, I also consider the impact of fast Internet on seasonality of employment and earnings type.[23] Treatment increases year-round employment for women by 5.5 percentage points on a 48.5% mean, but it does not have any significant impact on the seasonality of employment for men (Table 2.9). For both men and women, treatment has no effect on the type of earnings they receive: unpaid, cash only, in-kind only, or both cash and in-kind.[24]

For both genders, access to fast Internet has significant effects on sector of employment (Table 2.10).[25] The employment gains of treatment for men are concentrated in

[23]To determine seasonality of employment, both men and women are asked, "Do you usually work throughout the year, or do you work seasonally, or only once in a while?" To determine earnings type, both men and women are asked, "Are you paid in cash or kind for this work or are you not paid at all?"

[24]These results appear in Table A16 in the appendix.

[25]Both men and women are asked, "What is your occupation? That is, what kind of work do you mainly do?" The surveyor records the answer and then the occupation is grouped into one of

Table 2.9: Treatment Effect on Employment Seasonality

	Works Year-Round	Works Seasonally	Works Occasionally
Male*Treatment	-0.001	0.028	0.020
	(0.024)	(0.018)	(0.013)
Female*Treatment	0.055***	-0.015	-0.011
	(0.017)	(0.010)	(0.009)
Gender Difference	-0.056	0.043	0.031
Mean: Men	0.544	0.157	0.058
Mean: Women	0.485	0.117	0.051
Observations	117909	117909	117909
Clusters	1429	1429	1429
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

self-employed agriculture, where male employment increases by 7.2 percentage points off a 8.8% mean. For women, treatment increases both self-employed and employed agriculture, by 1.4 percentage points on a 6.9% mean and 1.1 percentage points on a 3.9% mean, respectively. In addition, fast Internet access increases skilled manual employment for women, by 1.4 percentage points on a 6.4% mean.

the following sectors: skilled manual, unskilled manual, professional, clerical, sales, self-employed agricultural, employed agricultural, domestic, and service. Since the distinction between self-employed and employed agricultural work is not made for every DHS wave in my sample, I construct my own self-employed and employed agriculture categories by grouping the DHS measures of self-employed and employed agriculture and then dividing the agricultural workers back into self-employed and employed based on their response on who they work for.

Table 2.10: Treatment Effect on Employment by Sector

	Manual, Skilled	Manual, Unskilled	Professional	Clerical	Sales	Agricultural, Self	Agricultural	Domestic	Services
Male*Treatment	0.001	0.005	-0.011	0.002	-0.011	0.072***	-0.018	0.005	-0.011
	(0.019)	(0.014)	(0.016)	(0.006)	(0.016)	(0.014)	(0.012)	(0.005)	(0.012)
Female*Treatment	0.014*	-0.000	0.002	-0.001	0.002	0.014**	0.011*	-0.000	-0.005
	(0.008)	(0.007)	(0.008)	(0.006)	(0.008)	(0.006)	(0.006)	(0.005)	(0.011)
Gender Difference	-0.013	0.005	-0.013	0.002	-0.013	0.058	-0.029	0.005	-0.006
Mean: Men	0.169	0.081	0.102	0.018	0.102	0.088	0.123	0.035	0.072
Mean: Women	0.064	0.027	0.065	0.019	0.065	0.069	0.039	0.025	0.081
Observations	117560	117560	117560	117560	117560	100640	117560	117560	117560
Clusters	1429	1429	1429	1429	1429	1429	1429	1429	1429
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Since the DHS interviews both women and men in households, it is possible to link couples who are in the same household and observe how this gender difference in treatment effect affects employment at the couple-level. To do this, I calculate the effect of fast Internet access on couple-level variables that indicate if both partners are employed, if neither partner is employed, if only the husband is employed, or if only the wife is employed.[26] The introduction of fast Internet reduces the likelihood that both partners in a couple are currently unemployed by 1.9 percentage points (Table 2.11). In addition, fast Internet access decreases the likelihood that both partners were unemployed over the past 12 months by 1.8 percentage points. As such, it appears that the impact on couples was to reduce the likelihood of couple-level unemployment. Although the point estimates are insignificant, likely due to the small couple sample size, it appears as though the gains in employment at the couple-level are to bring only the husband into employment: treatment increases the likelihood that only the husband is currently employed by 1.6 percentage points and also increases the likelihood that only the husband worked in the past year by 0.9 percentage points.

To better understand why a difference in treatment effect by gender is observed, I consider how education and household composition interact with access to fast Internet. As discussed in section 2.2, previous literature indicates that these two domains are important to fully understand why women may see different employment returns to fast Internet access. To study the interaction effects of education, I interact treatment with a number of education levels: no education, some primary, completed primary only, some secondary, completed secondary only, and higher education (Table 2.12).[27] These categories are mutually exclusive. In addition to the gender-by-

[26]These regressions use fixed effects that are not interacted with gender, as the fixed effects must be at the couple-level for this analysis.

[27]Educational attainment is a variable created from two questions, which are the same for both men

Table 2.11: Treatment Effect on Couple-Level Employment Status

	Both Currently Working	Neither Currently Working	Only Husband Currently Working	Only Wife Currently Working
Treatment	0.005	-0.019*	0.016	-0.002
	(0.044)	(0.010)	(0.043)	(0.014)
Mean	0.688	0.023	0.265	0.024
Observations	14851	14851	14851	14851
Clusters	1218	1218	1218	1218
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes
	Both Worked in Past 12 months	Neither Worked in Past 12 months	Only Husband Worked in Past 12 months	Only Wife Worked in Past 12 months
Treatment	0.005	-0.018**	0.009	0.004
	(0.040)	(0.008)	(0.039)	(0.007)
Mean	0.734	0.014	0.239	0.013
Observations	14911	14911	14911	14911
Clusters	1218	1218	1218	1218
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes

education-by-treatment interactions, I also include gender-by-education interactions, although these are suppressed in the table. Women with either some primary education or higher education experience positive employment gains from fast Internet access; men who have completed primary or secondary school or who have pursued higher education benefit most from the employment gains generated by fast Internet access.

In addition to education, I consider how household composition interacts with treatment. Single women experience a larger treatment effect on employment, than their married counterparts (Table 2.13).[28] In addition, women in households headed by a female see a positive treatment effect on employment status, while those in

and women: "What is the highest level of school you attended: primary, secondary, or higher?" and "What is the highest [grade/form/year] you completed at that level?"

[28]Men and women are asked, "Are you currently married or living together with a woman/man as if married?" to determine if the respondent is married or cohabitating. If the respondent answers "No", he or she is then asked "Have you ever been married or lived together with a woman/man as if married?" If the respondent answers "No", he or she is categorized as single. If the respondent answers "Yes", he or she is then asked "What is your marital status now: are you widowed, divorced, or separated?"

Table 2.12: Treatment Effect on Employment Status, with Interacted Education Levels

	Currently Working	Worked in Past 12 months
Male*No Education*Treatment	-0.024	0.025
	(0.042)	(0.038)
Female*No Education*Treatment	-0.003	0.003
	(0.023)	(0.024)
Male*Some Primary*Treatment	0.010	0.030
	(0.032)	(0.031)
Female*Some Primary*Treatment	0.057**	0.049**
	(0.023)	(0.022)
Male*Primary Only*Treatment	0.033	0.038
	(0.027)	(0.026)
Female*Primary Only*Treatment	0.040*	0.044*
	(0.024)	(0.023)
Male*Some Secondary*Treatment	0.052*	0.039
	(0.027)	(0.027)
Female*Some Secondary*Treatment	0.010	0.018
	(0.017)	(0.019)
Male*Secondary Only*Treatment	0.071***	0.065***
	(0.024)	(0.024)
Female*Secondary Only*Treatment	0.013	0.027
	(0.023)	(0.023)
Male*Higher*Treatment	0.059**	0.073***
	(0.029)	(0.027)
Female*Higher*Treatment	0.039	0.057**
	(0.028)	(0.027)
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117926	118471
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

Table 2.13: Treatment Effect on Employment Status, with Interacted Marital Status

	Currently Working	Worked in Past 12 months
Male*Treatment	0.042*	0.047**
	(0.022)	(0.023)
Female*Treatment	0.027	0.035**
	(0.016)	(0.017)
Male*Married	0.371***	0.355***
	(0.008)	(0.008)
Female*Married	0.233***	0.231***
	(0.007)	(0.007)
Male*Married*Treatment	-0.013	-0.025
	(0.019)	(0.018)
Female*Married*Treatment	-0.028*	-0.031*
	(0.016)	(0.017)
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117923	118468
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

male-headed household do not (Table 2.14).[29]

To further investigate the interaction of household composition with treatment,

[29]The name and sex of the household head is determined during the household survey when the household roster is created. The first question of the household survey is, "Please give me the names of the persons who usually live in your household and guests of the household who stayed here last night, starting with the head of the household." This is used to identify the name of the household head. The surveyor then asks, "Is (name) male or female?" to determine the gender of the household head. Although there is evidence that the DHS once had a definition of household head as "the person considered responsible for the household. This person may be appointed on the basis of age (older), sex (generally, but not necessarily male), economic status (main provider) or some other reason. It is up to the respondent to define who is the head" in the 1990s (Kishor and Neitzel 1996), this definition no longer appears in any DHS documentation nor in any of the relevant questionnaires for my sample, including those conducted in the late 1990s. Thus, it appears the respondent for the household survey chooses the household head without any guidance, allowing the choice to vary with local culture, local social norms, and individual preference.

Table 2.14: Treatment Effect on Employment Status, with Interacted Head of Household Gender

	Currently Working	Worked in Past 12 months
Male*Treatment	0.050**	0.050**
	(0.021)	(0.021)
Female*Treatment	0.006	0.015
	(0.017)	(0.017)
Male*Female HoH	-0.224***	-0.208***
	(0.009)	(0.009)
Female*Female HoH	0.019***	0.021***
	(0.005)	(0.005)
Male*Female HoH*Treatment	-0.034	-0.036
	(0.028)	(0.027)
Female*Female HoH*Treatment	0.035***	0.033**
	(0.013)	(0.013)
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117927	118472
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

I consider the role of household size and the number of children under the age of 5 in the household.[30] Women in single-member households experience a large positive effect of fast Internet access on employment status, and this effect diminishes as household size decreases (Table 2.15). Furthermore, women in households with 0 children experience a positive effect of fast Internet on employment, but this effect decreases with each additional child in the household (Table 2.16).

The mechanism by which fast Internet access increases employment in my sample is of significant economic interest. However, identifying and testing the various mechanisms that might explain the observed treatment effect is neither possible with the data I use nor the main purpose of the paper. Given this, I briefly discuss a potential mechanism and suggestive evidence supporting this mechanism, but I cannot exclude other possibilities.

Smallholder agriculture[31] makes up a large share of all agricultural farms in the countries within my sample. Around 80% of farms in these countries are smallholder (Food and Agriculture Organization of the United Nations 2017)[32], and the median size of agricultural land plots owned by households in my sample is only 6 hectares. Table 2.17 presents detailed descriptive statistics for both self-employed

[30]Household size is the number of household members listed in the household roster in response to the main prompt, "Please give me the names of the persons who usually live in your household and guests of the household who stayed here last night, starting with the head of the household," and 3 additional prompts: "Just to make sure that I have a complete listing: are there any other people such as small children or infants that we have not listed?", "Are there any other people who may not be members of your family, such as domestic servants, lodgers, or friends who usually live here?", and "Are there any guests or temporary visitors staying here, or anyone else who stayed here last night who have not been listed?" The number of children under 5 in the household is calculated from the answers to these same questions, but only children who are de jure residents, or usually live in the household, are included. This is determined for every child with the question, "Does (name) usually live here"?

[31]Smallholder agriculture consists of small-scale farmers, fishers, and other agricultural workers who manage areas of land that are smaller in size than the median national landholding, which is generally less than 10 hectares (Food and Agriculture Organization of the United Nations 2017).

[32]See Appendix Table A15 for detailed statistics on average farm size and percent of farms that are smallholder by country, for four countries in my sample: Ghana, Kenya, Nigeria and Tanzania.

Table 2.15: Treatment Effect on Employment Status, with Interacted Household Size

	Currently Working	Worked in Past 12 months
Male*Treatment	0.068***	0.057**
	(0.026)	(0.027)
Female*Treatment	0.057***	0.069***
	(0.018)	(0.019)
Male*HH Size	-0.022***	-0.021***
	(0.001)	(0.001)
Female*HH Size	-0.008***	-0.007***
	(0.001)	(0.001)
Male*HH Size*Treatment	-0.005	-0.002
	(0.004)	(0.004)
Female*HH Size*Treatment	-0.007***	-0.008***
	(0.002)	(0.002)
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117927	118472
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

Table 2.16: Treatment Effect on Employment Status, with Interacted Number of Children in HH

	Currently Working	Worked in Past 12 months
Male*Treatment	0.058***	0.062***
	(0.022)	(0.023)
Female*Treatment	0.035**	0.046***
	(0.016)	(0.017)
Male*Num Children in HH	0.038***	0.037***
	(0.003)	(0.003)
Female*Num Children in HH	0.011***	0.013***
	(0.003)	(0.003)
Male*Num Children in HH*Treatment	-0.014	-0.020**
	(0.010)	(0.010)
Female*Num Children in HH*Treatment	-0.017**	-0.021***
	(0.007)	(0.008)
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117927	118472
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

Table 2.17: Descriptive Statistics: Agriculture

	Self-Employed Agriculture	Employed Agriculture
Rural	0.852	0.828
Works for Family	0.000	0.754
Works for Someone Else	0.000	0.246
Works Year-Round	0.526	0.438
Works Seasonally	0.446	0.496
Works Occasionally	0.028	0.065
Works but is Currently Absent	0.308	0.231
Unpaid	0.351	0.435
Earns Cash Only	0.280	0.341
Earns In-Kind Only	0.061	0.060
Earns Cash and In-Kind	0.308	0.164
HH Owns Agricultural Land	0.841	0.765

and employed agricultural workers. A large percentage of agricultural workers live in households that own their own agricultural land: 84.1% and 76.5% for self-employed and employed agricultural workers, respectively. Employed agricultural workers are employed by family members 75.4% of the time. In addition, both self-employed and employed agricultural workers have relatively high rates of unpaid employment and in-kind payments: 35.1% and 43.5%, respectively, are unpaid, while 36.9% and 22.4%, respectively, receive at least some of their earnings in-kind.

Given this information, it is likely that the gains in self-employed agricultural employment due to fast Internet access are concentrated in entrance into smallholder agriculture. One potential mechanism for how fast Internet access increases smallholder agriculture is through an increase in information and a decrease in search costs that both ultimately increase agricultural productivity for smallholder agricultural workers. Previous literature has demonstrated that ICT in the form of mobile phones can substantially increase the productivity of smallholder farms when used to disseminate agricultural advice via SMS (Casaburi et al. 2014). By making agricultural

advice more accessible to farmers, fast Internet access likely increases agricultural productivity in a similar way. Internet access may also provide greater productivity gains than those seen with SMS, due to Internet users' ability to search for information using search engines and thereby tailor the agricultural information they view to their specific crops, location, season, and tools.

Casaburi et al. (2014) also note that the gains in productivity associated with SMS dissemination of agricultural advice are greatest for workers with little agricultural training. Similarly, fast Internet may be most influential for agricultural workers with little experience, thereby lowering the start-up costs associated with entering self-employment. Simply by connecting agricultural workers with agricultural information, the Internet likely increases productivity of smallholder farms and decreases the start-up costs of entering agriculture, incentivizing individuals to enter self-employed agriculture.

In both grain markets in Niger and fish markets in South India, ICT in the form of mobile phones decreases search costs associated with going to market and decreases price dispersion, improving both producer and consumer welfare (Aker 2010, Jensen 2007). It is likely that, insofar as Internet access allows agricultural workers and wholesalers to better search across markets and coordinate sales, fast Internet will have similar effects on search costs and market integration as mobile phones. Furthermore, if Internet allows individuals to conduct more targeted searches through the use of search engines, the revenue gains for agriculture might be amplified. In addition, if Internet access allows individuals to obtain information they could not obtain through their pre-established social networks, or if Internet access increases the size of users' social networks, it is possible that the Internet improves the ability of new self-employed agricultural workers to sell their goods, increasing the profitability of self-employed agriculture and incentivizing entrance into the sector.

2.5.2 Fast Internet and Household Bargaining

Previous literature suggests that any change in the relative earning power of household members also alters the relative bargaining power and, thus, decision-making power of these individuals. In addition, there is evidence that new technologies with the ability to disseminate new information may change individual attitudes and social norms and ultimately affect decision-making roles in the household, as demonstrated by the impact of cable television and soap operas on increasing women's autonomy and decreasing fertility (La Ferrara et al. 2012, Jensen and Oster 2009).

Given that fast Internet access increases male employment by a much greater degree than female employment, we may expect to see an average increase in male bargaining power and, as such, an increase in male decision-making.[33] However, we may expect that, by exposing individuals to new information and content, attitudes surrounding women's autonomy may change and increase women's decision-making power.

To see the actual impact of Internet access on household decision-making, I first estimate the treatment effect on female decision-making power, as perceived and reported by women themselves.[34] Fast Internet access increases women's sole decision-

[33]Since the DHS does not collect information on earnings or income outside of whether an individual is paid in cash, in kind, both in cash and in kind, or is unpaid, I make the assumption that any movement from unemployment to employment represents an increase in the value of time. Although the magnitude of this increase is not identified, we can assume it is positive and non-zero, as an individual would not move into employment if their time was better spent unemployed. This assumption relies on the value of unemployment not increasing significantly over time, which is not unreasonable given that unemployment benefits are minimal in all the relevant countries and all but one of the sampled countries kept laws on unemployment benefits constant over the sample timeframe (Social Security: Research, Statistics and Policy Analysis 2019). The only country to change any such law during the sample period was Namibia, which passed a new Labour Act in 2007; however, the only relevant change in this law affected severance pay, rather than unemployment insurance (Social Security: Research, Statistics and Policy Analysis 2019).

[34]Women's decision-making power is measured using questions on who usually makes various household decisions. To determine if the woman has any decision-making power in spending her own earnings or her husband's/partner's earnings, the woman is asked, "Who usually decides how (the money you earn/your (husband's/partner's) earnings) will be used: you, your (husband/partner),

Table 2.18: Treatment Effect on Female Sole Decision-Making Power, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.009	0.017	0.004	0.035*	0.047**
	(0.025)	(0.017)	(0.026)	(0.021)	(0.022)
Mean: Women	0.687	0.084	0.230	0.152	0.192
Observations	31266	36607	43583	43577	43581
Clusters	1225	1211	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

making power over major household purchases and visits to family and relatives by 3.5 percentage points on a 15.2% mean and 4.7 percentage points on a 67.9% mean, respectively (Table 2.18).[35] When we consider whether a woman has any say over various household decisions, we see a greater impact of fast Internet access on decision-making power over visits to family, but no significant effect on major household purchases (2.19). This suggests that, although women are more likely to partake in both sole and joint decision-making over family visits with treatment, there is not an increase in joint decision-making for major household purchases.

The increase in decision-making power over visits to family attributable to fast

or you and your (husband/partner) jointly?" To determine if the woman has any decision-making power in her own health care, major household purchases and visits to family, the woman is asked, "Who usually makes decisions about (health care for yourself/making major household purchases/visits to your family or relatives)?" When asking about health care decision-making, the question specifies 4 replies after the question, "you, your (husband/partner), you and your (husband/partner) jointly, or someone else"; although these options are not repeated after the questions on major household purchases and family visits, these questions appear directly after the question on health care, so the woman knows these are the relevant options. For all five domains of household decision-making, I define a woman as having some say if she says the decision is usually made by herself or herself and her husband/partner jointly; I define a woman as having sole say if she says the decision is usually made by herself.

[35] All decision-making results are presented using both a summary index and sharpened q-values in Appendix A.4, to account for multiple testing.

Table 2.19: Treatment Effect on Female Decision-Making Power, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.016	0.005	-0.004	0.027	0.069***
	(0.011)	(0.032)	(0.028)	(0.027)	(0.023)
Mean: Women	0.930	0.441	0.618	0.577	0.679
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Internet access is seen for women of all education levels (Table 2.20). This regression includes female-by-education interaction effects, although these are suppressed in the table. In addition, fast Internet access increases uneducated women's perceived ability to take part in the decision of how to spend her husband's earnings. Internet access also decreases a woman's decision-making power over spending her own earnings if she has completed primary school only; however, only the significance, not the direction of this effect, is unique to women of this education level.

When the interaction effects of household composition and treatment are considered, we see that unmarried women benefit more than their married counterparts from the increase in decision-making power regarding visits to family attributable to fast Internet access (Table 2.21). In addition, treatment is associated with an increase in female decision-making regarding spending her own earnings for unmarried women. These results are similar for women's sole decision-making power over these decisions.[36] The head of household's gender, the size of the household, and the number of children in the household do not affect the impact of fast Internet access on decision-making power.[37]

[36]This result is presented in Appendix Table A17

[37]These results are presented in the Appendix. Tables A18, A20, and A22 show how gender of the

Table 2.20: Treatment Effect on Female Decision-Making Power, with Interacted Education Level, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*No Education*Treatment	-0.029	0.092**	0.032	0.047	0.090***
	(0.022)	(0.040)	(0.033)	(0.036)	(0.032)
Female*Some Primary*Treatment	0.004	0.011	0.037	0.049	0.081***
	(0.015)	(0.040)	(0.031)	(0.035)	(0.028)
Female*Primary Only*Treatment	-0.031*	0.031	-0.004	-0.000	0.074**
	(0.018)	(0.036)	(0.032)	(0.031)	(0.031)
Female*Some Secondary*Treatment	-0.011	-0.014	-0.026	0.020	0.062**
	(0.015)	(0.035)	(0.030)	(0.029)	(0.026)
Female*Secondary Only*Treatment	-0.024	0.013	-0.023	0.029	0.055*
	(0.020)	(0.036)	(0.033)	(0.032)	(0.029)
Female*Higher*Treatment	-0.015	-0.010	-0.006	0.024	0.071**
	(0.019)	(0.042)	(0.032)	(0.034)	(0.032)
Mean: Women	0.930	0.441	0.618	0.577	0.679
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table 2.21: Treatment Effect on Female Decision-Making Power, with Interacted Marital Status, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	0.026*	0.039	0.007	0.044	0.120***
	(0.014)	(0.042)	(0.032)	(0.035)	(0.031)
Female*Married	0.018***	0.060***	0.051***	0.061***	0.061***
	(0.006)	(0.013)	(0.010)	(0.012)	(0.014)
Female*Married*Treatment	-0.055***	-0.043	-0.017	-0.024	-0.068**
	(0.018)	(0.033)	(0.025)	(0.031)	(0.032)
Mean: Women	0.930	0.441	0.618	0.577	0.679
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

To better understand how fast Internet access impacts decision-making within the household, I also consider the impact of Internet access on female decision-making power, according to men.[38] Access to fast Internet decreases female sole-decision making power over her husband's or partner's earnings, according to men, by 3.9 percentage points on a 5.1% mean (Table 2.22); however, the impact of fast Internet on whether a woman has any say in spending her husband's or partner's earnings is positive but insignificant, indicating that men may be more likely to make joint decisions over their earnings. In addition, fast Internet access decreases women's sole decision-making power over her husband's or partner's health care, according to men, by 3.4 percentage points on a 6.8% mean. There is no impact of fast Internet on women's decision-making power over major household purchases, according to men, although this could be due to a change in how the question was asked between earlier and later waves of the surveys used.[39]

To understand how these changes in reported female decision-making power affect the household, I estimate the treatment effect of fast Internet access on couple-level beliefs about women's power over household decisions. However, there are only two domains of household decisions for which men and women are both surveyed: major household purchases and spending the husband's earnings. The introduction of fast

household head, household size, and number of children in the household, respectively, impact the treatment effect on any decision-making power; Tables A19, A21, and A23 show how gender of household head, household size, and number of children in the household, respectively, impact the treatment effect on sole decision-making power.

[38]To measure men's perception of household decision-making power regarding his own earnings and health care, the DHS asks men, "Who usually (decides how the money you earn will be used/makes decisions about health care for yourself): you, your (wife/partner), or you and your (wife/partner) jointly?" Women are defined as having some decision making power according to the man if he states the decision is made jointly; similarly, women are defined as having sole decision-making power according to the man if he states the decision is made by his wife/partner alone.

[39]In early waves, men were asked "In a couple, who do you think should have the greater say in each of the following decisions: the husband, the wife or both equally: making large household purchases?" In later waves, however, men were asked "Who usually makes decisions about making major household purchases?"

Table 2.22: Treatment Effect on Female Sole Decision-Making Power, According to Men

	Spending His Earnings	His Health Care	Major HH Purchases
Male*Treatment	-0.039**	-0.034***	-0.015
	(0.016)	(0.013)	(0.031)
Mean: Men	0.051	0.068	0.096
Observations	12979	11265	21379
Clusters	1098	1000	1264
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

Table 2.23: Treatment Effect on Female Decision-Making Power, According to Men

	Spending His Earnings	His Health Care	Major HH Purchases
Male*Treatment	0.071	-0.011	0.013
	(0.058)	(0.046)	(0.039)
Mean: Men	0.414	0.351	0.469
Observations	12979	11265	21379
Clusters	1098	1000	1264
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

Table 2.24: Treatment Effect on Couple-Level Reports of Female Sole Decision–Making Power

	Both Report Wife Alone Decides Major HH Purchase	Neither Report Wife Alone Decides Major HH Purchase	Only Husband Reports Wife Alone Decides Major HH Purchase	Only Wife Reports Wife Alone Decides Major HH Purchase
Treatment	0.019*	-0.025	-0.026	0.031
	(0.010)	(0.043)	(0.032)	(0.030)
Mean	0.017	0.772	0.099	0.112
Observations	14096	14096	14096	14096
Clusters	1172	1172	1172	1172
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes
	Both Report Wife Alone Decides Spending Husband's Earnings	Neither Report Wife Alone Decides Spending Husband's Earnings	Only Husband Reports Wife Alone Decides Spending Husband's Earnings	Only Wife Reports Wife Alone Decides Spending Husband's Spending Earnings
Treatment	0.002	0.026	-0.040**	0.013
	(0.012)	(0.031)	(0.018)	(0.024)
Mean	0.006	0.879	0.049	0.067
Observations	11215	11215	11215	11215
Clusters	982	982	982	982
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes

Internet increases the likelihood that both members of a couple report that the woman is the sole decision-maker over large household purchases by 1.9 percentage points (Table 2.24). This suggests that treatment increases the likelihood that couples agree that the wife has sole decision-making power over this domain. Although the point estimate is not significant, perhaps due to small sample size, it appears as though treatment also increases the number of women who report being the sole decision-maker over major household purchases, even though their husband disagrees.

In addition, access to fast Internet reduces the likelihood that only the husband reports that the wife has sole decision-making power over spending his earnings. It also appears that treatment increases the likelihood that neither member of a couple report the women being the sole decision-maker in this domain, although this estimate is not significant. These results suggest that access to fast Internet decreases women's sole decision-making power over her husband's earnings, at least in the eyes of the husband.

Table 2.25: Treatment Effect on Couple-Level Reports of Female Decision-Making Power

	Both Say Wife Decides Major HH Purchase	Neither Say Wife Decides Major HH Purchase	Only Husband Says Wife Decides Major HH Purchase	Only Wife Says Wife Decides Major HH Purchase
Treatment	0.025	0.013	-0.018	-0.020
	(0.041)	(0.034)	(0.033)	(0.040)
Mean	0.337	0.248	0.155	0.261
Observations	14096	14096	14096	14096
Clusters	1172	1172	1172	1172
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes
	Both Report Wife Decides Spending Husband's Earnings	Neither Report Wife Decides Spending Husband's Earnings	Only Husband Reports Wife Decides Spending Husband's Earnings	Only Wife Reports Wife Decides Spending Husband's Spending Earnings
Treatment	0.057	-0.042	0.022	-0.038
	(0.061)	(0.044)	(0.041)	(0.041)
Mean	0.251	0.361	0.160	0.228
Observations	11138	11138	11138	11138
Clusters	976	976	976	976
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes

The introduction of fast Internet access also seems to increase the likelihood that both members of a couple report the wife as having any, rather than sole, decision-making power over major household purchases; however, this effect is not significant (Table 2.25). This effect, however, may not be uniform, as treatment also seems to increase the likelihood that neither member of a couple report the wife as having any decision-making power over this domain; again, however, this point estimate is not significant. In addition, treatment seems to increase the likelihood that couples agree that the wife has any decision-making power over spending her husband's earnings, although this effect is not significant. This seems to be driven by an increase in the likelihood that husbands report women as having some decision-making power in this domain.

The seeming increase in women's decision-making power may seem counter to the expectation that male bargaining power will increase if fast Internet access boosts male employment without affecting female employment. However, it is also important

to note that the gains in decision-making which are reported by women are in the domains of family visits and household purchases. It may be the case that these two domains of household decision-making are delegated to women when men spend less time in the household. This explanation is consistent with the decrease in female sole-decision-making, according to men, over male earnings and health care that is attributable to fast Internet access.

In addition to household decision-making roles, I also consider the impact of fast Internet access on the acceptability of domestic violence against women.[40] Previous literature presents the theory that, although men either gain no inherent enjoyment from domestic violence or actively dislike it, they may still engage in the behavior in order to secure control of household resources (Heath 2014, Haushofer et al. 2019). This theory suggests that when male bargaining power increases, domestic violence will decrease, as men have less incentive to resort to violence to control household resource allocation. As such, any differential impacts by gender of fast Internet access on employment may change domestic violence incidence through changes in household bargaining. Furthermore, past literature indicates that technologies that communicate new information and attitudes, such as television, can change opinions on domestic violence by exposing individuals to new ideas (Jensen and Oster 2009). As such, insofar as fast Internet exposes individuals to new ideas at a higher rate, we

[40]The DHS also includes data on the incidence of domestic violence. The sample size of this data is small, as the DHS only asks a select number of women about domestic violence and women can refuse to answer. In addition, the responses to these questions indicate that under-reporting is extremely common. This is not surprising, given that the survey is conducted by a stranger and that other household members are sometimes present during the domestic violence module. I have conducted analysis on the incidence data, and I find no treatment effect. Although it is possible that there is no actual treatment effect on domestic violence, it is also likely that any such treatment effect would not be detected due to the small sample sizes and tendency of women to severely under-report abuse. I choose to instead look at opinions on domestic violence. Although these are interesting in and of themselves, I also believe that opinions on domestic violence might move in line with actual experiences. That is, in environments when women may fear negative repercussions from reporting abuse, women may feel more comfortable reporting opinions on domestic violence, and these opinions reflect their own experiences with domestic violence.

may expect fast Internet to decrease the acceptability of domestic violence.

Fast Internet access increases the acceptability of domestic violence against women among both men and women (Table 2.26).[41][42] Fast Internet increases the acceptability of domestic violence to men if a wife neglects her children or refuses sex, by 4.6 percentage points on a mean of 18.6% and 2 percentage points on a 7.5% mean, respectively.[43] Fast Internet also increases the acceptability of domestic violence to women if a wife goes out without telling her husband or argues with her husband, by 3.8 percentage points on a 20.3% mean and 5.1 percentage points on a 19.3% mean, respectively.

To consider how education levels affect treatment effects, I consider the interaction of education with fast Internet access (Table 2.27). This regression includes gender-by-education interaction terms, although these are suppressed in the table. The increase in acceptability of domestic violence attributable to fast Internet is seen among women of all education levels, except women with no education. In addition, the impact of fast Internet access on the acceptability of domestic violence is strongest for men with no education, men who completed primary education only, and men with higher education. Interestingly, the treatment effects on acceptability of domestic violence are strongest among men with higher education.

[41] In order to obtain individuals', men's, and women's opinions on domestic violence, DHS asks, "In your opinion, is a husband justified in hitting or beating his wife in the following situations: if she goes out without telling him? If she neglects the children? If she argues with him? If she refuses to have sex with him? If she burns the food?" Importantly, the DHS frames the question about a hypothetical husband and wife, rather than asking about how the respondent would act in these situations. If the respondent responds "Yes" they are categorized as accepting of domestic violence in that situation.

[42] All results on approval of domestic violence are presented using both a summary index and sharpened q-values in Appendix A.4, to account for multiple testing.

[43] On average, men are less likely than women to respond that it is acceptable for a husband to beat his wife. This may be because men are less likely to report support of domestic violence for fear that the surveyor will assume they engage in this behavior themselves. For women, this fear is likely less influential, as there is less social stigma attached to being the victim of domestic abuse.

Table 2.26: Treatment Effect on Approval of Domestic Violence

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*Treatment	0.033	0.046*	0.026	0.020*	0.017
	(0.021)	(0.025)	(0.018)	(0.011)	(0.010)
Female*Treatment	0.038**	0.016	0.051***	0.012	0.010
	(0.019)	(0.022)	(0.019)	(0.014)	(0.013)
Gender Difference	-0.005	0.030	-0.025	0.009	0.007
Mean: Men	0.147	0.186	0.140	0.075	0.054
Mean: Women	0.203	0.252	0.193	0.131	0.096
Observations	105190	105359	105169	104606	105562
Clusters	1367	1367	1367	1367	1367
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

The impact of fast Internet on acceptability of domestic violence is larger for men in small households (Table 2.29), while it is larger for women in households with fewer children (Table 2.30). Marital status and household head's gender do not have significant interactions with fast Internet access when considering the impact on domestic violence acceptability.[44]

In order to see how these changes affect couples, I also estimate the impact of the introduction of fast Internet on couple-level approval of domestic violence in each of the 5 scenarios posed in the DHS. That is, I estimate the treatment effect on couple-level variables that indicate whether both partners, neither partner, the husband only, or the wife only believe a husband is justified in beating his wife in each situation.

Access to fast Internet decreases the likelihood that neither member of a couple approve of domestic violence if a wife goes out without telling her husband by 6.6 percentage points (Table 2.28). A large portion of this effect is driven by increase in husbands who approve of domestic violence in response to this scenario, even though

[44]The results are reported in Tables A24 and A24 for marital status and gender of household head, respectively, in the Appendix.

Table 2.27: Treatment Effect on Approval of Domestic Violence, with Interacted Education Level

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*No Education*Treatment	0.037	0.044	0.022	0.053*	-0.012
	(0.034)	(0.036)	(0.032)	(0.029)	(0.020)
Female*No Education*Treatment	0.013	-0.004	0.021	0.006	-0.016
	(0.030)	(0.029)	(0.027)	(0.022)	(0.019)
Male*Some Primary*Treatment	0.029	0.027	0.014	0.004	0.001
	(0.028)	(0.034)	(0.027)	(0.017)	(0.017)
Female*Some Primary*Treatment	0.046*	0.030	0.044*	0.002	0.008
	(0.025)	(0.026)	(0.025)	(0.019)	(0.017)
Male*Primary Only*Treatment	0.024	0.035	0.047*	0.019	0.009
	(0.027)	(0.033)	(0.028)	(0.018)	(0.015)
Female*Primary Only*Treatment	0.056**	0.022	0.049**	0.023	0.013
	(0.025)	(0.027)	(0.022)	(0.018)	(0.018)
Male*Some Secondary*Treatment	0.022	0.026	0.021	0.008	0.018
	(0.024)	(0.026)	(0.020)	(0.011)	(0.012)
Female*Some Secondary*Treatment	0.032*	0.008	0.052***	0.011	0.010
	(0.019)	(0.022)	(0.019)	(0.015)	(0.014)
Male*Secondary Only*Treatment	0.035	0.050	0.013	0.017	0.016
	(0.023)	(0.032)	(0.022)	(0.016)	(0.012)
Female*Secondary Only*Treatment	0.044**	0.026	0.064***	0.012	0.010
	(0.021)	(0.022)	(0.022)	(0.015)	(0.014)
Male*Higher*Treatment	0.036*	0.072***	0.028	0.040***	0.024*
	(0.021)	(0.025)	(0.019)	(0.015)	(0.013)
Female*Higher*Treatment	0.037*	0.005	0.048**	0.009	0.024*
	(0.021)	(0.026)	(0.020)	(0.017)	(0.013)
Mean: Men	0.147	0.186	0.140	0.075	0.054
Mean: Women	0.203	0.252	0.193	0.131	0.096
Observations	105190	105359	105169	104606	105562
Clusters	1367	1367	1367	1367	1367
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

their wives do not, by 4.1 percentage points. The remaining portion is driven by an increase in couples who agree domestic violence is justified in this situation by 2.1 percentage points.

Treatment also decreases the likelihood that neither member of a couple approves of domestic violence if the wife neglects the children; the magnitude of this effect is 7 percentage points. This is almost entirely driven by an increase of 5.8 percentage points in the likelihood that men approve of domestic violence in this situation, even though their wives do not. Similarly, the introduction of fast Internet increased the likelihood that only the husband approves of domestic violence if the wife argues with her husband by 4.3 percentage points. This increase comes primarily from a decrease in couples that both disapprove of domestic violence in this scenario; however, this decrease is not statistically significant.

The introduction of fast Internet also decreased the likelihood that neither member of a couple approves of domestic violence if the wife refuses sex by 6.3 percentage points. However, this effect is driven both by an increase in the likelihood that husbands alone report approval of domestic violence in this situation, by 2.5 percentage points, and by an increase in the likelihood that wives alone approve of domestic violence in this scenario, by 4.7 percentage points. The only scenario in which treatment did not impact couple-level approval of domestic violence was if the wife burned food.

Taken together, these couple-level results suggest that access to fast Internet caused a significant shift in the within-couple dynamic surrounding domestic violence. In particular, in all but one domain, fast Internet access increased the likelihood that a husband approved of domestic violence when his wife did not. This suggests that, although the treatment effect on approval of domestic violence was positive for both men and women individually, the effect within couples was primarily a decrease in joint disapproval of domestic violence and an increase in one-sided approval of the

practice on the part of husbands.

Given previous literature, it is surprising that fast Internet access is associated with increased acceptability of domestic violence among both men and women. Again, while the mechanism driving this effect is of economic interest, identifying and testing the various mechanisms that might explain the treatment effect is not possible with the data I use, nor is it within the purview of this paper. I will, however, briefly discuss a potential mechanism that might explain the treatment effect observed.

If opinions on domestic violence respond in the same direction to changes in incidence of domestic violence, the effect of fast Internet likely reflects an increase in physical domestic violence against women. To understand why this may occur, consider the theory of domestic violence set forth in the literature whereby a man either dislikes or is indifferent to domestic violence, but uses it as a tool to obtain control over household resources. In a household with few household resources, it may be the case that the value of domestic violence, in terms of the additional household resources secured, is small. As household income grows, however, the amount of household resources that can be extracted through domestic violence grows, and the relative value of engaging in the behavior increases. In this case, as men bring more resources into a household, they may be more likely to engage in physical abuse in order to secure a greater share of the additional wealth, even though their bargaining power increases.

Such conflict may be exacerbated if the increase in household resources coincides with an increase in female household decision-making power due to men spending less time in the household. If, for instance, women gain more control over quotidian household decisions as men spend more time working, but men also have a greater stake in controlling household resources due to their increased income, a power struggle may emerge. This might explain why women in the treatment group report an

Table 2.28: Treatment Effect on Couple-Level Approval of Domestic Violence

	Both Approve Beating if Wife Goes Out w/o Telling Husband	Neither Approve Beating if Wife Goes Out w/o Telling Husband	Only Husband Approves Beating if Wife Goes Out w/o Telling Husband	Only Wife Approves Beating if Wife Goes Out w/o Telling Husband
Treatment	0.021*	-0.066*	0.041**	0.004
	(0.011)	(0.037)	(0.020)	(0.028)
Mean	0.047	0.698	0.089	0.166
Observations	13859	13859	13859	13859
Clusters	1164	1164	1164	1164
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes
	Both Approve Beating if Wife Neglects Children	Neither Approve Beating if Wife Neglects Children	Only Husband Approves Beating if Wife Neglects Children	Only Wife Approves Beating if Wife Neglects Children
Treatment	0.009	-0.070*	0.058**	0.002
	(0.017)	(0.041)	(0.024)	(0.030)
Mean	0.064	0.634	0.104	0.198
Observations	13902	13902	13902	13902
Clusters	1168	1168	1168	1168
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes
	Both Approve Beating if Wife Argues w/ Husband	Neither Approve Beating if Wife Argues w/ Husband	Only Husband Approves Beating if Wife Argues w/ Husband	Only Wife Approves Beating if Wife Argues w/ Husband
Treatment	0.002	-0.050	0.043*	0.006
	(0.016)	(0.034)	(0.023)	(0.023)
Mean	0.040	0.710	0.083	0.167
Observations	13875	13875	13875	13875
Clusters	1167	1167	1167	1167
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes
	Both Approve Beating if Wife Refuses Sex	Neither Approve Beating if Wife Refuses Sex	Only Husband Approves Beating if Wife Refuses Sex	Only Wife Approves Beating if Wife Refuses Sex
Treatment	-0.008	-0.063**	0.025**	0.047**
	(0.007)	(0.028)	(0.011)	(0.023)
Mean	0.021	0.806	0.047	0.126
Observations	13902	13902	13902	13902
Clusters	1168	1168	1168	1168
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes
	Both Approve Beating if Wife Burns Food	Neither Approve Beating if Wife Burns Food	Only Husband Approves Beating if Wife Burns Food	Only Wife Approves Beating if Wife Burns Food
Treatment	-0.002	0.001	0.001	-0.000
	(0.007)	(0.023)	(0.014)	(0.020)
Mean	0.011	0.872	0.035	0.082
Observations	13973	13973	13973	13973
Clusters	1170	1170	1170	1170
Country x time FE	Yes	Yes	Yes	Yes
Grid-cell x connected FE	Yes	Yes	Yes	Yes

Table 2.29: Treatment Effect on Approval of Domestic Violence, with Interacted Household Size

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*Treatment	0.060***	0.066***	0.025	0.031**	0.015
	(0.020)	(0.025)	(0.020)	(0.013)	(0.011)
Female*Treatment	0.046**	0.026	0.063***	0.021	0.014
	(0.021)	(0.023)	(0.020)	(0.015)	(0.014)
Male*HH Size	0.003***	0.003***	0.001**	0.001*	0.001*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Female*HH Size	0.001**	0.002***	0.002***	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)
Male*HH Size*Treatment	-0.006***	-0.004**	0.000	-0.002*	0.000
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Female*HH Size*Treatment	-0.001	-0.002	-0.002	-0.002	-0.001
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Mean: Men	0.147	0.186	0.140	0.075	0.054
Mean: Women	0.203	0.252	0.193	0.131	0.096
Observations	105190	105359	105169	104606	105562
Clusters	1367	1367	1367	1367	1367
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table 2.30: Treatment Effect on Approval of Domestic Violence, with Interacted Number of Children in HH

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*Treatment	0.036*	0.046*	0.022	0.023**	0.015
	(0.020)	(0.026)	(0.019)	(0.011)	(0.010)
Female*Treatment	0.047**	0.025	0.060***	0.018	0.016
	(0.019)	(0.022)	(0.019)	(0.014)	(0.013)
Male*Num Children in HH	0.004*	0.003	0.000	-0.002	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Female*Num Children in HH	0.011***	0.011***	0.012***	0.007***	0.003***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)
Male*Num Children in HH*Treatment	-0.005	-0.000	0.006	-0.003	0.002
	(0.006)	(0.007)	(0.006)	(0.004)	(0.004)
Female*Num Children in HH*Treatment	-0.010**	-0.010**	-0.010**	-0.008*	-0.007**
	(0.004)	(0.005)	(0.005)	(0.004)	(0.003)
Mean: Men	0.147	0.186	0.140	0.075	0.054
Mean: Women	0.203	0.252	0.193	0.131	0.096
Observations	105190	105359	105169	104606	105562
Clusters	1367	1367	1367	1367	1367
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table 2.31: Treatment Effect on Access to Electricity

	Access to Electricity	Access to Roads
Male*Treatment	0.001	-0.000
	(0.001)	(0.001)
Female*Treatment	-0.000	-0.001
	(0.001)	(0.001)
Gender Difference	0.001	0.000
Mean: Men	0.103	0.013
Mean: Women	0.108	0.013
Observations	127967	127967
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

increase in decision-making power, at the same time as domestic violence becomes more acceptable in treatment areas.

2.5.3 Robustness

A primary concern when analyzing the impact of Internet access on employment status is that the analysis is instead capturing time trends in industrialization that coincide with connection to submarine cables. To check this is not the case, I first calculate the impact of fast Internet access on access to electricity. I use data from Hjort and Poulsen (2019) on electricity lines and road networks and designate an observation as having access to electricity if they are within 500m of an electricity cable and having access to roads if they are within 500m of a road. There is no treatment effect of Internet access on access to electricity or roads (Table 2.31), suggesting that a simple correlation of treatment and industrialization is not driving the results.

I then construct two placebo tests. First, I define a placebo electricity access treatment that considers an individual as treated if he or she lives within 500m of an electricity cable and lives in a country after it was first connected to a submarine cable. I calculate the treatment effect of both fast Internet access and the electricity placebo on employment status, within the same regression. If the effect of fast Internet access on employment is driven by a time trend in industrialization that coincided with the introduction of submarine cables, we would expect the coefficient on the treatment of fast Internet to attenuate. However, including the electricity placebo does not change the treatment effect (Table 2.32), suggesting that the treatment effect is not driven by industrialization.

Second, I define a placebo road access treatment that considers an individual as treated if he or she lives within 500m of a road and lives in a country after it was first connected to a submarine cable. I then calculate the treatment effect of fast Internet access, the electricity placebo, and the road placebo on employment status, in the same regression. Including the road placebo slightly increases female employment, but the change is small and the gender gap between men and women remains similar in magnitude (Table 2.33). Again, this suggests that the treatment effect of Internet access on employment is not driven by industrialization or other infrastructure expansion.

I next consider the sensitivity of my analysis to the 500m rule used to determine if an individual has access to the terrestrial Internet cables or not.[45] Figure 2.2 shows how the treatment effect on current employment is affected by the distance used to define whether an observation is "connected"; Panel A shows the sensitivity of the result for the male treatment effect, while Panel B shows the sensitivity of the result for the female treatment effect. Although the effect size and significance of the

[45]This sensitivity exercise holds the 10km sample restriction constant.

Table 2.32: Placebo Check: Electricity and Road Access

	Currently Working	Worked in Past 12 months
Male*Treatment	0.049**	0.049**
	(0.022)	(0.022)
Female*Treatment	0.018	0.027*
	(0.015)	(0.016)
Male*Electricity Placebo Treatment	-0.003	-0.005
	(0.013)	(0.013)
Female*Electricity Placebo Treatment	0.013	0.010
	(0.009)	(0.009)
Gender Difference	0.031	0.022
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117927	118472
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

treatment effect on male employment varies, the magnitude of the effect size remains large and positive no matter the distance used to define whether an observation is "connected" to a terrestrial backbone. Similarly, although the effect size of the treatment effect on female employment varies as the distance used to measure access to the backbone network changes, the magnitude of the effect size is always small and close to zero.

I next consider the sensitivity of my analysis to the 10km rule used to determine if an individual is included within the sample or not.[46] Figure 2.3 shows how the treatment effect on current employment is affected by the distance used to define whether an observation is within the sample; Panel A shows the sensitivity of the result for the male treatment effect, while Panel B shows the sensitivity of the result for the female treatment effect. The results are robust to expanding and shrinking

[46]This sensitivity analysis holds the definition of "connected" to the terrestrial cables fixed at 500.

Table 2.33: Placebo Check: Electricity and Road Access

	Currently Working	Currently Working	Worked in Past 12 months	Worked in Past 12 months
Male*Treatment	0.049**	0.052**	0.049**	0.052**
	(0.022)	(0.022)	(0.022)	(0.023)
Female*Treatment	0.018	0.024	0.027*	0.032**
	(0.015)	(0.015)	(0.016)	(0.016)
Male*Electricity Placebo Treatment	-0.003	-0.003	-0.005	-0.005
	(0.013)	(0.013)	(0.013)	(0.013)
Female*Electricity Placebo Treatment	0.013	0.014	0.010	0.011
	(0.009)	(0.009)	(0.009)	(0.009)
Male*Road Placebo Treatment		-0.009		-0.009
		(0.011)		(0.010)
Female*Road Placebo Treatment		-0.018**		-0.017**
		(0.008)		(0.009)
Gender Difference	0.031	0.028	0.022	0.020
Mean: Men	0.735	0.735	0.762	0.762
Mean: Women	0.618	0.618	0.653	0.653
Observations	117927	117927	118472	118472
Clusters	1429	1429	1429	1429
Country x time x gender FE	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes

Figure 2.2: Sensitivity to 500m Connected Rule

Panel A: Treatment Effect on Male Employment

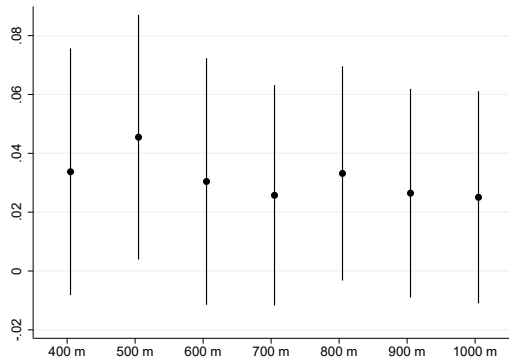

Panel B: Treatment Effect on Female Employment

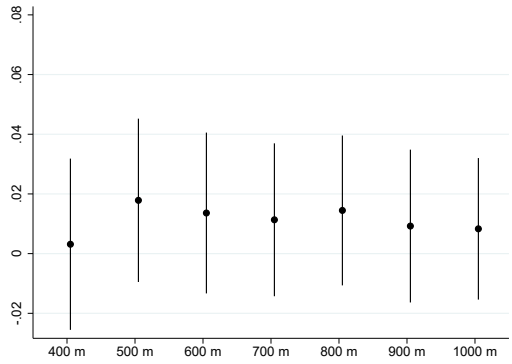

Figure 2.3: Sensitivity to 10km In-Sample Rule

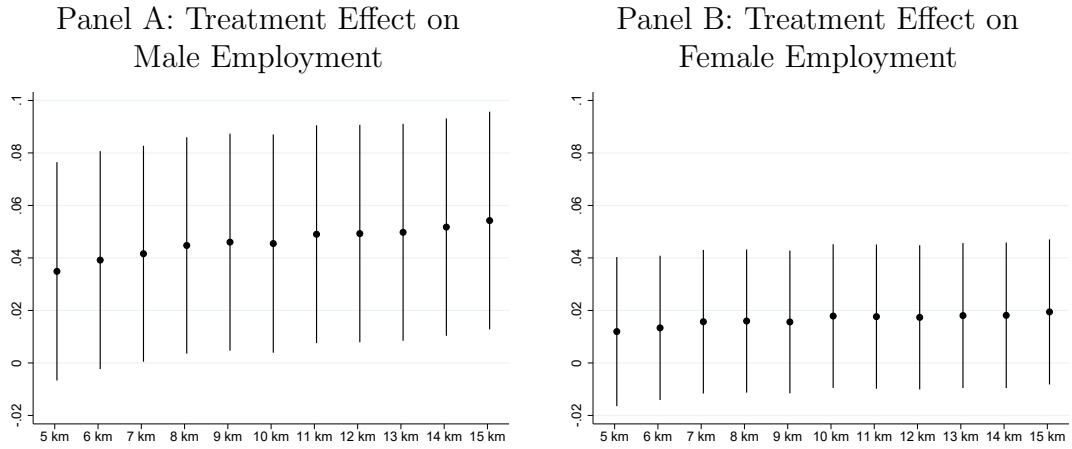

the sample restriction.

2.6 Conclusion

In this paper, I consider the impact of fast Internet access on employment and household bargaining in sub-Saharan Africa. I find that fast Internet has a large positive effect on current male employment, while it has a smaller and statistically insignificant impact on current female employment. In addition, fast Internet access increases both male and female employment over the past 12 months, although the treatment effect is much larger for males. I consider the impact of education and household composition on this gender difference, and I find that these factors partially explain the differential treatment effect by gender. I also find that the gains in male employment are primarily within self-employed agriculture, and that fast Internet access shifts men into self-employment and employment for family, and out of work for someone else. Internet access also increases female employment in self-employed and employed agricultural work. These employment effects are consistent with a mechanism

whereby Internet access increases access to agricultural information, decreases search costs associated with selling agricultural outputs, and increases producer welfare in the agricultural sector through increased market integration.

In addition to these findings, I also consider the impact of fast Internet access on household decision-making roles and opinions on domestic violence. I find that Internet access increases women's perceived decision-making power in regards to visits to family and major household purchases but decreases women's sole decision-making power, according to men, over their husband's or partner's earnings and health care. These effects may be explained by men increasingly delegating decisions regarding household functioning to women as their employment, and thus time spent outside the household, increases. In addition, I find that fast Internet access increases the acceptability of domestic violence against women among both men and women in a variety of situations. This treatment effect is consistent with a model where men dislike or are indifferent to domestic violence but use it to gain control of household resources. As male employment increases, the amount of household resources increases, and the value of using domestic violence, in terms of additional household resources secured, also increases; as such, domestic violence may increase as male employment increases, even if female employment remains constant.

The findings of this paper suggest that the expansion of Internet and other ICTs in developing countries may disproportionately benefit male employment over female employment. Insofar as employment gains from Internet access are concentrated among men, the spread of Internet technology may widen the gap between male and female economic empowerment in these settings. As demonstrated in this paper, differential employment effects spurred by Internet access may also have further repercussions on household decision-making roles and domestic violence against women in developing countries. Importantly, advances in access to this technology are not al-

ways associated with advances in female empowerment and instead actually increase the reported acceptability of domestic violence against women among both men and women. As such, policy-makers should take care to ensure that any advances in Internet access in developing contexts are accompanied by efforts to increase female employment, female economic empowerment, and female autonomy within society and the household.

There are several limitations to this study that should be noted. First, the DHS does not include any measures of income or productivity. Data on income would be the ideal outcome to consider both to assess the overall economic gains of individuals who gained access to fast Internet and to better understand the mechanism underlying the changes in household decision-making and acceptability of domestic violence. In addition, data on productivity would help shed light on whether the gains from fast Internet access to self-employed agriculture are due to the suggested mechanism. Future work could include geospatial data measuring agricultural productivity through satellite images to estimate the impact of fast Internet access on agricultural productivity.

Second, the DHS does not have any data on Internet use in the waves used for this analysis. As such, I cannot estimate the direct impact of proximity to terrestrial networks on Internet use for my sample. In addition, I cannot calculate the treatment effect of fast Internet access on Internet use or frequency of Internet use. Future work could further incorporate later waves of the DHS which include self-reported measures of whether an individual uses the Internet and the frequency of this use.

Chapter 3

Skill Certification and Labor Market Outcomes: Differential Effects by Gender

3.1 Introduction

Labor market interventions are an important part of economic policy, and interventions that assess and certify skills are a potentially impactful and cost-effective way to improve labor market outcomes for underemployed and unemployed participants. These programs improve labor market outcomes by addressing limited information about workseeker skills, which is faced by both workseekers and firms during the employment matching process. When firms hire in the face of limited information about workseeker skills, it can lead to poor match quality, creating distortions that result in reduced employment and reduced earnings conditional on employment (Aigner and Cain 1977; Altonji and Pierret 2001; Arcidiacono et al. 2010; Farber and Gibbons 1996; Kahn and Lange 2014; Pallais 2014). When workseekers conduct job search in the face of such limited information, it can lead to poorly targeted job search and uninformed job acceptance decisions, also resulting in reduced employment and earnings conditional on employment (Belot et al. 2019; Conlon et al. 2018). Labor market interventions that provide credible information about workseekers' skills ease these information frictions and improve labor market outcomes for participants

[0]I would like to thank Eliana Carranza, Robert Garlick, Kate Orkin, Neil Rankin, and Harambee Youth Employment Accelerator for their generosity in sharing the data used in this paper. In addition, I would like to thank Robert Garlick for providing additional feedback on the contents of this paper.

(Abebe et al. 2019; Abel et al. 2020; Bassi and Nansamba 2019; Carranza et al. 2020; Pallais 2014). Given the high potential of such interventions, it is important to understand whether and how these interventions affect men and women differently in order to develop a comprehensive understanding of how such interventions affect the labor market as a whole. This paper explores the gender differential effects of skill certification using an intervention conducted in South Africa.

A primary reason why skill certification may affect men and women differently is due to gender differences in baseline characteristics that are related to the skills being certified or the productivity of workseekers. In certain settings, researchers have measured gender differences in various cognitive and non-cognitive skill levels, which predict occupational choices and may affect worker productivity (Baker and Cornelson 2008, Dickerson et al. 2015, Fryer and Levitt 2010, Heckman et al. 2006, Hedges and Nowell 1995, Jacob 2002). If skill assessment reveals gender differences in skill levels, and skill certification is more beneficial to individuals of a certain skill, we would expect to see a gender difference in the impact of these interventions. Similarly, in certain settings, there is evidence of gender differences in quantity and timing of labor market experience, often due to gender differences in home production and workforce attachment (Albanesi and Olivetti 2009, Blau and Kahn 2017, Erosa et al. 2017, Light and Ureta 1995). Experience plays a role in revealing ability of workseekers to firms, and firms thus encounter more information frictions when considering inexperienced workers, leading to lower employment and lower wages for inexperienced workers (Pallais 2014). Skill certification is thus most valuable to inexperienced workers, since certification serves as an alternative way to communicate ability to potential employers (Carranza et al. 2020). Therefore, if there is a gender gap in experience, we may see a gender difference in how these interventions affect labor market outcomes.

In many settings, researchers have also found gender differences in levels of formal education (Pekkarinen 2012, van Broekhuizen and Spaull 2017). Formal education signals ability of workseekers to potential employers, suggesting that firms face greater information frictions when considering workers with only a secondary education (Alfonsi et al. 2017; Arcidiacono et al. 2010; Farber and Gibbons 1996). As such, skill certification is most valuable to workers without post-secondary education qualifications, since it provides an alternate way for workseekers to communicate ability (Carranza et al. 2020). As such, if there is a gender difference in rates of completion of higher education, we would expect to see a gender difference in how the intervention affects labor market outcomes. Finally, there is evidence that, in certain settings, women may engage in fewer career-enhancing actions, potentially due to concerns over marriage market consequences (Bursztyn et al. 2017). If skill certification serves as a substitute for such career-enhancing actions in communicating ability to potential employers, we would expect to see a greater impact of the intervention on women.

A second reason why skill certification interventions may affect men and women differently is due to gender differences in the way individuals search for jobs and the jobs they are willing to accept. Referral networks are influential in hiring decisions due to their role in decreasing information frictions, and workseekers with referrals are more likely to be hired than, experience an initial wage advantage over, and have longer tenure than workseekers without a referral (Brown et al. 2016; Heath 2018; Pallais and Sands 2016). Women are often excluded from referral networks, particularly in firms and industries with high proportions of male incumbent employees, potentially leading firms to face higher information frictions when considering female applicants (Beaman et al. 2018; Zeltzer 2020; Zhu 2019). If skill assessment and certification works to ease these information frictions, we may expect to see a greater

impact of these interventions on women.

There is also evidence that men and women differ in the confidence they have regarding their own ability. In certain settings, men demonstrate more overconfidence about their ability than women (Buser et al. 2014; Cho 2017; Niederle and Vesterlund 2007); however, within the intervention analyzed in this paper, women are actually more overconfident about their skills than men. Skill assessment reveals true skill levels to workseekers, allowing them to better target job search to jobs for which their skills are better matched (Ahn et al. 2020). Since female participants exhibit greater overconfidence than their male counterparts, and skill assessment is useful to workseekers due to its ability to correct beliefs and recalibrate search efforts, we may expect to see a greater impact on the intervention on women. Furthermore, there is evidence that men and women differ in how they update their beliefs about their own ability in response to new information, as well as in what instances and about which abilities they seek information (Möbius et al. 2011; Roberts and Nolen-Hoeksema 1989). As such, we may expect to see a gender difference in the impact of a skills assessment intervention if one gender is more receptive to updating their beliefs and changing search behavior in response to the information conveyed in the assessment. Similarly, if willingness to use certificates from the skill assessment varies with underlying confidence, it is possible a gender difference in treatment effect will appear due solely to differences in utilization of the certificate.

Both men and women demonstrate greater willingness-to-pay for jobs with tasks typically associated with their own gender, in large part due to past experience with these tasks (Gelblum 2020). In addition, men and women have different willingness-to-pay for a number of job attributes, including flexibility, stability, teamwork, earnings growth potential, and competition (Bursztyn et al. 2017; Croson and Gneezy 2010; Flory et al. 2015; Niederle and Vesterlund 2007; Wiswall and Zafar 2017).

If skill assessment and certification is more informative or valuable for employers requiring certain job tasks or offering certain job attributes, we might expect that occupational sorting by gender on these preferences might translate into gender differences in labor market effects of the intervention. Finally, it may be the case that women who take part in the intervention face statistical discrimination. Such statistical discrimination arises when firms have incomplete information about productivity and, as such, use easily observable characteristics, such as gender, to inform their expectations about productivity and, as such, their employment and wage offers (Aigner and Cain 1977; Arrow 1973; Altonji and Pierret 2001; Phelps 1972). If gender-based statistical discrimination exists in the market and skill certification reduces information frictions, potential employers may respond by relying more on certified skills and less on statistical discrimination, potentially generating a larger impact of the intervention for women.

In this paper, I consider whether labor market interventions that certify skills affect men and women differently. To do this, I use data from an intervention in South Africa that measures workseeker skills and provides workseekers with a certificate to credibly communicate these skills to potential employers. This intervention was designed and implemented by Carranza et al. (2020). The intervention assessed and certified the skills of 6,891 young unemployed or underemployed workseekers in urban South Africa, using a randomized control trial. The individuals in the sample have limited post-secondary education levels, limited work experience, and limited access to referral networks. Information frictions are an important feature of the labor market in general, and are even more relevant for the sample population. The skill assessment measures skills that are relevant in many workplaces, such as communication, numeracy, grit, and concept formation. The skill certification process provides participants with both electronic and physical certificates that give the in-

dividual's scores on the assessment and describe the assessments used. In addition, the certificates are branded by the local agency that conducts the assessments and the World Bank, in order to bolster their credibility. The intervention improved the overall labor market outcomes of participants, increasing employment rate by 17%, weekly earnings by 34%, and hourly wages by 20% (Carranza et al. 2020). These effects reflected workseeker use of the certificates in job applications to provide information to potential employers, as well as updates in workseekers' beliefs and search behavior in response to the information conveyed by the assessments.

First, I find that the effect of skill certification on labor market outcomes differs by gender. Although skill certification improves employment and earnings for both men and women, the effect is larger for men. While treatment increases employment for men by 6.9 percentage points (20%) and employment for women by 4.3 percentage points (15%); in addition, treatment increases male earnings by 51% and female earnings by 24%. Although the gender difference in these treatment effects is not statistically significant, the magnitude of the difference is appreciable. I also decompose the gender-specific treatment effect of earnings into an extensive margin effect, which represents the increase in earnings due to an increase in employment, and an intensive margin effect, which represents the increase in earnings among already employed individuals. Although the effect of certification on earnings occurs at both the extensive and intensive margin for men, it occurs only at the extensive margin for women. This suggests that, unlike for men, the earnings effect for women is driven entirely by an increase in the job-finding rate, rather than better firm-worker matches. In addition, although skill certification increases the rate of written contracts, Statistics South Africa's definition of a formal job, by 5 percentage points (38%) for men, it has no effect on the rate of written contracts for women.

Second, I find that although men and women update their beliefs similarly in

response to certification, men are more likely to use certificates in job applications. Certification causes both men and women to update their beliefs about their own abilities, increasing accuracy of these beliefs to rise by 14 percentage points for both men (38%) and women (43%). Although men and women update their beliefs in response to certification in the same way, certification decreases the percentage of men who have high self-esteem by 5 percentage points (8%) and increases the percentage of women who have high self-esteem by 3 percentage points (6%). Reflecting these changes in beliefs, certification increases the likelihood that both men and women target their search to match their strongest skills. Interestingly, while both men and women change their application behavior in response to certification and receive more interviews and offers in response to certification, men are both more likely to use the certificates in applications and more likely to receive interviews and offers when using the certificates.

Third, I find that the differential effect of certification on employment is almost fully explained by baseline gender differences in characteristics, while the differential effect on earnings is largely unexplained and the differential effect on written contracts is almost fully unexplained by these differences. To determine this, I decompose the gender difference in treatment effect using a Blinder-Oaxaca decomposition. This decomposition separates the gender differential into a portion that is explained by gender differences in baseline observed characteristics and a portion that is unexplained. The gender difference in how skill certification affects employment and hours is almost entirely explained by baseline characteristics, particularly numeracy skill, concept formation skill, and overconfidence. This suggests that treatment has differential employment effects for men and women due to underlying differences between the two groups, rather than differential responses; as such, if men and women had identical characteristics, we would expect skill certification to have similar ef-

fects on employment and hours for both groups. Less than a third of the gender difference in treatment effect on earnings and hourly wage, however, is explained by observed characteristics; and virtually none of the gender difference in how certification affects rates of written contracts is explained by these characteristics. These large unexplained effects may be indicative of discrimination or of gender differences in unobserved characteristics, such as preferences.

The primary contribution of this paper is twofold. First, I demonstrate that, while skill certification causes similar effects on belief updating and search targeting across genders, the effect of skill certification on employment, earnings, and written contracts is larger for men than women. Second, I demonstrate that the differential treatment effect on employment is explained not by differential response to certification but by baseline characteristics, and the differential treatment effects on earnings and written contracts are largely unexplained by these baseline characteristics. This adds to the existing literature on how labor market interventions designed to reduce relevant information frictions affect men and women differently. Abel et al. (2020) find that providing workseekers with reference letters from former employers doubled the employment rate of women with no impact on male employment, indicating that interventions may affect each gender differently due to gender differences in firm-side information frictions. Hicks et al. (2016) find that providing accurate information on the comparative economic gain of male- versus female-dominated occupations to workseekers caused women to increase enrollment in male-dominated occupational training but had no such impact for men, indicating that such interventions may result in differential effects by gender due to gender differences in workseeker-side information frictions. However, these papers do not consider the impact of an evaluation that addresses limited information on both the workseeker- and firm-side, nor do they decompose the difference in treatment effect to determine if it can be

attributed to baseline differences in characteristics.

This paper also adds to the broader literature on how general labor market interventions affect men and women differently. A number of papers have found that interventions designed to directly increase skills through technical training and internships, as well as programs that combine language training with labor market coaching and technical training, have larger employment effects for men than women (Andersson Joona and Nekby 2009; Cho et al. 2013; Walter et al. 2014); in addition, interventions that provide entrepreneurship training increase the number of start-ups created by men more than by women (Premand et al. 2016). However, other papers find that programs which combine technical training, soft-skills training and internships improved women's employment and earnings more than men's (Attanasio et al. 2011; Attanasio et al. 2015; González-Velosa et al. 2012); Ibarrarán and Rosas-Shady 2009). In a meta-analysis of over 200 active labor market program (ALMP) evaluations, Card et al. (2018) find larger effect sizes of ALMPs on women's labor market outcomes, but also find that this gender difference is driven by the upper tail of effect sizes for female participants. Although some of these interventions do reduce limited information on both the workseeker- and firm-side, none of these papers isolate this effect; instead, the interventions are designed to first increase skills correlated with worker productivity. This paper is, to the best of my knowledge, the first to consider and isolate the gender-specific impacts of information provision on both sides of the labor market.

Section 3.2 provides context on the local labor market and sample population. Section 3.3 describes the intervention and presents the empirical specification. Section 3.4 details the effects of skill certification on labor market outcomes by gender, the effects of skill certification on beliefs and job search, and the decomposition of the gender difference in treatment effects on labor market outcomes into explained and

unexplained portions. Section 3.5 concludes.

3.2 Economic Environment

3.2.1 Labor Market Context

The intervention takes place in Johannesburg, South Africa, a major metropolitan area. It is likely that firm-workseeker matching in this setting is affected by information frictions. Employers likely have difficulty inferring skills from secondary school performance and completion, as both progression and performance in primary and secondary schools are weakly correlated with measured skills (Lam et al. 2011; Taylor et al. 2011; van der Berg and Shepherd 2015). In addition, although workseekers typically report their post-secondary exit examination grades on job applications, these grades are not strong predictors of post-secondary performance, and firms report that these grades reveal only limited information about actual skills (Schöer et al. 2010). As such, it is likely that the intervention will reveal information to both workseekers and their potential employers, both of whom receive limited information about actual skills from prior school experience.

In addition, this setting is one in which such information is valuable to firms, as the cost of poorly matched hires is high. Labor regulations make firing an employee a complex and lengthy process that can be challenged in court, and business owners report such labor regulation as a constraint to growth (ILO 2016). Furthermore, firms are often ill-informed about the specifics of labor regulations and, as such, may perceive the costs of separation to be even higher; this is evidenced by the fact that providing consultation on labor regulation to these firms increases hiring rates (Bertrand and Crepón 2019). Although probationary periods for workers are allowed,

such labor is regulated and is limited in duration (Bhorat and Cheadle 2009).

Furthermore, the setting is one in which reservation and minimum wages exist, giving room for the intervention to affect employment rates. Compliance with minimum wage is high in South Africa's formal sector (Bhorat et al. 2016; ILO 2016). In addition, workseekers have reservation wages that are raised by several factors, including high commuting costs (Kerr 2017) and access to income from non-labor market activities, such as through pension payments to other members of the household (Abel 2019).

The setting of the intervention is characterized by high unemployment and significant gender differences in labor market outcomes. At the time of the intervention, overall unemployment was 28%, unemployment for youths aged 15-24 was 51%, and unemployment was 32% for individuals aged 25-34. In addition, there is also both a historical and current gender gap in relevant labor market outcomes in this setting (Rospabé 2001; Shepherd 2008). At the time of this intervention, women were 4.5 percentage points more likely than men to be unemployed, 0.8 percentage points less likely than men to have a written contract, and 5.4 percentage points less likely than men to have a job that contributed to a pension or retirement fund (Nishimwe-Niymbanira and Sabela 2019). There was also a significant difference in the types of employment in which men and women were engaged, with women more likely to work as a domestic worker, clerk, or technician and men more likely to work as a tradesman or plant/machine operator (Nishimwe-Niymbanira and Sabela 2019). Furthermore, women often face a disadvantage in conveying ability to potential employers, since women are often excluded from the referral networks upon which a large number of South African firms rely during the hiring process (Abel et al. 2020).

3.2.2 Sample Recruitment and Data Collection

This paper uses the data collected by and intervention conducted by Carranza et al. (2020). Carranza et al. conducted the intervention using a sample of 6,891 young, active workseekers. These workseekers were from low-income backgrounds and had limited work experience, university education and access to referral networks. As such, the individuals in the sample lacked the traditional ways to learn about their skills and credibly communicate these skills to potential employers.

Carranza et al. recruited the sample in collaboration with the Harambee Youth Employment Accelerator, a social enterprise that "builds solutions to address a mismatch of demand and supply in the South African youth labor market by connecting employers with inexperienced workseekers." The recruitment process used radio and social media advertising, as well as door-to-door recruitment in low-income areas of Johannesburg. Candidates register with the organization online and complete a phone-based questionnaire to determine eligibility.[1] Eligible candidates are invited to take part in 2 days of standardized skill assessments.[2] The final sample consists of all individuals who attend the second day of the two-day assessment, on 84 operational days.

These workseekers take part in 6 assessments designed to measure their skills in 6 domains: communication, concept formation, focus, grit, numeracy, and planning. Harambee has previously used these assessments to screen approximately 160,000 prospective workers for client firms, and firms have expressed interest in using the

[1]The eligibility requirements to work with Harambee are as follows: aged 18-19; has legal permission to work in South Africa; has completed secondary school; has 12 months or less of formal work experience; has no criminal record; and is from a disadvantaged background. This information is self-reported but verified using administrative data for some candidates.

[2]Less than 2% of candidates are later invited to job readiness training based on their assessment results and location, but Carranza et al. (2020) demonstrate that this is not relevant to the intervention effects.

assessment results during hiring. The assessments are led by 2-3 industrial psychologists, who manage a team of facilitators. The assessments are computer-based skill tests that are taken at Harambee and are self-administered but supervised by the facilitators. The assessments are conducted in English.

In addition to the assessment results, Carranza et al. conduct 3 surveys to measure participants' labor market outcomes, search behavior, and beliefs about their skills and the labor market. The baseline survey is a computer-based questionnaire that is taken at Harambee and is self-administered but supervised. This survey is administered after candidates have completed their skill assessments, but before they have learned the results. The endline survey is a phone-based survey that is administered 3-4 months after treatment. The endline has a response rate of 96%, resulting in an endline sample of 6,607 participants. For both genders, the response rate is balanced across treatment groups. A short SMS-based survey was also administered 2-3 days after treatment. Participants were incentivized to answer the phone- and SMS-based surveys with mobile phone airtime payments.

3.2.3 Sample Characteristics

The respondents in the sample are 99% Black African and 62% female. The average age of men in the sample is 23 years, while the average age of women in the sample is 24 years (Table 3.1 Panel A). 99% of both men and women in the sample have completed secondary education, but a higher percentage of women have completed post-secondary education. While only 16% of men have completed a university degree or diploma and 18% of men have completed some other post-secondary qualification, 17% of women have received a university degree or diploma and 23% of women have received a different post-secondary qualification.

Baseline employment during the past week is low for both genders, but higher for males: 42% of men are employed at baseline, but only 35% of women are (Table 3.1 Panel B). In addition, most participants have worked at some point, with 77% of men and 66% of women having previously worked. However, very few participants have ever held a long-time job: 10% of men and 9% of women reported having ever held a long-term job at baseline. Both men and women report low weekly earnings conditional on working, although men in the sample receive higher earnings than women: men earn an average of 97 USD PPP (609 South African Rands) weekly, while women earn an average of 85 USD PPP (533 South African Rands) weekly.

Despite the low levels of employment in the sample, participants are actively searching for employment (Table 3.1 Panel C). 96% of men and 97% of women in the sample searched for work during the week preceding baseline. Although men and women both reported submitting an average of 10 applications in the month prior to baseline, the men in the sample spent more time searching for jobs: men searched for an average 18 hours in the week before the baseline survey was conducted, and women reported searching for an average of 16 hours. Search costs, conditional on searching, were high for both men and women: men spent an average of 34 USD PPP (213 South African Rands) on job search in the past week and women spent an average of 42 USD PPP (260 South African Rands). As a result of these search efforts, both men and women reported receiving just over 1 job offer in the past week.

Men in the sample received higher scores than women in all of the assessments (Table 3.1 Panel D). The gender gap in scores was largest for concept formation (0.31 standard deviations), numeracy (0.30 standard deviations), and the other score category (0.35 standard deviations). The gender difference in scores was smaller for communication (0.10 standard deviations) and was small and insignificant for grit (0.03 standard deviations). Scores are weakly correlated across assessments.

Both men and women have inaccurate beliefs about their skills, but men on average have more accurate beliefs than women (Table 3.1 Panel E). Before they receive their results, candidates are asked to report in which tercile they believed they ranked for the communication, concept formation, and numeracy assessments. Only 11% of men and 7% of women had accurate perceptions of their skills in all three assessments. Furthermore, 24% of men and 31% of women had inaccurate beliefs about their skills in all three assessments. Both men and women are more likely to be overconfident than underconfident, but women are even more likely to overestimate their abilities than men are: on average, women are overconfident on 53% of assessments, whereas men are overconfident on 46% of assessments. Men, on the other hand, are more likely to be underconfident than women, as they are underconfident on 12% of assessments, while women are underconfident on 11% of assessments.

3.3 Labor Market Intervention

3.3.1 Intervention

The workseeker sample is randomly divided into control and 2 treatment groups: public certification and private certification. Treatment is randomized by assessment date to reduce the risk of spillovers between the control and treatment groups. 2,274 workseekers over 27 days were assigned to control, 2,247 workseekers over 27 days were assigned to public certification, and 2,114 workseekers over 27 days were assigned to private certification.

In the public certification intervention, Carranza et al. (2020) provide treated workseekers with information about their skills that they can share with potential

Table 3.1: Summary Statistics by Gender for Baseline Variables

	Male			Female			Difference	
Variable	Mean (St. Dev.)	Obs.	p:Balance	Mean (St. Dev.)	Obs.	p:Balance	Mean Diff. (Male-Female)	p:Diff=0
Panel A: Demographic Measures								
Age	23.43	2630	0.259	23.78	4261	0.437	-0.350	0.000
	(3.140)			(3.387)				
University degree/diploma	0.163	2630	0.771	0.170	4261	0.949	-0.007	0.502
	(0.370)			(0.375)				
Other post-secondary qualification	0.180	2630	0.839	0.231	4261	0.634	-0.051	0.000
	(0.384)			(0.422)				
Completed secondary education only	0.647	2630	0.738	0.587	4261	0.892	0.060	0.000
	(0.478)			(0.492)				
Panel B: Labor Market Measures								
Employed	0.416	2630	0.440	0.354	4261	0.412	0.062	0.000
	(0.493)			(0.478)				
Earnings	608.9	887	0.056	533.2	1229	0.332	75.70	0.023
	(791.1)			(699.3)				
Ever worked	0.768	2625	0.106	0.664	4252	0.685	0.104	0.000
	(0.422)			(0.472)				
Ever held a long-term job	0.095	2625	0.864	0.087	4252	0.280	0.008	0.272
	(0.293)			(0.282)				
Panel C: Job Search Measures								
Searched	0.960	2630	0.265	0.973	4261	0.221	-0.013	0.007
	(0.196)			(0.161)				
Applications submitted[a]	10.09	2598	0.255	9.763	4217	0.413	0.327	0.500
	(18.25)			(18.76)				
Search cost	213.4	2339	0.005	259.8	3808	0.141	-46.400	0.149
	(584.0)			(1876)				
Search hours	18.09	2533	0.189	16.28	4166	0.683	1.810	0.002
	(21.81)			(20.18)				
Offers received[a]	1.364	2591	0.565	1.092	4219	0.406	0.272	0.134
	(7.908)			(6.729)				
Panel D: Assessment Results								
Numeracy score	0.236	2630	0.848	-0.062	4261	0.688	0.298	0.000
	(1.008)			(0.959)				
Communication score	0.113	2630	0.177	0.011	4261	0.439	0.102	0.000
	(1.001)			(0.985)				
Concept formation score	0.238	2630	0.499	-0.071	4261	0.973	0.309	0.000
	(1.028)			(0.949)				
Grit score	0.046	2630	0.366	0.021	4261	0.191	0.025	0.324
	(0.991)			(0.992)				
Other scores	0.212	2630	0.908	-0.134	4261	0.730	0.346	0.000
	(1.035)			(1.071)				
Panel E: Belief Measures								
Planned applications[a]	94.69	2605	0.198	20.70	4235	0.135	73.99	0.153
	(2637)			(112.4)				
Fraction of assessments overconfident	0.456	2626	0.432	0.533	4249	0.554	-0.077	0.000
	(0.353)			(0.348)				
Fraction of assessments underconfident	0.122	2626	0.454	0.110	4249	0.428	0.012	0.034
	(0.217)			(0.202)				

Table shows summary statistics for selected baseline variables. Percentiles are omitted for binary variables. All monetary figures are reported in South Africa Rands. 1 Rand ≈ USD0.16 in purchasing power parity terms. Intensive-margin labor market measures are set to missing for non-workers. Intensive-margin search measures are set to missing for non-searchers. All assessment results are standardized to have mean zero and standard deviation one in the control group. Missing values reflect item non-response, mostly due to respondents reporting that they don't know the answer. All period-specific outcomes use a 7-day recall/forecast period unless marked with [a] (30-day recall/forecast period) or [b] (since treatment). The final column for each gender reports the p-value for testing equality of means of the baseline variables across all treatment groups, using heteroskedasticity-robust standard errors clustered by treatment date. The final column for the gender difference reports the p-value for testing equality of baseline means of the variables across both genders, using heteroskedasticity-robust standard errors clustered by treatment date.

employers. After completing the assessments and baseline survey, treated individuals are given an electronic copy and 20 high-quality printed color copies of a certificate that describes the assessments and the workseeker's individual performance on the assessment by tercile (Figure 3.1). The public certificate also includes the candidate's name and national identity number in order to link individuals to the certificate. In addition, the certificate describes Harambee and its placement and assessment work and identifies the assessed candidates as high school graduates from disadvantaged backgrounds, aged 18-34, in order to help potential employers understand the relative scores. The certificate also provides a link to a website with more information on the assessments, including sample questions and information on how the assessments were designed and evaluated, as well as additional information on Harambee itself. Finally, the certificate is branded with both the Harambee logo and the World Bank logo.

The public certificates are designed not only to provide workseekers with information about their skills, but also to allow workseekers to communicate these skills to potential employers in a credible certificate from a trusted source. Candidates receive these certificates during a group briefing with a psychologist, who explains what the assessments measure and how to interpret the results. Candidates are informed that they can, but are not required to, use the certificate in future job applications. In addition, they are informed that they can request additional copies of the certificate from Harambee.

In the private certification, on the other hand, Carranza et al. (2020) provide treated workseekers with information about their skills but do not give workseekers a credible way to convey this information to potential workseekers. Individuals in the private certification group receive one low-quality printed black-and-white certificate that is anonymous and unbranded (Figure 3.2). In addition, while those in the public

Figure 3.1: Example of Public Certificate

THE WORLD BANK

harambee
YOUTH EMPLOYMENT ACCELERATOR
WORK FOR WORK

REPORT ON CANDIDATE COMPETENCIES

name.. surname..
ID No. id..

This report provides information on assessments conducted by Harambee Youth Employment Accelerator (harambee.co.za), a South African organisation that connects employers looking for entry-level talent to young, high-potential work-seekers with a matric or equivalent. Harambee has conducted more than 1 million assessments and placed candidates with over 250 top companies in retail, hospitality, financial services and other sectors. Assessments are designed by psychologists and predict candidates' productivity and success in the workplace. This report was designed and funded in collaboration with the World Bank. You can find more information about this report, the assessments and contact details at www.assessmentreport.info. «name» was assessed at Harambee on 13 September, 2016.

«**name**» completed assessments on English Communication (listening, reading, comprehension), Numeracy, and Concept Formation:

1. The Numeracy tests measure candidates' ability to apply numerical concepts at a National Qualifications Framework (NQF) level, such as working with fractions, ratios, money, percentages and units, and performing calculations with time and area. This score is an average of two numeracy tests the candidate completed.
2. The Communication test measures a candidate's grasp of the English language through listening, reading and comprehension. It assesses at an NQF level, for example measuring the ability to recognise and recall literal and non-literal text.
3. The Concept Formation Test is a non-verbal measure that evaluates candidates' ability to understand and solve problems. Those with high scores are generally able to solve complex problems, while lower scores indicate an ability to solve less complex problems.

«**name**» also completed tasks and questionnaires to assess their soft skills:

4. The Planning Ability Test measures how candidates plan their actions in multi-step problems. Candidates with high scores generally plan one or more steps ahead in solving complex problems.
5. The Focus Test assesses a candidate's ability to distinguish relevant from irrelevant information in potentially confusing environments. Candidates with high scores are generally able to focus on tasks in distracting surroundings, while candidates with lower scores are more easily distracted by irrelevant information.
6. The Grit Scale measures whether candidates show determination when working on challenging problems. Those with high scores generally spend more time working on challenging problems, while those with low scores choose to pursue different problems.

«name»'s results have been compared to a large benchmark group of young (age 18-34) South Africans assessed by Harambee. All candidates have a matric certificate and are from socially disadvantaged backgrounds. The benchmark group is 5,000 for cognitive skills and 400 for soft skills.

«name» scored in the «tercile_num» THIRD of candidates assessed by Harambee for Numeracy, «tercile_lit» THIRD for Communication, «tercile_cft» THIRD for Concept Formation, «tercile_tol» THIRD for Planning Ability, «tercile_troop» THIRD for Focus and «tercile_grit» THIRD for the Grit Scale.

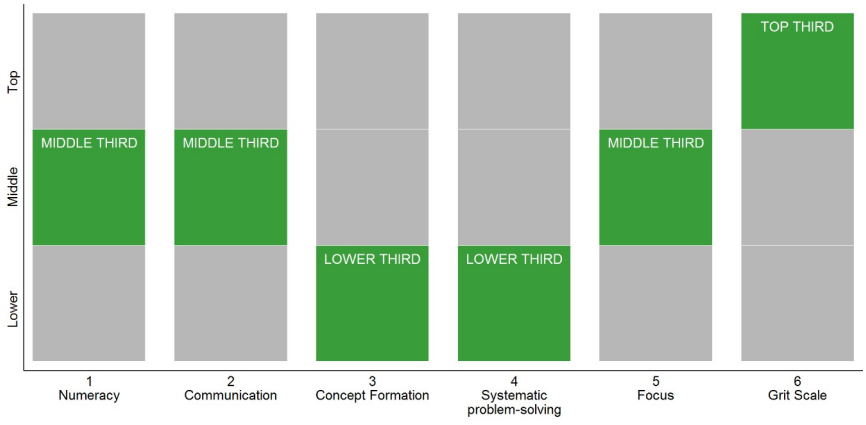

DISCLAIMER: This is a confidential assessment report for use by the person specified above. The information in the report should only be disclosed on a "need to know basis" with the prior understanding of the candidate. Assessment results are not infallible and may not be entirely accurate. Best practice indicates that any organisation's career management decisions should depend on factors in addition to these assessment results. Harambee cannot accept responsibility for decisions made based on the information contained in this report and cannot be held liable for the consequences of those decisions.

treatment arms are informed they can share the certificate with firms, those in the private arm are not informed that this is a possibility. Both types of certificate communicate the same information on the types of assessments administered and the participant's scores on these assessments. However, only the public certificate credibly conveys information about workseeker skills to firms. As such, it is expected that both the public and private certifications will have similar effects on workseeker beliefs and search behavior, but the private certificate will have a much smaller impact on firm-side information frictions.

3.3.2 Prior Findings

In their analysis, Carranza et al. (2020) find that this intervention improves labor market outcomes by easing both workseeker- and firm-side information frictions. Public certification increases employment by 5.2 percentage points, an increase of 17% from the control mean, and weekly hours by 20%. In addition, public certification increases earnings by 34% and hourly wages by 20%. Finally, the authors find that public certification increases the rate of having a written contract - a measure of formality - by 2 percentage points, or 17% of the control mean.

The treatment effects on hours, earnings, and written contracts could occur due to an increase in employment rate - the extensive margin - or an increase in match quality among those who were already employed - the intensive margin. As such, Carranza et al. (2020) decompose these treatment effects into an intensive and extensive margin effect using a method that extends upon Attanasio et al. (2011). This method defines the extensive margin as the treatment effect on employment multiplied by the mean of the relevant outcome for employed participants in the control group. In other words, the extensive margin effect is the treatment effect on employment, valued at

Figure 3.2: Example of Private Certificate

REPORT ON CANDIDATE COMPETENCIES
-Personal Copy-

This report contains results from the assessments you took at Harambee in Phase 1 and Phase 2. These results can help you learn about some of your strengths and weaknesses and inform your job search.

You completed assessments on English Communication (listening, reading and comprehension) and Numeracy today in Phase 2. In Phase 1, you completed a Concept Formation assessment which asked you to identify patterns.

1. The Numeracy tests measure various maths abilities. Your score is the average of the two maths tests you did today at Harambee.
2. The Communication test measures English language ability through listening, reading and comprehension.
3. The Concept Formation test measures the ability to understand and solve problems. Candidates with high scores can generally solve complex problems, while lower scores show an ability to solve less complex problems.

You also did some games and questionnaires to measure your soft skills:

4. The Planning Ability Test measures how you plan your actions in multi-step problems. Candidates with high scores generally plan one or more steps ahead in solving complex problems.
5. The Focus Test looks at your ability to pick out which information is important in confusing environments. Candidates with high scores are able to focus on tasks in distracting situations.
6. The Grit Scale measures candidates' determination when working on difficult problems. Candidates with high scores spend more time working on the problems rather than choosing to pursue different problems.

Your results have been compared to a large group of young South African job seekers who have a matric certificate, are from socially disadvantaged backgrounds and have been assessed by Harambee.

You scored in the MIDDLE THIRD of candidates assessed by Harambee for Numeracy, MIDDLE THIRD for Communication, LOWER THIRD for Concept Formation, LOWER THIRD for Planning Ability, MIDDLE THIRD for Focus and TOP THIRD for the Grit Scale.

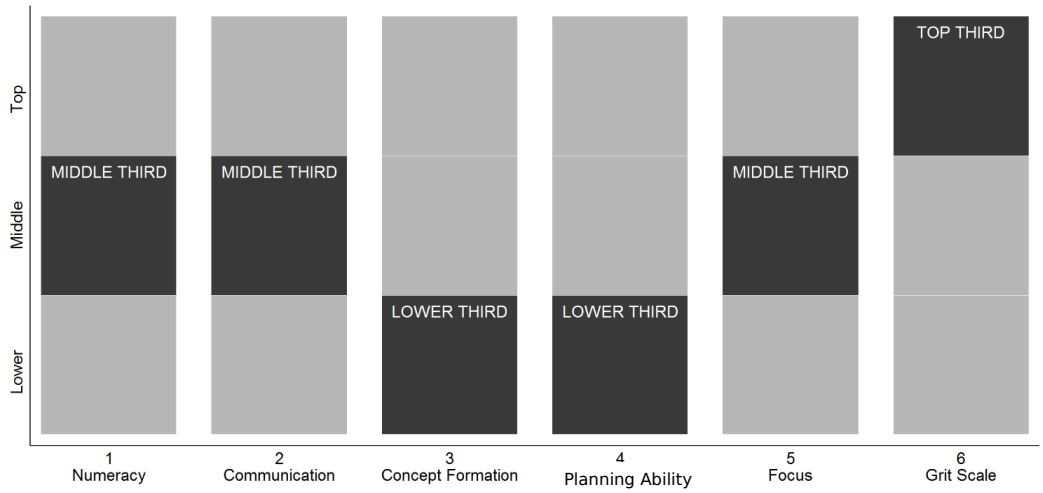

DISCLAIMER
Please note that this is a confidential assessment report and is intended for use by the person specified above. Assessment results are not infallible and may not be entirely accurate.

the mean of the outcome variable among employed members of the control group. The intensive margin is defined as the difference between the treatment effect and the extensive margin effect.

After decomposing the treatment effects of public certification on hours and written contracts, Carranza et al. (2020) find that these effects occur only at the extensive margin. The effects on earnings and hourly wages, however, occur at both the intensive and extensive margin, suggesting that there is both an increase in the job-finding rate and an increase in match quality among those who were already employed.

By comparing the treatment effects of the public and private certifications, Carranza et al. (2020) are able to determine that the intervention improves labor market outcomes by providing information to both workseekers and firms. Public and private certification affect beliefs about skills and search behavior in a similar manner, suggesting that certificates are valuable to workseekers for their ability to reveal new information about skills to the workseekers themselves. However, the employment and earnings impacts of public certification are much larger than those of private certification, suggesting that the ability of certificates to credibly communicate workseeker ability to firms is highly valuable. As such, the authors conclude that not only are workseeker- and firm-side information frictions present, but that public skill certification is effective due to abilities to address limited information on both sides of the labor market at once.

3.3.3 Empirical Specification

In order to determine gender specific treatment effects of the intervention on labor market outcomes, I use the following specification:

$$Y_{igd} = \boldsymbol{T}_{gd}\delta + \boldsymbol{X}_{igd}\Gamma + S_d + \epsilon_{igd}$$

where Y_{igd} is the outcome of workseeker i of gender g and assessed on date d, $\boldsymbol{T}_{gd}$ is a vector of treatment assignment dummies interacted with gender, $\boldsymbol{X}_{igd}$ is a vector of baseline covariates, and S_d is a stratification block by gender fixed effect. Standard errors are clustered by assessment date and are heteroskedasticity-robust.

All labor market outcomes and search measures use 7-day recall periods, unless otherwise specified. Right-skewed outcome variables, such as earnings, are transformed using inverse hyperbolic sine. This allows an interpretation of these treatment effects as percentage changes. In addition, zeros are assigned to job characteristics for respondents that are not working, and to search measures for respondents who are not searching. Both the IHS transformations and the zero replacement described above are consistent with Carranza et al. (2020). In addition, the choice of baseline covariates is consistent with the analysis performed by Carranza et al. (2020).

3.4 Results

3.4.1 Treatment Effects

The gender-specific treatment effects of public skill certification are calculated using the empirical specification presented in the previous section. Public certification improves both male and female employment, although the effect for men is somewhat larger: 6.9 percentage points, or 20% from the control mean, for male participants and 4.3 percentage points, or 15% from the control mean, for female participants (Table 3.2 Column 1). Similarly, certification increases weekly hours for both males and females, although the effect seems somewhat larger for men: a 27% increase for

males and a 16% increase for females (Table 3.2 Column 2).

Since the effect on hours could occur at the extensive margin, due to a change among those who shift into employment, or at the intensive margin, due to a change among those who were already employed, I decompose the effect into an intensive and extensive margin effect following Carranza et al. (2020). This method defines the extensive margin for hours as the treatment effect on employment multiplied by mean hours for employed participants in the control group. In other words, the extensive margin effect is the treatment effect on employment, valued at the mean number of hours among employed members of the control group. The intensive margin is defined as the difference between the treatment effect and the extensive margin effect. For both male and females, this effect on hours is driven by the extensive margin effect (Table 3.3 Column 1).

Public certification also increases earnings for both genders (Table 3.2 Column 3). The earnings effect for men, 51%, however, is somewhat larger than that for women, 24%. In addition, while treatment improves earnings at both the intensive and extensive margin for men - a 36.1% increase at the extensive margin and a 15.2% increase at the intensive margin - the only effect of certification on earnings for women is at the extensive margin (Table 3.3 Column 2). This suggests that the earnings effect for women is driven by an increase in job-finding rate, rather than better firm-worker matches. Similarly, certification increases hourly wage for both genders, although the effect is somewhat larger for men: a 28% increase for men and a 15% increase for women (Table 3.2 Column 4). For both genders, the effect on hourly wage comes at both the extensive and intensive margin (Table 3.3 Column 3).

Finally, public certification increases the probability of having a written contract for men by 5 percentage points, or 38% from the male control mean (Table 3.2 Column 5). This effect for men is at both the extensive and intensive margin: 2.6

Table 3.2: Gender-Specific Treatment Effects on Labor Market Outcomes

	(1) Employed	(2) Hours[c]	(3) Earnings[c]	(4) Hourly wage[c]	(5) Written contract
Treatment*Male	0.069***	0.272***	0.513***	0.278***	0.050***
	(0.022)	(0.091)	(0.127)	(0.072)	(0.016)
Treatment*Female	0.043***	0.160**	0.240**	0.152***	0.002
	(0.016)	(0.067)	(0.095)	(0.056)	(0.013)
Male mean outcome	0.349	11.178	198.985	10.770	0.132
Male mean outcome for employed		32.051	573.103	31.019	0.381
Female mean outcome	0.283	7.376	134.361	9.252	0.112
Female mean outcome for employed		26.329	475.949	33.281	0.399
p: treatment*male = treatment*female	0.397	0.347	0.105	0.215	0.019
# observations	6607	6598	6589	6574	6575
# clusters	84	84	84	84	84

Coefficients are from regressing each outcome on a vector of gender-specific treatment assignments, randomization block fixed effects interacted with gender, and prespecified baseline covariates interacted with gender (measured skills, self-reported skills, education, age, employment, discount rate, risk aversion). Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. Mean outcomes are for the control group. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation. The sample sizes differ across columns due to item non-response, mostly from respondents reporting that they don't know the answer.

percentage points at the extensive margin and 2.4 percentage points at the intensive margin (Table 3.3 Column 4). Certification, however, has no overall effect on the propensity to have a written contract for women. Interestingly, there does appear to be a positive effect for women at the extensive margin, but the intensive margin effect is negative, resulting in an overall neutral effect on written contracts.

3.4.2 Workseeker versus Firm Information Frictions

Certification changes both male and female workseekers' beliefs about their own skills (Table 3.4 Column 1). Candidates were asked if they believed they scored in the bottom, middle, or top third on each of the 6 assessments. The accuracy of these beliefs is then measured by calculating the fraction of assessments where self-reported beliefs about ability matched measured performance. Certification increased the accuracy of men's beliefs by 16 percentage points, or 38% from the male control mean. In addition, it also increased the accuracy of women's beliefs by 16 percentage points, or 43% from the female control mean.

Table 3.3: Gender-Specific Treatment Effects on Labor Market Outcomes at Extensive and Intensive Margins

	(1) Hours[c]	(2) Earnings[c]	(3) Hourly wage[c]	(4) Written contract
Panel A: Male treatment effect				
Total effect	0.272	0.513	0.278	0.050
	(0.090)	(0.127)	(0.072)	(0.016)
Extensive margin	0.260	0.361	0.183	0.026
	(0.082)	(0.114)	(0.058)	(0.008)
Intensive margin	0.012	0.152	0.094	0.024
	(0.036)	(0.075)	(0.054)	(0.015)
Treatment effect conditional on employment	0.029	0.377	0.233	0.060
	(0.088)	(0.186)	(0.133)	(0.038)
Panel B: Female treatment effect				
Total effect	0.160	0.240	0.152	0.002
	(0.067)	(0.095)	(0.055)	(0.013)
Extensive margin	0.152	0.222	0.119	0.017
	(0.057)	(0.084)	(0.045)	(0.006)
Intensive margin	0.009	0.017	0.033	-0.015
	(0.026)	(0.044)	(0.035)	(0.010)
Treatment effect conditional on employment	0.027	0.054	0.102	-0.047
	(0.080)	(0.138)	(0.109)	(0.031)

This table reports decompositions of gender-specific treatment effects on job characteristics into extensive and intensive margins. The extensive margins are the gender-specific treatment effects on job characteristics due to the gender-specific treatment effect on employment, evaluated at the mean job characteristics for the gender-specific control group. The intensive margins are the residual gender-specific treatment effects on job characteristics, which must be due to changes in job characteristics for the employed candidate in the gender-specific treatment group. The conditional effect is the implied mean change in job characteristics per employed gender-specific treatment group candidate. Employment in the male treatment group is 42%, so the male-specific conditional effects on all outcomes are roughly 2.4 times larger than the corresponding male intensive margin effect. Employment in the female treatment gorup is 33%, so the female-specific conditional effects on all outcomes are roughly 3 times larger than the corresponding female intensive margin effect. Heteroskedasticity-robust standard errors are shown in parentheses, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

Interestingly, certification decreased male participants' general self-esteem, reducing the probability that male participants had self-esteem above the sample median by 5 percentage points, or 8% of the male control mean (Table 3.4 Column 2). On the other hand, certification appeared to increase female participants' general self-esteem, increasing the probability that they had self-esteem above the median by 3 percentage points, or 6% of the female control mean.

Certification also changes the types of jobs that both male and female workseekers target during their search (Table 3.4 Column 3). Candidates were asked whether the types of jobs they are targeting in their search most value communication, numeracy, or concept formation. A measure of targeted search was then constructed by determining if this skill focus in job search matched the participants' highest scoring skill. Certification increased the probability that male participants targeted jobs that valued their strongest skill by 4 percentage points, or 29% of the male control mean. It also increased targeted search among female candidates to a somewhat larger degree, by 6 percentage points, or 36% of the female control mean.

Both males and females use the certificates in their applications, but males are more likely to do so (Table 3.5 Column 1). In the 3-4 months from certification to endline, 75% of male participants used the certificate in at least one job application, whereas 67% of women used the certificate. Furthermore, men used the certificate in an average of 1.9 job applications since treatment, while women used it in an average of 1.5 applications (3.5 Column 2). For both men and women, use of the report was significantly higher in the public treatment, suggesting there was not a substantial gender difference in understanding the value of a branded report for employees.

Using the report with applications resulted in interviews for both men and women, but the number of interviews was larger for men (3.5 Column 3). Men saw an average of 0.52 interviews with the public report, whereas women saw an average of 0.38

Table 3.4: Gender-Specific Public and Private Certification Effects on Beliefs and Search

	(1) Skill belief accurate	(2) > median self-esteem	(3) Targeted search
Public certification*Male	0.160***	-0.047**	0.038*
	(0.010)	(0.022)	(0.020)
Public certification*Female	0.158***	0.031*	0.060***
	(0.011)	(0.018)	(0.014)
Private certification*Male	0.132***	-0.031	0.036*
	(0.010)	(0.024)	(0.022)
Private certification*Female	0.118***	0.015	0.056***
	(0.009)	(0.020)	(0.016)
Male mean outcome	0.418	0.613	0.132
p: public*male = private*male	0.006	0.511	0.917
Female mean outcome	0.371	0.515	0.169
p: public*female = private*female	0.000	0.367	0.761
p: public*male = public*female	0.860	0.012	0.410
p: private*male = private*female	0.235	0.170	0.520
# observations	6607	6609	6609
# clusters	84	84	84

Coefficients are from regressing each outcome on a vector of gender-specific treatment assignments, randomization block-by-gender fixed effects, and prespecified baseline covariates interacted with gender (measured skills, self-reported skills, education, age, gender, employment, discount rate, risk aversion). Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. Mean outcome is for the control group. Skill belief accurate is the share of the six assessments where the candidate's perceived tercile matches their actual tercile. SAbove-median self-esteem is an indicator equal to one if the candidate's response on a shortened version of the Rosenberg (1965) self-esteem scale is above the sample median. Targeted search is an indicator equal to one if the candidate reports mainly applying for jobs that most value the skill in which the candidate scored highest. All outcomes use a 7-day recall/forecast period unless marked with [a] (30-day recall/forecast period) or [b] (since treatment). The sample sizes differ across columns due to item non-response, mostly from respondents reporting that they don't know the answer.

interviews. When the private report was used, both men and women saw a more modest number of interviews, and men again had more interviews. Similarly, using the report with applications resulted in job offers for both men and women, but men received more job offers on average (3.5 Column 4). Men saw an average of 0.14 offers with the public report, whereas women saw an average of only 0.10 offers. This suggests that use of the certificate in applications resulted in more interviews and offers for men and women, but this could be due to a number of reasons, including differential results by gender or differential patterns of use by gender.

Given the similarity in offers with the private report - 0.041 for men and 0.032 for women - it also appears as though the public certificate reduces firm-side information frictions that affect job offers more for male applicants than female applicants. Again, this difference may result from a number of factors that differ by gender. It does, however, suggest that the hypothesis that certificates might reduce statistical discrimination against women and, thereby, more greatly impact women's labor market outcomes is likely not a dominant force in this labor market.

Men and women expect to see a similar increase in job offers due to public certification (3.5 Column 5). Men expect to see a 12% increase, and women expect to see a 10% increase. Interestingly, however, men expect this entire increase to come from the public nature of the certificates: men with private certificates expect an increase in job offers of only 1%. Women, on the other hand, believe much of their expected increase in offers comes from assessment and private certification alone: women with private certificates expect an increase in job offers of 8%. This indicates that women believe workseeker-side frictions are a major impediment to receiving job offers.

As noted earlier, public certification improves employment rates, hours, earnings, and hourly wage for both men and women, although the magnitude of these effects is somewhat larger for men. Private certification for men has little impact on em-

ployment, but it does increase hours, earnings, and hourly wage to a lesser degree than public certification (3.5 Columns 6-9). This suggests that labor market effects of the certification intervention for men are due to easing both workseeker-side and firm-side information frictions. On the other hand, private certification has virtually no effect on any of these labor market outcomes for women. This suggests that, although women expect easing workseeker-side information frictions to substantially increase offers and although doing so does increase offers with report for women, this is not translated into greater employment outcomes.

Finally, although public certification improves the probability of having a written contract for men, it does not have any effect on this outcome for women (3.5 Column 10). Interestingly, private certification has a very similar effect on written contracts for men, suggesting that the major benefit of certification for finding a job with a written contract is easing workseeker-side information frictions. Like public certification, private certification has no effect on rates of written contracts for women.

3.4.3 Blinder-Oaxaca Decomposition

The effect of the intervention on labor market outcomes may differ between men and women for a number of reasons, but it is difficult to assess the root cause without a readily available counterfactual. Since men and women have different underlying characteristics that might interact with treatment, it is not possible to observe what the treatment effect on women would be if they were exactly the same as men except for their gender, and vice versa. As such, it is difficult to determine whether the gender difference in treatment effect is due to differences in underlying characteristics that interact with treatment or due simply to gender differences in response to

Table 3.5: Gender-Specific Public and Private Certification Effects on Application and Labor Market Outcomes

	(1) Used report[b]	(2) Applications with report[b,c]	(3) Interviews with report[b]	(4) Offers with report[b]	(5) Expected offers[a,c]
Public certification*Male	0.753***	1.916***	0.523***	0.139***	0.119***
	(0.018)	(0.056)	(0.037)	(0.020)	(0.032)
Public certification*Female	0.665***	1.533***	0.378***	0.095***	0.097***
	(0.015)	(0.047)	(0.034)	(0.011)	(0.025)
Private certification*Male	0.281***	0.607***	0.176***	0.041***	0.012
	(0.017)	(0.048)	(0.027)	(0.013)	(0.037)
Private certification*Female	0.292***	0.543***	0.122***	0.032***	0.080***
	(0.015)	(0.036)	(0.017)	(0.007)	(0.030)
Male mean outcome	0.000	0.000	0.000	0.000	4.472
p: public*male = private*male	0.000	0.000	0.000	0.000	0.002
Female mean outcome	0.000	0.000	0.000	0.000	4.025
p: public*female = private*female	0.000	0.000	0.000	0.000	0.595
p: public*male = public*female	0.000	0.000	0.011	0.048	0.614
p: private*male = private*female	0.531	0.177	0.054	0.442	0.156
# observations	6609	6598	6597	6597	6531
# clusters	84	84	84	84	84

	(6) Employed	(7) Hours[c]	(8) Earnings[c]	(9) Hourly wage[c]	(10) Written contract
Public certification*Male	0.069***	0.272***	0.513***	0.278***	0.050***
	(0.022)	(0.091)	(0.127)	(0.072)	(0.016)
Public certification*Female	0.043***	0.160**	0.240**	0.152***	0.002
	(0.016)	(0.067)	(0.095)	(0.056)	(0.013)
Private certification*Male	0.032	0.140*	0.307**	0.159*	0.040**
	(0.020)	(0.076)	(0.131)	(0.083)	(0.016)
Private certification*Female	-0.002	0.018	0.070	0.055	0.003
	(0.015)	(0.062)	(0.095)	(0.053)	(0.011)
Male mean outcome	0.349	11.178	198.985	10.770	0.132
p: public*male = private*male	0.098	0.158	0.138	0.146	0.563
Female mean outcome	0.283	7.376	134.361	9.252	0.112
p: public*female = private*female	0.011	0.048	0.094	0.091	0.971
p: public*male = public*female	0.397	0.347	0.105	0.215	0.019
p: private*male = private*female	0.189	0.226	0.147	0.285	0.053
# observations	6607	6598	6589	6574	6575
# clusters	84	84	84	84	84

Coefficients are from regressing each outcome on a vector of gender-specific treatment assignments, randomization block-by-gender fixed effects, and prespecified baseline covariates interacted with gender (measured skills, self-reported skills, education, age, gender, employment, discount rate, risk aversion). Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. Mean outcome is for the control group. All outcomes use a 7-day recall/forecast period unless marked with [a] (30-day recall/forecast period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation. The sample sizes differ across columns due to item non-response, mostly from respondents reporting that they don't know the answer.

treatment that are unrelated to observed characteristics.

In order to distinguish between these two channels that might contribute to the gender difference we observe, I use a Blinder-Oaxaca decomposition of the treatment effect. The Blinder-Oaxaca decomposition allows us to decompose mean differences in outcomes between groups based on linear regression models in a counterfactual manner (Blinder 1973; Oaxaca 1973). Although the Blinder-Oaxaca decomposition has traditionally been used for mean differences in outcomes outside of the context of RCTs, it can also be used to decompose differences in treatment effects. This is detailed in Appendix B. Intuitively, the treatment effect is the difference in expectation of the outcome variable between the treatment and control groups; the Blinder-Oaxaca decomposition can then be applied to the gender difference in this treatment effect.

The twofold Blinder-Oaxaca decomposition allows me to separate the gender difference in treatment effect into two parts. The first part is the "explained" effect, which is the part of the differential that is explained by group differences in observed characteristics. The second part is the "unexplained" effect, which is the residual. The "unexplained" effect is often attributed to discrimination in the literature, but it also includes the effects of unobserved characteristics. The twofold decomposition relies on the use of an assumed nondiscriminatory vector of observed characteristics in order to determine the "explained" effect. For this, I use the average of male and female observed characteristics, following Reimers (1983). These results are robust to instead using a weighted average of male and female observables, following Cotton (1988); this specification is presented in Appendix B.[3]

The twofold Blinder-Oaxaca decomposition reveals that the gender difference in

[3]The threefold decomposition, which divides the differential into an "endowments effect", a "coefficients effect", and an "interaction effect" is presented in Appendix B.

Table 3.6: Twofold Blinder-Oaxaca Decomposition of Gender Differential in Treatment Effect

E(Y|public)-E(Y|control)	Employed	Hours[c]	Earnings[c]	Hourly wage[c]	Written contract
Differential					
Prediction Male	0.055	0.213	0.472	0.259	0.039
	(0.019)	(0.072)	(0.107)	(0.056)	(0.014)
Prediction Female	0.038	0.140	0.218	0.138	0.001
	(0.012)	(0.045)	(0.072)	(0.044)	(0.010)
Difference	0.017	0.073	0.254	0.121	0.038
	(0.024)	(0.091)	(0.135)	(0.079)	(0.018)
Decomposition					
Explained	0.019	0.078	0.082	0.026	-0.002
	(0.009)	(0.033)	(0.050)	(0.028)	(0.006)
Unexplained	-0.002	-0.005	0.172	0.096	0.040
	(0.027)	(0.099)	(0.142)	(0.082)	(0.019)

The male and female predictions are generated by regressing each outcome on a vector of randomization block fixed effects and baseline covariates (measured skills, self-reported skills, self-esteem, overconfidence, education, age, baseline employment, baseline search, discount rate, and risk aversion) for each gender's public treatment and control group, and then predicting the treatment effect based on the difference. The gender difference in treatment effect is then calculated as the difference in the male and female treatment effect. Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

treatment effect on both employment and weekly hours is almost fully explained by gender differences in observed characteristics (Table 3.6 Columns 1-2). This suggests that the differential effect of treatment on employment and hours worked is due to preexisting differences between men and women that interact with treatment, rather than a difference in how treatment affects these outcomes for men and women with similar baseline characteristics.

For both of these treatment effects, the baseline characteristics that explain the largest portion of the gender difference in treatment effects are numeracy, concept formation, other skills, and the fraction of assessments for which the candidate was overconfident (Figure 3.3 Panels A-B). Men perform significantly better than women on these 3 assessments, and women are significantly more overconfident than men at baseline. The decomposition results reveal that the gender difference in mean levels of these skills and overconfidence play a large role in explaining why male and female

employment responds differently to treatment.

A much smaller portion of the gender difference in the treatment effect on earnings and hourly wage is explained by baseline differences in characteristics: only 32% of the difference in treatment effect on earnings and 21% of the difference in treatment effect on hourly wage (Table 3.6 Columns 3-4). Although the gender difference for both treatment effects are partially explained by candidates' numeracy and concept formation skills, as well as their beliefs about these skills, the majority of the gender difference remains unexplained by baseline observed characteristics (Figure 3.3 Panels C-D). This suggests that the earnings and wage effects of certification differ by gender for some reason other than observed characteristics. That is, treatment is likely to affect earnings of men and women differently, even if they share the same baseline characteristics.

Furthermore, the gender difference in the treatment effect on having a written contract is almost entirely unexplained by baseline characteristics (Table 3.6 Column 5). Although concept formation and overconfidence explain a small portion of the gender difference in treatment effect, almost all of it remains unexplained (Figure 3.3 Panel E). This reveals that the difference in how certification affects rates of written contracts between men and women is almost entirely unrelated to gender differences in baseline observables.

The traditional interpretation of large unexplained gender differences, as seen with the effect of certification on earnings, hourly wage, and written contracts, is discrimination. However, the unexplained portion of the Blinder-Oaxaca decomposition also includes the effect of gender differences in any unobserved characteristics, including preferences.

Figure 3.3: Twofold Blinder-Oaxaca Covariate Decomposition of Gender Differential in Treatment Effect

Panel A: Employed

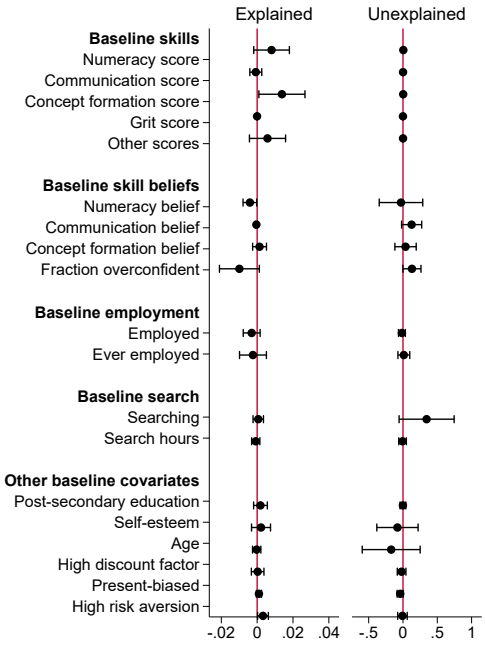

Panel B: Earnings[c]

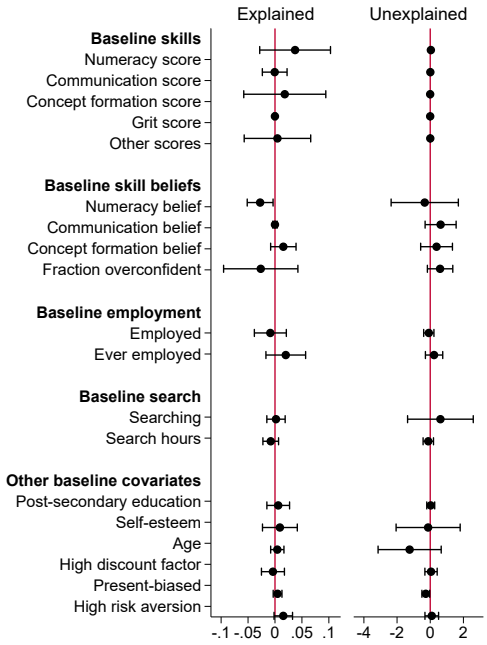

Panel C: Written Contract

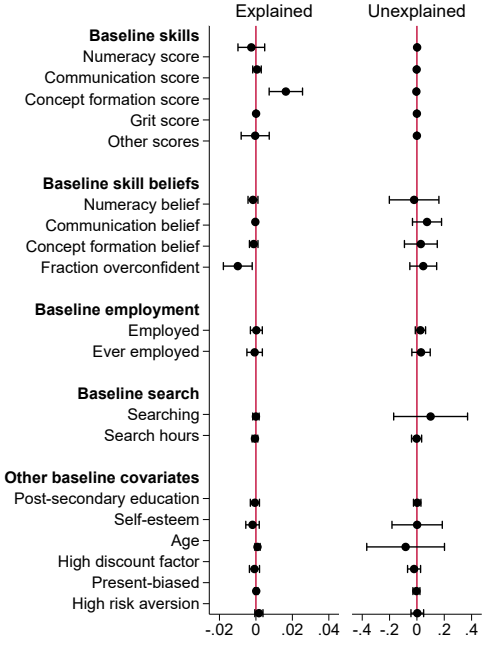

Note: The decomposition is generated using a vector of randomization block fixed effects and baseline covariates (measured skills, self-reported skills, self-esteem, overconfidence, education, age, baseline employment, baseline search, discount rate, and risk aversion). Effects for randomization block fixed effects are suppressed, as is the constant in the "unexplained" portion of the decomposition. Confidence intervals are generated using heteroskedasticity-robust standard errors, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

3.5 Conclusion

Both workseekers and firms face information frictions about workseeker skills when making matching decisions. Interventions, such as skill assessment and certification, that seek to ease these information frictions are a potentially impactful and cost-effective way to improve labor market outcomes. Skill certification, however, may have differential impacts by gender due to gender differences both in baseline characteristics that relate to the skills being certified or worker productivity and also in the ways individuals seek jobs and the types of jobs they seek. Understanding whether skill certification affects men and women differently is essential to understanding the full effect of such interventions on the labor market.

I use data from an intervention conducted by Carranza et al. (2020) that assessed and certified skills of young workseekers in South Africa who likely face high information frictions due to their low levels of post-secondary education, experience, and access to referral networks. These workseekers are also part of a labor market where skill certification is valuable to both workseekers and firms, as weak education-skill relationships provide little information about skills to either party and firms face high costs to hiring poorly matched employees. Furthermore, the labor market is one in which gender gaps already exist in labor market outcomes, such as employment levels and occupational sorting, and in which gender differences in labor market-relevant characteristics, such as skills and access to referral networks, are present.

First, I demonstrate that the effect of skill certification on labor market outcomes differs by gender. I find that, although skill certification improves employment and earnings for both men and women, the effect is larger for men; although this gender difference is statistically insignificant, it is large in magnitude. In addition, although skill certification increases the rate of written contracts for men, it has no effect on

this labor market outcome for women. I also find that the effect of certification on earnings occurs at both the extensive and intensive margin for men, but occurs only at the extensive margin for women. This suggests that the earnings effect for women is driven entirely by an increase in the job-finding rate, rather than better firm-worker matches.

Second, I find that men and women update their beliefs in response to certification in a similar manner, but men are more likely to use the certificates and are more likely to receive interviews and job offers when using the reports. Both men and women update their beliefs about their own abilities when they undergo certification, causing both genders to increase the accuracy of these beliefs; in addition, there is no gender difference in how certification affects this process of updating. However, while both men and women change their application behavior in response to certification and also receive more interviews and offers when using the reports, these effects are stronger for men.

Third, I decompose the gender difference in the effect of skill certification on labor market outcomes into a portion that is explained by gender differences in baseline observed characteristics and a portion that is unexplained. The gender difference in how skill certification affects employment is almost entirely explained by baseline characteristics, particularly numeracy, concept formation, and overconfidence. This suggests that treatment has differential employment effects for men and women due to underlying differences between the two groups; as such, if men and women had identical characteristics, we would expect skill certification to have similar effects on employment and hours for both groups. Less than a third of the gender difference in treatment effect on earnings, however, is explained by observed characteristics; and virtually none of the gender difference in how certification affects rates of written contracts is explained by these characteristics. These large unexplained effects may

be indicative of discrimination or of gender differences in unobserved characteristics.

These results suggest that, although skill certification improves labor market outcomes of both men and women by increasing information about workseeker skills for both workseekers and potential employers, men see greater effects. Furthermore, although the gender difference in treatment effects on employment is explained by differences in baseline characteristics, much of the gender difference in the effect of certification on earnings and rates of written contracts remains unexplained. These results motivate future work to examine the potential role of gender differences in other characteristics, such as preferences for job attributes and job tasks, in explaining such differences in treatment effect by gender.

Chapter 4

Female Genital Cutting in Mali

4.1 Introduction

Female genital cutting (FGC)[1] is a practice observed among many communities in Northern Africa, the Middle East, and Southeast Asia. During the procedure, a girl's genitalia are cut for non-medical reasons, usually by a traditional health provider. Cutting is performed almost exclusively on young girls, although the age of cutting varies by community, and the procedure is occasionally undergone by adult women. Recovery time and health effects depend on the extent of cutting performed, but it is common for girls to experience high levels of pain and to be immobile until scar tissue has formed. Even after a significant period of time, the scar tissue may remain highly sensitive, and it is not uncommon for women to experience genital pain during urination, intercourse, and at rest for the remainder of their lives.

In response to the various health consequences of FGC (Nour 2008, Banks et al. 2006, Berg et al. 2014a, Berg et al. 2014b), there has been global criticism of FGC and mounting outside pressure to outlaw the practice. In addition to being a public health problem, cutting is seen widely as a human rights violation, since it is performed on children. Despite the increased pressure to abandon FGC and the costs of cutting borne by the women who receive it, FGC remains widespread, particularly in

[1]Female genital cutting is alternately known as female genital mutilation and female circumcision. "Female genital mutilation" is often used in a political context, while the term "female circumcision" draws misleading parallels between male circumcision and FGC. For these reasons, the paper at hand adopts "female genital cutting" for the naming convention. It should be understood that this choice does not reflect a practical difference in the procedure in question.

Northern Africa. Furthermore, the transmission of FGC is often intergenerational: girls undergoing FGC are frequently the daughters of women who were cut themselves in childhood and, thus, have potentially experienced the personal costs of FGC first-hand.

Although FGC has been widely documented and described, there remain many questions about how the practice originated and why it continues. While the historical origins of FGC remain of high interest, the contemporary reasons for its continued existence may have long ago departed from those that sparked the practice.[2,3] As such, this paper will focus on the persistence, rather than origins, of FGC. Given the negative health and quality of life costs associated with FGC - and the psychological costs tied to increased morbidity, diminished quality of life, and the potentially traumatic experience of the procedure itself - a natural question is why the practice continues. In particular, this paper will ask why parents choose to exact a risky and medically unnecessary procedure upon their children. Put concisely: given the high health, psychic, and individual welfare costs; the absence of any health benefits; and the increasing social and legislative costs of FGC, why does the practice persist?

Previous accounts on FGC have documented the anecdotal reasons behind the practice's persistence. The oldest and most common reasons cited by researchers have been pure convention and marriageability. The reasoning behind these mechanisms is that "African women feel compelled by tradition to mutilate their own children so that they will be eligible for marriage, the only career option" (Hosken 1993).

[2]It should be noted that many scholars believe that uncovering the origins of FGC is currently a fruitless task, as "its origins appear to have been lost in the past" (Skaine 2005).

[3]What is known about the origins of FGC is that it predates the earliest written accounts and that it preexisted both Christianity and Islam (Rahman and Toubia 2000). It is believed that FGC probably originated in Egypt and the Nile valley, and that it spread first by Arab traders and then by slave traders who traveled the Nile in the 18th century (Lightfoot-Klein 1989). However, researchers emphasize that "historical resources and anthropological findings leave us with not much more than speculation as to how [FGC] began" (Skaine 2005).

The convention hypothesis suggests that the practice persists because parents either wish or feel the need to socially conform. Another reason cited by scholars is that FGC is linked in some way to femininity, and, as such, FGC persists because it is valuable to women as part of their identity. For instance, some scholars postulate that the persistence of FGC lies in an ancient belief that has become entrenched in modern society: that the masculine part of a woman's soul resides in her clitoris[4] (Lightfoot-Klein 1989, Epelboin and Epelboin 1979). Thus, the extraction of the clitoris is psychologically and emotionally linked to the affirmation of a woman's femininity, a crucial aspect of her own sense of self. Additional common reasons expressed by communities that practice FGC are hygiene, aesthetics, and the belief that FGC maintains fertility. Other explanations proposed by scholars include FGC as a rite of passage into adulthood, the desire or need to control female sexuality, reduction of female sex drive, prevention of women being unfaithful or unchaste, cultural tradition, and religious tradition.[5]

In response to the question of why FGC persists, this paper investigates several claims to elucidate the driving forces behind FGC's persistence. In particular, the paper considers the magnitude of a mothers' influence on the decision to cut her daughter and then focuses on three potential explanations for why mothers might choose to cut their daughters. The three potential explanations on which I focus are as follows: (1) the identity hypothesis, which states that performing FGC on daughters is internally valuable, as it contributes to a mother's self-worth and strengthens her identity, broadly defined; (2) the signaling hypothesis, in which FGC is externally valuable as a signal of familial quality, broadly defined, and mothers choose FGC in order to signal this quality; and (3) the marriage market hypothesis, which suggests

[4]And, in parallel, that the feminine portion of a man's soul resides in his prepuce.

[5]It should be noted that, although many writers reference religion as the reason for FGC's persistence, FGC is not required by any religion.

that demand for cut wives is high, and thus mothers choose to cut their daughters in order to improve their value in the future marriage market. These mechanisms are investigated using new data on child, rather than adult, FGC status. Importantly, this allows the paper to analyze variation in cutting decisions across daughters, where previously maternal decision-making was assumed to be uniform.

First, this paper rejects a theoretical model in which the father makes FGC decisions unilaterally. Mothers' stated preferences have a large and significant impact on the FGC status of their daughters, suggesting that mothers have at least some decision-making power in regards to FGC.[6] Thus, I conclude that the decision-making process behind FGC is a bilateral parental decision, in which the mother holds significant influence.[7]

Second, this paper rejects the identity hypothesis as the sole causal factor of FGC's persistence. There is significant variation in FGC status across daughters of the same mother, which suggests that identity (which would imply self-reinforcing FGC decisions, leading to homogeneous decisions across all daughters) is insufficient to explain the continuance of FGC.

Third, this paper rejects the marriage market hypothesis as the sole causal factor of FGC's persistence. Although we cannot rule out the importance of the marriage market in perpetuating FGC, we empirically observe that marital outcomes (such as propensity to marry and wife rank within polygynous families) are not significantly affected by the wife's FGC status.[8]

[6]The concern that mothers are reporting their husbands' preferences is discussed further in 4.6.

[7]Although this result raises interesting questions about household decision making power, such research is outside the purview of the paper at hand.

[8]It should be noted that the measures used to test this hypothesis, such as uxorial rank and propensity to marry, may not be the best measures of marital quality or match. If this is the case, the insignificant results we observe may not speak very strongly to the role of the marriage market in the continuance of FGC.

Given the claims above, this paper concludes that mothers (who, by all prior anthropological and sociological accounts, believe they are acting in their daughters' best interest) have significant decision-making power over FGC, and yet many still choose to have their daughters cut, despite the health and quality of life risks. This casts doubt on reasoning that places fathers or elder males as the primary or sole decision-making unit in FGC decisions. It also suggests that mothers' motives must be considered and that policy designed to alter FGC patterns may best be directed at mothers.

In addition, I argue that the marriage market hypothesis, which has widely been used to explain the persistence of FGC and to guide policy to eradicate the practice, is an insufficient - although not unimportant - motivation for mothers to cut their daughters. As such, I argue that the identity and signaling hypotheses play a large - although not all-encompassing - role in FGC's persistence. In particular, the variation across sisters in FGC status suggests that identity cannot be the sole driving force behind FGC patterns. Therefore, I conclude that a mixture of identity, signaling, and marriage market forces accumulate to influence the decisions surrounding FGC. As such, researchers must look beyond the marriage market when investigating FGC. Furthermore, policy makers must integrate explanations beyond the marriage market if they wish to effectively target cutting patterns. In particular, my results suggest that much of the policy suggested by researchers prior to this paper, such as having mothers commit to not marrying their sons to cut women (Macki 1996), will likely be ineffective.

The paper proceeds in the following order. Section 4.2 provides a basic exposition of FGC, including the various types and possible health risks of the procedure. Section 4.3 briefly reviews prior literature and discusses this paper's empirical improvements on this literature. Section 4.4 gives an overview of the data and a description

of the sample. Section 4.5 develops a conceptual model for understanding FGC, with a focus on three distinct mechanisms that could plausibly contribute to FGC's persistence: the marriage market, the use of FGC as a signal for socially desirable traits, and the role of FGC in personal identity. Section 4.6 describes the main empirical tests and results. Section 4.7 concludes.

4.2 Background on FGC

Cutting varies in degree depending on location, community and parental choice. The World Health Organization (WHO) officially recognizes four types of FGC (World Health Organization 2014). Clitoridectomy, or removal of the clitoris, is categorized as type I FGC. Removal of the labia minora (with or without the removal of the labia majora), or excision, falls under type II. Type III FGC involves narrowing the vaginal wall or suturing the vaginal opening, a process known as infibulation.[9] In this case, a girl's clitoris, labia minor, and the inside surface of her labia majora are cut such that the resulting scar tissue forms a barrier over the genitalia. This process of scar formation is often assisted with sutures and by binding the legs for several weeks. Prior to binding, the insertion of a foreign object is used to preserve a small hole in the scar tissue, which will be used for the excretion of both urine and menstrual blood. Upon marriage, the barrier is penetrated or cut to allow for intercourse. During labor, pregnant women will be defibulated by cutting through the scar tissue, creating a large enough hole for passing the infant. The final category of FGC - type IV - involves any other modifications to the female genitalia that are

[9]These categories are not mutually exclusive. Excision usually occurs along with clitoridectomy, and infibulation is often performed after excision.

non-medical in kind.[10]

Evidently, the process of FGC can have significant negative impacts. In addition to the immediate pain induced by cutting and recovery, residual pain and discomfort often remains throughout a woman's life. Although this pain may be acute during intercourse, it can also affect a woman in other areas of life; not surprisingly, FGC is often associated with a reduced quality of life in general and diminished sexual life in particular. In addition to pain, the procedure bears the immediate risk of infection, blood loss and shock. Although many long-term health risks of the procedure are associated with FGC, the medical literature is not conclusive on the risks and consequences of cutting (Browning et al. 2013). Studies often have access to a highly restricted sample pool, especially when data are collected from clinic or hospital patients. The use of midwives and traditional birth attendants (TBAs) is still common in developing countries, and those who attend health facilities for antenatal, delivery, or postnatal care are unlikely representative of the average mother. Furthermore, it seems likely that selection into formal health facilities might be biased by FGC status, especially in areas where FGC is outlawed. In addition to the sample selection problem researchers face, many studies in this area suffer from small sample sizes, particularly for infibulated patients.

Although the medical literature has not fully established the causal impact of FGC on health and quality of life outcomes, there is a consensus that the practice carries many risks and no health benefits (World Health Organization 2014,Nour 2008, Banks et al. 2006 Berg et al. 2014a). Among the many negative health outcomes believed to be caused or exacerbated by FGC are: malformed scar tissue, dyspareunia, cysts, increased risk of contracting sexually transmitted infections, recurrent

[10]Two examples of type IV FGC come from Northern Nigeria: the insertion of corrosive substances into the vagina, performed to narrow the orifice; and *Angruya*, which involves scraping the vaginal opening (Mandara 2004).

or chronic vaginal and bladder infections, hemorrhage, and infertility (Nour 2008). In addition, FGC - particularly infibulation - is associated with increased problems during childbirth, including prolonged labor, difficult delivery, caesarean section, and fetal and neonatal distress and death (Banks et al. 2006, Berg et al. 2014b, Berg et al. 2014a, Skaine 2005). In particular, a relationship between FGC and vaginal fistula - a common disease in developing contexts - has been emphasized, although not yet supported in the data[11]. Fistula, which occurs when the vaginal walls neighboring the bladder and rectum collapse, is of special importance in developing countries, due to the social ostracism often faced by fistula patients (Kristof and WuDunn 2010).

4.3 Prior Literature

4.3.1 Prior Findings

Researchers have spent much time and effort documenting FGC in various communities. In particular, anthropological and sociological literature has targeted FGC's persistence by recording what local women cite as the benefits - personal, familial, and community-wide - that stem from the tradition. While such anecdotal case studies are available, there is little empirical, data-driven research to determine which documented and reported reasons drive the persistence of FGC. Currently, there are two papers of which I know that perform this task. Bellemare, Novak and Steinmetz use repeated cross-sectional data in West Africa (Bellemare et al. 2015). They find that, controlling for observed characteristics, women who underwent FGC themselves

[11]The scar tissue left by FGC - and infibulation in particular - is often weaker and less plastic than tissue removed. Thus, vaginal tearing may be more likely to occur among women with extensive genital scar tissue. In addition, scar tissue does not dilate during labor, reducing the vaginal expansion required for safe childbirth (Skaine 2005, Browning et al. 2013, Shell-Duncan and Hernlund 2000)

are on average 16 percentage points more likely to favor the practice.[12] In addition, Bellemare, Novak and Steinmetz consider how much variation is absorbed by various levels of fixed effects - individual, household, village, etc. - and find that the vast majority (87% on average and 71-93% across countries) of heterogeneity in FGC preferences is attributable to individual- and household-level factors. Although attempting to attribute individual preferences to individual- versus village-level effects generates obvious concerns, the authors conclude that this result is suggestive evidence that social norms and customs may not be the primary driving force behind FGC.[13]

Wagner also uses cross-country data to analyze FGC in Africa (Wagner 2015). She finds that women who have undergone FGC are 38% more likely to marry and tend to marry at a younger age than women who are uncut, indicating that the marriage market might play an important role in perpetuating FGC[14]. An earlier version of Wagner's paper also considers the effect of reputation and peer pressure, as measured by the density of one's ethnic or religious group in the local area (Wagner 2011). Wagner finds that FGC outcomes are not impacted by the local density of an individual's group. This result is significant, as it suggests that local marriage market composition does not heavily affect the rate of cutting. If the marriage market is the primary mechanism in FGC persistence, we would expect that reputation, peer pressure, and marriage market composition play a central role in determining rates

[12] A previous version of this paper found that cut women were 40 percentage points more likely to support FGC in the Gambia, suggesting large variability by country.

[13] It is important to note that this result could also be evidence that villages (and potentially other geographical units) are not effective proxies for norm-sharing groups or communities.

[14] It should be noted that a negative correlation between years of education and undergoing FGC could explain the latter finding. In addition, demographic shifts in attitudes toward FGC could explain the former result. For example, if younger generations are less likely to undergo FGC, we would see a similar result, simply by virtue of the fact that these women are younger and thus have not yet entered the marriage market. Such correlations could cast doubt on the causal effect of FGC on marriage market performance.

of FGC.[15]

4.3.2 Improvements upon Prior Research

This paper expands upon the previous literature by using data on the reported FGC status of children, rather than the reported status and stated preferences of their mothers. This change has several empirical benefits. First, the Demographic and Health Survey (DHS) collects data on adult women aged 15 to 49. Although the age of cutting varies by region and ethnicity, most girls are cut between the ages of 4 and 8 (Jones et al. 1999).[16] Given this information, the FGC status of adult women reflects trends in the practice from 10 years prior to data collection. Although this information is useful in determining the long-term trends of FGC, it does not reflect the current state of the practice. Using child data allows us to describe and analyze the state of FGC in recent years. Recent data for determining the current state of FGC, the recent trends of the practice, and the impact of various laws and interventions designed to reduce FGC is thus valuable.

Second, using reported child status ensures access to relevant maternal and household characteristics that are likely to impact the decision to cut a daughter. Studies using adult status do not have access to any information about a woman's parents or family of origin. To the extent that this information is simply ignored, omitted variable bias (OVB) will heavily bias results. In addition, when current observed characteristics of the women, such as wealth and education, are used to proxy for

[15] In the case that marriage markets are widely dispersed, local marriage market composition will be poor predictors of the role of marriage market forces in the occurrence of FGC.

[16] In countries where FGC prevalence is high, many girls are cut at birth or in infancy. In countries where the practice is viewed as an initiation - such as Liberia, where FGC is associated with entrance into a secret society (Sande) - the age of cutting may be later (Yoder and Wang 2013). For a thorough overview of age of cutting in different countries, please consult USAID's 2013 DHS Comparative Reports, No. 33.

those of her family of origin, results will suffer from concerns over reverse causality. The use of child status, however, allows us to eliminate much of this OVB and eliminate concerns of reverse causality, as the FGC status of a daughter is unlikely to have a direct effect on her mother's education, household wealth, etc.

To demonstrate the importance of access to parental characteristics, consider the following concrete example: researchers observe that women who have undergone FGC have, on average, more wealth once married. Given this observation, the researchers argue that cut women perform better in the marriage market and that there is thus a causal influence of FGC on marriageability. Now, consider the argument that wealthier families cut their daughters more frequently (for unspecified reasons) but also invest more in non-FGC-related methods to improve their daughters' marriageability. Now, we cannot disentangle whether it is the positive FGC status or the non-FGC inputs (and, therefore, parental wealth) that are driving the increase in marriageability. Controlling for parental wealth, however, will enable us to disentangle this relationship.

Third, where previous studies used a woman's stated preferences to proxy for her decisions regarding FGC and her daughters, this paper employs actual outcomes. Not only does this choice better reflect the real state of the world, but it also eliminates the implicit assumption that women act uniformly across all daughters and always act in accordance with their stated preference. Allowing for such seemingly inconsistent behavior is particularly useful in the context of FGC, as it is not clear a priori that women have sufficient decision-making power to act upon their preferences. Furthermore, where we do believe that preferences are influential, using actual child outcomes allows us to corroborate stated preferences with outcomes and, thereby, reduce the natural concerns that arise surrounding stated preferences in survey data.[17]

[17]It is important to note that this method of corroboration will be effective only if outcomes are

To demonstrate the importance of using actual outcomes over stated preferences, we will consider another concrete example: researchers observe that FGC preferences are not correlated with marriageability of daughters. Thus, the researchers argue that FGC does not have a causal impact on marriage market outcomes. Now, consider that this occurs in an environment where FGC is a highly valued trait in the marriage market and a high percentage of women support FGC, but the procedure itself (i.e. hiring a traditional health care provider, losing a daughter's labor during her potentially long recovery, etc.) is very costly. In this environment, many families that support FGC cannot act on their preferences by cutting their daughters. The regression of FGC preferences on marriage match quality does indeed show no significant relationship; however, if the researchers instead regressed child FGC status on marriage match quality, they would observe a positive and significantly large effect of being cut on marriageability. Here, preferences are an ineffective proxy for FGC status and thus controlling directly for this variable will significantly improve accuracy.

Fourth, collecting information on adult FGC status involves asking women to accurately remember and report experiences from early childhood. In contexts where FGC is common, it may be the case that the topics of female genitalia and sexuality are taboo. In such an environment, expecting a young girl to understand what was happening and to accurately remember that experience from childhood seems unreasonable. Anecdotally, it seems to be the case that that many women do not know what was performed.[18] Furthermore, data indicates that women often misreport the form of FGC they underwent as children, with a high number underestimating

not misreported to better align with preferences, or vice versa. It will be effective in cases where women misreport preferences in contrast with their actions (perhaps because, as in the case of FGC, outside communities disapprove of the practice).

[18]For a recent anecdotal journalistic account of FGC that received widespread attention, please see Karimjee 2015.

the severity of the procedure they received (Elmusharaf et al. 2006). Since mothers typically authorize the procedure, accompany their daughter to the provider, and watch or participate in the procedure, it is reasonable to believe that women will have more accurate and detailed knowledge about their daughters' FGC experiences than their own.

In addition to these empirical benefits, using child status rather than female preferences for FGC allows us to test various theories of FGC persistence. Most importantly, child-level cutting status allows us to utilize within-mother variation and to observe the dynamics of the FGC decision-making process for families. First, we can determine where cutting decisions differ by daughter. Second, we can see where cutting decisions deviate from opinions regarding the practice. This variation among daughters of the same mother is of first-order importance because it sheds light on the various mechanisms that could reinforce the practice of FGC. If, for instance, we observe mothers acting uniformly across daughters, this lends credence to the identity hypothesis, since a positive feedback loop between major decisions (such as FGC) and a strengthened sense of self produces continuity of decisions through time. Alternately, if we observe mothers acting differently across daughters, this detracts from the identity hypothesis. In addition, such a pattern strengthens the argument for strategic decision making when cutting daughters, which not only speaks less to the identity hypothesis, but also suggests that the marriage market hypothesis - in which parents may choose to cut their daughter in order to compensate for other characteristics that reduce her marriageability - may be more plausible.

4.4 Data

The data used in this paper were collected by the Demographic and Health Survey (DHS), funded by the United States Agency for International Development (USAID). Within each country, the DHS is nationally representative and is designed to be repeated at 5 year intervals. Several developing countries, however, do not have such frequent or consistent data available. The survey randomly selects households, collecting both individual and household level data on demographic characteristics and health histories. Although the survey is aimed at women, a random sub-sample of households are selected for male interviews as well.[19]

In addition to the standard questionnaire, the DHS has several modules and country-specific questions and modifications that are used when relevant. In particular, the DHS has a standardized FGC module that collects information relevant to the practice. In 2010, the DHS overhauled the FGC questionnaire to expand the information collected on child cutting.[20] Prior to this change, the module focused on FGC as it related to adults: questions included stated preferences on FGC for both men and women and FGC status for adult women. In some versions, the DHS collected data on whether a woman's first- or last-born daughter had been cut, but the questions were inconsistent across countries and waves. In 2010, the FGC module began collecting information on FGC status for all entries in a woman's birth history. As such, the most recent survey data include FGC information on every daughter whose mother was surveyed for the module.

The new data allows for a more salient and in-depth analysis of FGC practices in recent years. There are several empirical benefits to exploiting the new data on child

[19]The most common form of randomization is every second or every third household.

[20]The DHS redesigned the FGC module in conjunction with UNICEF's Multiple Indicator Cluster Survey (MICS) in order to create a standardized questionnaire for use in both surveys.

FGC status. These include a better understanding of the practice in recent years, a distinction between opinions of FGC and outcomes, the analysis of variation among daughters of the same mother, and access to natal household covariates.

As a case study, this paper will focus on FGC practices within Mali. Although few countries have yet surveyed residents under the revised FGC module, Mali used the questionnaire in its 2012-2013 DHS survey. In addition, Mali has 4 previous waves of DHS with the older FGC module. As global perceptions and continental laws shift toward FGC abolition, prior data on FGC prevalence becomes increasingly important in order to control for changes in reporting habits. Cohort comparisons in the repeated cross-section allow researchers to minimize bias produced purely by reporting-based changes and to isolate actual changes in FGC rates.[21] In the case of Mali, no legislation against FGC has been introduced, so the expected reporting bias is likely stable over time; however, several community-level initiatives to end cutting have been launched.[22]

Furthermore, Mali is a prime candidate for FGC analysis, as the practice is extremely common but not universal, providing significant variation in individual characteristics (Figure 4.1). In addition, FGC is geographically widespread, with variation inside survey clusters and ethnicities. This variation allows a close analysis of the practice of FGC and the possible reasons for its persistence, which can speak to the persistence of FGC on a wider geographical scale. It should be noted, however, that there is a key drawback to the data: due to conflict in North Mali, the sample excludes 3 regions (Figure 4.2) and, thus, a small percent of the population (approximately 10.3% of the population in 2009, prior to the current conflict (République de

[21] Cohort comparisons are not utilized in this paper, but this is an area for future research.

[22] The availability of data suggest that Mali is an especially conducive environment to analyze the effect of localized efforts to abandon FGC. The paper at hand does not complete this task, but the use of GPS data and reports of community-level campaigns to end FGC could be utilized in future research.

Mali: Institut National de la Statistique 2010).

Figure 4.1: FGC Prevalence among Girls

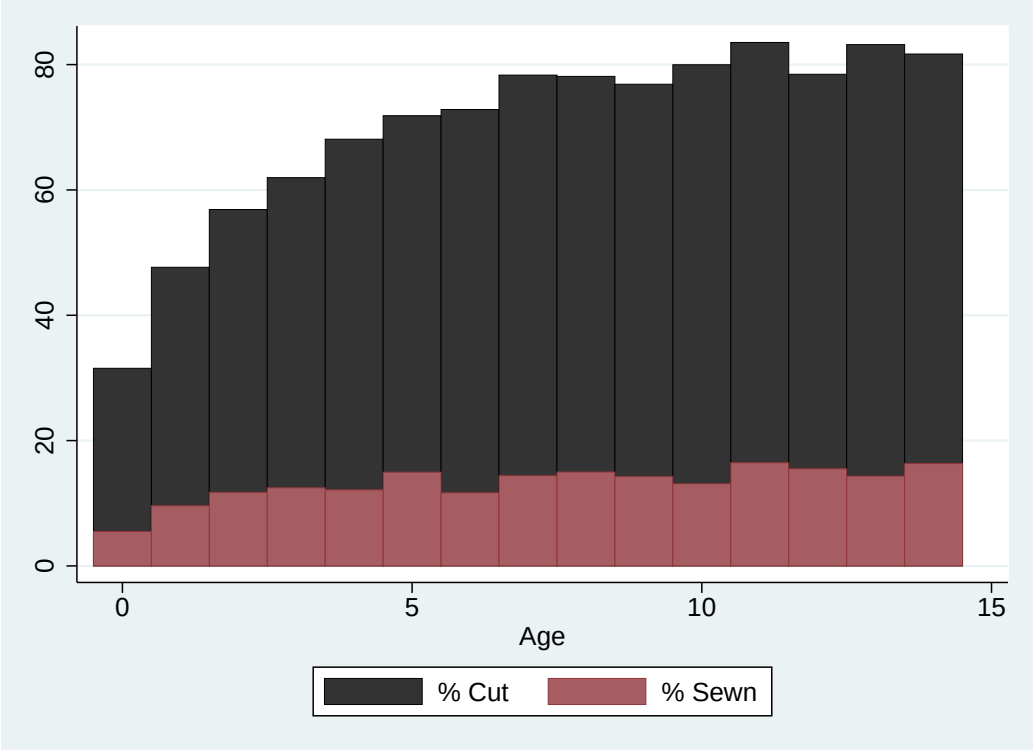

Note: This figure shows the reported rates of FGC among girls ages 0-14.

The data on Mali includes a wide variety of women and targets women between the ages of 15 and 49.[23] A description of the average traits of surveyed women is presented in Table 4.1, and more detailed information about the distribution of certain traits is presented in Table 4.2.[24]. As can be seen, the average age is relatively young: 28.5 years old.[25] In addition, the average level of schooling is minimal. The mean number of children borne to Malian women is just over three, and the first birth occurs early in a woman's life. Most households are rural, and just over a third

[23]It should be noted that both never- and ever-married women are surveyed in Mali, as opposed to the ever-married only DHS sample of some countries.

[24]In addition, descriptive statistics by cutting status are presented in Table C1

[25]Note that this average excludes all observations under 15 and above 49 years of age.

Figure 4.2: Location of DHS clusters within Mali, 2012-2013

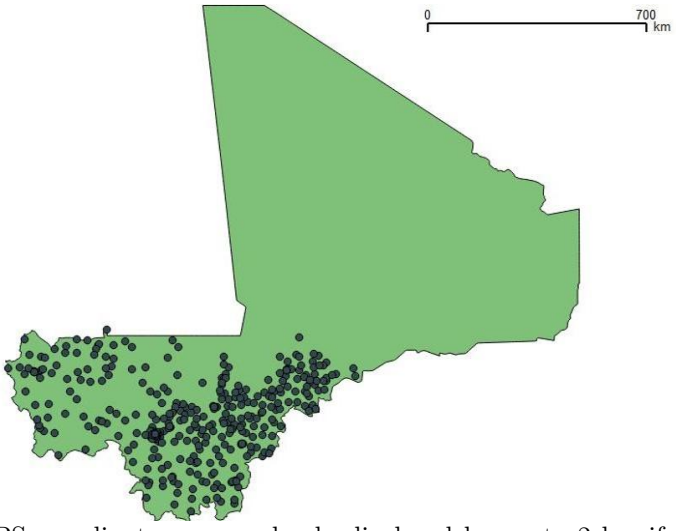

Notes: The GPS coordinates are randomly displaced by up to 2 km if urban and up to 5 km if rural. The sample is representative of 5 regions (Kayes, Koullikoro, Sikasso, Ségou, Mopti) and the district of Bamako. 3 regions (Gao, Tombouctou, Kidal) are excluded.

of households have electricity. A minority of women hold their own job,[26] but the number of working women is perhaps higher than might be expected in a patriarchal society.

Table 4.1: Summary Statistics, Female Adults

Variable	Mean	Variable	Percent
Age	28.503	Rural	68.707
Years Schooling	2.003	Working	42.556
Children	3.243	Pregnant	11.176
Living sons	1.465	Has Electricity	36.138
Deceased sons	0.217	Has Television	41.970
Living daughters	1.391	Has Radio	79.106
Deceased daughters	0.169	Has Fridge	13.008
Age of 1^{st} birth	18.822		

Although a variety of age groups and religious groups are represented in the data,

[26]This statistic presumably includes only women who earn an income at a job outside of the family agricultural plot, but the DHS survey is not specific.

Table 4.2: Detailed Summary Statistics, Female Adults

Variable	Mean	St Dev	Min	Max
Age	28.503	8.896	15	49
Years Schooling	2.003	3.837	0	18
Children	3.243	2.677	0	15
Living sons	1.465	1.449	0	9
Living daughters	1.392	1.403	0	9
Age of 1^{st} birth	18.822	4.201	12	43

the vast majority of women are Muslim (Table 4.3). In addition, the age distribution is skewed in favor of the young. Ages 15-29 are relatively evenly populated. However, around the age of 30, the number of women observed in the data set drops significantly. As is seen in many developing countries, the number of elderly individuals is low and the population density decreases steadily with age.

Table 4.3: Age and Religion of Adult Females

Age		*Religion*	
Group	Percent	Group	Percent
Age 15-19	18.400	Muslim	93.198
Age 20-24	18.035	Catholic	2.427
Age 25-29	19.906	Methodist	0.192
Age 30-34	15.896	Evangelical	0.844
Age 35-39	12.759	Other Christian	0.451
Age 40-44	8.653	Animist	0.844
Age 45-49	6.351	No Religion	1.976

Around 82% of women surveyed are married. Two-thirds of these women are in a monogamous marriage, while the remaining one-third are one of multiple wives to a single husband (Table 4.4). Since men are randomly sampled from the initial sample of households, we do not have information collected directly from every woman's husband. However, we do have both wife and husband questionnaires from many couples.

Although it is common for multiple household members to be surveyed, DHS

Table 4.4: Marital Status of Female Adults

All Women		*Married Women*	
Status	Percent	Status	Percent
Now Married	82.099	Monogamous	65.319
Never Married	14.208	Polygynous	34.681
Co-habitating	1.717		
Divorced	0.489		
Widowed	0.940		
Separated	0.547		

interviewers attempt to isolate the interviewee to prevent misreporting and encourage honesty. However, it is not always possible to conduct the interview alone, which raises concerns about the reliability of responses, particularly to questions which are sensitive. In order to clarify the environment in which the interviews were conducted, the DHS includes variables that report whether others were present at the time of the interview. Since interviews may be conducted in more than one session and since other individuals may enter and leave the interview location during the course of the survey, the variables are specified for certain questions: questions about the social acceptability of wife-beating[27] and questions about sexual activity.

We will use the presence of other individuals during the wife-beating questions as an imperfect proxy for the presence of other individuals during the questions about FGC. Since these questions are similar in nature to the questions on FGC - sensitive and likely to be influenced by the presence of others, particularly the husband - this measure helps shed light on misreporting. Table 4.5 demonstrates the percentage of women who were interviewed in the absence of other people, in the presence of others who were actively listening, and in the presence of others who were not listening (as

[27]These questions about wife-beating are separate from the domestic violence questionnaire, which is not administered to all interviewees. The questions in the wife-beating category concern when the respondent thinks it is justified for a husband to beat his wife, including the following scenarios: leaving the house without notifying her husband, neglecting the children, arguing with her husband, refusing sex with her husband, and burning food.

determined by the interviewer). The categories of other people is broken down into the following categories: children under the age of 10, the respondent's husband, men other than the respondent's husband, and women. Table 4.6 shows the percentage of women who were completely alone during the wife-beating questions and those who were not accompanied by another person who was listening.

Table 4.5: Presence of Non-Respondents During Abuse Questions

	Children	Husband	Other Men	Other Women
% Absent	94.023	97.477	97.688	96.182
% Present, Listening	1.804	0.470	0.489	1.410
% Present, Not Listening	4.173	2.053	1.823	2.408

Table 4.6: Percent of Interviews Without Listeners

% Completely Alone	89.917
% No One Listening	96.585

As is evident in these two tables, most women were alone when interviewed and the majority of those who were not alone were not actively being monitored. Perhaps most importantly, very few women were interviewed while their husband was present. This eases some concerns about reporting bias. In particular, it suggests that women were unlikely to report their husband's preferences for FGC as their own due simply to his presence. It should be noted, however, that this does not preclude the women misreporting their preferences for other reasons, such as social desirability bias.

Furthermore, Table 4.7 is a balance table of certain characteristics for women who had their husband present, had anyone present, had no-one present, and had no-one listening.[28] Although there are differences among the groups in reported preferences, these disparities are small. Furthermore, women seem less likely to report that they support FGC if their husbands are present. This is comforting if we are concerned

[28]Women who are alone tend to be from wealthier families, perhaps because the opportunity cost of staying at home during the survey is higher for their husband and other relatives.

about women reporting the preferences of their husbands, who on average express greater support for the practice, due to their presence. Women who are accompanied by their husband do seem less likely to express a well-defined preference for FGC (i.e. support or oppose). However, this cannot be ascribed a causal interpretation, and it may be due not to reporting bias but to the possibility that women who are monitored by their husbands develop weaker opinions in general.

Table 4.7: Comparison of Women by Presence of Others

	Husband present	Someone present	No-one present	No-one listening
% Cut	87.833	89.744	91.006	90.902
% Support FGC	68.821	69.231	70.628	70.481
% Oppose FGC	17.490	21.795	18.148	18.067
Education	0.281	0.308	0.451	0.443
Wealth index	2.821	2.718	3.244	3.228
No. co-wives	0.807	0.361	0.690	0.685

4.4.1 FGC Prevalence

FGC is very common in Mali. Over 90% of all adult women have undergone the procedure (Table 4.8). The prevalence is relatively constant across cohorts, and the minor fluctuations between age groups are not statistically significant.

When interviewing women, the DHS survey also inquires about what type of cutting the women received. An adult woman can report one of three specific types: being nicked, being sewn, and having flesh removed. However, these types of cutting do not correspond perfectly to the four classifications defined by WHO, and the interpretation of these types likely varies by woman. Being sewn falls under infibulation, or type III FGC, but may exclude infibulation that is achieved without sutures (e.g. through leg binding alone). Flesh removal could correspond to types I-IV, but it is was most likely designed as a catchall for clitoridectomy and excision, types I and II.

Table 4.8: FGC Prevalence by Age Group, Adults

	% Cut	% Flesh Removed	% Nicked	% Sewn
15-19 years	90.04	51.67	11.21	9.91
20-24 years	90.96	52.55	12.45	8.99
25-29 years	91.90	54.84	12.58	9.45
30-34 years	90.65	52.81	11.59	8.99
35-39 years	90.23	53.46	13.01	10.38
40-44 years	91.91	54.99	13.97	10.53
45-40 years	91.39	53.02	13.44	6.95
All Adults	90.94	53.24	12.38	9.43

Finally, being nicked roughly corresponds to type IV. However, it should be noted that a clitoridectomy is often referred to as a small cut in communities that practice FGC; as such, women may misidentify their clitoridectomy as "being nicked" when the interviewer poses the question. More generally, the percentages for different types of FGC should be taken as rough estimates, since women may not know the details of the childhood procedure and often underestimate the extent of cutting performed (Elmusharaf et al. 2006).

Even so, certain trends can be identified in these rough estimates. Flesh removal is, by far, the most commonly reported procedure. Being nicked is reported more often than being sewn, but the margin is small and the probability of reporting error is large.[29]

Rates of FGC among underage girls are not reported to be quite as high as among adults, but over 80% of girls are reportedly cut by the age of 10 (Table 4.9). Mothers are not asked if their daughters had flesh removed or were nicked, but they are asked whether their daughter's genitalia were sewn. Rates of infibulation appear to be higher among young girls, and at least 15% of girls are sewn by the time they are 11

[29]These patterns align with continental patterns in FGC. Flesh removal (i.e. clitoridectomy and excision) is estimated to account for at least 85% of all reported FGC cases, and infibulation is rare, confined to only a handful of countries, including Mali (Shell-Duncan and Hernlund 2000).

years of age. It is not clear why there are differences in prevalence between the oldest children and the youngest adults. It could be the result of differential reporting, informational differences or possibly even a result of sampling bias or measurement error. What is clear, however, is that reported rates of FGC among young girls are still extremely high.

Table 4.9: FGC Prevalence, Girls Under 15

Age	% Cut	% Sewn	Age	% Cut	% Sewn
< 1 yr	31.54	5.49	8 yrs	78.12	14.99
1 yr	47.66	9.60	9 yrs	76.87	14.26
2 yrs	56.88	11.74	10 yrs	79.98	13.11
3 yrs	61.96	12.47	11 yrs	83.54	16.46
4 yrs	68.10	12.14	12 yrs	78.46	15.51
5 yrs	71.85	14.96	13 yrs	83.19	14.34
6 yrs	72.84	11.67	14 yrs	81.69	16.34
7 yrs	78.33	14.40			

The rates of FGC among young girls can also be seen in Figure 4.1. A feature of this graph is that it visually demonstrates the near-quadratic structure of the relationship between age and probability of having been cut. This quadratic relationship indicates that age can be controlled for relatively cleanly in the empirical analysis.[30]

4.4.2 FGC Preferences

The DHS FGC module collects information from adult respondents on their opinions surrounding FGC. After being asked "Have you ever heard of female circumcision?", all adults are asked, "Should this practice of female genital circumcision continue or be stopped?" The question is open ended, but four answers are coded by interviewers: "continued", "stopped", "depends" and "don't know".

Table 4.10 shows the opinions of men and women regarding FGC. Men are more

[30]Robustness checks are also performed as if this relationship is logarithmic.

supportive of the practice than women, with about 82% of men and 75% of women supporting the practice. In addition, the table shows how couples match in reference to their opinions on FGC. Just over 65% of couples are perfectly aligned in their beliefs (rates of perfect positive assortment can be read directly from the diagonal). However, a surprisingly high portion of couples - almost 20% - have well-defined and opposing preferences (i.e. one partner wishes FGC to be stopped and the other wishes it to continue). Among these couples, the man opposes FGC in approximately one-third of relationships.

Table 4.10: All Couples: FGC Preferences

		Husband				
		Support	Oppose	Depends	Unsure	Total
	Support	62.92	6.28	4.33	1.11	74.64
	Oppose	11.89	2.86	1.81	0.07	16.63
Wife	Depends	4.43	0.81	0.24	0.03	5.51
	Unsure	1.11	0.07	0.03	0.10	3.22
	Total	81.79	10.38	6.52	1.31	

Table 4.11 distinguishes female preferences for those who live in urban versus rural areas. Although there are small differences, it does not seem that there is a large rural-urban split in FGC preferences in Mali.

Table 4.11: Female Preferences by Location

%	Urban	Rural
Support	0.724	0.700
Oppose	0.192	0.173
Depends	0.056	0.070
Unsure	0.020	0.035
No. obs	3262	7162

4.5 Conceptualizing FGC

In order to understand why FGC continues, we must consider the decision problem faced by parents of young girls, which rests on the utility each parent derives from having his or her daughter cut. We can assume that the parent considers the personal costs and benefits of cutting the daughter, as well as the effects of cutting experienced by the daughter. Thus, the costs and benefits accrued by the daughter enter the parent's utility function; conceptually, we can think of this as love or altruism. We will allow weights on the direct and indirect components of utility to vary by the decision-maker, thus allowing for high flexibility in parental preferences.

Parental costs There are obvious immediate costs: the wage of the local circumcisor, time costs of the procedure, lost productivity associated with the daughter's recuperation time, and the cost of immediate medical complications. If a parent is opposed to FGC, there may also be personal psychic costs associated with the procedure. In addition, there are long-term costs that would be considered in expectation: long-term medical costs, possible negative fertility impacts, and the resulting marriage market costs associated with the realization of any medical or fertility risks. There are possibly also both immediate and enduring social and reputation costs of having a daughter undergo FGC. In cases where FGC is very common, these costs may be insignificant at the status quo but potentially malleable.

Parental benefits If an individual personally supports FGC or lives in an environment where it is common, the immediate benefits likely include psychic benefits from adhering to a personal value, strengthening a valued identity, or conforming to a socially valued norm. In addition to these, there are also likely to be social and reputation benefits associated with FGC in societies where the prac-

tice is widespread. In cases where a parent believes - correctly or incorrectly - that FGC fulfills specific purposes, such as maintaining chastity, or enhances other goods, such as health or cleanliness, there are additional psychic benefits. In the short term, there may be pecuniary gains, as parents often receive money and presents as part of the celebration around the procedure. In the longer term, parents may benefit directly by improving their daughter's marriage prospects. Some of these marriage benefits may be monetary: the bride price received at marriage and long-term financial support from a daughter and her husband. However, many of these marriage benefits are likely non-pecuniary: expanded insurance and credit networks, improved social standing through marriage, and the benefits of progeny from the marriage union.

Daughter's costs The most obvious of these costs are health costs that result from FGC. These health costs could be acute, temporary shocks or long-term, chronic conditions that impede the health and quality of life of the individual undergoing the procedure. Furthermore, there are likely psychological costs that can result from the surgery itself, societal expectations surrounding gender and FGC, and long-term pain or discomfort from the initial trauma. In addition to the myriad health risks of FGC, there are also fertility and maternity risks that could harm a girl, both in her own experience and in the marriage market. Any costs that affect a girl's marriage market prospects likely also have long-term effects on her quality of life and economic prospects.

Daughter's benefits Immediate benefits include the psychic and social benefits of acceptance and approval, which, along with reputation and identity benefits, likely persist throughout her life. In the long-term, the girl may accrue benefits from increased status in the marriage market, which include higher match

quality and improved standard of living. To the extent that the girl can "marry up," there will also be benefits associated with increased social status for herself and her children. In the same vein, better network connections associated with a higher quality husband may also result from undergoing FGC.

It is likely that all of the factors mentioned above influence the parental decision regarding FGC. However, in order to understand why and how FGC persists we will sort the above incentives into a handful of categories: marriage market concerns, signaling concerns (outside of the marriage market), and identity concerns. Each of these explanations can affect parental utility both through his or her own direct experiences or through considerations of the daughter's welfare. It should also be noted that these hypotheses are applicable to a range of household decision-making models. Although the following discussion will talk of how each mechanism affects one parent, we should be cognizant that both parents are indeed affected and that the resulting decision is based on some aggregation process (which may or may not place zero weight on one parent's preferences) that is likely influenced by bargaining power.[31]

4.5.1 Marriage Market

Assuming that utility is, in part, derived from matching a daughter with a "good" husband, a parent will choose his or her actions to improve the match quality expected in the marriage market. That is, he will act to maximize his daughter's marriageability, given his constraints and the expected actions of others. As discussed above, there are a myriad of costs (social, psychological, physical, pecuniary,

[31]The exact decision-making process is not explored in this paper. I will consider only if maternal preferences have non-zero weight in the process.

etc.) associated with choosing FGC. To compensate for these costs of FGC, the marriage market hypothesis predicts that there will be large differential gains in match quality and/or familial transfers (i.e. bride price) by cutting status. Anecdotally, this theory has been widely cited. For instance, in her exposition of FGC, Hughes claims that "the smaller the artificial passage is, the greater the value of the girl and the higher the brideprice" (Hughes 1995). In the data, this hypothesis can be explored through a variety of measures: probability of marriage, wealth of household, rank among co-wives, and other indicators of match quality.

A second prediction of the marriage market hypothesis is how FGC might respond to shifting marriage market composition or demographics, specifically in terms of sex ratio. First, consider a scenario where the proportion of the population that is female decreases, and, consequently, the proportion that is male increases (e.g. female emigration to a distant garment factory). Since the quantity of marriageable women has decreased, the number of potential and, thus, realized matches decreases. In addition, since men must compete more intensely to match with one of the limited number of women, the price of matching (i.e. the cost of searching for a spouse in the marriage market) increases for men. Similarly, the costs associated with the marriage market decrease for women. We can think of this as the costs shifting such that men bear a higher proportion of these costs. That is, the bargaining power of women in the marriage market increases relative to the prior equilibrium. From here, we can see the marriage market hypothesis's prediction for FGC: if FGC is fueled by a marriage market in which men prefer wives who have been cut and women do not have a preference for FGC, the rate of FGC will decrease as men's bargaining power decreases due to the reduced supply of potential wives.

Second, consider a scenario in which the proportion of men in the marriage market decreases, and, as such, the proportion of women increases (e.g. male mortality

due to a war or other violent conflict). Again, the total quantity of marriages will decrease, as the number of men eligible to match has decreased. However, unlike in the first scenario, women must now compete more intensely to find a match, since the supply of potential husbands has diminished. As such, the costs of finding a spouse have increased for women but decreased for men. That is, the burden has shifted slightly onto the women, and men now have increased bargaining power in the market. Therefore, the marriage market hypothesis predicts that the rate of FGC among women will increase as they compete for the now-limited supply of men. Note that, "in many corners of Africa, previously because of wars and today due to immigration, there are double the women than men of marriageable age," (Octapixx Worldwide 2004), a fact which might help explain the prevalence of FGC, if the marriage market hypothesis presented above holds.

4.5.2 Signaling

To understand how signaling could drive parents to cut their daughters, we will review the basic framework of signaling. There are an array of people and a number of types (for simplicity, we can think of the case of two types: high and low). Type is socially important, and certain types are valued above other types by the rest of the community. In the case of FGC, this type might be religiosity, honorableness, social conformism, or some other characteristic that is held in high esteem in the relevant social environment. For example, if a community is highly religious - as seems to be the case in many FGC-practicing communities (Skaine 2005) - then individuals who care deeply about practicing and upholding their faith (i.e. their type is religiously devoted) will be valued higher than individuals who are less devoted. In our general framework, type is intrinsic and known by the individual himself, but it is publicly unobservable.

Since type is socially valued, the utility derived from social interactions naturally depends upon the type of others. Individuals gain greater utility by interacting with the preferred type. As such, individuals not only want to know others' types but also want to signal that they themselves are of the preferred type - whether or not they actually are this type. For instance, scholars have suggested that "a mother will impose FGC on her daughter because the mother wants social respect" (Skaine 2005). When the signaling convention arises, the community recognizes it as a way to judge others' type, and it becomes a common way to communicate an individual's unverifiable type to his neighbors. Importantly, while type is publicly unobservable, the signal itself must be observable by individuals other than the agent himself.

In our example, then, villagers wish to judge the type of others, but doing so directly is impossible (in this case, because one's religiosity is internal). At first, we might think that religious participation would be a natural signal. However, as a rule, signals must be costly to fake, to deter non-preferred types from lying about their types via sending an inauthentic signal. In our example, since the stakes are relatively high, the cost of attending a mosque, church, or other house of worship once a week and partaking in special religious holidays and events is relatively low. As such, it is likely that all individuals - regardless of type - will signal their devotion by attending their house of worship. Here, we will observe a pooling equilibrium in which all types signal and, as such, the signal is uninformative.

In order to unravel the pooling equilibrium, signals must be costly, and the signal must be less costly for the socially preferred type.[32] This encourages a separating equilibrium where individuals signal in line with their true type. In our example, we will consider FGC - a practice with unknown origins that spread and fused with cultural, societal and religious expectations - as a costly way to signal religious devo-

[32]However, even when these two conditions are met, a pooling equilibrium might still emerge.

tion.[33] Intuitively, we can see that asking parents to undergo the dangerous procedure might test parents' commitment to signal (particularly if they are not especially devoted to the religious and cultural values to which FGC is connected). In addition, there are clearly high costs to FGC: the procedure is irreversible, painful, and potentially harmful to health. In line with the general rule that signaling must be very costly, FGC (i.e. the method of signaling religious devotion) is, reportedly, ensured to be "always deliberately intensely painful" (Magesa 1997).

Conceptually, the cost differential between types could appear in numerous ways. For instance, religious families may find it less costly to cut their daughter than non-religious families if FGC is rationalized as an expression of religious devotion. Similarly, families who emphasize chastity and honor may find it less costly to cut their daughter than other families because the relative importance of demonstrating these values is higher among the former group. Note, however, that it is not readily obvious just how large the cost differential is. In cases where the difference is large, we will likely see a separating equilibrium; however, in cases where the difference is small or the potential gain from the signal is large, we may see a pooling equilibrium, in which all agents signal.

If the benefit of signaling is high enough - as is likely the case when dealing with marriage and reproduction - this pooling equilibrium may occur, resulting in a convention or social norm.[34] If such a pooling equilibrium evolves, types are indistinguishable on the basis of the signal, yet the costs of the signal will still be exacted. Although the society may be better off eliminating the signal as a whole, it is too costly for any one individual to not signal and thereby suffer the consequences of

[33]Empirically, such a process has been observed: for instance, among Mandingas in Guinea-Bissau, "the practice has come to be seen as a prerequisite for the ritual purity necessary to pray as well as a marker of belonging to an Islamic community" (Shell-Duncan and Hernlund 2000).

[34]Social conventions are sometimes understood as one equilibrium in a coordination game (Macki 1996). Here, I use the term more generally.

being perceived as the non-preferred type. As such, an important prediction of signaling is that there must be a critical mass of people willing to commit to not signal in order to convince the rest to do the same.[35] Empirically, we have seen support for this implication, since "experience shows that large-scale abandonment can only be expected when FGC is no longer an all-dominant social norm and families can abandon the practice without the risk of stigmatization and exclusion" (Johansen et al. 2013).

It should be noted that testing for signaling as a mechanism will be difficult, particularly if we are in a world where the pooling equilibrium has evolved. Ideally, we would want a measure of underlying type that is observed by the econometrician but not by the rest of society. Unfortunately, social desirability bias makes it incredibly hard to find such a measure.

Another way we could investigate the signaling possibility is to look for multiple markets within Mali, where the value of a signal (e.g. the information communicated by the signal or the consequences of a signal becoming uninformative) varies across markets. Since the value of signaling as a tool to communicate unobserved type depends to some extent on how public it is made, it is possible that geographical variation in the medicalization of FGC might help us test this theory. If FGC is more publicly observable when it is performed by a local circumcisor or a traditional health care provider than when it is performed by a medical professional, then we might expect differences in the value - and thus prevalence - of the signal between these locations.

Unfortunately, the data for such empirical tests are not included in the DHS survey. As such, I cannot speak to signaling in this paper.

[35] Informally, we can think of this as a tipping point, where it is individually too costly to not signal until a large number of community members will do so simultaneously.

4.5.3 Identity

In order to consider the identity hypothesis, I will adopt the framework proposed by Akerlof and Kranton (Akerlof and Kranton 2000). In this general framework, individuals' sense of identity enters their utility function, which also depends on their actions and the actions of others. When a change in identity improves the agent's utility, we consider this as a gain in identity. Similarly, a loss in identity is any change in an agent's sense of self that decreases his overall utility.

The identity component of utility (i.e. the identity function) is a function of numerous components. First, there are societal categories to which agents relate. In the case of FGC, this category might be as narrow as "a virtuous and chaste woman" or as broad as "feminine." These categories have social import and carry certain expectations, images or prescriptions for proper behavior. The expectations associated with womanhood in societies with FGC, for instance, might be virginity, modesty, and cut genitalia, among other characteristics. These requirements and expectations also enter the identity function. Finally, an individual's characteristics enter the identity component of utility, such that the discrepancy between the identity ideal, or expectations, and the individual's actual characteristics determine the overall utility contributed by the identity function.

As a woman's characteristics near those prescribed by society, she will generate gains in identity that improve her utility. Similarly, as more distance between her own characteristics and the ideal emerges, losses in identity occur. Thus, an individual will act to change her characteristics so that they better align with the expectations associated with her social category. It is important to note that, unlike intrinsic type in the signaling model, individual characteristics are malleable, can be changed by an individual's own actions, and may be publicly observable.

Since individuals can increase their utility by acting in accordance with the ideal with which they socially identify, they will choose their actions to maximize their identity, given the social environment and the actions of others. For example, a woman may strengthen her utility by remaining chaste until marriage, bringing her closer to the societal ideal of the group with which she is associated (i.e. females).

In regards to the specific action of FGC, choosing to cut one's own or one's daughters' genitals would change the individual characteristics associated with femininity and, as such, generate gains in identity that affect utility (assuming the individual is a woman who is expected to conform to FGC). In line with this idea, Shell-Duncan and Hernlund have argued that "the procedure of genital cutting is seen by many as essential to the creation of femininity and full adult status" (Shell-Duncan and Hernlund 2000). More explicitly, Johnson argues that among certain groups, FGC "marks the first step in initiating the segregation of the sexes and the construction of gender identity, removing a girl's 'hard male parts' and making a woman forever soft and feminine" (Johnson 2000). The identity hypothesis, then, suggests that women choose FGC because it is internally valuable to their own sense of self. As Kawira, the medical officer in charge of the local hospital in Sharati, Tanzania, expresses, "You see something which seems to be harmful to women and yet women are the ones who do it. The women are the ones who perpetrate it and girls run away from their parents in order to have it done" (Kawira in Skaine 2005).

Since girls are mostly cut when they are young, we must consider the role that identity plays in affecting the parents' decision to cut a daughter. This could be as direct as parents identifying with the category of "good parents," with which performing FGC upon daughters is the associated prescription. However, the link could also be more abstract, where a mother strengthens her identity as a woman by ensuring that her daughters meet the same prescriptions to which she was held,

perpetuating the practice. In addition to advancing one's own identity, some scholars suggest that women engage in FGC to simultaneously attach themselves to a group and then advance that group's identity as a whole: "the bulk of Kono women [in Sierra Leone] who uphold these rituals do so because they want to...embrace the legitimacy of female authority" (Coomaraswamy 2002).

The identity hypothesis generates certain testable implications surrounding FGC. The most significant is that, if FGC induces gains in identity and a strengthened identity is positively associated with FGC rates, decisions surrounding FGC will be self-reinforcing. That is, if a woman cuts one daughter she will be more likely to cut subsequent daughters. As such, we can analyze variation in FGC decisions among daughters of the same mothers to present evidence about the role of identity in FGC's persistence.

4.5.4 Comparing the Hypotheses

Distinguishing between the various hypotheses is, at times, challenging, as the categories likely overlap somewhat in the real world. Here, I will try to make the distinction clearer to the reader. The first difference I will tackle is that between signaling and identity.[36] The clearest distinction between these two hypotheses is that actions taken to improve identity are associated with one's *own* self-image. As such, individuals act to affect their own perception of their characteristics. We can think of this, in a sense, as self-signaling one's identity; that is, demonstrating to oneself that you meet the expectations of your socially assigned category.[37] On the other hand, signaling involves affecting *others'* perception of oneself. Thus, agents

[36]For a more detailed theoretical exploration of this topic in a specific setting, please see Grossman 2015.

[37]Please see Bénabou and Tirole 2011 for the full model of self-signaling.

act in order to affect their community's perception of their characteristics, or type.

This distinction is extremely significant, as the two hypotheses (differentiated by concern for one's own understanding of self versus others' perception of self) will result in different predictions for various situations. For instance, if innovations occur that allow for the dispersion of information about type, the signaling hypothesis would predict a decrease in the rate of FGC, while the identity hypothesis would predict no change.

The second distinction I will clarify is that between the marriage market and signaling hypotheses. The marriage market is a complex process, and the quality of women and men may be communicated in numerous ways. We might expect parents to invest in their child's marriageability in ways that communicate, indirectly, the quality of the child. These investments might well be seen as signals sent to potential mates in the marriage market. For this paper, we will separate such signals from the signaling hypothesis by intention. If a signal is sent in order to directly improve a child's marriage market prospects, this signal will fall under the marriage market hypothesis, which considers how behavior changes in response to supply and demand of potential spouses. On the other hand, signals that are used for purposes other than improving marriage market prospects (e.g. improving the entire family's image in the community to expand networks) will fall under the signaling hypothesis.

Again, this distinction is significant, as in many cases the hypotheses predict different responses to the same event. For instance, the marriage market hypothesis implies that a change in sex-ratio would alter the rates of FGC in response to the shift in market supply and demand. However, if the signaling hypothesis prevails, rates of FGC should not respond similarly to the change, unless it is associated with other underlying changes in the community.

4.6 Empirical Results

4.6.1 How do stated preferences relate to reported outcomes?

The first empirical analysis I perform explores how reported child status relates to the mother's stated preferences. The primary goal of this exercise is to determine whether using reported child status in place of maternal preferences is a useful improvement upon prior literature. To determine the answer, we will first look at the determinants of a daughter's FGC status. We will compare the results from (1) using maternal support as a proxy for daughter's status, as prior literature has done, and (2) using the actual reported status of the daughter (Table 4.12). These regressions will employ the specifications used by Bellemare et al. 2015, to maintain continuity with prior literature.

The controls used in this specification include maternal age, education level and religion; household wealth and assets, including television, radio, and electricity; a full set of maternal ethnicity dummies (suppressed in the table); household location in either an urban or rural area; maternal health knowledge, as assessed by her beliefs on whether HIV is transmitted supernaturally; and two measures of domestic violence. The first measure of domestic violence - labeled "domestic violence, behavior" - is a binary variable that equals one if the mother answered that a man is justified in beating his wife for one or more of the following reasons: burning food, neglecting children, or arguing with her husband. The second measure - labeled "domestic violence, property" - is also a binary variable, and it equals one if the mother believes a husband is justified in beating his wife for one or both of the following reasons: going out without permission or refusing sex.[38]

The first model using this specification considers the relationship between a

[38]These categories were created by (Bellemare et al., 2015) to separate patterns of abuse.

Table 4.12: Determinants of Daughter's FGC Status, Bellemare and Steinmetz Specification

	(1) Mother supports FGC	(2) Mother supports FGC	(3) Daughter underwent FGC	(4) Daughter underwent FGC
Mother underwent FGC	.626***	.529***	.520***	.440***
	(.029)	(.034)	(.020)	(.024)
Mother's age	-.003***	-.001	-.001*	-.001
	(.001)	(.001)	(.001)	(.001)
Mother has primary education	-.021	-.025	-.001	.001
	(.027)	(.028)	(.018)	(.018)
Mother has secondary education	-.060*	-.051	.006	.003
	(.031)	(.033)	(.018)	(.019)
Mother is Muslim	.118***	.065*	.133***	.101***
	(.038)	(.038)	(.031)	(.025)
HH wealth index	.002	.012	.002	.003
	(.009)	(.011)	(.006)	(.006)
HH has TV	-.044*	-.047*	-.018	-.001
	(.024)	(.028)	(.013)	(.013)
HH has radio	.003	.005	-.000	-.001
	(.019)	(.021)	(.011)	(.011)
HH has electricity	.009	-.016	.036**	.019
	(.026)	(.031)	(.018)	(.018)
Domestic violence, behavior	.077***	.083***	.028**	.009
	(.025)	(.026)	(.013)	(.013)
Domestic violence, property	.046*	.014	.007	.026*
	(.024)	(.025)	(.015)	(.014)
Urban	.042		.025	
	(.027)		(.017)	
Mother knows HIV is not supernatural	.022	.014		
	(.019)	(.020)		
Mother unsure if HIV is supernatural	.033	.059*		
	(.028)	(.030)		
Age			.057***	.057***
			(.003)	(.003)
Age^2			-.003***	-.003***
			(.000)	(.000)
Ethnicity FEs	Yes	Yes	Yes	Yes
Cluster FEs	Yes	No	Yes	No
Observations	3039	3033	14506	14506
Clusters	412	406	413	413

Heteroskedasticity-robust standard errors shown in parentheses, clustering by DHS cluster.

mother's support for FGC and her FGC status (model (1) in Table 4.12). The coefficient and standard error estimates are very close to those calculated by Bellemare, Novak and Steinmetz using data from the Gambia. Estimates are generated using a linear probability model.

The second model using this specification considers the relationship between a daughter's FGC status and her mother's status (model (3) in Table 4.12). There are, however, two differences in the covariates. First, variables concerning a mother's beliefs about supernatural transmission of HIV are dropped in this equation because these responses are not linked to the birth record and are insignificant in the first model. Second, a quadratic control for age is included in order to account for the relationship between age and FGC status observed in data (Figure 4.1). As in the previous model, estimates are generated using a linear probability model.

The comparison between these initial results reveals that using reported child status as the outcome variable generates different results than when using maternal preferences. Although the general pattern is similar, the coefficient estimates differ for several variables. For instance, the effect of the mother's own FGC status decreases significantly when considering her daughter's outcome. In addition, measures of domestic violence - which seemed important explanatory variables in the first model - lose magnitude when child status is the outcome variable. The results differ significantly when using reported child status as the outcome variable, suggesting that a more accurate understanding of FGC can be gained from using data from the new FGC module, particularly when analyzing covariates that could influence whether or not the daughter is cut.

In addition to this analysis, we can also assess the relationship between stated preferences and reported outcomes by regressing the latter on a mother's preferences (Table 4.13). This will reveal the basic association between maternal preferences

and her daughter's FGC status. In this analysis, the outcome variable is binary and equals 1 if a daughter has been cut. The covariates include dummy variables for both the mother's and the father's reported preferences. The regression is run with no additional controls, except for age in the second specification.

Table 4.13: Determinants of Daughter's FGC Status, No Controls

	(1) Daughter underwent FGC	(2) Daughter underwent FGC	(3) Daughter underwent FGC	(4) Daughter underwent FGC
Mother supports FGC	0.120***	0.098***	0.137***	0.113***
	(0.017)	(0.018)	(0.019)	(0.019)
Mother opposes FGC	-0.054***	-0.041**	-0.061***	-0.050**
	(0.018)	(0.020)	(0.021)	(0.022)
Father supports FGC	0.017	0.024	0.032**	0.029*
	(0.016)	(0.016)	(0.016)	(0.017)
Father opposes FGC	-0.014	-0.003	-0.009	-0.004
	(0.021)	(0.020)	(0.022)	(0.021)
Daughter's age			0.029***	0.029***
			(0.002)	(0.002)
Daughter's age^2			-0.002***	-0.002***
			(0.000)	(0.000)
Cluster FEs	No	Yes	No	Yes
Observations	11842	11842	10345	10345
Clusters	411	411	411	411

Heteroskedasticity-robust standard errors are reported in parentheses, clustered by DHS cluster.

Table 4.13 reports the results of this regression.[39] The coefficient estimate for the variable indicating that a mother supports FGC ranges from 0.098-0.137, depending on age controls and cluster fixed effects. Maternal support for FGC is evidently a strong but indefinitive predictor of intergenerational transmission of FGC. This is significant, as it indicates that using maternal support as a proxy for FGC outcomes - the primary focus of prior literature - is insufficient. As such, it is likely that previous

[39]This result is robust to using the log of age in place of squared age (Table C2).

research may have overlooked important aspects of the practice.

Given this and the earlier findings, we can conclude that results will differ depending on whether we use reported outcomes or stated preferences and that the results will be more accurate when using child FGC status itself, rather than a proxy. As such, there is room for improvement upon prior literature by using this new measure.

4.6.2 Are fathers the unilateral decision-makers?

The next empirical analysis I perform examines the role that maternal preferences have in the final FGC outcome. In particular, I wish to assess whether a mother's preferences have a significant impact on her daughter's FGC status. This will reveal whether or not mothers have an influence on FGC decisions, separate and apart from their husbands. Such influence was assumed by prior literature that used maternal preferences as a proxy for child status. However, this is an assumption that has been widely questioned in popular media, which often portrays FGC as a unilateral male decision.[40] As such, it is wise to empirically test whether or not mothers have influence in the decision-making process.

To explore the influence mothers have, I consider the relationship between a daughter's FGC status and her parents' preferences (Equation 4.1). The outcome variable, Y_i, is a binary variable equaling one if the child has undergone FGC. X_i is a vector of individual controls, including observable characteristics of the mother, father and household as a whole. P_i is a vector of dummy variables representing parental preferences for FGC. Individual idiosyncratic error is represented by ϵ_i. The coefficients of interest - Π - indicate the influence of parental preferences in a linear probability model.

[40]Karimjee 2015 is a specific journalistic account that strongly supports the model of an all-male decision-making process.

$$Y_i = \alpha + X_i\beta + P_i\Pi + \epsilon_i \tag{4.1}$$

The estimates of this model are displayed Table 4.14.

Table 4.14: Determinants of Daughter's FGC Status

	(1) Daughter underwent FGC	(2) Daughter underwent FGC
Mother supports FGC	.266***	.232***
	(.030)	(.030)
Mother opposes FGC	-.138***	-.110***
	(.039)	(.039)
Father supports FGC	.046	.015
	(.029)	(.029)
Father opposes FGC	-.012	-.022
	(.036)	(.035)
Cluster FEs	No	Yes
Observations	5058	5055
Clusters	411	408

Estimates include controls for maternal and paternal education, maternal and paternal religion, maternal prenatal care, maternal health care in the past year, maternal contraceptive use, maternal employment, age of household head, urban/rural location of household, electricity availability of household, household wealth index, daughter's age and daughter's age squared. Heteroskedasticity-robust standard errors are reported in parentheses clustered by DHS cluster.

The results of this model demonstrate that maternal preferences are much more closely related to child FGC status than paternal preferences. This result is inconsistent with a model where a father is the unilateral FGC decision-maker. It also suggests that mechanisms for which mothers and fathers are likely to have similarly strong incentives might be insufficient.

Although this model indicates that maternal preferences are important determinants of a daughter's cutting status, there is evidently concern about the validity of stated preferences. Maternal stated preferences, for instance, might be systemati-

cally biased. This is especially important in the case of FGC, as it seems reasonable to believe that women might report opinions that align with past behavior, social perception of FGC, and, most likely, spousal preferences. As a robustness check, I repeat the estimation, restricting the regression to couples with strongly incompatible preferences (Equation 4.2).

$$i \in \{i \mid (p_m = 1 \;\&\; p_f = 0) \; or \; (p_m = 0 \;\&\; p_f = 1)\} \tag{4.2}$$

The results of this estimation are displayed in Table 4.15. A father's support for FGC - which, in this subset, equals a mother's opposition to FGC - has a significant negative effect on the probability a daughter will be cut. Thus, even when parents have opposite views on FGC, the mother's opinion seems to dominate. Although it cannot eliminate all concerns about the use of stated preferences, it does provide additional evidence that mothers can have a decisive impact on their daughters' FGC status.

Special Case: Polygynous Households

As was demonstrated in Table 4.4, a large number of Malian families are polygynous. This fact is the basis of my next empirical strategy, which restricts attention to polygynous families. Doing this allows me to exploit variation in maternal FGC preferences and child FGC status, while holding the father's influence fixed. In addition, because the father's preferences are absorbed by fixed effects, I do not need to control for them directly. Thus, I can use observations for which the father questionnaire was not completed. As such, even though the restriction of polygyny restricts the sample significantly, the results can still achieve a relatively high level of precision.

Table 4.15: Determinants of Daughter's FGC Staus, Parents who Disagree

	(1) Daughter underwent FGC	(2) Daughter underwent FGC
Father supports FGC	-.370***	-.334***
	(.045)	(.086)
Cluster FEs	No	Yes
Observations	903	858
Clusters	232	187

The variable indicating that the father supports FGC is the equivalent of a variable indicating that the mother opposes FGC in this specification, as the sample is restricted to couples who disagree. Estimates include controls for maternal and paternal education, maternal and paternal religion, maternal prenatal care, maternal health care in the past year, maternal contraceptive use, maternal employment, age of household head, urban/rural location of household, electricity availability of household, household wealth index, daughter's age and daughter's age squared. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

The model is specified by Equation 4.3. Observations are indexed by individual i, mother m, and father f. The outcome variable, Y_{imf} is again a binary variable equaling one if the daughter is cut. X_{im} is a vector of individual controls. P_m is a binary variable equaling one if the mother supports FGC. The specification also includes father fixed effects: μ_f. Individual error is represented by ϵ_i. The coefficient of interest - π - indicates the association of maternal support for FGC with daughter's cutting status.

$$Y_{imf} = \alpha + X_{im}\beta + P_m\pi + \mu_f + \epsilon_i \tag{4.3}$$

The results of this estimation are presented both with and without controls in Table 4.16.[41] The estimates indicate that, even when we hold a father's unobserved

[41] In this regression, the relevant reference group consists of women who oppose or hold ill-defined preferences on FGC. If the reference group is changed to women without clear preferences on FGC, the results remain consistent, but it appears that women who oppose FGC are also more

influence and other household characteristics fixed, the mother's preferences still have a significant effect on the decision to cut her daughter. Thus, the major result of this empirical model is that uxorial tastes still have a statistical effect when paternal influence is fully absorbed. This suggests we can reject a model where the father makes FGC decisions unilaterally.

Table 4.16: Determinants of Daughter's FGC Status, in Polygynous Families

	(1) Daughter underwent FGC	(2) Daughter underwent FGC
Mother supports FGC	.132***	.136***
	(.046)	(.045)
Maternal controls	No	Yes
Father FEs	Yes	Yes
Observations	5230	4972
Fathers	1531	1431
Clusters	384	376

Both columns control for daughter's age and daughter's age squared. Column 2 controls for the following maternal characteristics: rank among co-wives, religion, education, and employment. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

4.6.3 Can the identity hypothesis fully account for observed behavior?

An important benefit of using child-level FGC status rather than maternal preferences is that it allows us to determine if there is variation in cutting status among daughters of the same mother. As is demonstrated in Figure 4.3, there is a significant amount

likely to cut their daughter. This indicates that women with ill-defined preferences are less likely to report cutting their daughter. This could be due to a difference in reporting or in actual behavior, such as women delaying the FGC decision if they have not yet reached a strong opinion. It could also reflect a pattern where women form strong preferences after observing the outcomes of FGC among their daughters (e.g. a woman strongly opposes FGC only after observing negative health effects).

Table 4.17: Transition Matrix: Cutting Status, Daughters 1 and 2

		Daughter 2	
		Uncut	Cut
Daughter 1	Uncut	23.405%	15.217%
	Cut	1.779%	59.599%

of within-mother variation. In addition, Table 4.17 shows that a significant number of women cut their second daughter after leaving their first daughter uncut.[42]

Figure 4.3: Percent of a Mother's Daughters, 8 and Older, Who Are Cut

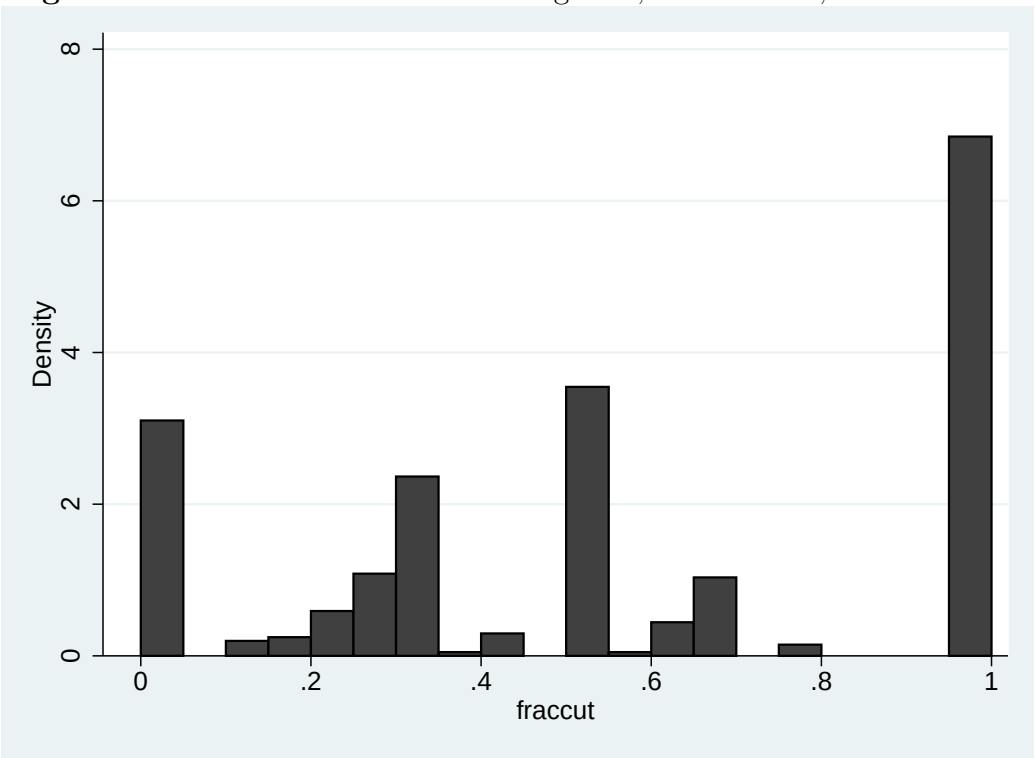

Note: This figure shows the distribution of a variable reporting the fraction of a mother's daughters (8 years and older) who are cut.

The high variation among daughters of the same mother helps shed light on why FGC persists. First and foremost, it casts doubt on explanations, such as simple

[42]This table does not control for age. However, since cutting is linked to age, we can expect the elder child (child 1) to have passed the decision point for cutting by the time the younger child (child 2) is cut. As such, the number of mothers who cut their second daughter but not their first daughter is likely underestimated.

convention, that imply mothers would treat their daughters similarly. Furthermore, these results can help us analyze the mechanism of identity. Although identity seems the mechanism most likely to explain the differential impact of maternal and paternal preferences on their daughters' FGC status, these results suggest that identity - a mechanism in which the decisions to act in accordance with one's identity reinforce one's identity and thus likely lead to similar decisions in the future - might be an insufficient explanation.

4.6.4 Can the marriage market hypothesis fully account for observed behavior?

In order to assess how important the marriage market is in maintaining the practice of FGC, we can analyze marriage market outcomes by FGC status. This will not give us a definitive answer about the importance of marriage markets in the decision-making process, but it should give us a very rough idea of how large a role marriage plays in the persistence of FGC.

In order to see if there are significant differences in marriage outcomes between girls who are cut and those who are not, I regress a variety of plausible measures of match quality on cutting status. I control for age and height-for-age[43] quadratically. Height-for-age is included in order to account, in some part, for the stock of health and the standard of living to which the woman was exposed prior to marriage - two factors which might affect marriageability. In addition, I control for geographical variation in marriage market conditions in two ways: first, by including a binary variable indicating if the woman currently resides in an urban area; second, by including cluster fixed effects. In all cases, the results of these two regressions were similar and,

[43]This is a variable that reports the number of standard deviations a woman is from the worldwide median for women of similar age and height.

as such, I report the former[44].

The first measure I use to analyze marriage market outcomes is a binary variable indicating whether a woman is currently married. Table 4.18 reports the results of this regression.[45] The first column of this table reports the results without controlling for education, while the second controls for whether or not the woman has completed primary and/or secondary education. Note that controlling for education will produce problems if it is co-determined with marriage. The large negative coefficient on secondary education in the second column is consistent with an environment in which this is the case: a girl leaves school upon marriage, those who complete a secondary education are those who were not selected for or did not select into marriage earlier, and, as such, women with a secondary education are less likely to marry.[46] The coefficient on FGC status is positive but marginally insignificant. As such, it does not appear that cutting status is associated with a large improvement in the probability of entering a marriage union.

The second measure I use to analyze marriage market outcomes is the rank a woman has among her co-wives. If we regress wife rank on FGC status for the entire sample, the estimates are likely to be biased, as societies that engage in polygyny may systematically engage in more or less FGC than monogamous societies.[47] As such, I will restrict the sample to women who are already in polygynous unions. It is important to note that this specification excludes the first wives of men who will, at

[44]The regression tables with cluster fixed effects are available upon request.

[45]The regression is performed as a linear probability model, but it is robust to using probit and logit (Tables C6 and C7, respectively).

[46]In addition to regressing a woman's current marital status on her cutting status, I also regress whether a woman has ever been married on her FGC status. The results are very similar to those above, except that the effect of undergoing FGC becomes marginally significant at the 10% level. Results are available upon request.

[47]In addition, as explained by the marriage market hypothesis in Section 5.1, any marriage market conditions that cultivate polygyny (e.g. a skewed sex-ratio) will also likely alter the rates of FGC.

Table 4.18: Regression of Marriage on Cutting Status

	(1) Currently married	(2) Currently married	(3) Currently married	(4) Currently married
Underwent FGC	0.017	0.025	0.014	0.022
	(0.017)	(0.019)	(0.016)	(0.019)
Education Controls	No	No	Yes	Yes
Cluster FEs	No	Yes	No	Yes
Observations	5346	5346	5346	5346
Clusters	413	413	413	413

All columns include controls for the age and age squared, the height-for-age and height-for-age squared, and the urban/rural location of the household. Columns 3 and 4 control for education level. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

some time in the future, marry one or more additional wives. Table 4.19 shows the result of this regression for the restricted sample.[48] The coefficient on cutting status is statistically insignificant[49], indicating that FGC status does not seem to have a significant effect on a wife's rank. To the extent that rank reflects match quality or quality of life in the marriage union[50], it does not seem that FGC status is associated with a large improvement in marriage market outcomes.

The third measure I use to analyze marriage market outcomes is the woman's wealth index. Evidently, there are major problems with using wealth index as the dependent variable, since the wealth a woman experiences in childhood might well influence parental opinions on FGC and marriage market outcomes. However, we will use this as an imperfect proxy for the husband's income, since the wealth index reports the wealth of the household in which the woman now lives. Table 4.20 shows

[48]These results are robust to using negative binomial and Poisson regressions (Tables C8 and C9, respectively).

[49]When cluster fixed effects are used, the point estimate is larger, but the coefficient remains insignificant.

[50]This is plausible, as the primary co-wife in polygynous families often has a disproportionate amount of power over her co-wives, and domestic violence by women against their lower-ranked co-wives is considered prevalent.

Table 4.19: Regression of Wife Rank on Cutting Status, Polygynous Wives Only

	(1) Wife's Rank	(2) Wife's Rank	(3) Wife's Rank	(4) Wife's Rank
Underwent FGC	-0.413	-1.724	-0.439	-1.677
	(0.975)	(1.636)	(0.955)	(1.560)
Education Controls	No	No	Yes	Yes
Cluster FEs	No	Yes	No	Yes
Observations	1567	1506	1567	1506
Clusters	366	305	366	305

All columns include controls for the age and age squared, the height-for-age and height-for-age squared, and the urban/rural location of the household. Columns 3 and 4 control for education level. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

the regression results. There is no significant association between FGC status and wealth in adulthood. Again, this is not a definitive measure of the impact of FGC on husband quality, as determined by the husband's income. Even so, the lack of a result is consistent with the lack of large marriage market effects found in the regressions above.

Table 4.20: Regression of Wealth Index on Cutting Status

	(1) Wealth index	(2) Wealth index	(3) Wealth index	(4) Wealth index
Underwent FGC	0.040	-0.030	0.044	-0.028
	(0.074)	(0.053)	(0.072)	(0.052)
Education controls	No	No	Yes	Yes
Cluster FEs	No	Yes	No	Yes
Observations	5346	5346	5346	5346
Clusters	413	413	413	413

All columns include controls for the age and age squared, the height-for-age and height-for-age squared, and the urban/rural location of the household. Columns 3 and 4 control for education level. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

Although we cannot rule out the importance of the marriage market in the persistence of FGC, it seems unlikely from these estimates that the marriage market is the main driving force behind the practice.

4.7 Conclusion

The empirical tests in this paper do not provide a definitive answer to why FGC persists. However, the results do reveal significant information about the practice of FGC. First, the results suggest that using maternal preferences to proxy for child FGC status may lead us to false conclusions about the practice. There is evidence of a significant amount of within-mother variation in whether daughters undergo FGC. Thus, characterizing daughters' experiences unilaterally by their mother's opinion overlooks important differences in how children are exposed to FGC. In addition, the within-mother variation suggests that decisions regarding FGC depend on the individual characteristics of the child or other conditions not necessarily shared by siblings. This may indicate a level of strategic decision-making around the practice, suggesting that FGC might persist for reasons other than social convention. Furthermore, although maternal opinions and child-cutting status are positively correlated, the relationship is far from perfect, suggesting that using maternal opinions in place of actual cutting outcomes may bias the analysis. As such, this paper uses reported child status to better measure the practice and suggests the use of this measure in future research.

Second, the results suggest that mothers do have some decision-making power in choosing FGC for their daughters, even when accounting for paternal preferences. This is important descriptive information as it is not a priori clear that women would have significant influence in this decision in societies where FGC is common. This information casts doubt on work that characterizes FGC as a decision made unilaterally by fathers or elder males. In addition, it has potential policy implications: policy designed to reduce or eliminate the incidence of FGC may be best directed to mothers. Furthermore, the information that women have influence above and beyond their

husbands helps shed light on the possible mechanisms that keep FGC common. This information might suggest that models where mothers and fathers have equal stake in the daughter's cutting status - such as marriage market models where bride prices are either shared or dominated by males - might be insufficient to explain FGC's persistence. That is, the results suggest that there is room for other mechanisms to explain residual variation in cutting status.

Third, the results suggest that the difference in marriage market outcomes due to FGC is relatively small. This is suggestive evidence that, although the marriage market likely plays a role in maintaining FGC, it cannot fully account for the persistence of the practice. Furthermore, these results suggest that policy which unilaterally targets the marriage market will likely fail. A popular policy suggestion has been to follow the program used to end footbinding in China: in each community, convince a small group of mothers to promise that they will not allow their sons to marry women who have undergone FGC (Macki 1996). Cultivating demand for uncut wives in this way may affect FGC rates if the marriage market is the dominant force. However, in environments where other mechanisms - such as identity and signaling - are also influential, such a program will likely fail. If the identity mechanism is strong, women value FGC for their own sense of self and, as such, a small change in marriage market demand will unlikely eliminate FGC. In addition, if signaling is strong, creating a small pocket of opposition to FGC will likely be ineffective, as it is too costly to deviate from the practice until the signal is no longer the dominant social means to communicate extrinsically valuable information.

Fourth, variation in cutting status among daughters of the same mother casts doubt on a model where identity is the sole driving force behind FGC. In cases of identity and self-signaling, we expect FGC decisions to be self-reinforcing. As such, we would expect continuity in FGC decisions across daughters. However, there is

evidence that a high percentage of women make different FGC choices for different daughters. As such, this is evidence that the role of identity cannot fully explain observed FGC outcomes. Given this result, policy which targets FGC solely through its role in identity will also likely fail. Instead, the findings suggest that a multilateral policy, which targets a full range of possible mechanisms - including marriage, identity, and signaling - for FGC's persistence is required.

Chapter 5

Conclusion

This dissertation considers the role gender plays in labor markets, household decision-making, and health in sub-Saharan Africa.

The first chapter considers the impact of fast Internet access on employment outcomes and household dynamics. I find the introduction of fast Internet to sub-Saharan Africa significantly increased employment for males, but had little impact on female employment. Most of the effect on male employment came from an increase in self-employed agriculture. In addition, the introduction of fast Internet significantly increased perceived acceptability, among both genders, of domestic violence against women.

The second chapter considers the differential impact, by gender, of an experimental labor market intervention in South Africa, which measured skills of workseekers and provided a mechanism for workseekers to communicate their results to potential employers. I find that men experienced a larger effect of the intervention on employment outcomes than did women. This difference is largely explained by pre-existing differences between genders, rather than differential responses to treatment.

The third chapter considers the factors that contribute to female genital cutting (FGC) in Mali and tests various hypotheses to explain the persistence of the tradition. I find that maternal preference is pivotal in the decision to cut daughters; in fact, maternal preference is more predictive of a daughter's FGC status than is paternal preference. In addition, I find variation in FGC status among daughters of the same mother, suggesting that the identity hypothesis alone is insufficient to explain FGC

decision-making. Finally, I find that undergoing FGC as a child increases neither the probability of ever being married nor wife rank in polygynous households, suggesting that the marriage market hypothesis alone is also insufficient to explain FGC decision-making.

Appendix A

Appendices to Chapter 2

A.1 Internet Use

The introduction of submarine cables to the African continent increased the speed at which African users could access the Internet. This introduction may have had effects at both the extensive margin, by increasing the number of Internet users, and the extensive margin, by increasing the amount of content accessed by those who were already using the Internet. It is not possible to directly measure the impact of the submarine cables on either margin for the individuals in the sample, as the DHS did not collect information on Internet access or use in the relevant survey waves.

Figure A1 presents country-level Internet use statistics for the period of analysis. These measures are constructed using data from the International Telecommunication Union and World Bank (2019). The rate of Internet use increased significantly in the time period after the arrival of submarine cables. It is also important to note that, although the rate of Internet use was low even after the introduction of fast Internet, these statistics are at the country level; as such, the rate of Internet use in connected areas is likely significantly higher.

Another pattern that is informative for this study is how men and women differ in their Internet use. Again, this data is not available for the individuals in the sample, as the DHS did not collect information on Internet use during the period of analysis. However, subsequent waves of the DHS did collect individual information on Internet use and frequency of this use. This data is available only for 2 countries used in this

Figure A1: Internet Use Over Time, by Country

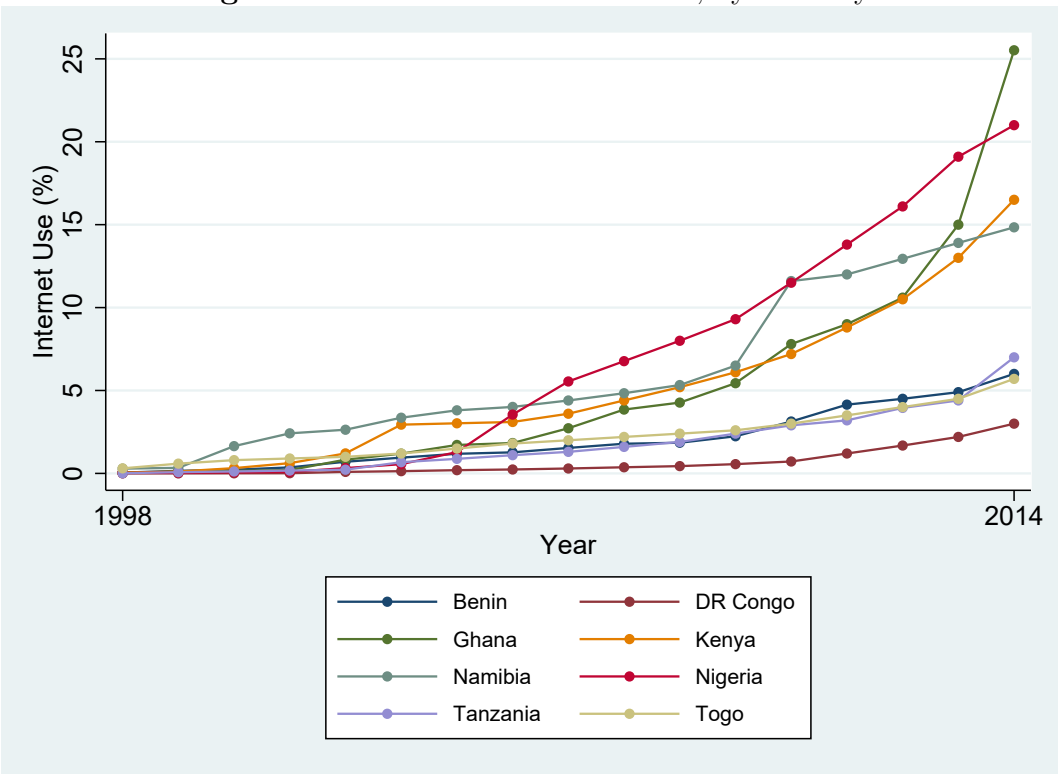

This figure uses data from International Telecommunication Union, World Bank Development Indicators.

analysis, Benin and Tanzania, and the data cannot be tied to the household results used in this survey, since the DHS is a repeated cross-section.

For these two waves of the DHS, it is possible to calculate Internet use statistics for connected individuals, or those who were within 500 meters of the nearest terrestrial backbone cable, using the cable maps from the main analysis. Table A1 presents these Internet use statistics, separated by gender. There is a large gender gap in Internet use: while 42.8% of men in connected areas have used the Internet in the past year, only 16.5% of women in these areas have. Among both men and women, those who use the Internet do so frequently. 24.1% of men and 10.8% of women use the Internet almost every day.

Table A1: Internet Use of Connected Individuals

	Male	Female
Has Used Internet Ever	0.450	0.186
Has Used Internet in Past Year	0.428	0.165
Has Used Internet Over a Year Ago	0.022	0.021
Uses Internet Almost Every Day	0.241	0.108
Uses Internet At Least Once a Week	0.115	0.035
Uses Internet Less than Once a Week	0.053	0.017

This table uses data from the DHS, Benin 2017 and Tanzania 2015.

The primary way for Internet users in developing countries to access the Internet is via mobile phones. Table A2 shows country-level statistics on mobile phone subscriptions per 100 people, for the period of analysis. These measures are constructed using data from the International Telecommunication Union and World Bank (2019). The number of mobile phone subscriptions per 100 people increased significantly in the time period after the submarine cables were connected to the national backbones. It is important to note that these statistics on mobile phones include both smartphones and basic mobile phones that do not connect to the Internet.

Smartphone ownership is only a fraction of total mobile phone ownership in these

Figure A2: Mobile Phone Subscriptions Over Time, by Country

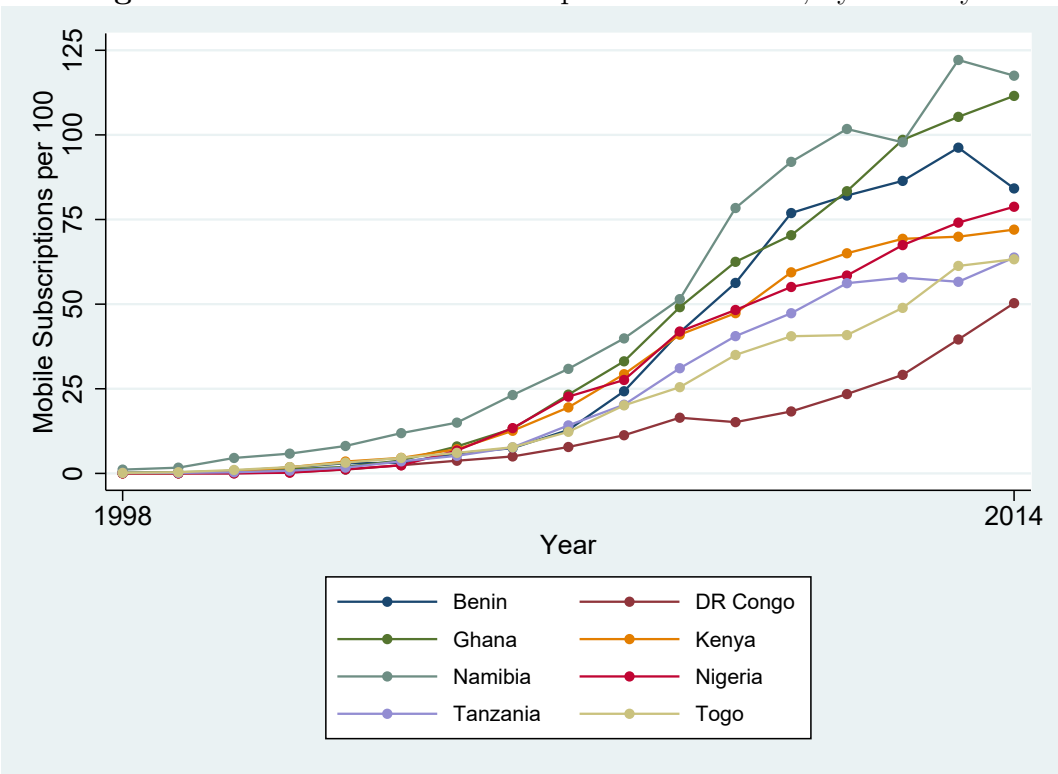

This figure uses data from International Telecommunication Union, World Bank Development Indicators

countries. The Pew Research Center (2015) collected data on rates of smartphone ownership among individuals living in 4 of the countries included in this study; this data was for the year 2014, the last year included in this study's time frame. Table A2 presents these statistics at the country-level. Although smartphone ownership was nowhere near universal in any of these countries at this time, it was certainly present, particularly among younger individuals. Furthermore, as these statistics are at the country-level, it is likely that smartphone ownwership in areas connected to the terrestrial backbones was higher, as use of a smartphone would be more valuable in these locations.

Table A2: Smartphone Ownership, by Country

	Total	Ages 18-34
Ghana	14%	21%
Kenya	15%	18%
Nigeria	27%	34%
Tanzania	8%	10%

This table uses data from Pew Research Center (2015).

As part of this study, the Pew Research Center (2015) also asked mobile phone users about the activities they conduct using their phones. Table A3 shows the percentage of individuals surveyed in each country that used their mobile phone for various activities in the past 12 months. In all countries, mobile phone owners used their devices for SMS messaging and to take pictures and videos. In both Kenya and Tanzania, countries with high adoption of M-Pesa, mobile phones were often used to make or receive payments. In addition, although such use was less common, mobile phone users were actively using their devices to connect with social networks; search and apply for jobs; and seek out information, including political news, health information, and consumer information. As such, it is reasonable to

believe that an increase in Internet speed may affect how individuals search for work, which information individuals consume, and how individuals connect with their social network, even if they only access this content from their phone.

Table A3: Uses of Mobile Phones, by Country

	SMS	Pictures & Video	Make/Receive Payments	Politcal News	Social Networks	Health Information	Search/Apply For a Job	Consumer Info
Ghana	51%	48%	15%	18%	15%	11%	9%	8%
Kenya	88%	54%	61%	28%	28%	19%	26%	14%
Nigeria	80%	57%	15%	28%	35%	23%	15%	14%
Tanzania	92%	53%	39%	21%	19%	11%	14%	12%

This table uses data from Pew Research Center (2015).

A.1.1 Top Websites by Country

To understand how the introduction of submarine cables to the African backbone affected the speed and content African Internet users were accessing, it is important to know which sites were frequented by these users. Since the speed increase was most dramatic for content hosted in Europe and North America, knowing the server locations of popular websites provides insight into whether and to what extent the submarine cables affected content viewed by African users. In addition, as the type of content accessed may influence the labor market, household decision-making, and domestic violence impacts of fast Internet, it is important to know the categories of websites that were accessed most frequently.

To construct lists of the top 15 websites in each country, I used Wayback Machine (2019) to view archived versions of the top sites lists for each country from Alexa (2019). The list of top websites is based on page views over the 3-months prior to the date archived. If information was available on a country's top sites during the second survey wave for that country, I used the archived version closest to the end of that survey wave. If this information was not available during the second survey wave,

I instead chose the first archived version of the list that appeared after the second survey wave; for many countries, this occurred years after the second survey wave concluded. Once I obtained the top 15 sites for each country, I then used Checker (2019) to determine the server location of each site. Although the server location is that of the website on the date Checker (2019) was accessed (Oct. 14, 2019), the server is likely in the same country as it when the top site information was collected. I determined the type of content of each site manually. Information was available on top sites for every sampled country except Togo.

Tables A1-7 present information for the top 15 websites in each country, except Togo. In addition, the table shows the server location for each website, as well as the type of content each site provides.

Table A4: Top 15 Websites in Benin (Dec 5, 2016)

Rank	Website	Category	Server Country
1	Google.com	Search Engine	USA
2	Youtube.com	Video	USA
3	Google.bj	Search Engine	USA
4	Yahoo.com	Search Engine	USA
5	Google.fr	Search Engine	USA
6	Cpasbien.cm	Video	USA
7	Wikipedia.org	Information	USA
8	Live.com	Search Engine	USA
9	Msn.com	News	USA
10	Ask.com	Search Engine	USA
11	CommentCaMarche.net	Information	France
12	OnClickAds.net	Advertising	Netherlands
13	SaveFrom.net	Video	Germany
14	Blogspot.com	Information	USA
15	Alibaba.com	Shopping	USA

Table A5: Top 15 Websites in Democratic Republic of Congo (Dec 5, 2016)

Rank	Website	Category	Server Country
1	Google.com	Search Engine	USA
2	Google.cd	Search Engine	USA
3	Youtube.com	Video	USA
4	Yahoo.com	Search Engine	USA
5	Google.fr	Search Engine	USA
6	Wikipedia.org	Information	USA
7	MediaCongo.net	News	Switzerland
8	RadioOkapi.net	News	USA
9	Cpasbien.cm	Video	USA
10	Ask.com	Search Engine	USA
11	OnClickAds.net	Advertising	Netherlands
12	Msn.com	News	USA
13	MyWay.com	Search Engine	USA
14	Live.com	Search Engine	USA
15	CommentCaMarche.net	Information	France

Table A6: Top 15 Websites in Ghana (Oct 20, 2014)

Rank	Website	Category	Server Country
1	Facebook.com	Social Media	USA
2	Google.com.gh	Search Engine	USA
3	Google.com	Search Engine	USA
4	Yahoo.com	Search Engine	USA
5	Youtube.com	Video	USA
6	Ask.com	Search Engine	USA
7	GhanaWeb.com	News	USA
8	MyJoyOnline.com	News	Canada
9	Wikipedia.org	Information	USA
10	Live.com	Search Engine	USA
11	Tonaton.com	Shopping	USA
12	Twitter.com	Social Media	USA
13	Amazon.com	Shopping	USA
14	Goal.com	Sports	UK
15	LinkedIn.com	Job Search	USA

Table A7: Top 15 Websites in Kenya (Oct 15, 2014)

Rank	Website	Category	Server Country
1	Google.com	Search Engine	USA
2	Facebook.com	Social Media	USA
3	Youtube.com	Video	USA
4	Yahoo.com	Search Engine	USA
5	Google.co.ke	Search Engine	USA
6	El-Balad.com	News	USA
7	Twitter.com	Social Media	USA
8	StandardMedia.co.ke	News	USA
9	Wikipedia.org	Information	USA
10	Blogspot.com	Information	USA
11	Nation.co.ke	News	UK
12	Ask.com	Search Engine	USA
13	LinkedIn.com	Job Search	USA
14	Amazon.com	Shopping	USA
15	Olx.co.ke	Shopping	USA

Table A8: Top 15 Websites in Namibia (May 19, 2017)

Rank	Website	Category	Server Country
1	Google.com.na	Search Engine	USA
2	Google.com	Search Engine	USA
3	Youtube.com	Video	USA
4	Facebook.com	Social Media	USA
5	Yahoo.com	Search Engine	USA
6	UNAM.na	Education	-
7	Wikipedia.org	Information	USA
8	NUST.na	Education	Namibia
9	Onclkds.com	Advertising	Netherlands
10	Polytechnic.edu.na	Education	-
11	UNAM.edu.na	Education	Namibia
12	Instagram.com	Social Media	USA
13	Ask.com	Search Engine	USA
14	AliExpress.com	Shopping	USA
15	FNBNamibia.com.na	Banking	South Africa

Note: UNAM.na and Polytechnic.edu.na no longer exist, but appear to be former versions of UNAM.edu.na and NUST.na, respectively.

Table A9: Top 15 Websites in Nigeria (May 6, 2013)

Rank	Website	Category	Server Country
1	Facebook.com	Social Media	USA
2	Google.com.ng	Search Engine	USA
3	Google.com	Search Engine	USA
4	Yahoo.com	Search Engine	USA
5	Youtube.com	Video	USA
6	Blogspot.com	Information	USA
7	Nairaland.com	News	USA
8	MyWebSearch.com	Search Engine	USA
9	Twitter	Social Media	USA
10	Amazon.com	Shopping	USA
11	LinkedIn.com	Job Search	USA
12	Wikipedia.org	Information	USA
13	VanguardNGR.com	News	USA
14	PunchNG.com	News	USA
15	Jumia.com.ng	Shopping	USA

Table A10: Top 15 Websites in Tanzania (Feb 2, 2016)

Rank	Website	Category	Server Country
1	Google.co.tz	Search Engine	USA
2	Google.com	Search Engine	USA
3	Yahoo.com	Search Engine	USA
4	Facebook.com	Social Media	USA
5	Youtube.com	Video	USA
6	Ask.com	Search Engine	USA
7	Blogspot.com	Information	USA
8	ZoomTanzania.com	Shopping	USA
9	JamiiForums.com	Information	USA
10	Wikipedia.org	Information	USA
11	MyStart.com	Search Engine	USA
12	Amazon.com	Shopping	USA
13	MillardAyo.com	Sports	France
14	XVideos.com	Pornography	Netherlands
15	Alibaba.com	Shopping	USA

A.1.2 Agricultural Searches

To understand whether individuals were taking advantage of faster Internet to search for agricultural information, it would be ideal to know how frequently informative agricultural websites were accessed. However, although these agricultural websites may be influential, they are not necessarily popular enough to be included in Alexa's list of top websites by country. However, if a website receives a substantial number of visits from search engines, Alexa catalogs the top search terms that lead to the site, even if the site does not appear in the top website list. These search terms allow us to see how individuals are finding agricultural websites through search terms, as well as the type of agriculture information individuals look for on the Internet.

Table A11 lists a number of popular agricultural websites in Kenya and the top search terms that led Kenya Internet users to that site. A number of types of searches lead users to these websites. Searches for specific crops, such as "soursop," "French beans," and "green grams," are a large portion of the most popular search terms, as are searches for crop-specific farming, such as "chili farming," "sugar beet growing," and "how to grow coffee." Another key type of search is for specific agricultural techniques and technology, such as "walking tractor," "sunken boosters," and "solar water pumping." Internet users are also searching for specific problems that are encountered in agriculture, such as "bacterial wilt," "compatibility chart of insecticides and fungicides," and "coryza." Finally, some Internet users are searching for information about markets for their products, as evidenced for searches for "flower consolidators."

It is important to note that search engines are only one way in which Internet users access agricultural information. Users may also obtain information via the Internet through email, social networking sites, forums, and by directly accessing

agricultural websites without the use of a search engine. In addition, Internet users may access sites, such as news pages or online encyclopedias, that offer agricultural information within a broader set of knowledge and, as such, are not easily identified as key agricultural sites through tools such as Alexa.

Nevertheless, these search engine results suggest that individuals are searching for a wide range of agricultural information on the Internet and are able to obtain this information through relevant websites. Not only do users seek generalized information about crops and farms, but they also look for specialized information on technology to improve their agricultural productivity, solutions to problems that would otherwise impede their productivity, and information on how to take their product to market.

Table A11: Agricultural Websites and Search Terms, Kenya

Website	Keywords
selinawamucci.com	Soursop, French Beans, African Horned Melon
farmlinkkenya.com	Brachiaria Grass, Coryza, Ground Nuts
mfarm.co.ke	Green Grams, Chili Farm, Bacterial Wilt
farmerstrend.co.ke	Hass Avocado, Silage, Chilli Farming, Dairy Farming
topfarmer.co.ke	Flower Consolidators, Chili Farm, Grevillea Tree
farmbizafrica.com	Chilli Farming, Walking Tractor, Napier Grass, Flower Consolidators, Sunken Boosters
graduatefarmer.co.ke	Dairy Farming, Cabbage Farm, Wambugu Apples, Solar Water Pumping
oxfarmorganic.com	Hass Avocado, Wambugu Apples, Paw Paw Farm, Apple Tree
yagrein.blogspot.com	Chilli Farming, Garlic Farming, Agro Farms, Dairy Farming
realipm.com	Metharhizium Anisopliae Products, Compatibility Chart of Insecticides and Fungicides, Bacillus Subtillus Fungicide
smartfarmerkenya.com	Sugar Beet Growing, Onion Farming, How to Grow Coffee

The data in this table is from Alexa for 2019. Alexa returns the 5 top search terms that lead a country's Internet users to the website. In the case of redundant searches (e.g. "French beans" and "French bean" or "onion farming" and "onion farm"), the more popular of the redundant searches is included in the table.

A.2 Additional Descriptive Statistics

A.2.1 Sample Selection

The DHS surveys both male and female individuals, but its focus is on female respondents. As such, most countries conduct the male survey at only a random subsample of households that are eligible for the female survey. Table A12 shows the size of the male-eligible subsample, relative to the female-eligible sample, for each country and wave used in this paper. The designation of a household as female-eligible simply means that a female survey will be conducted along with the household survey, if at least one female over the age of 15 is present; however, a household need not have any female members to be female-eligible. As such, male surveys can be conducted in household without females, as long as they are chosen as part of the random subsample of female-eligible households. For many, but not all, waves and countries, the female-eligible sample is the entire household-eligible sample.

The male and female response rates differ in most countries and waves of the DHS. Table A13 shows the urban and rural response rates for both men and women in every country and wave used in this analysis. In all cases, the male response rate is lower, since men have higher rates of employment outside the household than women. The magnitude of this gender difference, however, differs significantly across countries.

A.2.2 Balance

Table A14 shows the balance on demographics between unconnected and connected observations prior to the arrival of any submarine cables. Column 1 shows the unconditional mean of the relevant variable among unconnected male observations before

Table A12: Relative Size of DHS Male Subsample

Country	Year	Male-Eligible HHs (Fraction of Female-Eligible HHs)
Benin	2001	1/2
	2012	1/3
DR Congo	2007	1/2
	2013	1/2
Ghana	2008	1
	2014	1/2
Kenya	2008	1/2
	2014	1/2
Namibia	2006	1/2
	2013	1/2
Nigeria	2008	1/2
	2013	1/2
Tanzania	1999	1
	2010	1/3
Togo	1998	1/2
	2013	1/2

Table A13: Response Rates by Gender

		Urban		Rural	
Country	Year	Male	Women	Men	Women
Benin	2001	94.0***	95.8	94.2***	96.9
	2012	92.0***	94.7	96.2***	96.2
DR Congo	2007	93.9**	96.4	96.8**	97.0
	2013	95.6**	97.8	98.5**	99.0
Ghana	2008	95.0**	96.6	96.3**	96.4
	2014	93.7**	96.8	96.6**	97.8
Kenya	2008	85.4*	95.6	90.2*	96.6
	2014	86.6*	94.8	92.5*	97.1
Namibia	2006	83.9	92.9	91.5	96.2
	2013	81.7***	90.9	88.5***	93.9
Nigeria	2008	91.7**	96.5	93.1**	96.5
	2013	94.6	97.3	95.7	97.8
Tanzania	1999	91.4**	98.1	94.5**	97.7
	2010	88.3	95.8	91.3	96.2
Togo	1998	86.3**	90.6	94.1**	97.0
	2013	91.1**	96.6	97.4**	98.5

Note: * indicates response rates is for ages 15-54; ** indicates response rates is for ages 15-59; *** indicates response rate is for ages 15-64

connection, while column 2 shows the difference for males in connected areas, conditional on grid-cell and country-by-year fixed effects. Columns 3 and 4 follow the same pattern, but for females. Overall, the sample is relatively well-balanced on demographic characteristics, but this balance is not perfect. Men in connected areas are slightly younger and have higher rates of primary and secondary education than those in unconnected areas, while women in connected areas are less likely to be married than those in unconnected areas. However, since this paper relies on a difference-in-difference approach, these baseline differences should not bias the results.

Table A14: Demographic Balance Table: Unconnected vs Connected at Baseline

	Male		Female	
	Unconnected Mean	Difference	Unconnected Mean	Difference
	(st dev)	(st err)	(st dev)	(st err)
Age	28.619	-0.560***	28.496	-0.172
	(9.506)	(0.278)	(9.352)	(0.171)
Marital Rate	0.394	-0.014	0.491	-0.023***
	(0.489)	(0.016)	(0.500)	(0.011)
Primary Completion Rate	0.724	0.065***	0.599	0.019
	(0.447)	(0.016)	(0.490)	(0.014)
Secondary Completion Rate	0.284	0.035**	0.203	0.021**
	(0.451)	(0.018)	(0.402)	(0.012)
Male HH Head	0.848	-0.026*	0.685	-0.021
	(0.359)	(0.016)	(0.464)	(0.017)
HH Size	5.754	-0.108	6.455	-0.181
	(3.869)	(0.203)	(3.899)	(0.163)
Chidren in HH	0.876	-0.047	1.148	-0.065*
	(1.116)	(0.048)	(1.234)	(0.043)

A.2.3 Smallholder Agriculture

Smallholder agriculture is a large portion of total agriculture in the countries studied. Table A15 shows descriptive statistics, from the Food and Agriculture Organization of the United Nations (FAO), for smallholder farming in four of the countries studied: Ghana, Kenya, Nigeria and Tanzania. The table details, for each country and

year specified, the average farm size in the country, the country-specific largeholder threshold, and the percentage of farms in the country that are classified as smallholder according to this threshold. In all four countries, the average farm size is small and more than three-quarters of all farms are classified as smallholder.

Table A15: Smallholder Farming

	Year	Average Farm Size (Ha)	Largeholder Threshold (Ha)	Smallholder Farms (% of Total Farms)
Ghana	2013	2.56	3.64	79.10
Kenya	2005	0.86	1.21	80.88
Nigeria	2013	0.85	1.70	87.76
Tanzania	2013	1.89	3.31	83.23

Note: Data on smallholder farming from Food and Agriculture Organization of the United Nations (2017).

A.3 Additional Treatment Effects

This section contains additional treatment effects that were not included in the main paper.

A.3.1 Earnings

Table A16 presents the effect of fast Internet on the type of earnings received for work. There are no significant treatment effects on type of earnings for either men or women.

A.3.2 Decision-Making

Table A17 shows the effect of fast Internet on the likelihood a woman reports being the sole decision-maker in 5 domains - spending her earnings, spending her partner's

Table A16: Treatment Effect on Earnings Type

	Unpaid	Earns Cash Only	Earns Cash and In-Kind	Earns In-Kind
Male*Treatment	0.013	0.019	0.000	0.006
	(0.012)	(0.027)	(0.012)	(0.005)
Female*Treatment	0.009	0.017	-0.011	0.001
	(0.012)	(0.020)	(0.009)	(0.002)
Gender Difference	0.003	0.002	0.011	0.005
Mean: Men	0.081	0.562	0.062	0.011
Mean: Women	0.081	0.444	0.055	0.011
Observations	110126	110126	110126	110126
Clusters	1422	1422	1422	1422
Country x time x gender FE	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes

or husband's earnings, her health care, major household purchases, and visits to family - including an interaction effect of marital status and treatment. Internet access increases sole-decision making power for unmarried women more than it does for their married counterparts.

Table A18 reports the effect of fast Internet on the likelihood a woman reports having any decision-making power in these same 5 domains, including an interaction effect of treatment and having a female head of household; Table A19 reports the effect on the likelihood a woman reports having sole decision-making power, including the interaction effect of treatment and having a female head of household. Having a female head of household does not affect the treatment effect.

A.4 Multiple Testing Adjustments

The results in this paper test the effect of Internet access on multiple outcomes. This is problematic, as testing multiple hypotheses simultaneously increases the probability that coefficients will erroneously appear as significant. In order to account for

Table A17: Treatment Effect on Female Sole Decision-Making Power, with Interacted Marital Status, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	0.027	0.045*	0.018	0.031	0.100***
	(0.035)	(0.026)	(0.034)	(0.031)	(0.032)
Female*Married	-0.064***	-0.017**	-0.025***	-0.018***	-0.052***
	(0.014)	(0.007)	(0.008)	(0.007)	(0.008)
Female*Married*Treatment	-0.042	-0.036	-0.018	0.005	-0.067***
	(0.035)	(0.022)	(0.021)	(0.021)	(0.025)
Mean: Women	0.687	0.084	0.230	0.152	0.192
Observations	31266	36607	43583	43577	43581
Clusters	1225	1211	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A18: Treatment Effect on Female Decision-Making Power, with Interacted Head of Household Gender, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.015	-0.002	-0.002	0.030	0.071***
	(0.012)	(0.032)	(0.029)	(0.029)	(0.024)
Female*Female HoH	0.019***	-0.034***	0.066***	0.072***	0.060***
	(0.004)	(0.008)	(0.008)	(0.009)	(0.006)
Female*Female HoH*Treatment	-0.007	0.034	-0.018	-0.021	-0.013
	(0.016)	(0.025)	(0.021)	(0.020)	(0.018)
Mean: Women	0.930	0.441	0.618	0.577	0.679
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A19: Treatment Effect on Female Sole Decision-Making Power, with Interacted Head of Household Gender, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.005	0.012	0.002	0.036*	0.050**
	(0.025)	(0.018)	(0.025)	(0.020)	(0.021)
Female*Female HoH	0.115***	0.046***	0.151***	0.154***	0.161***
	(0.009)	(0.005)	(0.009)	(0.008)	(0.010)
Female*Female HoH*Treatment	-0.027	0.021	-0.005	-0.017	-0.023
	(0.027)	(0.018)	(0.025)	(0.020)	(0.019)
Mean: Women	0.687	0.084	0.230	0.152	0.192
Observations	31266	36607	43583	43577	43581
Clusters	1225	1211	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A20: Treatment Effect on Female Decision-Making Power, with Interacted Household Size, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.009	0.001	-0.008	0.028	0.066**
	(0.014)	(0.034)	(0.031)	(0.029)	(0.029)
HH Size	-0.001	-0.003***	-0.002**	-0.004***	-0.003***
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Female*HH Size*Treatment	-0.001	0.001	0.001	-0.000	0.001
	(0.002)	(0.003)	(0.003)	(0.002)	(0.004)
Mean: Women	0.930	0.441	0.618	0.577	0.679
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A21: Treatment Effect on Female Sole Decision-Making Power, with Interacted Household Size, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.000	0.028	0.013	0.043**	0.056**
	(0.028)	(0.020)	(0.027)	(0.022)	(0.023)
HH Size	0.002***	-0.000	-0.000	-0.002**	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Female*HH Size*Treatment	-0.002	-0.002	-0.002	-0.002	-0.002
	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
Mean: Women	0.687	0.084	0.230	0.152	0.192
Observations	31266	36607	43583	43577	43581
Clusters	1225	1211	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A22: Treatment Effect on Female Decision-Making Power, with Interacted Number of Children in HH, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.009	0.009	0.003	0.039	0.070***
	(0.013)	(0.033)	(0.028)	(0.028)	(0.027)
Num Children in HH	0.000	-0.006***	-0.014***	-0.019***	-0.011***
	(0.001)	(0.002)	(0.003)	(0.003)	(0.003)
Female*Num Children in HH*Treatment	-0.007	-0.004	-0.007	-0.010	-0.000
	(0.007)	(0.009)	(0.010)	(0.010)	(0.010)
Mean: Women	0.930	0.441	0.618	0.577	0.679
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A23: Treatment Effect on Female Sole Decision-Making Power, with Interacted Number of Children in HH, According to Women

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.000	0.023	0.016	0.044**	0.051**
	(0.027)	(0.018)	(0.028)	(0.022)	(0.024)
Num Children in HH	0.002	-0.001	-0.009***	-0.010***	-0.009***
	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)
Female*Num Children in HH*Treatment	-0.008	-0.006	-0.011*	-0.008	-0.004
	(0.009)	(0.005)	(0.006)	(0.005)	(0.007)
Mean: Women	0.687	0.084	0.230	0.152	0.192
Observations	31266	36607	43583	43577	43581
Clusters	1225	1211	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A24: Treatment Effect on Approval of Domestic Violence, with Interacted Marital Status

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*Treatment	0.029	0.041	0.022	0.022**	0.017
	(0.022)	(0.026)	(0.018)	(0.011)	(0.011)
Female*Treatment	0.037*	0.013	0.056***	0.012	0.009
	(0.020)	(0.022)	(0.020)	(0.014)	(0.013)
Male*Married	-0.015***	-0.024***	-0.025***	-0.019***	-0.016***
	(0.004)	(0.005)	(0.005)	(0.003)	(0.003)
Female*Married	0.009**	0.001	0.008**	0.013***	-0.004
	(0.004)	(0.005)	(0.004)	(0.003)	(0.003)
Male*Married*Treatment	0.011	0.014	0.013	-0.003	0.000
	(0.012)	(0.016)	(0.014)	(0.010)	(0.008)
Female*Married*Treatment	0.002	0.008	-0.013	-0.003	0.002
	(0.010)	(0.011)	(0.009)	(0.008)	(0.006)
Mean: Men	0.147	0.186	0.140	0.075	0.054
Mean: Women	0.203	0.252	0.193	0.131	0.096
Observations	105186	105355	105165	104602	105558
Clusters	1367	1367	1367	1367	1367
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A25: Treatment Effect on Approval of Domestic Violence, with Interacted Head of Household Gender

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*Treatment	0.034*	0.048*	0.025	0.021*	0.016
	(0.020)	(0.025)	(0.018)	(0.012)	(0.011)
Female*Treatment	0.036*	0.015	0.044**	0.009	0.007
	(0.020)	(0.023)	(0.020)	(0.014)	(0.013)
Male*Female HoH	0.010	0.013*	0.006	0.005	0.012***
	(0.007)	(0.007)	(0.006)	(0.004)	(0.004)
Female*Female HoH	-0.005	-0.003	-0.008**	0.002	0.001
	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)
Male*Female HoH*Treatment	-0.004	-0.009	0.004	0.000	0.005
	(0.020)	(0.019)	(0.018)	(0.018)	(0.011)
Female*Female HoH*Treatment	0.006	0.003	0.017	0.007	0.008
	(0.010)	(0.011)	(0.011)	(0.007)	(0.007)
Mean: Men	0.147	0.186	0.140	0.075	0.054
Mean: Women	0.203	0.252	0.193	0.131	0.096
Observations	105190	105359	105169	104606	105562
Clusters	1367	1367	1367	1367	1367
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

this, this section includes 2 types of multiple testing adjustments. The first multiple testing adjustment is to calculate sharpened q-values, following Anderson (2008). These q-values adjust the p-value according to the false discovery rate (FDR). This is done by iteratively testing whether the hypothesis is rejected at every value of q and recording the minimum q-value at which the hypothesis is rejected.

The second multiple testing adjustment is to use a summary index test that pools related outcomes into one measure (Anderson 2008). Summary index tests eliminate the problem of multiple inference by collapsing all outcomes into one test and provide a treatment effect on a general measure of an outcome of interest. In addition, summary index tests can be more powerful than individual tests: if the treatment effect on multiple outcomes are all marginally significant, the treatment effect on an index of these outcomes may reach significance.

Summary index tests are most useful when there is an a priori reason to expect

that treatment will impact a group of outcome measures in a consistent direction. This is not the case for the various employment measures used in this paper, as we may expect, for example, an increase in agricultural work to come at the expense of other sectors or to increase seasonal work but decrease work for outside employers. However, summary index tests are useful for both household decision-making and approval of domestic violence, insofar as questions about various domains of household decision-making reflect overall household bargaining power and questions about domestic violence in various scenarios reflect overall acceptability of domestic violence. As such, I use a summary index test for reported decision-making power of women and for approval of domestic violence.

A.4.1 Employment Q-Values

Table A26 presents the sharpened q-values for the treatment effect of fast Internet access on employment status. The results for males are robust to using q-values: Internet access increases both current employment and employment over the past 12 months by a large and significant amount. The impact on current female employment is also robust to using q-values: Internet access has a positive but insignificant impact. In addition, after adjusting for multiple testing, the small impact of Internet access on female employment over the past 12 months is no longer significant.

Table A27 includes the sharpened q-values for the treatment effect on employer type. The results for men are robust to using q-values: among men, fast Internet access significantly increases self-employment and employment for family members, but has a significant negative impact on employment for an outside employer. In addition, the results for women are robust: among women, Internet access has no impact on employer type.

Table A26: Treatment Effect on Employment Status (Q-Values)

	Currently Working	Worked in Past 12 months
Male*Treatment	0.049	0.048
	(0.021)	(0.022)
Female*Treatment	0.021	0.028
	(0.015)	(0.016)
Male Q-Value	0.030	0.030
Female Q-Value	0.172	0.172
Observations	117927	118472
Clusters	1429	1429
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

Table A27: Treatment Effect on Employer Type (Q-Values)

	Works for Family	Works for Someone Else	Works for Self
Male*Treatment	0.046	-0.054	0.041
	(0.010)	(0.025)	(0.019)
Female*Treatment	0.005	-0.002	0.018
	(0.008)	(0.013)	(0.016)
Male Q-Value	0.001	0.021	0.021
Female Q-Value	1.000	1.000	1.000
Observations	101488	101488	101488
Clusters	1429	1429	1429
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

Table A28 presents the sharpened q-values for the treatment effect on seasonality of employment. The results for both men and women are robust to using q-values: fast Internet access causes a significant increase in year-round employment for women, but it has no impact on the seasonality of work in which men are employed.

Table A28: Treatment Effect on Employment Seasonality (Q-Values)

	Works Year-Round	Works Seasonally	Works Occasionally
Male*Treatment	-0.001	0.028	0.020
	(0.024)	(0.018)	(0.013)
Female*Treatment	0.055	-0.015	-0.011
	(0.017)	(0.010)	(0.009)
Male Q-Value	0.481	0.274	0.274
Female Q-Value	0.005	0.149	0.185
Observations	117909	117909	117909
Clusters	1429	1429	1429
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

Table A29 shows the sharpened q-values for the treatment effect of fast Internet access on sector-specific employment. The male results are robust to including q-values: fast Internet access causes a large and significant increase in self-employed agricultural work. In addition, the treatment effect on female self-employed agricultural work is robust: women are significantly more likely to work in this sector, although the size of this effect is smaller than that for men. However, the positive treatment effects on female employment in skilled manual labor and agricultural work for others are no longer significant when adjusting for multiple testing.

Table A29: Treatment Effect on Employment by Sector (Q-Values)

	Manual, Skilled	Manual, Unskilled	Professional	Clerical	Sales	Agricultural, Self	Agricultural	Domestic	Services
Male*Treatment	0.001	0.005	-0.011	0.002	-0.011	0.072	-0.018	0.005	-0.011
	(0.019)	(0.014)	(0.016)	(0.006)	(0.016)	(0.014)	(0.012)	(0.005)	(0.012)
Female*Treatment	0.014	-0.000	0.002	-0.001	0.002	0.014	0.011	-0.000	-0.005
	(0.008)	(0.007)	(0.008)	(0.006)	(0.008)	(0.006)	(0.006)	(0.005)	(0.011)
Male Q-Value	1.000	1.000	0.931	1.000	0.931	0.001	0.523	0.931	0.931
Female Q-Value	0.219	1.000	1.000	1.000	1.000	0.077	0.219	1.000	1.000
Observations	117560	117560	117560	117560	117560	100640	117560	117560	117560
Clusters	1429	1429	1429	1429	1429	1429	1429	1429	1429
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

A.4.2 Decision-Making Index

The main treatment effects on women's self-reported sole decision-making power used 5 questions from the DHS that asked women who usually made decisions in each of 5 domains: spending her earnings, spending her husband's earnings, her health care, large household purchases, and visits to family. The summary index combines all 5 of these outcome measures into one index, representing overall female sole decision-making power, according to women.

Table A30 presents the treatment effect of fast Internet access on the sole decision-making index, according to women. The results show a slightly positive but insignificant effect on the index. This aligns with the results calculated using each of the 5 components: women saw an increase in sole decision-making power in some, but not all, domains. As such, women's self-reported sole decision-making power did not significantly increase across all domains.

Table A30: Treatment Effect on Female Sole Decision-Making Power Index, According to Women

	Female Sole Decision-Making Power Index
Female*Treatment	0.055
	(0.051)
Mean: Women	0.048
Observations	47389
Clusters	1383
Country x time x gender FE	Yes
Grid-cell x connected x gender FE	Yes

A summary index for women's self-reported decision-making power was calculated in the same manner as the sole decision-making power index discussed above.

However, this index reflects whether women reported having any, rather than sole, decision-making power across the 5 domains. Table A31 presents the treatment effect of fast Internet access on this index. Again, although there is a positive sign on the treatment effect, the effect is insignificant, suggesting there was no impact on women's overall decision-making power within the household. This reflects the individual treatment effects in each of the 5 domains: although women saw a significant increase in decision-making power in 1 of the 5 domains, they did not see any increase in the other 4. As such, women's self-reported decision-making power did not significantly increase across all domains.

Table A31: Treatment Effect on Female Decision-Making Power Index, According to Women

	Female Decision-Making Power Index
Female*Treatment	0.042
	(0.048)
Mean: Women	0.001
Observations	47389
Clusters	1383
Country x time x gender FE	Yes
Grid-cell x connected x gender FE	Yes

The main treatment effects on female sole decision-making power, as reported by men, used 3 questions from the DHS that asked men who usually made decisions in each of 5 domains: spending his earnings, his health care, and large household purchases. The summary index combines all 3 of these outcome measures into one index, representing overall female sole decision-making power, according to men.

Table A32 presents the treatment effect of fast Internet access on the sole decision-

making index, according to men. The results show a negative but insignificant effect on the index. This result reflects the individual treatment effects in each of the 3 domains - in which men reported a significant decrease in female sole decision-making in 2 domains and no effect in the third domain - but the decrease in precision reflects the adjustment made for multiple testing.

Table A32: Treatment Effect on Female Sole Decision-Making Power Index, According to Men

	Female Sole Decision-Making Power Index
Male*Treatment	-0.099
	(0.088)
Mean: Men	-0.030
Observations	21797
Clusters	1268
Country x time x gender FE	Yes
Grid-cell x connected x gender FE	Yes

A summary index for female decision-making power, as reported by men, was calculated in the same manner as the decision-making power index discussed above. However, this index reflects whether men reported that women had any, rather than sole, decision-making power in the 3 domains. Table A33 presents the treatment effect of fast Internet access on this index. The results show that, according to men, fast Internet access had no impact on women's overall decision-making power in the household. This result is consistent with the individual results in the 3 domains, all of which showed no impact on female decision-making power, as reported by men.

Table A33: Treatment Effect on Female Decision-Making Power Index, According to Men

	Female Decision-Making Power Index
Male*Treatment	0.036
	(0.085)
Mean: Men	-0.017
Observations	21797
Clusters	1268
Country x time x gender FE	Yes
Grid-cell x connected x gender FE	Yes

A.4.3 Decision-Making Q-Values

While the summary indices are useful in testing for broad changes in female decision-making power across all domains, they have a key drawback: if the treatment effect on decision-making as a whole is insignificant, we cannot determine if any of the treatment effects on individual domains of decision-making would be significant. Sharpened q-values allow us to adjust for multiple testing, while also allowing us to understand the impact of fast Internet access in each of the relevant domains.

Table A34 presents the sharpened q-values for the treatment effect of fast Internet access on female sole decision-making power, according to women. After adjusting for multiple testing, there is no significant impact of fast Internet on women's self-reported sole decision-making power. This differs from the main results since, although there is a positive effect in the domains of large household purchases and visits to family, these effects are either not strong enough or not precisely estimated enough to be significant after accounting for multiple testing.

Table A35 shows the sharpened q-values for the treatment effect of fast Internet

Table A34: Treatment Effect on Female Sole Decision-Making Power, According to Women (Q-Values)

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.009	0.017	0.004	0.035	0.047
	(0.025)	(0.017)	(0.026)	(0.021)	(0.022)
Female Q-Value	1.000	0.462	1.000	0.233	0.167
Observations	31266	36607	43583	43577	43581
Clusters	1225	1211	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

access on the probability that a woman has any decision-making power in various domains, as reported by women. The results indicate that the treatment effects are robust to accounting for multiple testing: fast Internet access significantly increases women's self-reported decision-making power over visits to family but has no impact on the other 4 domains.

Table A35: Treatment Effect on Female Decision-Making Power, According to Women (Q-Values)

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.016	0.005	-0.004	0.027	0.069
	(0.011)	(0.032)	(0.028)	(0.027)	(0.023)
Female Q-Value	0.459	1.000	1.000	0.648	0.017
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A36 presents the sharpened q-values for the treatment effect of fast Internet access on female sole decision-making power, according to men. The results indicate

that the treatment effects are robust to accounting for multiple inference with sharpened q-values: fast Internet access reduces women's sole decision-making power over spending her partner's earnings and his health care, according to men.

Table A36: Treatment Effect on Female Sole Decision-Making Power, According to Men (Q-Values)

	Spending His Earnings	His Health Care	Major HH Purchases
Male*Treatment	-0.039	-0.034	-0.015
	(0.016)	(0.013)	(0.031)
Male Q-value	0.021	0.021	0.270
Observations	12979	11265	21379
Clusters	1098	1000	1264
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

Table A37 shows the sharpened q-values for the treatment effect of fast Internet access on the probability that a woman has any decision-making power in various domains, according to men. The results indicate that the treatment effects are robust to accounting for multiple testing: fast Internet access has no significant impact on the probability that women have any decision-making power over the 3 domains.

A.4.4 Approval of Domestic Violence Index

The main treatment effects on approval of domestic violence used 5 questions from the DHS that asked both men and women whether it was acceptable for a man to beat his wife in 5 different scenarios: if the wife goes out without telling her husband, if the wife neglects the children, if the wife argues with her husband, if the wife refuses sex with her husband, and if the wife burns food. The summary index combines all 5 of these outcome measures into one index, representing overall approval of domestic

Table A37: Treatment Effect on Female Decision-Making Power, According to Men (Q-Values)

	Spending His Earnings	His Health Care	Major HH Purchases
Male*Treatment	0.071	-0.011	0.013
	(0.058)	(0.046)	(0.039)
Male Q-value	1.000	1.000	1.000
Observations	12979	11265	21379
Clusters	1098	1000	1264
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

violence.

Table A38 presents the gender-specific treatment effects of fast Internet access on approval of domestic violence. The results reveal that fast Internet access significantly increases the probability that men support domestic violence, across 5 domains. This result reflects the increase in male approval of domestic violence in each of the 5 domains, as reported in the main results. The index also reveals that, although the increase in male approval of domestic violence was only significant in 2 of the 5 scenarios when tested individually, the overall effect on male approval over the 5 domains is large and significant. This suggests that the general acceptability of domestic violence against women increased substantially among men as a result of access to fast Internet.

Table A38 also shows that fast Internet had a positive but insignificant effect on women's approval of domestic violence across all 5 domains. This is likely driven by the fact that the female treatment effect on acceptability of domestic violence was large and positive for only 2 scenarios - when a wife goes out without telling her husband and when a wife argues with her husband - and small for the other 3. As such, acceptability of domestic violence among women did not increase significantly

across all domains.

Table A38: Treatment Effect on Approval of Domestic Violence Index

	Approval of Domestic Violence Index
Male*Treatment	0.090**
	(0.046)
Female*Treatment	0.083
	(0.052)
Gender Difference	0.007
Mean: Men	-0.147
Mean: Women	0.052
Observations	106416
Clusters	1367
Country x time x gender FE	Yes
Grid-cell x connected x gender FE	Yes

A.4.5 Approval of Domestic Violence Q-Values

While the summary index is helpful for considering the impact of fast Internet access on broad approval for domestic violence across all situations, it cannot reveal the impact on approval in specific scenarios or reveal if the treatment effect on men and women's approval of domestic violence differs within each domain. Sharpened q-values allow us to adjust for multiple testing, while also allowing us to understand the impact of fast Internet access on approval of domestic violence in specific scenarios.

Table A39 presents the sharpened q-values for the gender-specific treatment effects of fast Internet access on approval of domestic violence. The results indicate that, although fast Internet access has a positive effect on female approval of domestic violence in each situation, this effect is only significant for one domain, when a wife argues with her husband.

The treatment effects on male approval of domestic violence indicate that, although fast Internet access has positive effects on acceptability of domestic violence among men, these effects are not individually significant after adjusting for multiple testing. These results should not be taken as proof that there was no impact on male approval of domestic violence, however, as the male treatment effect on the summary index of domestic violence was both positive and significant. This reminds us of the benefit of a summary indices: summary index tests are often more powerful.

Table A39: Treatment Effect on Approval of Domestic Violence (Q-Values)

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*Treatment	0.033	0.046	0.026	0.020	0.017
	(0.021)	(0.025)	(0.018)	(0.011)	(0.010)
Female*Treatment	0.038	0.016	0.051	0.012	0.010
	(0.019)	(0.022)	(0.019)	(0.014)	(0.013)
Male Q-Value	0.163	0.163	0.163	0.163	0.163
Female Q-Value	0.104	0.381	0.040	0.381	0.381
Observations	105190	105359	105169	104606	105562
Clusters	1367	1367	1367	1367	1367
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

A.5 Sample-Weighted Results

The DHS collects data according to a two-stage probability sampling of a country's population, using a recent census to define the population. The population is first stratified into subgroups, typically by geographic region and by urban/rural areas within each region. Within each stratum, sampling clusters are selected independently. Within each clusters, households are then selected. As such, the probability that a household is sampled is dependent both on the probability that a cluster is

chosen from a stratum and that a household is chosen from a cluster. Although the probability that a household is chosen is constant within a cluster, it differs across clusters.

As such, the DHS provides household sampling weights that reflect the probability a household is selected, as well as the household response rate in the stratum. The DHS also provides individual sampling weights that reflect both the household sample weight and the response rate for the respondent's gender in the stratum.

In effect, the household and individual sample weights provided by the DHS ensure that, when weighted, the sample is nationally representative within each country. However, the importance of using sampling weights for this paper's analysis, in which both the treatment and control group are determined by distance to infrastructure that is only available in select areas, is less clear. As such, the main analysis chooses not to use sample weights. However, the results of the main analyses using the individual sample weights are reported in this appendix.

A.5.1 Employment

Table A40 presents the gender-specific treatment effect of fast Internet access on employment status, after incorporating individual sample weights. Although the estimated magnitude of the effects are smaller for men when including weights, these effects are still significant. In addition there is still a gender difference, albeit a smaller gender difference, in the effect of Internet access on employment status, after incorporating for the sample weighs.

Table A41 presents the treatment effect of fast Internet access on employer type for both men and women, with the use of sample weights. The results are consistent with the unweighted results, but the magnitudes are slightly different: incorporating

Table A40: Treatment Effect on Employment Status (Sample-Weighted)

	Currently Working	Worked in Past 12 months
Male*Treatment	0.038*	0.042*
	(0.021)	(0.022)
Female*Treatment	0.021	0.024
	(0.017)	(0.017)
Gender Difference	0.017	0.018
Mean: Men	0.735	0.762
Mean: Women	0.618	0.653
Observations	117927	118472
Clusters	1429	1429
Sample weights	Yes	Yes
Country x time x gender FE	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes

sample weights slightly reduces the estimated magnitude of the positive treatment effect on male employment for family members and the estimated magnitude of the negative treatment effect on male employment for outside employers.

Table A42 presents the sample-weighted treatment effect on seasonality of employment. The results are consistent with those presented in the main analysis, although use of individual sample weights slightly reduces the estimated magnitude of the positive treatment effect on year-round employment for women.

Table A43 presents the treatment effect of fast Internet access on sector-specific employment, when incorporating individual sample weights. Among men, the sample-weighted treatment effects are broadly consistent with the unweighted effects, and the primary result that fast Internet access increases agricultural self-employment does not change. Interestingly, however, the inclusion of sample weights does affect the estimated magnitude of the treatment effect on male agricultural employment for an employer: when individual sample weights are included, the treatment effect

Table A41: Treatment Effect on Employer Type (Sample-Weighted)

	Works for Family	Works for Someone Else	Works for Self
Male*Treatment	0.032***	-0.046**	0.042**
	(0.009)	(0.022)	(0.020)
Female*Treatment	-0.002	0.005	0.018
	(0.008)	(0.014)	(0.021)
Gender Difference	0.033	-0.051	0.024
Mean: Men	0.065	0.227	0.234
Mean: Women	0.060	0.150	0.407
Observations	101488	101488	101488
Clusters	1429	1429	1429
Sample weights	Yes	Yes	Yes
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

Table A42: Treatment Effect on Employment Seasonality (Sample-Weighted)

	Works Year-Round	Works Seasonally	Works Occasionally
Male*Treatment	0.009	0.018	0.013
	(0.026)	(0.018)	(0.013)
Female*Treatment	0.049***	-0.010	-0.014
	(0.018)	(0.010)	(0.012)
Gender Difference	-0.039	0.028	0.027
Mean: Men	0.544	0.157	0.058
Mean: Women	0.485	0.117	0.051
Observations	117909	117909	117909
Clusters	1429	1429	1429
Sample weights	Yes	Yes	
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

appears to be both negative and significant.

Among women, the sample-weighted treatment effects are also broadly consistent with the unweighted results. However, inclusion of individual sample weights increases the estimated magnitude of the positive treatment effect on skilled manual labor; in addition, it reduces the estimated magnitude of the positive treatment effect on agricultural employment for women and, thus, renders this treatment effect insignificant.

Table A43: Treatment Effect on Employment by Sector (Sample-Weighted)

	Manual, Skilled	Manual, Unskilled	Professional	Clerical	Sales	Agricultural, Self	Agricultural	Domestic	Services
Male*Treatment	-0.010	0.013	-0.002	0.000	-0.002	0.075***	-0.029***	0.003	-0.013
	(0.023)	(0.015)	(0.017)	(0.008)	(0.017)	(0.015)	(0.010)	(0.004)	(0.013)
Female*Treatment	0.020**	0.001	0.009	-0.004	0.009	0.014**	0.008	-0.004	-0.007
	(0.010)	(0.007)	(0.010)	(0.006)	(0.010)	(0.007)	(0.005)	(0.006)	(0.011)
Gender Difference	-0.030	0.012	-0.012	0.004	-0.012	0.060	-0.037	0.007	-0.006
Mean: Men	0.169	0.081	0.102	0.018	0.102	0.088	0.123	0.035	0.072
Mean: Women	0.064	0.027	0.065	0.019	0.065	0.069	0.039	0.025	0.081
Observations	117560	117560	117560	117560	117560	100640	117560	117560	117560
Clusters	1429	1429	1429	1429	1429	1429	1429	1429	1429
Sample weights	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

A.5.2 Decision-Making Power

Table A44 presents the treatment effect of fast Internet access on female sole decision-making power in 5 domains, as reported by women, after weighting using the individual sample weights. The results are broadly consistent with the unweighted treatment effects, although the magnitude of the effects does differ. Incorporating sample weights increases the estimated magnitude of the positive treatment effects on women's self-reported sole decision-making power over major household purchases and over visits to family.

Table A44: Treatment Effect on Female Sole Decision-Making Power, According to Women (Sample-Weighted)

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	0.007	0.031	0.030	0.056***	0.060***
	(0.027)	(0.022)	(0.027)	(0.019)	(0.023)
Mean: Women	0.687	0.084	0.230	0.152	0.192
Observations	31266	36607	43583	43577	43581
Clusters	1225	1211	1353	1353	1353
Sample weights	Yes	Yes	Yes	Yes	Yes
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

Table A45 shows the sample-weighted treatment effect on the probability that a woman reports having any decision-making power over various household decisions. These results are consistent with the unweighted results.

Table A46 presents the sample-weighted treatment effect of fast Internet access on female sole decision-making power, according to men. The treatment effects in the domains of spending the man's earnings and his health care are consistent with the unweighted results, but using individual sample weights increases the estimated magnitude of both these negative treatment effects. The inclusion of individual sample

Table A45: Treatment Effect on Female Decision-Making Power, According to Women (Sample-Weighted)

	Spending Her Earnings	Spending Husband's Earnings	Her Health Care	Major HH Purchases	Visits to Family
Female*Treatment	-0.014	-0.018	0.011	0.026	0.073***
	(0.011)	(0.041)	(0.031)	(0.031)	(0.024)
Mean: Women	0.930	0.441	0.618	0.577	0.679
Observations	31264	36194	43583	43577	43581
Clusters	1225	1210	1353	1353	1353
Sample weights	Yes	Yes	Yes	Yes	Yes
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

weights, however, has a larger impact on the estimated treatment effect on women's sole decision-making power over major household purchases, according to men: when weighted, we see that this treatment effect is both negative and significant.

Table A46: Treatment Effect on Female Sole Decision-Making Power, According to Men (Sample-Weighted)

	Spending His Earnings	His Health Care	Major HH Purchases
Male*Treatment	-0.048***	-0.038***	-0.059*
	(0.017)	(0.011)	(0.031)
Mean: Men	0.051	0.068	0.096
Observations	12979	11265	21379
Clusters	1098	1000	1264
Sample weights	Yes	Yes	Yes
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

Table A47 shows the sample-weighted treatment effect on the probability that a woman has any decision-making power, according to men. When using individual sample weights, the estimated treatment effect of fast Internet access on women's decision-making power over spending her partner's earnings, as reported by men, is

positive and significant. This is a deviation from the unweighted results, in which this treatment effect is estimated to be positive but not statistically significant; in addition, the weighted estimate of the treatment effect is much larger than the unweighted estimate.

Table A47: Treatment Effect on Female Decision-Making Power, According to Men (Sample-Weighted)

	Spending His Earnings	His Health Care	Major HH Purchases
Male*Treatment	0.135**	-0.014	-0.000
	(0.059)	(0.037)	(0.044)
Mean: Men	0.414	0.351	0.469
Observations	12979	11265	21379
Clusters	1098	1000	1264
Sample weights	Yes	Yes	Yes
Country x time x gender FE	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes

A.5.3 Approval of Domestic Violence

Table A48 presents the gender-specific treatment effect of fast Internet access on approval of domestic violence in 5 scenarios, after weighting using individual sample weights. The male sample-weighted treatment effects are broadly consistent with the unweighted results, although the estimated magnitudes of some treatment effects differ. In particular, the estimated magnitude of the positive treatment effect of Internet access on male approval of domestic violence if a wife neglects the children is larger when individual sample weights are incorporated.

Similarly, the sample-weighted treatment effect on female approval of domestic violence is largely consistent with the unweighted results. However, inclusion of sample weights decreases the estimated magnitude of the treatment effect on women's

approval of domestic violence if a wife goes out without telling her husband or she argues with her husband; in addition, the sample-weighted estimate of the former treatment is insignificant.

Table A48: Treatment Effect on Approval of Domestic Violence (Sample-Weighted)

	Wife Goes Out w/o Telling Husband	Wife Neglects Children	Wife Argues with Husband	Wife Refuses Sex	Wife Burns Food
Male*Treatment	0.023	0.054**	0.024	0.023**	0.013
	(0.021)	(0.024)	(0.017)	(0.012)	(0.010)
Female*Treatment	0.029	0.005	0.039**	0.008	0.015
	(0.021)	(0.023)	(0.018)	(0.013)	(0.012)
Gender Difference	-0.007	0.049	-0.015	0.016	-0.002
Mean: Men	0.147	0.186	0.140	0.075	0.054
Mean: Women	0.203	0.252	0.193	0.131	0.096
Observations	105190	105359	105169	104606	105562
Clusters	1367	1367	1367	1367	1367
Sample weights	Yes	Yes	Yes	Yes	Yes
Country x time x gender FE	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes

A.6 Country-Level Results

The value of using multiple countries in this analysis is that it both provides variation in the timing of connection to submarine cables and also increases the power to detect changes by boosting the sample size. The latter is particularly important due to the high-definition fixed effects - both country by year by gender and also 10km*10km grid cell by connected by gender - employed in the analysis. Despite this, the individual country-level results may be of interest and, as such, are presented in this section.

A.6.1 Employment Status

Table A49 shows the effect of fast Internet access on current employment, for each country. The results are imprecise but consistent with the multinational results - men see a greater employment boost than women - for 6 countries: Benin, Kenya, Namibia, Nigeria, Tanzania and Togo. In the Democratic Republic of Congo, however, men and women both see a large increase in employment in response to Internet access, and women see a slightly larger boost. In Ghana, it appears that women may see a greater employment effect than men, but the results are not statistically significant.

Table A49: Treatment Effect on Current Employment, by Country

	Currently Working							
	Benin	DR Congo	Ghana	Kenya	Namibia	Nigeria	Tanzania	Togo
Male*Treatment	0.044	0.156***	-0.031	0.068	0.029	0.131*	0.036	0.057
	(0.058)	(0.013)	(0.038)	(0.044)	(0.048)	(0.074)	(0.103)	(0.051)
Female*Treatment	0.016	0.174***	0.019	0.010	-0.021	0.050	-0.018	0.041
	(0.043)	(0.008)	(0.058)	(0.032)	(0.032)	(0.036)	(0.102)	(0.025)
Gender Difference	0.028	-0.017	-0.049	0.058	0.050	0.082	0.054	0.016
Mean: Men	0.717	0.790	0.768	0.796	0.609	0.741	0.827	0.693
Mean: Women	0.681	0.505	0.731	0.576	0.459	0.616	0.729	0.726
Observations	14207	4281	9723	23272	19418	23625	6331	17070
Clusters	83	18	155	443	238	206	164	128
Time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A50 shows the effect of fast Internet access on employment over the past year, for each country. The results are consistent with the multinational results - men see a greater employment effect than women - for 7 countries: Benin, Democratic Republic of Congo, Kenya, Namibia, Nigeria, Tanzania and Togo. However, in Ghana, it appears that women may see a greater employment effect than men; again, the results are not statistically significant.

Table A50: Treatment Effect on Employment in Past 12 Months, by Country

	Worked in the Past 12 Months							
	Benin	DR Congo	Ghana	Kenya	Namibia	Nigeria	Tanzania	Togo
Male*Treatment	0.021	0.163***	-0.021	0.084*	0.045	0.154**	0.048	0.025
	(0.065)	(0.018)	(0.035)	(0.043)	(0.049)	(0.068)	(0.106)	(0.056)
Female*Treatment	0.017	0.146***	0.024	0.025	0.022	0.062*	-0.030	0.021
	(0.047)	(0.010)	(0.053)	(0.029)	(0.044)	(0.033)	(0.083)	(0.023)
Gender Difference	0.004	0.016	-0.045	0.059	0.023	0.092	0.078	0.004
Mean: Men	0.750	0.699	0.796	0.818	0.662	0.767	0.857	0.725
Mean: Women	0.706	0.541	0.764	0.616	0.520	0.631	0.771	0.761
Observations	14212	4569	9748	23305	19491	23729	6337	17081
Clusters	83	18	155	443	238	206	164	128
Time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

A.6.2 Decision-Making Power

Table A51 presents the effect of fast Internet access on women's self-reported sole decision-making power over household decisions. In this specification, sole decision-making power is measured using an index. This index is generated from women's responses to 5 questions that ask who usually makes decisions in each of 5 domestic areas: spending her earnings, spending her husband's earnings, her health care, major household purchases, and visits to family. Tanzania and Togo are omitted because there is not enough variation in the index after accounting for grid-cell-by-connected fixed effects to estimate a treatment effect.

The results are consistent with the multinational results - Internet access generates an insignificant but small positive effect on female's self-reported sole decision-making power as a whole - for 3 countries: Benin, Namibia, Nigeria. In the Democratic Republic of Congo, Internet access has a strongly negative impact on women's self-reported sole decision-making power. In Ghana, Internet access has a strongly positive impact, perhaps due to the fact that Internet access appears to improve female employment more than male employment in this country. In Kenya, Internet

access has a mildly negative impact, but this estimate is very imprecise.

Table A51: Treatment Effect on Female Sole Decision-Making Power Index, According to Women, by Country

	Female Sole Decision-Making Power Index					
	Benin	DR Congo	Ghana	Kenya	Namibia	Nigeria
Female*Treatment	0.092	-0.166***	0.249***	-0.075	0.105	0.132
	(0.100)	(0.054)	(0.077)	(0.091)	(0.156)	(0.138)
Mean: Women	0.036	0.170	0.087	0.040	0.224	-0.235
Observations	8326	2315	3389	8204	4899	10566
Clusters	83	18	155	439	228	206
Time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes

Table A52 presents the effect of fast Internet access on whether a woman reports having any decision-making power in household decisions. In this specification, decision-making power is measured using an index. This index is generated in the same manner as the sole decision-making power, but using women's responses to having any decision-making power in the 5 domains. Again, Tanzania and Togo are excluded due to a lack of variation after applying fixed effects.

The results are broadly consistent with the multinational results - Internet access generates an insignificant but small positive effect on female's self-reported decision-making power as a whole. However, in Benin, the effect of Internet access on decision-making power is strongly positive.

Table A53 presents the effect of fast Internet access on whether a man reports his wife or partner has sole decision-making power over household decisions. In this specification, sole decision-making power is measured using an index. This index is generated from men's responses to 3 questions that ask who usually makes decisions in each of 3 domestic areas: spending her husband's earnings, her health care, and

Table A52: Treatment Effect on Female Decision-Making Power Index, According to Women, by Country

	Female Decision-Making Power Index					
	Benin	DR Congo	Ghana	Kenya	Namibia	Nigeria
Female*Treatment	0.248***	0.066*	-0.027	0.026	0.014	-0.056
	(0.086)	(0.037)	(0.195)	(0.105)	(0.103)	(0.139)
Mean: Women	-0.141	-0.233	0.270	0.278	0.442	-0.270
Observations	8326	2315	3389	8204	4899	10566
Clusters	83	18	155	439	228	206
Time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes

major household purchases. Tanzania and Togo are omitted because there is not enough variation in the index after accounting for grid-cell-by-connected fixed effects to estimate a treatment effect.

The results are broadly consistent with the multinational results - Internet access decreases female sole decision-making power, as reported by men. However, the effect in Ghana is strongly negative, although ultimately insignificant; this result may relate to the finding that Internet access seems to boost women's employment more than that of men in Ghana.

Table A54 presents the effect of fast Internet access on whether a man reports that his wife or partner has any decision-making power in household decisions. In this specification, decision-making power is measured using an index. This index is generated in the same manner as the sole decision-making power, but using men's responses about their spouse having any decision-making power in the 3 domains. Again, Tanzania and Togo are excluded due to a lack of variation after applying fixed effects.

The results are broadly consistent with the multinational results - Internet ac-

Table A53: Treatment Effect on Female Sole Decision-Making Power Index, According to Men, by Country

	Female Sole Decision-Making Power Index					
	Benin	DR Congo	Ghana	Kenya	Namibia	Nigeria
Male*Treatment	-0.216***	-0.462***	0.417	-0.249	0.095	-0.478*
	(0.080)	(0.061)	(0.260)	(0.180)	(0.188)	(0.249)
Mean: Men	-0.183	0.062	-0.046	-0.016	0.042	0.047
Observations	2369	628	2541	5644	3505	5306
Clusters	80	18	152	411	194	205
Time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes

cess has no impact on women's decision-making power on the whole, as reported by men. However, in the Democratic Republic of Congo, the effect of Internet access on women's decision-making power, as reported by men, is strongly negatively. In addition, in Ghana, the effect is strongly positive, perhaps because Internet access seems to boost female employment more than male employment in this country.

Table A54: Treatment Effect on Female Decision-Making Power Index, According to Men, by Country

	Female Sole Decision-Making Power Index					
	Benin	DR Congo	Ghana	Kenya	Namibia	Nigeria
Male*Treatment	-0.107*	-0.069	0.438***	-0.014	-0.046	-0.037
	(0.061)	(0.053)	(0.151)	(0.128)	(0.216)	(0.334)
Mean: Men	-0.588	-0.137	0.083	0.298	0.551	-0.404
Observations	2369	628	2541	5644	3505	5306
Clusters	80	18	152	411	194	205
Time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes

A.6.3 Approval of Domestic Violence

Table A55 presents the effect of fast Internet access on self-reported approval of domestic violence against women. In this specification, approval of domestic violence is measured using an index. This index is generated from individuals' responses to 5 questions that ask whether the respondent believes it is justified for a husband to hit his wife in response to 5 situations: she goes out without telling her husband, she neglects the children, she argues with her husband, she refuses sex with her husband, and she burns the food. Tanzania and Togo are omitted because there is not enough variation in the index after accounting for grid-cell-by-connected fixed effects to estimate a treatment effect. Similarly, the treatment effect for men in the Democratic Republic of Congo is omitted because there is not enough variation in the index among male respondents after including these fixed effects.

The impact of fast Internet access on male approval of domestic violence is consistent with the multinational results - that Internet access increases reported acceptability of domestic violence among men - but imprecisely estimated in all countries. However, the effect on women's reported approval of domestic violence is not as consistent with the multinational results of a positive but ultimately insignificant effect. In the Democratic Republic of Congo, fast Internet access decreases reported acceptability of domestic violence among female respondents. In addition, although the effect is positive in Namibia, the impact of Internet access on reported approval of domestic violence is actually higher for women than for men.

Table A55: Treatment Effect on Approval of Domestic Violence Index, According to Men, by Country

	Approval of Domestic Violence Index					
	Benin	DR Congo	Ghana	Kenya	Namibia	Nigeria
Male*Treatment	0.059		0.075	0.141	0.133	0.046
	(0.080)		(0.095)	(0.101)	(0.100)	(0.090)
Female*Treatment	0.189	-0.279***	-0.013	0.028	0.164**	0.004
	(0.124)	(0.062)	(0.082)	(0.107)	(0.070)	(0.163)
Gender Difference	-0.131	0.279	0.088	0.113	-0.031	0.042
Mean: Men	-0.271	0.161	-0.343	0.013	-0.160	-0.199
Mean: Women	-0.037	0.555	-0.190	0.187	-0.014	0.004
Observations	14114	3890	9731	23162	19233	23486
Clusters	83	18	155	443	238	206
Time x gender FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell x connected x gender FE	Yes	Yes	Yes	Yes	Yes	Yes

Appendix B

Appendices to Chapter 3

B.1 Endline Summary Statistics

Table B1 reports summary statistics for the endline workseeker sample, separately for men and women. Column 1 reports the mean and standard deviation of each measure for men, while Column 2 reports the number of non-missing observations for men. Column 3 reports the mean and standard deviation of each measure for women, while Column 3 reports the number of non-missing observations for men. Intensive-margin employment variables are set to 0 if the workseeker is not employed. Column 4 presents the difference in the mean between men and women, while Column 5 tests whether this difference is significant.

B.2 Decomposition of Gender Differential

B.2.1 Decomposition of Gender Differential in Treatment Effect

Labor market outcomes between two groups of people may differ for a number of reasons, but it is difficult to assess the root cause without a readily available counterfactual. In the case of this intervention, men and women experience different treatment effects. However, because men and women have different underlying characteristics that might interact with treatment, it is not possible to observe what the treatment effect on women would be if they were exactly the same as men except for

Table B1: Summary Statistics by Gender for Endline Variables

	Male		Female		Difference	
Variable	Mean (St. Dev.)	Obs.	Mean (St. Dev.)	Obs.	Mean Diff. (Male-Female)	p:Diff=0
Panel A: Labor Market Measures						
Employed	0.372	2499	0.293	4108	0.079	0.000
	(0.483)		(0.455)			
Earnings	650.5	924	601.2	1188	49.30	0.354
	(1202)		(1168)			
Hours worked	31.57	930	26.13	1191	5.440	0.000
	(22.22)		(20.81)			
Hourly wage	30.70	923	35.06	1174	-4.360	0.202
	(74.58)		(70.48)			
Wage employment	0.900	918	0.874	1184	0.026	0.086
	(0.300)		(0.332)			
Self employment	0.100	918	0.124	1184	-0.024	0.104
	(0.300)		(0.330)			
Panel B: Job Search Measures						
Any search	0.669	2499	0.706	4109	-0.037	0.001
	(0.471)		(0.456)			
Applications submitted[a]	13.25	2493	12.52	4084	0.730	0.177
	(21.66)		(21.34)			
Hours searched	10.57	2499	9.549	4102	1.021	0.005
	(15.35)		(13.43)			
Search cost	115.2	2495	117.0	4104	-1.800	0.680
	(181.8)		(156.5)			
Responses[a]	0.865	2495	0.858	4098	0.007	0.882
	(2.100)		(2.176)			
Offers[a]	0.232	2495	0.191	4097	0.041	0.019
	(0.722)		(0.652)			
Panel C: Belief Measures						
Fraction of assessments overconfident	0.317	2498	0.362	4109	-0.045	0.000
	(0.231)		(0.239)			
Fraction of assessments underconfident	0.173	2498	0.178	4109	-0.005	0.239
	(0.165)		(0.166)			
Targeted search	0.149	2630	0.191	4261	-0.042	0.000
	(0.356)		(0.393)			
Planned applications[a]	15.94	2494	16.15	4097	-0.210	0.764
	(25.13)		(32.14)			
Expected offers[a]	4.743	2475	4.330	4056	0.413	0.002
	(5.831)		(5.617)			

Table shows summary statistics for selected endline variables. Percentiles are omitted for binary variables. All monetary figures are reported in South Africa Rands. 1 Rand ≈ USD0.16 in purchasing power parity terms. Intensive-margin labor market measures are set to missing for non-workers. Intensive-margin search measures are set to zero for non-searchers. Missing values reflect item non-response, mostly due to respondents reporting that they don't know the answer. All period-specific outcomes use a 7-day recall/forecast period unless marked with [a] (30-day recall/forecast period) or [b] (since treatment). The final column for the gender difference reports the p-value for testing equality of endline means of the variables across both genders, using heteroskedasticity-robust standard errors clustered by treatment date.

their gender, and vice versa.

The Blinder-Oaxaca decomposition allows us to decompose mean differences in outcomes between groups based on linear regression models in a counterfactual manner (Blinder 1973; Oaxaca 1973). This method divides the differential in outcome into one part that is "explained" by group differences in observable characteristics and the residual that is "unexplained". The "unexplained" portion of the differential is often attributed to discrimination but is also comprised of the effects of any group differences in unobserved characteristics. This method has both a threefold and a twofold decomposition.

The threefold decomposition divides the differential in outcome into 3 parts. The first component is the "endowments effect," which is the part of the differential that is explained by group differences in observed characteristics. The second component is the "coefficients effect," which is the part of the differential that is explained by group differences in coefficients. The third component is the residual "interaction effect," which is the part of the differential that is explained by the interaction of differences in coefficients and differences in observed characteristics.

To construct the threefold decomposition for this paper, there are two groups: males, M, and females, F. The outcome variable is the treatment effect of public certification on a labor market outcome, which is the difference in expectation of the labor market outcome between the treatment group, $T = 1$, and control group, $T = 0$: $E[Y|T = 1] - E[Y|T = 0]$. After incorporating a set of predictors into a linear regression model, we are then interested in the difference in this outcome between men and women:

$$\Delta_Y = (E[Y_M|T=1] - E[Y_M|T=0]) - (E[Y_F|T=1] - E[Y_F|T=0]) \tag{B.1}$$
$$= (E[Y|M=1,T=1] - E[Y|M=1,T=0]) - (E[Y|M=0,T=1] - E[Y_F|M=0,T=0])$$

Based on the linear model:

$$Y_{M,T} = X'_{M,T}\beta_{M,T} + \epsilon_{M,T}\ , \qquad E(\epsilon_{M,T}) = 0 \qquad M \in (0,1) \qquad T \in (0,1)$$

where X is a vector of observable characteristics and a constant, the differential can be rewritten as the difference in the linear prediction of the group-specific means of the regressors:

$$\Delta_Y = (E[X|M=1,T=1]'\beta_{M=1,T=1} - E[X|M=1,T=0]'\beta_{M=1,T=0}) \tag{B.2}$$
$$- (E[X|M=0,T=1]'\beta_{M=0,T=1} - E[X|M=0,T=0]'\beta_{M=0,T=0})$$

Since treatment is randomly assigned, this can be rewritten as:

$$\Delta_Y = E[X|M=1]'(\beta_{M=1,T=1} - \beta_{M=1,T=0}) - E[X|M=0]'(\beta_{M=0,T=1} - \beta_{M=0,T=0}) \tag{B.3}$$

This can then be rearranged into the form for threefold decomposition:

$$\Delta_Y = \underbrace{(E[X|M=1] - E[X|M=0])'(\beta_{M=0,T=1} - \beta_{M=0,T=0})}_{Endowments\ Effect} \tag{B.4}$$

$$+ \underbrace{E[X|M=0]'((\beta_{M=1,T=1} - \beta_{M=1,T=0}) - (\beta_{M=0,T=1} - \beta_{M=0,T=0}))}_{Coefficients\ Effect}$$

$$+ \underbrace{(E[X|M=1] - E[X|M=0])'((\beta_{M=1,T=1} - \beta_{M=1,T=0}) - (\beta_{M=0,T=1} - \beta_{M=0,T=0}))}_{Interaction\ Effect}$$

This can be estimated using:

$$\hat{\Delta_Y} = \underbrace{(\bar{X}_M - \bar{X}_F)'(\hat{\beta}_{F,T=1} - \hat{\beta}_{F,T=0})}_{Endowments\ Effect} + \underbrace{\bar{X}'_F((\hat{\beta}_{M,T=1} - \hat{\beta}_{M,T=0}) - (\hat{\beta}_{F,T=1} - \hat{\beta}_{F,T=0}))}_{Coefficients\ Effect} \tag{B.5}$$

$$+ \underbrace{(\bar{X}_M - \bar{X}_F)'((\hat{\beta}_{M,T=1}) - \hat{\beta}_{M,T=0} - (\hat{\beta}_{F,T=1} - \hat{\beta}_{F,T=0}))}_{Interaction\ Effect}$$

The twofold decomposition divides the differential in outcome into 2 parts, using a nondiscriminatory coefficient vector to determine the effect of group differences in observed characteristics. The first component is the "explained effect" or "quantity effect," which is the part of the differential that is explained by group differences in observed characteristics. The second component is the "unexplained effect", which is the residual part of the differential, which could be attributed to discrimination or group differences in unobserved characteristics.

Let β^* be a nondiscriminatory coefficient vector, and we can rewrite the differential as:

$$\Delta_Y = \underbrace{(E[X|M=1] - E[X|M=0])'\beta^*}_{Explained\ Effect} \tag{B.6}$$
$$+ \underbrace{E[X|M=1]'((\beta_{M=1,T=1} - \beta_{M=1,T=0}) - \beta^*) + E[X|M=0]'(\beta^* - (\beta_{M=0,T=1} - \beta_{M=0,T=0}))}_{Unexplained\ Effect}$$

This can be estimated using:

$$\Delta_Y = \underbrace{(\bar{X}_M - \bar{X}_F)'\hat{\beta}^*}_{Explained\ Effect} + \underbrace{\bar{X}'_M((\hat{\beta}_{M,T=1} - \hat{\beta}_{M,T=0}) - \hat{\beta}^*) + \bar{X}'_F(\hat{\beta}^* - (\hat{\beta}_{F,T=1} - \hat{\beta}_{F,T=0}))}_{Unexplained\ Effect} \tag{B.7}$$

The choice of $\hat{\beta}^*$ is often chosen to be equal to one group's estimated coefficient vector. There is, however, perhaps no reason to assume that one group's coefficient vector is non-discriminatory. Using Reimers's (1983) suggestion of using the average of the two group's coefficient vectors, $\hat{\beta}^*$ can be estimated as:

$$\hat{\beta}^* = 0.5(\hat{\beta}_{M,T=1} - \hat{\beta}_{M,T=0}) + 0.5(\hat{\beta}_{F,T=1} - \hat{\beta}_{F,T=0})$$

Table B2 presents the threefold decomposition of the gender differential in treatment effects.

Section IV presents the twofold decomposition of the gender difference in treatment effects using Reimer's (1983) specification, as well as this detailed decomposition for employment, earnings, and written contracts by each baseline characteristics. Figure B1 presents the detailed decomposition for hours and hourly wage by baseline characteristic.

While this paper relies on Reimer's (1983) specification of the nondiscriminatory

Table B2: Threefold Blinder-Oaxaca Decomposition of Gender Differential in Treatment Effect

E(Y|public)-E(Y|control)	Employed	Hours[c]	Earnings[c]	Hourly wage[c]	Written contract
Differential					
Prediction Male	0.055	0.213	0.472	0.259	0.039
	(0.019)	(0.072)	(0.107)	(0.056)	(0.014)
Prediction Female	0.038	0.140	0.218	0.138	0.001
	(0.012)	(0.045)	(0.072)	(0.044)	(0.010)
Difference	0.017	0.073	0.254	0.121	0.038
	(0.024)	(0.091)	(0.135)	(0.079)	(0.018)
Decomposition					
Endowments	0.009	0.036	0.059	0.021	0.003
	(0.010)	(0.035)	(0.059)	(0.033)	(0.007)
Coefficients	-0.013	-0.046	0.149	0.091	0.044
	(0.030)	(0.108)	(0.152)	(0.087)	(0.019)
Interaction	0.021	0.083	0.045	0.010	-0.009
	(0.015)	(0.059)	(0.093)	(0.054)	(0.012)

The male and female predictions are generated by regressing each outcome on a vector of randomization block fixed effects and baseline covariates (measured skills, self-reported skills, self-esteem, overconfidence, education, age, baseline employment, baseline search, discount rate, and risk aversion) for each gender's public treatment and control group, and then predicting the treatment effect based on the difference. The gender difference in treatment effect is then calculated as the difference in the male and female treatment effect. Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

Figure B1: Twofold Blinder-Oaxaca Covariate Decomposition of Gender Differential in Treatment Effect

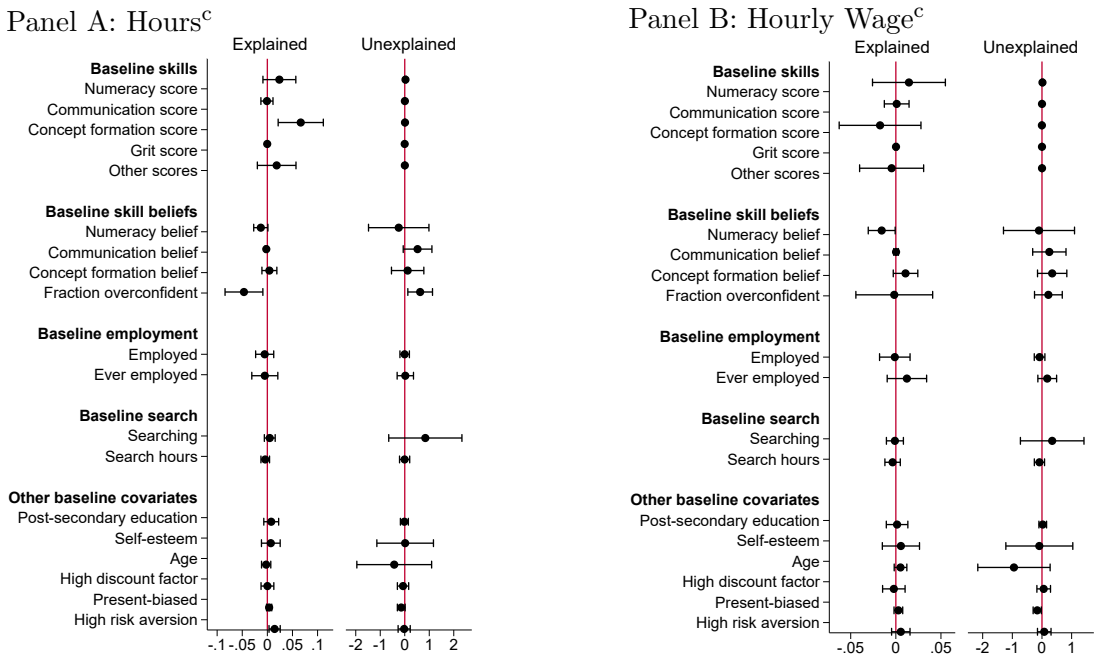

Note: The decomposition is generated using a vector of randomization block fixed effects and baseline covariates (measured skills, self-reported skills, self-esteem, overconfidence, education, age, baseline employment, baseline search, discount rate, and risk aversion). Effects for randomization block fixed effects are suppressed, as is the constant in the "unexplained" portion of the decomposition. Confidence intervals are generated using heteroskedasticity-robust standard errors, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

Table B3: Twofold Blinder-Oaxaca Decomposition of Gender Differential in Treatment Effect (Cotton)

E(Y|public)-E(Y|control)	Employed	Hours	Earnings	Hourly wage	Written contract
Differential					
Prediction Male	0.055	0.213	0.482	0.266	0.039
	(0.019)	(0.079)	(0.114)	(0.060)	(0.014)
Prediction Female	0.038	0.145	0.224	0.141	0.001
	(0.012)	(0.045)	(0.075)	(0.046)	(0.010)
Difference	0.017	0.068	0.258	0.125	0.038
	(0.024)	(0.098)	(0.146)	(0.086)	(0.018)
Decomposition					
Explained	0.017	0.076	0.084	0.026	-0.001
	(0.009)	(0.033)	(0.050)	(0.028)	(0.006)
Unexplained	0.000	-0.008	0.175	0.099	0.038
	(0.026)	(0.106)	(0.156)	(0.090)	(0.019)

The male and female predictions are generated by regressing each outcome on a vector of randomization block fixed effects and baseline covariates (measured skills, self-reported skills, self-esteem, overconfidence, education, age, baseline employment, baseline search, discount rate, and risk aversion) for each gender's public treatment and control group, and then predicting the treatment effect based on the difference. The gender difference in treatment effect is then calculated as the difference in the male and female treatment effect. Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

coefficient vector, Cotton (1988) suggests sample-weighted nondiscriminatory coefficient vector. Using Cotton's suggestion of using the sample-weighted average of the two group's coefficient vectors, $\hat{\beta}^*$ can be estimated as:

$$\hat{\beta}^* = \frac{n_M}{n_M + n_F}(\hat{\beta}_{M,T=1} - \hat{\beta}_{M,T=0}) + \frac{n_F}{n_M + n_F}(\hat{\beta}_{F,T=1} - \hat{\beta}_{F,T=0})$$

The twofold decomposition using Cotton's (1988) specification is presented in Table B3. These results are consistent with the results found using Reimer's (1983) specification.

B.2.2 Decomposition of Gender Differential in Outcome

The Blinder-Oaxaca decomposition can also be applied to the gender difference in outcome, rather than the gender difference in the treatment effect on the outcome.

$$\Delta_Y = E[Y_M] - E[Y_F] = E[Y|M=1] - E[Y|M=0] \tag{B.8}$$

The threefold decomposition is then estimated as follows:

$$\hat{\Delta_Y} = \underbrace{(\bar{X}_M - \bar{X}_F)'\hat{\beta}_F}_{Endowments\ Effect} + \underbrace{\bar{X}'_F(\hat{\beta}_M - \hat{\beta}_F)}_{Coefficients\ Effect} + \underbrace{(\bar{X}_M - \bar{X}_F)'(\hat{\beta}_M - \hat{\beta}_F)}_{Interaction\ Effect} \tag{B.9}$$

Using the nondiscriminatory coefficient vector, β^*, the twofold decomposition is estimated as follows:

$$\Delta_Y = \underbrace{(\bar{X}_M - \bar{X}_F)'\hat{\beta}^*}_{Explained\ Effect} + \underbrace{\bar{X}'_M(\hat{\beta}_M - \hat{\beta}^*) + \bar{X}'_F(\hat{\beta}^* - \hat{\beta}_F)}_{Unexplained\ Effect} \tag{B.10}$$

I perform this decomposition for the gender differential for both the public certification treatment group and the control group. Table B4 shows this decomposition for candidates in the public certification group, and Table B5 shows this decomposition for those in the control group. For both groups, the gender differential in all outcomes remains mostly unexplained.

Table B4: Twofold Blinder-Oaxaca Decomposition of Gender Differential in Outcome for Public Certification Group (Reimer)

E(Y)	Employed	Hours	Earnings	Hourly wage	Written contract
Differential					
Prediction Male	0.404	1.459	2.169	1.128	0.168
	(0.018)	(0.073)	(0.088)	(0.040)	(0.014)
Prediction Female	0.321	1.081	1.624	0.890	0.113
	(0.009)	(0.036)	(0.058)	(0.036)	(0.008)
Difference	0.083	0.378	0.545	0.238	0.055
	(0.022)	(0.087)	(0.108)	(0.056)	(0.015)
Decomposition					
Explained	0.031	0.111	0.133	0.059	0.003
	(0.007)	(0.025)	(0.042)	(0.023)	(0.004)
Unexplained	0.052	0.267	0.412	0.178	0.052
	(0.023)	(0.091)	(0.115)	(0.060)	(0.015)

The male and female predictions are generated by regressing each outcome on a vector of randomization block fixed effects and baseline covariates (measured skills, self-reported skills, self-esteem, overconfidence, education, age, baseline employment, baseline search, discount rate, and risk aversion) for each gender, and then predicting the outcome. The gender difference is then calculated as the difference in the male and female predicted outcome. Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

Table B5: Twofold Blinder-Oaxaca Decomposition of Gender Differential in Outcome for Control Group (Reimer)

E(Y)	Employed	Hours	Earnings	Hourly wage	Written contract
Differential					
Prediction Male	0.351	1.253	1.730	0.880	0.133
	(0.009)	(0.025)	(0.058)	(0.039)	(0.006)
Prediction Female	0.283	0.950	1.400	0.741	0.112
	(0.006)	(0.025)	(0.040)	(0.023)	(0.005)
Difference	0.068	0.303	0.330	0.139	0.021
	(0.012)	(0.037)	(0.079)	(0.055)	(0.008)
Decomposition					
Explained	0.008	0.019	0.040	0.028	0.006
	(0.007)	(0.023)	(0.033)	(0.019)	(0.005)
Unexplained	0.059	0.284	0.290	0.110	0.014
	(0.014)	(0.046)	(0.082)	(0.055)	(0.010)

The male and female predictions are generated by regressing each outcome on a vector of randomization block fixed effects and baseline covariates (measured skills, self-reported skills, self-esteem, overconfidence, education, age, baseline employment, baseline search, discount rate, and risk aversion) for each gender, and then predicting the outcome. The gender difference is then calculated as the difference in the male and female predicted outcome. Heteroskedasticity-robust standard errors shown in parentheses, clustering by treatment date. All outcomes use a 7-day recall period unless marked with [a] (30-day recall period) or [b] (since treatment). Outcomes marked with [c] use the inverse hyperbolic sine transformation.

Appendix C

Appendices to Chapter 4

C.1 Descriptive Statistics by Cutting Status

It is likely that the types of women who have undergone FGC are different to those who did not. Table C1 presents various descriptive statistics for adult women, by their cutting status. Women who did not undergo FGC are more likely to live in urban areas and less likely to be currently married or have been married in the past. Cut women also have better health stocks and flows, as measured by height-for-age and weight-for-height.

C.2 Robustness Checks

This section reviews the robustness checks performed for Chapter 4.

C.2.1 Age Curvature

The main analysis in this paper includes age and the square of age as covariates in order to account for any non-linearity in the impact of age. However, it is also possible to control for this non-linearity using the natural log of age.

Table C2 reports the determinants of a daughter's FGC status with no controls except daughter's age and the natural log of her age. The results of this specification do not differ from the specification presented in the main analysis.

Table C3 reports the determinants of a daughter's FGC status including a vector

Table C1: Comparison of Adult Women by FGC Status

	Had FGC	No FGC
% Urban	31.297	35.263
% Married	82.321	78.289
% Never married	13.966	17.895
Education	3.572	3.728
	(1.873)	(1.791)
Wealth index	3.224	3.216
	(1.457)	(1.507)
Spousal education	4.178	4.045
	(1.809)	(1.698)
No. co-wives	0.708	0.639
	(5.493)	(5.607)
Rank among wives	2.514	2.800
	(9.110)	(10.438)
Height-for-age	-24.700	-20.697
	(106.692)	(108.135)
Weight-for-height	-50.701	-44.575
	(112.802)	(116.004)
No. obs	9480	760

Standard deviations are presented in parentheses.

Table C2: Determinants of Daughter's FGC Status, No Controls

	(1) Daughter underwent FGC	(2) Daughter underwent FGC	(3) Daughter underwent FGC	(4) Daughter underwent FGC
Mother supports FGC	0.120***	0.098***	0.136***	0.113***
	(0.017)	(0.018)	(0.020)	(0.020)
Mother opposes FGC	-0.054***	-0.041**	-0.066***	-0.054**
	(0.018)	(0.020)	(0.022)	(0.023)
Father supports FGC	0.017	0.024	0.035**	0.030*
	(0.016)	(0.016)	(0.017)	(0.018)
Father opposes FGC	-0.014	-0.003	-0.008	-0.007
	(0.021)	(0.020)	(0.022)	(0.022)
Daughter's age			-0.038***	-0.037***
			(0.002)	(0.002)
Daughter's ln(age)			0.217***	0.217***
			(0.013)	(0.013)
Cluster FEs	No	Yes	No	Yes
Observations	11842	11842	10345	10345
Clusters	411	411	411	411

Heteroskedasticity-robust standard errors are reported in parentheses, clustered by DHS cluster.

of covariates, the daughter's age and the natural log of her age. The estimates of the impact of a mother's support or opposition of FGC are broadly consistent with the main specification presented in the paper; however, the magnitude of the point estimate associated with a mother opposing the practice are slightly larger when the natural log of age is accounted for. In addition, when the natural log of the daughter's age is included rather than her squared age, the point estimate related to a father supporting FGC is significant.

Table C3: Determinants of Daughter's FGC Staus

	(1) Daughter underwent FGC	(2) Daughter underwent FGC
Mother supports FGC	.263***	.227***
	(.032)	(.032)
Mother opposes FGC	-.154***	-.125***
	(.041)	(.041)
Father supports FGC	.055*	.018
	(.030)	(.029)
Father opposes FGC	-.003	-.015
	(.035)	(.035)
Cluster FEs	No	Yes
Observations	4702	4695
Clusters	411	404

Estimates include controls for maternal and paternal education, maternal and paternal religion, maternal prenatal care, maternal health care in the past year, maternal contraceptive use, maternal employment, age of household head, urban/rural location of household, electricity availability of household, household wealth index, daughter's age and daughter's age squared. Heteroskedasticity-robust standard errors are reported in parentheses clustered by DHS cluster.

Table C4 presents the determinants of a daughter's FGC status for daughters that are the children of parents who disagree about FGC, including a vector of covariates, the daughter's age, and the natural log of her age. These results are consistent with the main specification.

Table C4: Determinants of Daughter's FGC Staus, Parents who Disagree

	(1) Daughter underwent FGC	(2) Daughter underwent FGC
Father supports FGC	-.383***	-.350***
	(.045)	(.091)
Cluster FEs	No	Yes
Observations	839	795
Clusters	227	183

The variable indicating that the father supports FGC is the equivalent of a variable indicating that the mother opposes FGC in this specification, as the sample is restricted to couples who disagree. Estimates include controls for maternal and paternal education, maternal and paternal religion, maternal prenatal care, maternal health care in the past year, maternal contraceptive use, maternal employment, age of household head, urban/rural location of household, electricity availability of household, household wealth index, daughter's age and daughter's age squared. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

Finally, Table C5 presents the determinants of a daughter's FGC status for daughters that are the children of polygynous parents, including a vector of covariates, the daughter's age, and the natural log of her age. These results are consistent with the main specification presented in the paper.

C.2.2 Alternate Regression Models

The main specifications in this paper use OLS regressions. However, this may not be appropriate for all the regressions. For instance, when considering the marriage market hypothesis, a key outcome variable is whether a woman is married or not, which is a binary variable. For these regressions, the use of a probit or logit regression may be more suitable. Similarly, another key outcome variable used in the main specification is a polygynous wife's rank among her co-wives, a count variable. For these regressions, the use of a Poisson or negative binomial might be better suited.

Table C5: Determinants of Daughter's FGC Status, in Polygynous Families

	(1) Daughter underwent FGC	(2) Daughter underwent FGC
Mother supports FGC	.133***	.135***
	(.045)	(.044)
Maternal controls	No	Yes
Father FEs	Yes	Yes
Observations	4946	4695
Fathers	1522	1422
Clusters	384	376

Both columns control for daughter's age and daughter's age squared. Column 2 controls for the following maternal characteristics: rank among co-wives, religion, education, and employment. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

Table C6 presents the relationship between cutting status and marriage, using a probit regression. Although the point estimate on FGC status gets slightly larger, the results are consistent with the main OLS specification: cutting status does not affect the likelihood that a woman is married.

Table C6: Regression of Marriage on Cutting Status: Probit

	(1) Currently married	(2) Currently married
Underwent FGC	0.047	0.040
	(0.092)	(0.094)
Observations	5346	5346
Clusters	413	413

Both columns include controls for the age and age squared, the height-for-age and height-for-age squared, and the urban/rural location of the household. Column 2 controls for education level. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

Table C7 presents this same relationship between cutting status and marriage,

using a logit regression. Again, although the point estimate on FGC status appears slightly larger, the results do not change from the main specification: FGC status does not affect the likelihood that a woman is currently married.

Table C7: Regression of Marriage on Cutting Status: Logit

	(1) Currently married	(2) Currently married
Underwent FGC	0.091	0.075
	(0.168)	(0.172)
Observations	5346	5346
Clusters	413	413

Both columns include controls for the age and age squared, the height-for-age and height-for-age squared, and the urban/rural location of the household. Column 2 controls for education level. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

Table C8 presents the relationship between cutting status and wife rank, using a negative binomial specification. The results are consistent with the OLS specification: FGC status is not related to wife rank for wives in polygynous unions.

Table C8: Regression of Wife Rank on Cutting Status: Negative Binomial

	(1) Wife's Rank	(2) Wife's Rank
Underwent FGC	0.017	-0.439
	(0.193)	(0.955)
Observations	1567	1567
Clusters	366	366

Both columns include controls for the age and age squared, the height-for-age and height-for-age squared, and the urban/rural location of the household. Column 2 controls for education level. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

Finally, Table C9 shows the relationship between FGC status and wife rank for polygnous wives, using a Poisson specification. Again, these results are consistent

with those presented in the main specification: cutting status does not have a significant impact on a wife's rank within a polygynous family structure.

Table C9: Regression of Wife Rank on Cutting Status: Poisson

	(1) Wife's Rank	(2) Wife's Rank
Underwent FGC	-0.144	-0.439
	(0.352)	(0.955)
Observations	1567	1567
Clusters	366	366

Both columns include controls for the age and age squared, the height-for-age and height-for-age squared, and the urban/rural location of the household. Column 2 controls for education level. Heteroskedasticity-robust standard errors in parentheses, clustered by DHS cluster.

Bibliography

Abebe, G., Caria, S., Fafchamps, M., Falco, P., Franklin, S., and Quinn, S. (2019). Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City. *The Review of Economic Studies*, forthcoming.

Abel, M. (2019). Unintended Labor Supply Effects of Cash Transfer Programs: New Evidence from South Africa's Pension. *Journal of African Economies*, 28:558–581.

Abel, M., Burger, R., and Piraino, P. (2020). The Value of Reference Letters: Experimental Evidence from South Africa. *American Economic Journal: Applied Economics*.

Ahn, S. Y., Dizon-Ross, R., and Feigenberg, B. (2020). Improving Job Matching Among Youth. *Working Paper*.

Aigner, D. J. and Cain, G. G. (1977). Statistical Theories of Discrimination in Labor Markets. *Industrial and Labor Relations Review*, 30(2):175–187.

Aker, J. C. (2010). Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger. *American Economic Journal: Applied Economics*, 2(3):46–59.

Akerlof, G. A. and Kranton, R. E. (2000). Economics and Identity. *The Quarterly Journal of Economics*, 115(3):715–753.

Albanesi, S. and Olivetti, C. (2009). Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations. *Review of Economic Dynamics*, 12(1):80–107.

Alexa (2019). The Top 500 Sites on the Web. *http://alexa.com/topsites/countries*.

Alfonsi, L., Bandiera, O., Bassi, V., Burgess, R., Rasul, I., Sulaiman, M., and Vitali, A. (2017). Tackling Youth Unemployment: Evidence from a Labor Market Experiment in Uganda. *Working Paper*.

Altonji, J. G. and Pierret, C. R. (2001). Employer Learning and Statistical Discrimination. *Quarterly Journal of Economics*, 116(1):313–350.

Anderson, M. L. (2008). Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association*, 103(484):1481–1495.

Anderson, S. and Eswaran, M. (2009). What Determines Female Autonomy? Evidence from Bangladesh. *Journal of Development Economics*, 90(2):179–191.

Andersson Joona, P. and Nekby, L. (2009). TIPping the Scales towards Greater Employment Chances? Evaluation of a Trial Introduction Program (TIP) for Newly-Arrived Immigrants Based on Random Program Assignment. *IZA Discussion Paper No. 2072*.

Anipah, K., Mboup, G., Mawuéna Ouro-Gnao, A., Boukpessi, B., Messan, P. A., and Salami-Odjo, R. (1999). *Enquête Démographique et de Santé, Togo 1998*. Direction de la Statistique et Macro International Inc., Calverton, Maryland, USA.

Arcidiacono, P., Bayer, P., and Hizmo, A. (2010). Beyond Signaling and Human Capital: Education and the Revelation of Ability. *American Economic Journal: Applied Economics*, 2(4):76–104.

Arrow, K. (1973). The Theory of Discrimination. In Ashenfelter, O. and Rees, A., editors, *Discrimination in Labor Markets*, pages 3–33. Princeton University Press.

Atasoy, H. (2013). The Effects of Broadband Internet Expansion on Labor Market Outcomes. *Industrial and Labor Relations Review*, 66(2):315–345.

Attanasio, O., Guarín, A., Medina, C., and Meghir, C. (2015). Long Term Impacts of Vouchers for Vocational Training: Experimental Evidence for Colombia. *NBER Working Paper 21390*.

Attanasio, O., Kugler, A., and Meghir, C. (2011). Subsidizing Vocational Training for Disadvantaged Youth in Colombia: Evidence from a Randomized Trial. *American Economic Journal: Applied Economics*, 3(3):188–220.

Baker, M. and Cornelson, K. (2008). Gender-Based Occupational Segregation and Sex Differences in Sensory, Motor, and Spatial Aptitudes. *Demography*, 55:1749–1775.

Banks, E., Meirik, O., Farley, T., Akande, O., Bathija, H., and Ali, M. (2006). Female Genital Mutliation and Obstetric Outcome: WHO Collaborative Prospective Study in Six African Countries. *Lancet*, 367:1835–1841.

Bassi, V. and Nansamba, A. (2019). Screening and Signaling Non-Cognitive Skills: Experimental Evidence from Uganda. *Working Paper*.

Beaman, L., Keleher, N., and Magruder, J. (2018). Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi. *Journal of Labor Economics*, 36(1):121–157.

Bellemare, M. F., Novak, L., and Steinmetz, T. L. (2015). All in the Family: Explaining the Persistence of Female Genital Cutting in West Africa. *Journal of Development Economics*, 116:252–265.

Belot, M., Kircher, P., and Muller, P. (2019). Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice. *The Review of Economic Studies*, 86(4):1411–1477.

Bénabou, R. and Tirole, J. (2011). Identity, Morals, and Taboos: Beliefs as Assets. *Quarterly Journal of Economics*, 126(2):805–855.

Berg, R. C., Odgaard-Jensen, J., Fretheim, A., Underland, V., and Vist, G. E. (2014a). An Updated Systematic Review and Meta-Analysis of the Obstetric Consequences of Female Genital Mutilation/Cutting. *Obstetrics and Gynecology International*, 2014.

Berg, R. C., Underland, V., Odgaard-Jensen, J., Fretheim, A., and Vist, G. E. (2014b). Effects of Female Genital Cutting on Physical Health Outcomes: A Systematic Review and Meta-Analysis. *British Medical Journal Open*, 4(11).

Bertrand, M. and Crepón, B. (2019). Teaching Labor Laws: Evidence from a Randomized Trial in South Africa. *Working Paper*.

Bhorat, H., Caetano, T., Jourdan, B., Kanbur, R., Rooney, C., Stanwix, B., and Woolard, I. (2016). Investigating the Feasibility of a National Minimum Wage for South Africa. *Working Paper*.

Bhorat, H. I. and Cheadle, H. (2009). Labour Reform in South Africa: Measuring Regulation and a Synthesis of Policy Suggestions. *Development Policy Research Unit Working Paper 09/139*.

Bhuller, M., Kostol, A. R., and Vigtel, T. C. (2019). How Broadband Internet Affects Labor Market Matching. *Working Paper*.

Blau, F. D. and Kahn, L. M. (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, 55(3):789–865.

Blinder, A. S. (1973). Wage Discrimination: Reduced Form and Structural Estimates. *Journal of Human Resources*, 8(4):436–455.

Bound, J., Brown, C., and Mathiowetz, N. (2001). Measurement Error in Survey Data. In Heckman, J. and Leamer, E., editors, *Handbook of Econometrics*, volume 5(59), pages 3703–3843. Elsevier Science B.V.

Brown, M., Setren, E., and Topa, G. (2016). Do Informal Referrals Lead to Better Matches? Evidence from a Firm's Employee Referral System. *Journal of Labor Economics*, 34(1):161–209.

Browning, A., Allsworth, J. E., and Wall, L. L. (2013). The Relationship Between Female Genital Cutting and Obstetric Fistulas. *Obstetrics and Gynecology*, 115(3):578–583.

Bursztyn, L., Fujiwara, T., and Pallais, A. (2017). 'Acting Wife': Marriage Market Incentives and Labor Market Investments. *American Economic Review*, 107(11):3288–3319.

Buser, T., Niederle, M., and Oosterbeek, H. (2014). Gender, Competitiveness, and Career Choices. *The Quarterly Journal of Economics*, 129(3):1409–1447.

Carranza, E., Garlick, R., Orkin, K., and Rankin, N. (2020). Job Search and Hiring with Two-Sided Limited Information about Workseekers' Skills. *Working Paper*.

Casaburi, L., Kremer, M., Mullainathan, S., and Ramrattan, R. (2014). Harnessing ICT to Increase Agricultural Production: Evidence from Kenya. *Working Paper*.

Champion, S., Kosec, K., and Stanton, C. (2012). The Effects of Internet Access on Labor-Supply Decisions. *Time Warner Cable Research Program on Digital Communications Report Series*.

Cho, S.-Y. (2017). Explaining Gender Differences in Confidence and Overconfidence in Math. *MAGKS Papers on Economics 201701*.

Cho, Y., Kalomba, D., Mobarak, A. M., and Orozco, V. (2013). Gender Differences in the Effects of Vocational Training: Constraints on Women and Drop-Out Behavior. *IZA Discussion Paper 7408*.

Clarke, G. R. G. and Wallsten, S. J. (2006). Has the Internet Increased Trade? Developed and Developing Country Evidence. *Economic Inquiry*, 44(3):465–484.

Commander, S., Harrison, R., and Menezes-Filho, N. (2011). ICT and Productivity in Developing Countries: New Firm-Level Evidence from Brazil and India. *Review of Economics and Statistics*, 93(2):528–541.

Conlon, J. J., Pilossoph, L., Wiswall, M., and Zafar, B. (2018). Labor Market Search wtih Imperfect Information and Learning. *NBER Working Paper 24988*.

Coomaraswamy, R. (2002). Identity Within: Cultural Relativism, Minority Rights, and the Empowerment of Women. *George Washington International Law Review*, 34:483–492.

Crespi, G., Criscuolo, C., and Haskel, J. (2007). Information Technology, Organisational Change, and Productivity Growth: Evidence from UK Firms. *CERP Discussion Paper 6105*.

Croson, R. and Gneezy, U. (2010). Gender Differences in Preferences. *Journal of Economic Literature*, 47(2):448–474.

Czernich, N., Falck, O., Kretschmer, T., and Woessmann, L. (2011). Broadband Infrastructure and Economic Growth. *Economic Journal*, 121(552):505–532.

Dettling, L. J. (2015). Broadband in the Labor Market: The Impact of Residential High-Speed Internet on Married Women's Labor Force Participation. *Federal Reserve Board Finance and Economics Discussion Series 2013-065*.

Dickerson, A., McIntosh, S., and Valente, C. (2015). Do the Maths: An Analysis of the Gender Gap in Mathematics in Africa. *Economics of Education Review*, 46:1–22.

Duflo, E. (2000). Child Health and Household Resources in South Africa: Evidence from the Old Age Pension Program. *American Economic Review*, 90(2):393–398.

Efobi, U. R., Voufo, B. T., and Asongu, S. (2018). Female Economic Participation with Information and Communication Technology (ICT) Advancement: Evidence from Sub-Saharan Africa. *Working Paper*.

Elmusharaf, S., Elhadi, N., and Almroth, L. (2006). Reliability of Self Reported Form of Genital Mutilation and WHO Classification: Cross Sectional Study. *British Medical Journal*, 333(7559):124–127.

Epelboin, S. and Epelboin, A. (1979). Special Report: Female Circumcision. *People*, page 28.

Erosa, A., Fuster, L., Kambourov, G., and Rogerson, R. (2017). Hours, Occupations, and Gender Differences in Labor Market Outcomes. *NBER Working Paper 23636*.

Farber, H. S. and Gibbons, R. (1996). Learning and Wage Dynamics. *Quarterly Journal of Economics*, 111(4):1007–1047.

Federal Communications Commission (2019). Broadband Speed Guide. *https://www.fcc.gov/reports-research/guides/broadband-speed-guide*.

Flory, J. A., Leibbrand, A., and List, J. A. (2015). Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions. *The Review of Economic Studies*, 82(1):122–155.

Food and Agriculture Organization of the United Nations (2017). Smallholders Dataportrait. *http://www.fao.org/family-farming/data-sources/dataportrait*.

Forden, E. (2015). The Undersea Cable Boom in Sub-Saharan Africa. *USITC Executive Briefing on Trade*.

Forman, C., Goldfarb, A., and Greenstein, S. (2012). The Internet and Local Wages: A Puzzle. *American Economic Review*, 102(1):556–575.

Fryer, R. G. and Levitt, S. D. (2010). An Empirical Analysis of the Gender Gap in Mathematics. *American Economic Journal: Applied Economics*, 2(2):210–240.

Gelblum, M. (2020). Preferences for Job Tasks and Gender Gaps in the Labor Market. *Working Paper*.

Ghana Statistical Service, Ghana Health Service, and ICF International (2015). *Ghana Demographic and Health Survey 2014*. Ghana Statistical Service, Ghana Health Service, and ICF International, Rockville, Maryland, USA.

Ghana Statistical Service, Ghana Health Service, and ICF Macro (2009). *Ghana Demographic and Health Survey 2008*. Ghana Statistical Service, Ghana Health Service, and ICF Macro, Accra, Ghana.

González-Velosa, C., Ripani, L., and Rosas-Shady, D. (2012). How Can Job Opportunities for Young People in Latin America Be Improved? *Inter-American Development Bank Technical Notes IDB-TN-345*.

Grossman, Z. (2015). Self-Signaling and Social-Signaling in Giving. *Journal of Economic Behavior and Organization*, 117:26–39.

Haushofer, J., Ringdal, C., Shapiro, J. P., and Wang, X. Y. (2019). Income Changes and Intimate Partner Violence: Evidence from Unconditional Cash Transfers in Kenya. *NBER Working Paper 25627*.

Heath, R. (2014). Women's Access to Labor Market Opportunities, Control of Household Resources, and Domestic Violence: Evidence from Bangladesh. *World Development*, 57(C):32–46.

Heath, R. (2018). Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garment Factories. *Journal of Political Economy*, 126(4):1691–1724.

Heckman, J. J., Stixrud, J., and Urzua, S. (2006). The Effects Of Cognitive and Noncognitive Abilities On Labor Market Outcomes and Social Behavior. *Journal of Labor Economics*, 24(3):411–482.

Hedges, L. V. and Nowell, A. (1995). Sex Differences in Mental Test Scores, Variability, and Numbers of High-Scoring Individuals. *Science*, 269(5220):41–45.

Hjort, J. and Poulsen, J. (2019). The Arrival of Fast Internet and Employment in Africa. *American Economic Review*, 109(3):1032–1079.

Hosken, F. P. (1993). *The Hosken Report: Genital and Sexual Mutilation of Females*. Women's International Network News, 4 edition.

Hosting Checker (2019). Find Out Who Is Hosting Any Website. *http://hostingchecker.com*.

Hughes, K. (1995). Note: The Criminalization of Female Genital Mutilation in the United States. *Journal of Law and Policy*, 4:321.

Ibarrarán, P. and Rosas-Shady, D. (2009). Evaluating the Impact of Job Training Programmes in Latin America: Evidence from IDB Funded Operations. *Journal of Development Effectiveness*, 1:195–216.

ICF (1998-2017). *Demographic and Health Surveys (Various) [Datasets]*. ICF [Distributor], Rockville, Maryland. Funded by USAID.

ILO (2016). Enabling Environment for Sustainable Enterprises in South Africa. *International Labour Office - Geneva*.

Institut National de la Statistique et de l'Analyse Économique et ICF International (2013). *Enquête Démographique et de Santé du Bénin 2011-2012*. Institut National de la Statistique et de l'Analyse Économique ICF International, Calverton, Maryland, USA.

Institut National de la Statistique et de l'Analyse Économique et ORC Macro (2002). *Enquête Démographique et de Santé au Bénin 2001*. Institut National de la Statistique et de l'Analyse Économique et ORC Macro, Calverton, Maryland, USA.

International Telecommunications Union and World Bank. Individuals Using the Internet (% of Population). *World Bank Development Indicators*.

International Telecommunications Union and World Bank. Mobile Cellular Subscriptions (Per 100 People). *World Bank Development Indicators*.

Jacob, B. (2002). Where the Boys Aren't: Non-Cognitive Skills, Returns to School and the Gender Gap in Higher Education. *Economics of Education Review*, 21:589–598.

Jensen, R. (2007). The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector. *Quarterly Journal of Economics*, 122(3):879–924.

Jensen, R. and Oster, E. (2009). The Power of TV: Cable Television and Women's Status in India. *Quarterly Journal of Economics*, 124(3):1057–1094.

Johansen, R. E. B., Diop, N. J., Laverack, G., and Leye, E. (2013). What Works and What Does Not: A Discussion of Popular Approaches for the Abandonment of Female Genital Mutilation. *Obstetrics and Gynecology International*.

Johnson, M. C. (2000). Becoming a Muslim, Becoming a Person: Female "Circumcision," Religious Identity, and Personhood in Guinea-Bissau. In Shell-Duncan, B. and Hernlund, Y., editors, *Female "Circumcision" in Africa: Culture, Controversy, and Change*, pages 215–233. Lynne Rienner Publishers.

Jones, H., Diop, N., Askew, I., and Kaboré, I. (1999). Female Genital Cutting Practices in Burkina Faso and Mali and Their Negative Health Outcomes. *Studies in Family Planning*, 30(3):219–230.

Kahn, L. B. and Lange, F. (2014). Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures. *The Review of Economic Studies*, 81(4):1575–1613.

Karimjee, M. (2015). Damage. *The Big Roundtable*.

Kendall, T. D. (2011). The Relationship Between Internet Access and Divorce Rate. *Journal of Family and Economic Issues*, 32(3):449–460.

Kenya National Bureau of Statistics and ICF Macro (2010). *Kenya Demographic and Health Survey 2008-09*. Kenya National Bureau of Statistics and ICF Macro, Calverton, Maryland, USA.

Kenya National Bureau of Statistics, Ministry of Health, National AIDS Control Council, Kenya Medical Research Institute, National Council for Population and Development, and ICF International (2015). *Kenya Demographic and Health Survey 2014*. Kenya National Bureau of Statistics, Ministry of Health, National AIDS Control Council, Kenya Medical Research Institute, National Council for Population and Development, and ICF International, Rockville, Maryland, USA.

Kerr, A. (2017). Tax(i)ing the Poor? Commuting Costs in South African Cities. *Journal of Labor Economics*, 85:321–340.

Kishor, S. and Neitzel, K. (1996). *The Status of Women: Indicators for Twenty-Five Countries, DHS Comparative Studies No. 21*, chapter Gender and Household Headship, pages 13–27. Macro International, Calverton, Maryland, USA.

Kristof, N. D. and WuDunn, S. (2010). *Half the Sky: Turning Oppression into Opportunity for Women Worldwide*. Vintage Books.

La Ferrara, E., Chong, A., and Duryea, S. (2012). Soap Operas and Fertility: Evidence from Brazil. *American Economic Journal: Applied Economics*, 4(4):1–31.

Lam, D., Ardington, C., and Leibbrandt, M. (2011). Schooling as a Lottery: Racial Differences in School Advancement in Urban South Africa. *Journal of Development Economics*, 95(2):121–136.

Light, A. and Ureta, M. (1995). Early-Career Work Experience and Gender Wage Differentials. *Journal of Labor Economics*, 13(1):121–154.

Lightfoot-Klein, H. (1989). *Prisoners of Ritual: An Odyssey into Female Genital Circumcision in Africa*. Haworth Press.

Lundberg, S. J., Pollak, R. A., and Wales, T. J. (1997). Do Husband and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit. *The Journal of Human Resources*, 32(3):463–480.

Macki, G. (1996). Ending Footbinding and Infibulation: A Convention Account. *American Sociological Review*, 61:999–1017.

Madamombe, I. (2007). New Cable to Connect Eastern Africa: NEPAD Pushes for Cheaper, Faster Telecommunications. *Africa Renewal Magazine*.

Magesa, L. (1997). *African Religion: The Moral Traditions of Abundant Life*. Orbis Books.

Mahlknecht, G. (2019). Greg's Cable Map. *https://www.cablemap.info*.

Majlesi, K. (2013). Labor Market Opportunities and Sex-Specific Investment in Children's Health: The Role of Women's Decision Making Power. *Working Paper*.

Mandara, M. U. (2004). Female Genital Mutilation in Nigeria. *International Journal of Gynecology and Obstetrics*, 84(3):291–298.

Ministère de la Planification, du Développement et de l'Aménagement du Territoire, Ministère de la Santé et ICF International (2015). *Enquête Démographique et de Santé au Togo 2013-2014*. Ministère de la Planification, du Développement et de l'Aménagement du Territoire, Ministère de la Santé et ICF International, Rockville, Maryland, USA.

Ministère du Plan et et Suivi de la Mise en œuvre de la Révolution de la Modernité, Ministère de la Santé Publique et ICF International (2014). *Enquête Démographique et de Santé, République Démocratique du Congo 2013-2014*. Ministère du Plan et et Suivi de la Mise en œuvre de la Révolution de la Modernité, Ministère de la Santé Publique et ICF International, Rockville, Maryland, USA.

Ministère du Plan et Macro International (2008). *Enquête Démographique et de Santé, République Démocratique du Congo 2007*. Ministère du Plan et Macro International, Calverton, Maryland, USA.

Möbius, M. M., Niederle, M., Neihause, P., and Rosenblat, T. S. (2011). Managing Self-Confidence. *NBER Working Paper 17014*.

Namibia Minstry of Health and Social Services and ICF International (2014). *Namibia Demographic and Health Survey 2013*. Namibia Minstry of Health and Social Services and ICF International, Windhoek, Namibia and Rockville, Maryland, USA.

Namibia Minstry of Health and Social Services and Macro International Inc. (2008). *Namibia Demographic and Health Survey 2006-07*. Namibia Minstry of Health and

Social Services and Macro International Inc., Windhoek, Namibia and Calverton, Maryland, USA.

New Media and Development Communication (2008). EASSY Wikipage. *http://www.columbia.edu/itc/sipa/nelson/newmediadev/files/EASSYWikipage.doc*.

Niederle, M. and Vesterlund, L. (2007). Do Women Shy Away from Competition? Do Men Compete Too Much? *Quarterly Journal of Economics*, 122(3):1067–1101.

Nigeria National Population Commission and ICF International (2014). *Nigeria Demographic and Health Survey 2014*. Nigeria National Population Commission and ICF International, Abuja, Nigeria and Rockville, Maryland, USA.

Nigeria National Population Commission and ICF Macro (2009). *Nigeria Demographic and Health Survey 2008*. Nigeria National Population Commission and ICF Macro, Abuja, Nigeria.

Nishimwe-Niymbanira, R. and Sabela, P. T. (2019). Gender Dynamics in Employment and Labour Force Trends in South Africa. *International Journal of Economics*, 11(2):35–52.

Nour, N. M. (2008). Female Genital Cutting: A Persisting Practice. *Reviews in Obstetrics and Gynecology*, 1(3):135–139.

Oaxaca, R. (1973). Male-Female Wage Differentials in Urban Labor Markets. *International Economic Review*, 14(3):693–709.

Octapixx Worldwide (2004). Will You Marry Me?: Marriage Customs in Ethiopia, Mali, Niger, and Senegal. Film.

Pallais, A. (2014). Inefficient Hiring in Entry-Level Labor Markets. *American Economic Review*, 104(11):3565–3599.

Pallais, A. and Sands, E. G. (2016). Why the Referential Treatment? Evidence from Field Experiments on Referrals. *Journal of Political Economy*, 124(6):1793–1828.

Pekkarinen, T. (2012). Gender Differences in Education. *IZA Discussion Paper 6390*.

Pew Research Center (2015). Cell Phones in Africa: Communication Lifeline. *Pew Research Center Report*.

Phelps, E. (1972). The Statistical Theory of Racism and Sexism. *American Economic Review*, 62:659–661.

Premand, P., Brodmann, S., Almeida, R., Grun, R., and Barouni, M. (2016). Entrepreneurship Education and Entry into Self-Employment Among University Graduates. *World Development*, 77:311–327.

Rahman, A. and Toubia, N., editors (2000). *Female Genital Mutilation: A Guide to Laws and Policies Worldwide*. St. Martin's Press.

Rahman, L. and Rao, V. (2005). The Determinants of Gender Equity in India: Examining Dyson and Moore's Thesis with New Data. *Population and Development Review*, 30(2):239–268.

République de Mali: Institut National de la Statistique (2010). Resultats Provisoires RGPH 2009.

Roberts, T.-A. and Nolen-Hoeksema, S. (1989). Sex Differences in Reactions to Evaluative Feedback. *Sex Roles*, 21:725–747.

Rosenberg, M. (1965). *Society and the Adolescent Self-Image*. Princeton University Press.

Rospabé, S. (2001). An Empirical Evaluation of Gender Discrimination in Employment, Occupation Attainment and Wage in South Africa in the Late 1990s. *Working Paper*.

Rubalcav, L., Teruel, G., and Thomas, D. (2009). Investments, Time Preferences and Public Transfers Paid to Women. *Economic Development and Cultural Change*, 57(3):507–538.

Schöer, V., Ntuli, M., Rnakin, N., Sebastio, C., and Hunt, K. (2010). A Blurred Signal? The Usefulness of National Senior Certificate (NSC) Mathematics Marks as Predictors of Academic Performance at University Level. *Perspectives in Education*, 28(2):9–18.

Schumann, R. and Kende, M. (2013). Lifting Barriers to Internet Development in Africa: Suggestions for Improving Connectivity. *Report for the Internet Society*, 32729-502d.

Shell-Duncan, B. and Hernlund, Y. (2000). 'Female "Circumcision" in Africa: Dimensions of the Practice and Debate'. In Shell-Duncan, B. and Hernlund, Y., editors, *Female "Circumcision" in Africa: Culture, Controversy, and Change*, pages 1–40. Lynne Rienner Publishers.

Shepherd, D. (2008). Post-Apartheid Trends in Gender Discrimination in South Africa: Analysis through Decomposition Techniques. *Working Paper*.

Skaine, R. (2005). *Female Genital Mutilation: Legal, Cultural and Medical Issues*. McFarland & Company, Inc., Publishers.

Social Security: Research, Statistics and Policy Analysis (2019). Social Security Programs Throughout the World: Africa. *https://www.ssa.gov/policy/docs/progdesc/ssptw*.

Tanzania National Bureau of Statistics and ICF Macro (2011). *Tanzania Demographic and Health Survey 2010*. Tanzania National Bureau of Statistics and ICF Macro, Dar es Salaam, Tanzania.

Tanzania National Bureau of Statistics and Macro International Inc. (2000). *Tanzania Reproductive and Child Health Survey 1999*. Tanzania National Bureau of Statistics and Macro International Inc., Calverton, Maryland, USA.

Taylor, S., van der Berg, S., Reddy, V., and Janse van Rensburg, D. (2011). How Well Do South African Schools Convert Grade 8 Achievement into Matric Outcomes? *Working Paper*.

The National Center for Victims of Crime (2018). 2018 National Crime Victims' Rights Week Resource Guide: Crime and Victimization Fact Sheets. *https://ovc.ncjrs.gov/ncvrw2018/info_flyers/fact_sheets*.

Thomas, D. (1990). Intra-Household Resource Allocation: An Inferential Approach. *The Journal of Human Resources*, 25(4):635–664.

Valenzuela, S. (2014). Social Network Sites, Marriage Well-Being and Divorce: Survey and State-Level Evidence from United States. *Computers in Human Behavior*, 36:96–101.

van Broekhuizen, H. and Spaull, N. (2017). 'The Martha Effect': The Compounding Female Advantage in South African Higher Education. *Working Paper*.

van der Berg, S. and Shepherd, D. (2015). Signalling Performance: Continuous Assessment and Matriculation Examination Marks in South African Schools. *South African Journal of Childhood Education*, 5:78–94.

Wagner, N. (2011). Why Female Genital Cutting Persists. *Working Paper*.

Wagner, N. (2015). Female Genital Cutting and Long-Term Health Consequences: Nationally Representative Estimates Across 13 Countries. *Journal of Development Studies*, 51(3):226–246.

Walter, T., Bonin, H., Butschek, S., Schütz, H., Schröder, H., Knerr, P., Steinwede, J., and Thomsen, S. L. (2014). Evaluation "Programm zur berufsbezogenen Sprachförderung für Personen mit Migrationshintergrund (ESF-BAMF-Programm)". *Abschlussbericht, Bundesministerium für Arbeit und Soziales, Berlin*.

Wayback Machine (2019). Internet Archive. *http://archive.org/web*.

Wiswall, M. and Zafar, B. (2017). Preference for the Workplace, Investment in Human Capital, and Gender. *Quarterly Journal of Economics*, 133(1):457–507.

World Health Organization (2014). Female Genital Mutilation: Fact Sheet No. 241. http://www.who.int/mediacentre/factsheets/fs241/en/.

Yoder, P. S. and Wang, S. (2013). Female Genital Cutting: The Interpretation of Recent DHS Data. In *DHS Comparative Reports*, 33. United States Agency for International Development.

Zeltzer, D. (2020). Gender Homophily in Referral Networks: Consequences for the Medicare Physician Earnings Gap. *American Economic Journal: Applied Economics*, 12(2):169–197.

Zhu, M. (2019). Job Networks through College Classmates: Effects of Referrals for Men and Women. *Working Paper*.

Biography

Rachel Sayers pursued her undergraduate studies at Loyola University New Orleans in Louisiana, where she graduated in 2013 from the Honors program with a B.A. in Economics, Political Science, and Environmental Science and minors in Mathematics and African and African American Studies. She pursued her graduate studies at Duke University, where she obtained her M.A. in Economics in 2019 and her Ph.D. in Economics in 2020. She joined AidData at William & Mary in 2020 as a Senior Research Analyst.