

Gendered Language

Pamela Jakiela and Owen Ozier*

Preliminary draft. Do not circulate without permission.

September 28, 2017

Abstract

Languages use different systems for classifying nouns. *Gender languages* assign many — sometimes all — nouns to distinct sex-based categories, masculine and feminine. Drawing on a broad range of historical and linguistic sources, we estimate the proportion of each country's population whose native language is a gender language. At the cross-country level, we document a robust negative relationship between prevalence of gender languages and women's labor force participation. We also show that traditional views of gender roles are more common in countries with more native speakers of gender languages. In African countries where indigenous languages vary in terms of their gender structure, educational attainment and female labor force participation are lower among those whose native languages are gender languages. Cross-country and individual-level differences in labor force participation are large in both absolute and relative terms (when women are compared to men), suggesting that the observed patterns are not driven by development or some unobserved aspect of culture that affects men and women equally. Following the procedures proposed by Altonji, Elder, and Taber (2005) and Oster (forthcoming), we show that the observed correlations are unlikely to be driven by unobservables. Gender languages appear to reduce women's labor force participation and perpetuate support for unequal treatment of women and men in society.

Keywords: grammatical gender, language, labor force participation, culture

*Jakiela: University of Maryland and IZA, email: pjakiela@umd.edu; Ozier: World Bank Development Research Group, BREAD, and IZA, email: oozier@worldbank.org. All errors are our own.

The findings, interpretations and conclusions expressed in this paper are entirely those of the authors, and do not necessarily represent the views of the World Bank, its Executive Directors, or the governments of the countries they represent.

1 Introduction

Language structures thought. All human beings use language to articulate their ideas and communicate them to others. Yet, the world's languages show tremendous diversity in terms of their structure and vocabulary. Different languages obviously use different words to describe the same concept, but they also organize the relationships between concepts in remarkably different ways. Because languages are so diverse and language is so fundamental to thought, some scholars have argued that the language we speak may limit the scope of our thinking. Benjamin Lee Whorf, one of the original proponents of this theory of *linguistic determinism*, famously argued that it was difficult for humans to think about ideas or concepts for which there was no word in their language (Whorf 2011[1956]a).

Though specious anecdotes about obscure languages abound, cognitive scientists have largely refuted the strongest forms of Whorf's hypothesis (Boroditsky, Schmidt, and Phillips 2003). Nonetheless, there is mounting evidence for weaker forms of linguistic determinism: the languages we speak shape our thoughts in subtle, subconscious ways. For example, implicit association tests show that bilinguals display different subconscious attitudes when tested in their different languages (Ogunnaike, Dunham, and Banaji 2010, Danziger and Ward 2010). Differences in language structure also influence our behavior in the economic realm. Keith Chen (2013) demonstrates that speakers of languages that demarcate the future as separate from the present (e.g. English) save for the future less than those whose languages make no such distinction (e.g. German).

Several recent papers explore the link between language and gender roles. As Alesina, Giuliano, and Nunn (2013) note, views of the appropriate role for women in society differ markedly across cultures. Languages also vary in their treatment of gender distinctions. At one extreme, languages such as Finnish and Swahili do not mark gender distinctions in any systematic way: nouns are not categorized as either masculine, feminine, or neuter; and the same first, second, and third person pronouns are used for males and females. Many languages mark gender by using different pronouns for males and females: for example, "he"

and “she” in English. Some languages go even further, extending the gender distinction to inanimate nouns through a system of *grammatical gender*. For example, languages such as Spanish and Italian partition all nouns — even inanimate objects — into distinct gender categories. This feature of language forces gender into every aspect of life: for a speaker of a *gender language*, gender distinctions are both inherent and salient in every thought and utterance; every object is either masculine or feminine because it is intrinsically linked to a word that carries a grammatical gender.

Does grammatical gender shape (non-grammatical) gender norms? Does it impact women’s participation in economic life? Evidence from psychology suggests that grammatical gender impacts children’s perceptions of human gender differences (Guiora 1983). In the economic realm, one recent study of immigrants to the United States shows that those who grew up speaking a gender language are more likely to divide household tasks along gender lines (Hicks, Santacreu-Vasut, and Shoham 2015). Female labor supply is also lower among immigrants who speak a gender language at home (Gay, Hicks, Santacreu-Vasut, and Shoham forthcoming).

We provide new evidence to support the hypothesis that grammatical gender shapes views of women’s role in society and directly impact women’s labor force participation. To do this, we construct a data set that characterizes the grammatical gender structure of 3,930 living languages. Taken together, these languages account for 6.37 billion people, or 98 percent of the world population.¹

We use these data in two ways. First, we calculate — for every country in the world — an estimate of the proportion of each country’s population whose native language is a gender language. We are able to account for more than 90 percent of the estimated population in all but six countries. This allows us to explore the cross-country relationship between gender languages and female labor force participation. We find that a 10 percentage point increase in the share of a country’s population whose native language is a gender language is associated with a 1.68 percentage point decline in women’s labor force participation and

¹This calculation is based on *Ethnologue* estimates of the total number of native speakers in the world.

a 1.19 percentage point decline increase in the magnitude of the gender gap in labor force participation. These associations are robust to the inclusion of a wide range of geographic controls (including suitability for the plough) that could not plausibly have been impacted by language. Following the approach suggested by Altonji, Elder, and Taber (2005) and Oster (forthcoming), we estimate that unobservable country-level characteristics would need to be 1.19 times more correlated with treatment than observed covariates to fully explain the apparent impact of grammatical gender on the level of female labor force participation; unobserved factors would need to be 2.6 times more closely linked to treatment to explain the impact of grammatical gender on the gender gap in labor force participation.

Using data from the World Values Survey, we also show that grammatical gender also predicts support for traditional gender roles. However, the prevalence of gender languages does not explain cross-country differences in women’s educational attainment.

To further assess the likelihood of a causal link between gender languages and female labor force participation, we use our data to explore the link between grammatical gender and women’s labor force participation in parts of sub-Saharan Africa where both gender and non-gender languages are indigenous and widely spoken. Combining our language data with Afrobarometer surveys from Kenya, Nigeria, Niger, and Uganda, we show that grammatical gender is associated with drastically reduced female labor force participation. Women whose native language is a gender language are 18 percentage points less likely to be in the labor force than women whose native language is not a gender language. Women from groups who speak gender languages are also less likely to be in the labor force relative to men from their ethnic group. At the individual level, we also find a robust negative relationship between grammatical gender and women’s educational attainment. Thus, gender languages appear to reduce women’s labor force participation and shape gender attitudes.

The rest of this paper is organized as follows. Section 2 introduces the concept of grammatical gender and surveys recent research on its impacts. Section 3 provides an overview of our data sources, including the data we have compiled on the grammatical structure of nearly four thousand languages. Section 4 presents our analysis, and Section 5

concludes.

2 Grammatical Gender

Many languages partition the set of all nouns into mutually exclusive categories. Membership in these categories, which are typically referred to as either *genders* or *noun classes* (Corbett 1991, Aikhenvald 2003), can be manifest in several ways. First, members of a noun class may be semantically related. For example, Tamil has three noun classes: nouns representing male humans and gods, nouns representing female humans and goddesses, and all other nouns (Corbett 1991). However, such strictly semantic systems of noun classification (where a word’s meaning always determines its categorization) are relatively rare. In many languages, members of a class are linked by morphology rather than semantics. For example, members of the KI-/VI- class in Swahili often begin with *ki-* in the singular and *vi-* in the plural — e.g. “chair” is *kiti* and “chairs” is *viti*. However, though semantic and morphological commonalities are a common characteristic of noun classes, they are not required.

Instead, membership in a specific noun class is defined based on agreement: class must be reflected in the conjugation of associated words within the noun phrase or the predicate in grammatically correct speech (Aikhenvald 2003).² In Swahili, for example, the noun class determines the prefixes used to modify adjectives, verbs, demonstratives, and other parts

²There is some debate amongst linguists as to whether agreement rules that do not involve elements of the noun phrase or the predicate can form the basis of a noun class system — specifically, linguists disagree as to whether requiring “anaphoric agreement” between nouns and associated pronouns constitutes a system of grammatical gender (Corbett 1991, Aikhenvald 2003). Corbett (1991) argues that there is no fundamental distinction between pronominal agreement and other forms of grammatical agreement; he consequently classifies languages that (only) require pronominal agreement (e.g. English) as gender languages in his work (Corbett 2013a, Corbett 2013b, Corbett 2013c). Aikhenvald (2003) agrees that there is no fundamental distinction between pronominal agreement and other forms of grammatical concordance, but advocates the use of the traditional definition of grammatical gender to avoid confusion. She also suggests restricting the use of the term “grammatical gender” to systems of noun classification involving a relatively small number of categories that include masculine and feminine. Since our focus is on the links between grammatical gender and non-grammatical gender norms, we adopt her terminology to avoid confusion. Employing the traditional definition of grammatical gender also facilitates the use of data from a wide range of linguistic and anthropological sources. Grammatical systems that require pronominal agreement but do not require agreement within the noun phrase or between the noun and the predicate are relatively rare; English is the only prominent example.

of speech. So, “this new chair” is *kiti kipya hiki*, while “this new house” is *nyumba mpya hii* because the word “house” is part of the N-/N- rather than the KI-/VI- noun class.³ Nouns are said to belong to the same agreement class if, “given the same conditions, they will take the same agreement form” (Corbett 1991, p. 148), where the relevant “conditions” are linguistic and typically relate to number and case.

Systems of noun classification differ widely across languages, and not all languages have such a system. One of the most common bases for a system of noun classification is biological sex: female humans and some other nouns are assigned to one category, while male humans and some other nouns are assigned to a different category (Corbett 1991, Aikhenvald 2003, Hellinger and Bußman 2003).⁴ We refer to such systems, which assign nouns to agreement classes that are based on biological sex as *grammatical gender*; we refer to languages characterized by such systems of grammatical gender as *gender languages* (Aikhenvald 2003, Hellinger and Bußman 2003). Spanish is a prominent example of a gender language: all Spanish nouns are either masculine or feminine, and both definite articles and adjectives must be consistent with a noun’s gender. So, for example, “the white house” is

<i>la</i>	<i>casa</i>	<i>blanc-a</i>
the.FEM	house	white-FEM,

because “house” is feminine, but “the white car” is

<i>el</i>	<i>coche</i>	<i>blanc-o</i>
the.MASC	car	white-MASC

because “car” is masculine. A Spanish speaker must therefore maintain a mental map that assigns each noun to one of these two distinct gender categories.

³Corbett (1991) states: “The existence of gender can be demonstrated only by agreement evidence. . . Evidence taken only from the nouns themselves, such as the presence of markers on the nouns, does not of itself indicate that a language has genders (or noun classes); if we accepted this type of evidence, then we could equally claim that English had a gender comprising all nouns ending in *-ion*.” Thus, though many nouns within a class may share particular prefixes or suffixes, it is the requirement that other parts of speech (particularly elements of the noun phrase or the predicate) conjugate or inflect appropriately that distinguishes noun classes from other phonological or orthographic partitions of the set of all nouns.

⁴Almost all languages also distinguish between singular and plural, but this is not typically treated as a system of noun classification because the singular and plural forms are treated as two variants of the same noun.

Systems of grammatical gender differ along several dimensions. Languages differ in terms of the extent of agreement across parts of speech, and the extent to which the gender distinction represents a complete partition of the set of all nouns. Languages such as Spanish — with only two sex-based noun classes — are at one end of this spectrum. In such languages, every inanimate noun *must* be classified as either feminine or masculine. Languages such as German display a weaker form of grammatical gender because some objects are classified as neither feminine nor masculine. Intuitively, one might think that the partition of nouns into two dichotomous genders suggests that other aspects of the universe should also be so organized (for example, into male and female household tasks). In systems that assign objects (i.e. nouns) without natural gender to gender categories, there is also question of what the observed grouping signals about the relative status of women and men. Though the rules used to assign nouns to different classes are often phonological (e.g. Spanish nouns that end in “o” are typically masculine), many languages assign some nouns to the feminine gender using semantic guidelines that have a certain cultural intelligibility. For example, dangerous objects are feminine in the Australian language Dyirbal (Lakoff 1987), while one linguist studying the Siberian language Ket suggested that certain small animals were feminine “because they are of no importance to the Kets” (Corbett 1991, p. 19).

Even among languages that do not assign objects lacking a natural (i.e. biological) gender to grammatical gender categories, there is variation in the extent to which a grammatical distinction is made between males and females.⁵ For instance, though typically not classified as a gender language, English displays a limited system of pronominal agreement — different third-person singular pronouns are used for male and female humans and, in some cases, male and female animals (Aikhenvald 2003, Boroditsky, Schmidt, and Phillips 2003, Hellinger and Bußman 2003, Kilarski 2013).⁶ Such systems of limited pronom-

⁵Of course, among humans, biological sex and gender may not always be in alignment. However, since we are studying language-level variation in the assignment of (primarily inanimate) nouns to gender categories, issues that arise in cases where an individual human needs to be assigned to a sex-based gender category in line with either his or her gender or his or her biological sex are not particularly relevant.

⁶Female pronouns have traditionally been used to refer to ships and other large transportation vessels.

inal agreement based on biological gender are also present in a handful of Niger-Congo languages that show no other form of gender inflection (Aikhenvald 2003, Creissels 2005). This contrasts with languages such as Finnish, Hungarian, and Swahili that make no grammatical distinction between males and females.⁷

Whether grammatical gender distinctions influence (non-grammatical) gender attitudes is an empirical question, but the idea that they might is not new. Whorf, for example, argued that gender distinctions in language might make a gendered division of labor seem more natural, arguing that viewing the world through the lens of a gender language would create “a sort of habitual consciousness of two sex classes as a standing classificatory fact in our thought-world” (Whorf 2011[1956]b, p. 69). A number of studies in psychology suggest that Whorf was (at least partially) correct. For example, Guiora (1983) found that children who grow up speaking Hebrew, English, or Finnish come to understand their own biological genders at different ages; those who grow up using different pronouns for males and females become aware of their own natural gender earlier.⁸ Boroditsky, Schmidt, and Phillips (2003) conducted a study — in English — of native speakers of Spanish and German; participants in the study were asked to provide adjectives to describe nouns that had been chosen because they had opposite grammatical genders in Spanish and German.⁹ Results were consistent with the hypothesis that grammatical gender shapes the way we think about inanimate objects without inherent biological gender: participants tended to list English adjectives with gender connotations that matched the gender of the nouns in their native language.

Recent work by economists suggests that the influence of grammatical gender extends

As discussed above, because pronouns agree with the natural gender of animate nouns, Corbett (1991) classifies English as a gender language with a limited and strictly semantic system of noun classification (i.e. a system of grammatical gender based only on biological gender).

⁷Swahili uses a system of noun classes that functions as a grammatical gender system, but it is not based sex-based — i.e. all humans are in the same noun class; there is not distinction between males and females (Wilson 1985, Corbett 1991).

⁸As discussed above, English has a system of pronominal gender while Finnish does not. Hebrew also uses a dichotomous system of grammatical gender (all nouns are either masculine or feminine), so male and female Hebrew speakers must use grammatically correct verb forms, for example, that reflect their natural gender. Hebrew also uses different second-person pronouns for males and females.

⁹Sun, for example, is masculine in Spanish but feminine in German, while moon is feminine in Spanish but masculine in German (Boroditsky, Schmidt, and Phillips 2003).

past childhood and outside the psychology laboratory. Countries where the national language uses a sex-based system of grammatical gender are less likely to implement gender quotas for political office and have relatively fewer women in corporate leadership positions, relative to countries where the dominant national language is not a gender language (Santacreu-Vasut, Shoham, and Gay 2013, Santacreu-Vasut, Shenkar, and Shoham 2014). More recently, Hicks, Santacreu-Vasut, and Shoham (2015) show that immigrants to the United States assign tasks within the household along gendered lines if they grew up speaking a gender language; no such difference is found among immigrants who came to the U.S. before the age of language acquisition, or among the children of immigrants.¹⁰ Importantly, their findings suggest that one’s native language plays a particularly crucial role in shaping one’s views on the appropriate role for women in society.

3 Data

More than four thousand of the world’s living languages have at least one thousand native speakers, and 382 languages have more than one million native speakers (Lewis, Simons, and Fennig, eds., 2016). Moreover, very few countries are entirely devoid of linguistic heterogeneity, and many countries have two or more widely spoken languages that differ in terms of their gender structure.¹¹ Developing countries, in particular, tend to be characterized by relatively high ethnolinguistic fractionalization (Easterly and Levine 1997).

From an empirical perspective, within-country linguistic heterogeneity is both a problem and an opportunity. Empirically, linguistic heterogeneity is a problem for cross-country analysis if it introduces measurement error and attenuation bias, leading to underestimates of the impact of grammatical gender. Many developed countries — e.g. Canada — have high levels of linguistic heterogeneity and widely spoken languages that differ in terms of their grammatical gender structures. Among less-developed countries, linguistic heterogeneity is

¹⁰More recently, Gay, Hicks, Santacreu-Vasut, and Shoham (forthcoming) find that female immigrants to the United States exhibit lower labor market participation (working fewer hours, fewer weeks, etc.) if they speak a gender language at home.

¹¹Only four countries are linguistically homogeneous (i.e. have only one native language that is common to all citizens): Cape Verde, Maldives, North Korea, and San Marino (Lewis, Simons, and Fennig, eds., 2016).

the norm rather than the exception.¹² Consequently, analysis that focuses on countries with a single dominant national language will miss much of the world’s gender inequality, since Sub-Saharan Africa and South Asia have both high levels of linguistic heterogeneity and high levels of gender inequality.¹³ Thus, to understand the relationship between language and gender norms worldwide, one needs to employ an empirical approach that does not assume each country can be represented by a single national language.

However, within-country linguistic heterogeneity also creates an opportunity to estimate the impacts of grammatical gender while holding many cultural and institutional factors constant — in countries where both gender languages and non-gender languages are indigenous. Indeed, recent work by Gay, Hicks, Santacreu-Vasut, and Shoham (forthcoming) shows that speaking a gender language (at home) is negatively correlated with female labor supply among immigrants to the U.S. — even after controlling for country of origin fixed effects. Unfortunately, much of the existing work on grammatical gender and economic outcomes has been limited by the availability of data on the gender structure of individual languages. This has prevented scholars from extending cross-country analyses beyond a set of mostly wealthy countries, and has made analysis of languages within many countries impossible.

We compile a new data set characterizing the gender structure of nearly 4,000 living languages. Together, the languages that we classify account for approximately 98 percent of the world’s population. and historical sources. The downside of this approach is that, because the data set was not compiled by a single linguist, there may be measurement error at the language level. Specifically, while many historical sources explicitly state that languages either do or do not have a system of grammatical gender, we cannot be sure that the same precise definition of grammatical gender is being used across sources. On the plus

¹²This pattern is not accidental. Many developing countries — particularly those in Africa and Asia — were European colonies until the middle of the 20th century. In many developing countries (e.g. Paraguay and Senegal), most people’s mother tongue is an indigenous language that is not the primary language of government or the educational system.

¹³The UN’s Gender Inequality Index ranks the Arab World, South Asia, and Sub-Saharan Africa substantially higher than the rest of the world in terms of gender inequality. Of the world’s 15 most ethnolinguistically fractionalized countries, 14 are in Sub-Saharan Africa and the other is India (Easterly and Levine 1997).

side, this approach allows us code the grammatical structure of thousands of languages accounting for almost all of the world’s population.

3.1 Sources of Data

3.1.1 Data on Native Languages

Data on the world’s native languages comes from the *Ethnologue*, a comprehensive database of over 7,000 living languages (Lewis, Simons, and Fennig, eds., 2016). For every known language, the *Ethnologue* provides an estimate of the number of native speakers (if any) in every country. Data are drawn from a range of sources including national censuses and surveys compiled by linguists. Combining the *Ethnologue* data with information on the grammatical gender structure of the world’s languages allows us to construct an estimate of the fraction of each country’s people who speak a gender language as their native language.

Of the 7,457 languages included in the *Ethnologue* database, we drop languages that are extinct or have no native speakers, sign languages, and dying languages that had fewer than 100 native speakers when last assessed by *Ethnologue* researchers. This leaves 6,189 languages. Together, these languages account for an estimated 6.50 billion native speakers. Of these, we are able to identify at least one credible source characterizing the gender structure of 3,930 languages which together account for 6.37 billion native speakers.

3.1.2 Grammatical Gender Data

Data on the gender structure of languages comes from a range of sources. One of the best known is the World Atlas of Language Structures (WALS), which characterizes the noun classification system of 257 languages (Corbett 2013a, Corbett 2013b, Corbett 2013c). Unfortunately (from our perspective), many of these are indigenous languages from Papua New Guinea, Australia, and the Americas with very few living speakers.¹⁴ Additional

¹⁴Data from the WALS must be used with caution because they were compiled by the linguist Greville Corbett; as discussed above, Corbett advocates the use of a non-standard definition of grammatical gender that includes systems of anaphoric pronominal agreement (Corbett 1991). This is particularly problematic when one combines the WALS with other data sources that do not classify systems of pronominal agreement based on the gender of the referent as examples of grammatical gender. We address this by excluding WALS

sources of data include: the UCLA Language Materials Project (UCLA Language Materials Project 2014), which provides detailed descriptions and learning materials for 116 languages; George L. Campbell’s *Compendium of the World’s Languages* (Campbell 1991), which characterizes several hundred of the world’s most widely spoken languages; and George Abraham Grierson’s eleven-volume *Linguistic Survey of India* (Grierson 1903a, 1903b, 1904, 1905, 1907, 1908, 1909, 1916, 1919, 1921), which was compiled between 1891 and 1928 and covers more than 300 South Asian languages and dialects. Additional data on the grammatical gender structures of languages comes from academic articles and teaching materials focused on individual languages. Detailed information on the full range of sources used to code the world’s languages is provided in the Online Appendix.

We successfully classify the grammatical gender structure of 3,930 languages. Though this is only 64 percent of the world’s living languages, those that we classify together account for 98 percent of the native speakers counted in the *Ethnologue* database. As we discuss below, this allows us to construct reasonable estimates of the proportion of individuals who speak gender languages as their native languages for the overwhelming majority of countries.

3.1.3 Other Sources of Data

Additional data for our cross-country analysis comes from several sources. Data on labor force participation, income, and population come from the World Bank’s World Development Indicators database. We use data on labor force participation in 2011, which is available for 177 countries (i.e. member states of the United Nations) plus Hong Kong, Macao, Puerto Rico, and the West Bank and Gaza. We also make use of data on schooling attainment from the Barro-Lee Educational Attainment Data Set (Barro and Lee 2013), which is available for 142 countries plus Hong Kong and Macao. Data on gender attitudes comes from the World Values Survey; it is available for 56 countries plus Puerto Rico (World

data on languages that are classified as “strictly semantic” (i.e. agreement class can always be inferred from the meaning of the noun) since Corbett considers pronominal agreement an example of such a system. We rely on other sources to classify those languages. Languages that are classified in the WALS as either lacking a grammatical gender system or having a system that is “semantic and formal” are unambiguous.

Values Survey Association 2015). Finally, we take several country-level geographic controls (average precipitation and rainfall plus suitability for the plough) from Alesina, Giuliano, and Nunn (2013). These data are available for 173 countries plus Puerto Rico.

Data for our individual-level analysis comes from the Afrobarometer Surveys (Afrobarometer Data 2016). Afrobarometer surveys have been conducted in 36 African countries and are representative of the voting age population within each country. Given the salience of ethnolinguistic identities in many African societies, the Afrobarometer collects data on respondents’ native languages. We use data from five countries where both gender and non-gender languages are indigenous: Kenya, Niger, Nigeria, and Uganda. Data for Niger is only available in Round 5 of the Afrobarometer (2011–2013). For the other three countries, four rounds of data are available: 2002–2003, 2005–2006, 2008–2010, and 2011–2013.¹⁵ Across all countries and rounds, our data set includes 26,547 respondents who speak 292 different native languages.¹⁶

4 Cross-Country Analysis

4.1 Empirical Strategy

In our cross-country analysis, we examine the association between women’s labor force participation and the proportion of a country’s population whose native language is a gender language, $Gender_c$. Our main empirical specification is an OLS regression of the form:

$$LFP_c = \alpha + \beta Gender_c + \delta_{continent} + \lambda X_c + \varepsilon_c \quad (1)$$

where LFP_c is women’s labor force participation in country c (in 2011), $Gender_c$ is an estimate of the proportion of the population of country c whose native language is a gender language, $\delta_{continent}$ is a vector of continent fixed effects, X_c is a vector of of country level

¹⁵Kenya, Nigeria, and Uganda were also included in the first round of the Afrobarometer. However, that data set does not contain detailed information on native languages.

¹⁶We are forced to drop 0.93 percent of the sample because we are unable to establish the gender structure of the native language.

geography controls, and ε_c is a conditionally mean-zero error term.¹⁷ Standard errors are clustered at the level of the most widely spoken language within each country. Our cross-country analysis of the relationship between women’s labor force participation and grammatical gender includes data on 178 countries: all the independent states for which data on women’s labor force participation is available from the World Bank development indicators database.

Our main outcome of interest is women’s labor force participation. However, we do not wish to conflate gender differences in labor market participation with structural factors that impact labor force participation among both men and women. To rule out this possibility, we include specifications where the outcome variable is the gender difference in labor supply, i.e. women’s labor force participation minus men’s labor force participation.

We also examine two other outcome variables related to gender norms: women’s educational attainment and gender attitudes. As discussed above, data on women’s educational attainment comes from the Barro-Lee data set, and is available for 144 countries (Barro and Lee 2013). Our analysis of educational outcomes parallels our analysis of labor force participation: we consider both average levels of education among women and differences between men’s and women’s (average) educational attainment. Data on gender attitudes comes from the World Values Survey (WVS) and is available for 56 countries plus Puerto Rico. In our main analysis, we construct an index of gender attitudes by taking the first principal component of the eight WVS questions on gender roles.

4.2 Labor Force Participation

Panel A of Figure 2 presents a scatter plot of the relationship between women’s labor force participation and the proportion of a country’s population that speaks a gender language as a native language. The figure depicts a negative relationship: among countries where

¹⁷Our results are also robust to the inclusion of additional contemporaneous controls such as log GDP per capita and population. However, such controls might be directly impacted by gender norms and women’s involvement in the labor force, creating a “bad controls” problem and biasing the coefficient of interest (Angrist and Pischke 2008, Acharya, Blackwell, and Sen 2016). We therefore focus on geographic controls — proportion tropical, precipitation, temperature, suitability for the plough, and an indicator for being landlocked — which are plausibly exogenous (conditional on national boundaries).

no one speaks a gender language as a native language, the average rate of women’s labor force participation is 57 percent; it is only 40 percent in countries where everyone’s native language is a gender language. Panel B of 2 documents the positive association between the prevalence of gender languages and the gender difference in labor force participation. Indeed, women in countries where none of the native languages is a gender language are 18 percentage points less likely to participate in the labor force than men; this difference rises to 32 percentage points in countries where all of the native languages are gender languages.

We confirm the statistical significance of this relationship in a regression framework in Table 1. In the first three columns, the outcome variable is the average level of female labor force participation in country c in 2011. We report a parsimonious specification with no controls in Column 1. Gender languages are negatively and significantly associated with lower levels of female labor force participation. The coefficient estimate suggests that a 10 percentage point increase in the proportion of the population whose native language is a gender language is associated with a 1.7 percentage point decline in female labor force participation (p-value 1.78×10^{-8}). Column 2 of Table 1 reports a specification that includes continent fixed effects; Column 3 also includes geographic controls. The coefficient of interest is negative and statistically significant in both specifications. Moreover, it remains reasonably similar in magnitude: when all of our geographic controls are included, the coefficient suggests that a 10 percentage point increase in the proportion of the population whose native language is a gender language is associated with a 1.4 percentage point decrease in women’s labor force participation (p-value 3.63×10^{-5}).

In Columns 4 through 6 of Table 1, we replicate our analysis using the gender difference in labor force participation as the dependent variable. Gender languages are also associated with robust differences in women’s labor force participation relative to men. In a parsimonious specification with no controls (Column 4), we find that a 10 percentage point increase in the prevalence of gender languages is associated with a 1.2 percentage point increase in the gender difference in labor force participation (p-value 7.31×10^{-6}). When we include continent fixed effects and country-level geography controls, the coefficient rises to suggest

a 1.5 percentage point increase in the gender difference in labor force participation for each 10 percentage point increase in the prevalence of gender languages (p-value 9.61×10^{-6}). Thus, the proportion of a country’s population whose native language is a gender language is a robust predictor of gender differences in labor force participation.

The regressions presented in Table 1 document the association between grammatical gender and women’s labor force participation. Of course, gender languages are not randomly assigned, and the observed correlation may be driven by some unobserved casual factor that is correlated with language and labor force participation. To assess whether the observed correlation is likely to represent a causal link between language and labor force participation, we follow the approach suggested by Altonji, Elder, and Taber (2005) and further refined by Oster (forthcoming). Under the assumption that the relationship between the outcome variable, treatment, and the observed controls is similar to the relationship between the outcome, treatment, and unobserved controls, this approach relates changes in coefficient magnitudes as controls are added to changes in the observed R^2 . Analysis following the procedures outline by Oster (forthcoming) suggests that the causal effect of our measure of the prevalence of gender languages on women’s labor force participation is likely to lie between -13.99 and -5.91 percentage points, while the causal impact of gender languages on the gender difference in labor force participation is likely to fall between 15.34 and 19.25 percentage points. Thus, the analysis suggests that gender languages reduce women’s labor force participation in both absolute and relative terms.

4.3 Other Outcomes

We consider two other outcomes related to women’s labor force participation: education and gender attitudes. Education is a key determinant of wages; in many countries, gender differences in educational attainment translate into gender gaps in wages and economic empowerment (Grant and Behrman 2010). In Table 2, we show that the prevalence of gender languages does not explain either women’s educational attainment (measured in terms of years of schooling) or gender differences in educational attainment.

Our main measure of gender norms is a GENDER ATTITUDES INDEX which we construct by taking the first principal component of the eight World Values Survey (WVS) questions related to gender. In Figure 3, we plot the cross-country relationship between each of these questions and the proportion of a country whose native language is a gender language. The prevalence of gender languages predicts responses to seven of the eight WVS questions. For example, WVS respondents from countries with more gender languages are more likely to agree with the statement “When a mother works for pay, the children suffer” (p-value 1.20×10^{-6}) and “When jobs are scarce, men should have more right to a job than women” (p-value 2.20×10^{-3}). The country-level prevalence of gender languages is also associated with an increased likelihood of believing that men make better business executives and political leaders than women.

In Table 3, we confirm the association between the prevalence of gender languages and our summary index of gender attitudes in a regression framework. After controlling for continent fixed effects and country-level geography, the coefficient estimate suggests that a 10 percentage point increase in the prevalence of gender languages is associated with a decline in support for gender equality that is equivalent to 0.1 standard deviations. Moreover, the bounding procedures suggested by Oster (forthcoming) indicates that this correlation is very unlikely to be driven by omitted variables. At the cross-country level, our analysis suggests that gender languages are causally related to gender inequality.

5 Within-Country Analysis

5.1 Empirical Strategy

Next, we explore the relationship between gender languages and women’s labor force participation at the individual level in a cultural and institutional context where both gender and non-gender languages are indigenous. As shown in Figure 1, there are six countries in Africa where the proportion of the population that speaks a gender language as a native language is between 0.1 and 0.9: Chad, Kenya, Mauritania, Niger, Nigeria, and Uganda. In

these countries, both gender and non-gender languages are indigenous — in contrast to, for example, several countries in South America where non-gender indigenous languages and a colonial language that is a gender language are both widely spoken. Of the six African countries listed above, we focus on the four that have been included in at least one round of the Afrobarometer survey: Kenya, Niger, Nigeria, and Uganda. Four rounds of data are available for Kenya, Nigeria, and Uganda, while only one round of data is available for Niger.¹⁸ Our sample includes 26,547 Afrobarometer respondents who speak 292 different languages.

Our individual-level analysis parallels our cross-country analysis. We consider two main classes of outcomes: labor force participation (an indicator equal to one if a respondent is either does some type of income-generating activity or is actively looking for a job) and education (indicators for having completed primary and secondary school). We report two regression specifications. First, we estimate the association between gender languages and labor force participation in a sample of (only) women, estimating the OLS regression equation:

$$Y_{icr} = \alpha + \beta \text{Gender}_{icr} + \nu_{cr} + \gamma Z_{icr} + \varepsilon_{icr} \quad (2)$$

where Y_{icr} is the outcome of interest for woman i in country c who was interviewed in Afrobarometer Round r , Gender_{icr} is an indicator for speaking a gender language as native language, ν_{cr} is a vector of country-round fixed effects, Z_{icr} is a vector of controls (age, age², and a set of religion dummies), and ε_{icr} is a mean-zero error term. As in our cross-country analysis, we wish to avoid confounding the impact of gender on women’s labor force participation (and education) with other cultural factors that might impact both men’s and women’s likelihood of working. To do this, we also report pooled OLS regressions that

¹⁸The first round of the Afrobarometer surveys did not include sufficiently detailed data on native languages for inclusion in our analysis. Our analysis includes data from Afrobarometer Rounds 2 through 5 for Kenya, Nigeria, and Uganda. Niger was only added to the Afrobarometer in Round 5; that round is included in our analysis.

include data on both men and women. These take the form:

$$Y_{icr} = \alpha + \beta Gender_{icr} + \zeta Female_{icr} + \mu Gender \times Female_{icr} + \nu_{cr} + \gamma Z_{icr} + \varepsilon_{icr} \quad (3)$$

where $Gender \times Female_{icr}$ is an interaction between a female dummy and the indicator for being a native speaker of a gender language. Throughout our analysis, we cluster standard errors by language.

5.2 Labor Force Participation

We report the results of our regressions of individual-level labor force participation on the indicator for being a native speaker of a gender language in Table 4. Here, the sample is restricted to women. We find a robust negative association between speaking a gender language and labor force participation. After controlling for country-round fixed effects, age, and religion, coefficient estimates suggest that women who speak gender languages as their native languages are 18 percentage points less likely to be in the labor force (p-value 1.32×10^{-6}). Using the approach proposed by Oster (forthcoming) to bound the likely unbiased causal impact of grammatical gender on labor force participation suggests that gender languages reduce women’s labor force participation by at least 13 percentage points (or, equivalently, unobservable factors would need to be 2.11 times more related to treatment than observable controls).

Next, we examine differences in women’s labor force participation relative to men. Table 5 reports the results of pooled OLS regressions that include Afrobarometer data on both men and women; the coefficient of interest is now the interaction term FEMALE×GENDER LANGUAGE. Results suggest that speaking a gender native language is associated with lower female labor force participation relative to men from the same ethnolinguistic group. The coefficient on the FEMALE×GENDER LANGUAGE interaction is -0.17 whether or not controls (for age and religion, plus country-round fixed effects) are included (p-values 9.06×10^{-4} and 1.35×10^{-3} , respectively). The Oster (forthcoming) approach suggests that

unobservable covariates would need to be more than one thousand times more closely correlated with treatment than observables to explain the observed effect. Thus, our individual-level estimates suggest that grammatical gender has a negative causal impact on women’s labor force participation.

5.3 Educational Attainment

Next we consider the within-country association between grammatical gender and women’s educational attainment. In our cross-country analysis, we did not find a relationship between the country-level prevalence of gender languages and either women’s educational attainment or gender differences in educational attainment. However, rates of educational attainment are quite high (for both men and women) in many countries, limiting the statistical power of cross-country analysis. Moreover, many countries have compulsory schooling laws in place, and these may attenuate the impacts of both beliefs about labor market returns (which will differ by gender if women are less likely to participate in the labor force) and cultural values in decisions about girls’ enrollment in school.

Average levels of education are still quite low in many African countries. In the 37 African countries included in the Barro-Lee data set, the average level of educational attainment among adult males is only 5.6 years. Moreover, gender differences in educational attainment persist throughout Africa. Among countries in the Barro-Lee data set, the average level of education among women is only 4.3 years, and women obtain less schooling than men in all but five African countries. Primary school has only recently been made free in many African countries, and compulsory schooling laws are still relatively rare.

We estimate the association between speaking a gender language and the likelihood of completing primary school (Table 6, Columns 1 through 3) and secondary school (Table 6, Columns 1 through 3) in Kenya, Niger, Nigeria, and Uganda using the Afrobarometer data described above. Coefficient estimates suggest a very strong relationship between grammatical gender and educational attainment. After controlling for country-round fixed effects and individual characteristics (age and religion), speaking a gender native language

is associated with a 22 percentage point decline in the likelihood that a woman completed primary school and a 16 percentage point decline in the likelihood that a woman completed secondary school (Table 6, Columns 3 and 6, respectively). Both coefficients are negative and significant at the 99 percent confidence level (p-values, 4.18×10^{-5} and 4.66×10^{-4} , respectively.) Following the approach suggested by Oster (forthcoming) suggests that the true impact of grammatical gender on primary school completion is likely to fall between -0.10 and -0.15 , while the causal impact of gender language on secondary school completion is likely to fall between -0.11 and -0.13 .

Again, we are able to rule out the possibility that differences in levels of education among women are driven by cultural factors (across linguistic groups) that impact both men and women. We report pooled OLS specifications that include men and women in Table 7. We do find that cultural factors matter: the indicator for having a gender native language is negative and significant in all specifications. After controlling for age and religion (along with country-round fixed effects), *men* whose native language is a gender language are 9 percentage points less likely to finish primary school and 10 percentage points less likely to finish secondary school than men whose native language is non-gender. Nonetheless, the interaction between FEMALE and the indicator for gender native languages is also negative and significant in all specifications; moreover, after (implicitly) controlling for the level of education observed among men in each language group, the estimated FEMALE \times GENDER LANGUAGE interaction is largely impervious to controls. In a specification with no controls (Table 7, Column 1), the estimated coefficient on the interaction term of interest is -0.12 (p-value 4.24×10^{-19}). When we include controls for country-round fixed effects, age, and religion, the estimated coefficient of interest is still -0.12 (p-value 5.39×10^{-14}). The pattern is similar for secondary school: without controls, the coefficient on FEMALE \times GENDER LANGUAGE is -0.06 (p-value 1.28×10^{-5} ; with controls, the coefficient is still -0.06 (p-value 258×10^{-4}). Because the estimated coefficients show almost no change as controls are added (though the controls more than triple the R^2), the approach suggested by Altonji, Elder, and Taber (2005) and Oster (forthcoming) suggests that the true causal impact of

grammatical gender on the women’s educational attainment is very close to the estimated regression coefficients report in Table 7

6 Conclusion

Using a new data set on the grammatical gender structure of nearly 4,000 languages, we document a robust negative association between gender languages and women’s labor force participation. At the country level, an increase in the proportion of the population whose native language is a gender language is associated with lower female labor force participation and — perhaps more importantly — larger gender differences in labor force participation. Using data from the World Values Survey, we show that grammatical gender also predicts support for traditional gender roles. However, the prevalence of gender languages does not explain cross-country differences in women’s educational attainment.

Focusing on four African countries where both gender and non-gender languages are indigenous, we show that a similar pattern holds within countries. Speaking a gender native language is associated with labor force participation and educational attainment among women, both in absolute terms and relative to men from the same ethnolinguistic group. Both our cross-country and our individual-level regressions are robust to the inclusion of controls that could not plausibly have been impacted by treatment; if one is willing to assume that the relationship between unobserved omitted factors, treatment, and the outcomes of interest is similar to the observed relationship between controls, treatment, and the outcomes of interest, our estimates suggest that grammatical gender has a large negative impact on women’s labor force participation.

Our results are consistent with research in psychology, linguistics, and anthropology suggesting that languages shape patterns of thought in subtle and subconscious ways. Languages are a critical part of our cultural heritage, and it would be inappropriate to suggest that some languages are detrimental to development or women’s rights. However, languages do evolve over time; the direction of their evolution is shaped by both individual choices

(for example, whether to use gendered pronouns like “he” or “she” or gender-neutral alternatives such as “they”) and conscious decisions by government agencies (e.g. the Académie française) and other thought leaders (e.g. major newspapers and magazines). Our results suggest that individuals should reflect upon the social consequences of their linguistic choices, as the nature of the language we speak shapes the ways we think, and the ways our children will think in the future.

References

- ACHARYA, A., M. BLACKWELL, AND M. SEN (2016): “Explaining Causal Findings without Bias: Detecting and Assessing Direct Effects,” *American Political Science Review*, 110(3), 512–529.
- AFROBAROMETER DATA (2016): “Kenya, Niger, Nigeria, Tanzania, Uganda, Rounds 2 through 5, 2002–2013,” available at <http://www.afrobarometer.org>.
- AIKHENVALD, A. Y. (2003): *Classifiers: A Typology of Noun Categorization Devices*. Oxford University Press, Oxford, UK.
- ALESINA, A., P. GIULIANO, AND N. NUNN (2013): “On the Origins of Gender Roles: Women and the Plough,” *Quarterly Journal of Economics*, 128(2), 469–530.
- ALTONJI, J. G., T. E. ELDER, AND C. R. TABER (2005): “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools,” *Journal of Political Economy*, 113(1), 151–184.
- ANGRIST, J. D., AND J.-S. PISCHKE (2008): *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- BARRO, R., AND J.-W. LEE (2013): “A New Data Set of Educational Attainment in the World, 1950–2010,” *Journal of Development Economics*, 104, 184–198.
- BORODITSKY, L., L. A. SCHMIDT, AND W. PHILLIPS (2003): “Sex, Syntax, and Semantics,” in *Language in Mind: Advances in the Study of Language and Thought*, ed. by S. Goldin-Meadow, and D. Gentner, pp. 61–79. MIT Press, Cambridge, MA.
- CAMPBELL, G. L. (1991): *Compendium of the World’s Languages*. Routledge, London, UK.
- CORBETT, G. G. (1991): *Gender*. Cambridge University Press, Cambridge, UK.
- (2013a): “Number of Genders,” in *The World Atlas of Language Structures Online*, ed. by M. S. Dryer, and M. Haspelmath. Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany.
- (2013b): “Sex-based and Non-sex-based Gender Systems,” in *The World Atlas of Language Structures Online*, ed. by M. S. Dryer, and M. Haspelmath. Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany.
- (2013c): “Systems of Gender Assignment,” in *The World Atlas of Language Structures Online*, ed. by M. S. Dryer, and M. Haspelmath. Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany.
- CREISSELS, D. (2005): “Typology,” in *African Languages: An Introduction*, ed. by B. Heine, and D. Nurse, chap. 21, pp. 231–258. Cambridge University Press, Cambridge, UK.
- DANZIGER, S., AND R. WARD (2010): “Language Changes Implicit Associations between Ethnic Groups and Evaluation in Bilinguals,” *Psychological Science*, 21(6), 799–800.
- EASTERLY, W., AND R. LEVINE (1997): “Africa’s Growth Tragedy: Policies and Ethnic Divisions,” *The Quarterly Journal of Economics*, 112(4), 1203–1250.
- GAY, V., D. L. HICKS, E. SANTACREU-VASUT, AND A. SHOHAM (forthcoming): “Decomposing culture: An Analysis of Gender, Language, and Labor Supply in the Household,” *Review of Economics of the Household*.
- GRANT, M. J., AND J. R. BEHRMAN (2010): “Gender Gaps in Educational Attainment in Less Developed Countries,” *Population and Development Review*, 36(1), 71–89.

- GRIERSON, G. A. (1903a): *Linguistic Survey of India: Volume V, Indo-Aryan Family Eastern Group, Part I, Specimens of the Bengali and Assamese Languages*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1903b): *Linguistic Survey of India: Volume V, Indo-Aryan Family Eastern Group, Part II, Specimens of the Bihari and Oriya Languages*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1904): *Linguistic Survey of India: Volume VI, Indo-Aryan Family Mediate Group, Specimens of the Eastern Hindi Language*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1905): *Linguistic Survey of India: Volume VII, Indo-Aryan Family Southern Group, Specimens of the Marathi Language*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1907): *Linguistic Survey of India: Volume IX, Indo-Aryan Family Central Group, Part III, The Bhil Languages*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1908): *Linguistic Survey of India: Volume IX, Indo-Aryan Family Central Group, Part II, Specimens of Rajastani and Gujarati*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1909): *Linguistic Survey of India: Volume III, Tibeto-Burman Family, Part I, General Introduction, Specimens of the Tibetan Dialects, the Himalayan Dialects, and the North Assam Group*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1916): *Linguistic Survey of India: Volume IX, Indo-Aryan Family Central Group, Part I, Specimens of Western Hindi and Panjabi*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1919): *Linguistic Survey of India: Volume VIII, Indo-Aryan Family Northwestern Group, Part II, Specimens of the Dardic or Pisacha Languages*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- (1921): *Linguistic Survey of India: Volume X, Specimens of Languages of the Eranian Family*. Superintendent Government Printing, Calcutta, India, <http://dsal.uchicago.edu/books/lsi/index.html> [accessed 13 July 2016].
- GUIORA, A. Z. (1983): “Language and Concept Formation: A Cross-Lingual Analysis,” *Cross-Cultural Research*, 18(3), 228–256.
- HELLINGER, M., AND H. BUSSMAN (2003): “The Linguistic Representation of Women and Men,” in *Gender Across Languages: The Linguistic Representation of Women and Men, Volume 3*, ed. by M. Hellinger, and H. Bußman, pp. 1–26. John Benjamins Publishing Company, Amsterdam, The Netherlands.
- HICKS, D. L., E. SANTACREU-VASUT, AND A. SHOHAM (2015): “Does Mother Tongue Make for Women’s Work? Linguistics, Household Labor, and Gender Identity,” *Journal of Economic Behavior and Organization*, 110(2), 19–44.
- KILARSKI, M. (2013): *Nominal Classification: A History of Its Study from the Classical Period to the Present*. John Benjamins Publishing Company, Amsterdam, The Netherlands.
- LAKOFF, G. (1987): *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press, Chicago, IL.

- LEWIS, M. P., G. F. SIMONS, AND C. D. FENNIG, EDs., (2016): “Ethnologue: Languages of the World, nineteenth edition.” Dallas, Texas: SIL International, Online version: <http://www.ethnologue.com> [accessed 4 June 2016].
- OGUNNAIKE, O., Y. DUNHAM, AND M. R. BANAJI (2010): “The Language of Implicit Preferences,” *Journal of Experimental Social Psychology*, 46(6), 999–1003.
- OSTER, E. (forthcoming): “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business and Economic Statistics*.
- SANTACREU-VASUT, E., O. SHENKAR, AND A. SHOHAM (2014): “Linguistic Gender Marking and Its International Business Ramifications,” *Journal of International Business Studies*, 45, 1170–1178.
- SANTACREU-VASUT, E., A. SHOHAM, AND V. GAY (2013): “Do female/male distinctions in language matter? Evidence from gender political quotas,” *Applied Economics Letters*, 20(5), 495–498.
- UCLA LANGUAGE MATERIALS PROJECT (2014): “Teaching Resources for Less Commonly Taught Languages,” <http://www.lmp.ucla.edu/> [accessed 4 June 2016].
- WHORF, B. L. (2011[1956]a): “Language, Mind, and Reality,” in *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*, ed. by J. B. Carroll, pp. 246–270. Martino Publishing, Mansfield Centre, CT.
- (2011[1956]b): “A Linguistic Consideration of Thinking in Primitive Communities,” in *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*, ed. by J. B. Carroll, pp. 65–86. Martino Publishing, Mansfield Centre, CT.
- WILSON, P. M. (1985): *Simplified Swahili*. Longman, Harlow, UK.
- WORLD VALUES SURVEY ASSOCIATION (2015): “World Values Survey Wave 6 2010–2014 Official Aggregate v.20150418,” available online at www.worldvaluessurvey.org, aggregate file producer: Asep/JDS, Madrid SPAIN.

Figure 1: The Distribution of Gender Languages

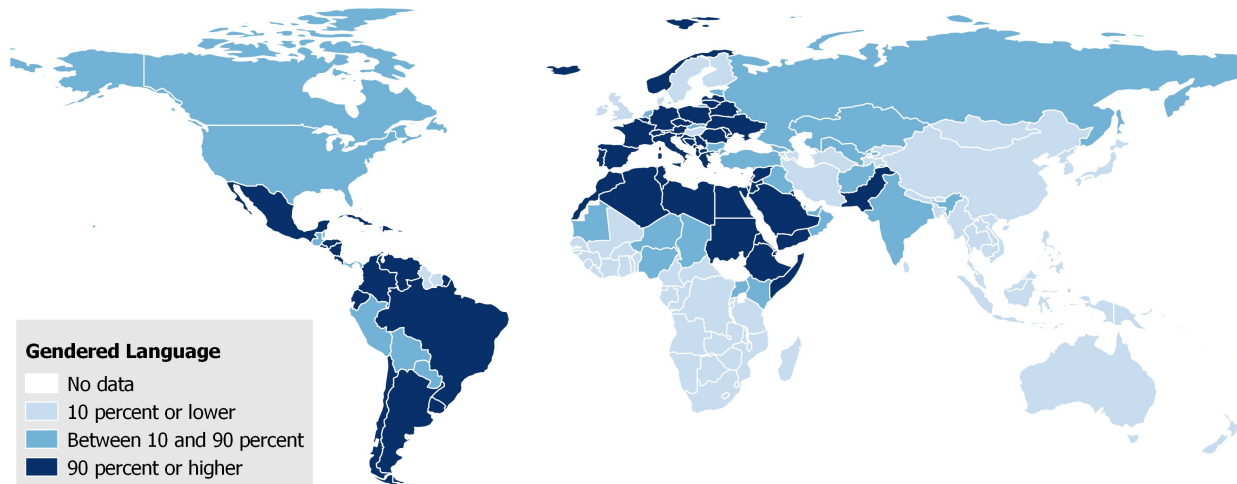
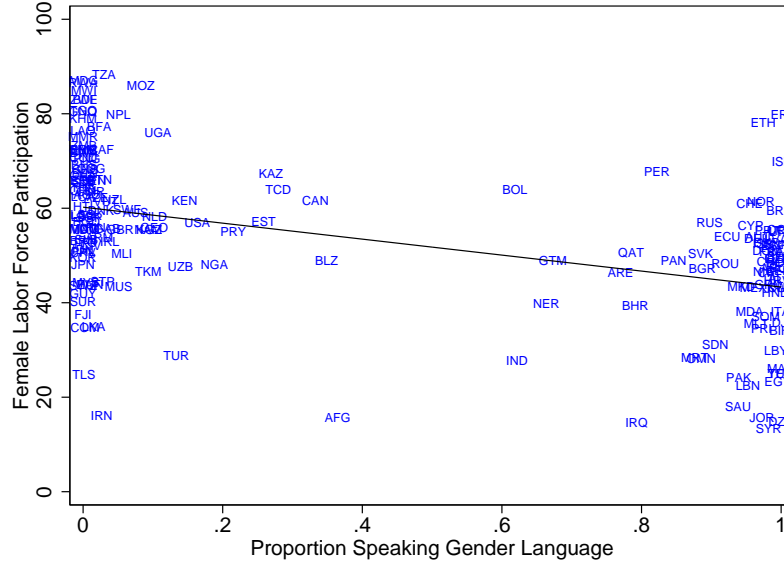


Figure 2: Women's Labor Force Participation and Proportion Speaking Gender Languages

Panel A: Scatter Plot of Women's LFP (Level) and Proportion Speaking Gender Languages



Panel B: Scatter Plot of Gender Difference in LFP and Proportion Speaking Gender Languages





Figure 3: Cross-Country Variation in Gender Attitudes



The figure summarizes the results from a series of regressions of (country-level averages of) responses to World Values Survey (WVS) questions on the proportion of a country's population whose native language is a gender language. We present the results for all eight WVS questions related to gender attitudes. Responses to all eight questions are coded so that the answer most consistent with traditional gender norms (involving separate roles for men and women) is equal to 1 and the response most consistent with gender equality is equal to 0. Each regression is estimated via OLS and includes continent fixed effects. The outcome in the first row is the average response to the question "When a mother works for pay, the children suffer" (agreement is coded as a 1, disagreement as a 0). The outcome variable in the second row is the average response to the statement "When jobs are scarce, men should have more right to a job than women." In the third row, the outcome variable is based on the statement "On the whole, men make better political leaders than women do." In the fourth row, the outcome variable is based on the statement "On the whole, men make better business executives than women do." In the fifth row, the outcome variable is based on the statement "Being a housewife is just as fulfilling as working for pay;" agreement was coded as 0 and disagreement was coded as 1. In the sixth row, the outcome variable is based on the statement "If a woman earns more money than her husband, it's almost certain to cause problems." In the seventh row, the outcome variable is based on the statement "A university education is more important for a boy than for a girl." In the last row, the outcome variable is based on the statement "Having a job is the best way for a woman to be an independent person;" in this case, disagreement was coded as 1 and agreement was coded as 0.

Table 1: Cross-Country OLS Regressions of Labor Force Participation

<i>Dependent variable:</i>	FEMALE			GENDER DIFFERENCE IN		
	LABOR FORCE PARTICIPATION			LABOR FORCE PARTICIPATION		
<i>Specification:</i>	OLS	OLS	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Proportion speaking gender language	-16.91*** (2.78)	-19.82*** (3.44)	-13.99*** (3.24)	-12.17*** (2.47)	-18.87*** (3.15)	-15.34*** (3.37)
Constant	60.23*** (1.47)	64.09*** (2.22)	64.83*** (5.98)	-16.63*** (1.20)	-14.37*** (1.88)	-5.58 (4.83)
Continent Fixed Effects	No	Yes	Yes	No	Yes	Yes
Geography Controls	No	No	Yes	No	No	Yes
Observations	174	174	174	174	174	174
R^2	0.22	0.32	0.39	0.13	0.39	0.44

Robust standard errors clustered by most widely spoken language in all specifications. FEMALE LABOR FORCE PARTICIPATION is the percentage of women in the labor force, measured in 2011. The GENDER DIFFERENCE IN LABOR FORCE PARTICIPATION is the difference between male and female labor force participation. Additional controls are log GDP per capita (PPP-adjusted 2011 US dollars); current population; and the percentage of land area in the tropics or subtropics. GDP data is missing from the World Development Indicators database for Argentina, Myanmar, North Korea, Somalia, and Syria.

Table 2: Cross-Country OLS Regressions of Educational Attainment

<i>Dependent variable:</i>	FEMALE			GENDER DIFFERENCE IN		
	EDUCATIONAL ATTAINMENT			EDUCATIONAL ATTAINMENT		
<i>Specification:</i>	OLS	OLS	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Proportion speaking gender language	1.84** (0.77)	-0.58 (0.6)	-0.8 (0.6)	0.25 (0.26)	-0.2 (0.27)	-0.32 (0.28)
Constant	6.83*** (0.68)	4.45*** (0.41)	9.56*** (1.05)	-0.83*** (0.19)	-1.23*** (0.21)	-0.59 (0.48)
Continent Fixed Effects	No	Yes	Yes	No	Yes	Yes
Geography Controls	No	No	Yes	No	No	Yes
Observations	140	140	140	140	140	140
R^2	0.06	0.51	0.62	0.01	0.16	0.18

Robust standard errors clustered by most widely spoken language in all specifications. FEMALE LABOR FORCE PARTICIPATION is the percentage of women in the labor force, measured in 2011. The GENDER DIFFERENCE IN LABOR FORCE PARTICIPATION is the difference between male and female labor force participation. Additional controls are log GDP per capita (PPP-adjusted 2011 US dollars); current population; and the percentage of land area in the tropics or subtropics. GDP data is missing from the World Development Indicators database for Argentina, Myanmar, North Korea, Somalia, and Syria.

Table 3: Cross-Country OLS Regressions of Gender Attitudes

<i>Dependent variable:</i>	GENDER ATTITUDES INDEX		
<i>Specification:</i>	OLS	OLS	OLS
	(1)	(2)	(3)
Proportion speaking gender language	-0.03 (0.05)	-0.11*** (0.03)	-0.12*** (0.04)
Constant	0.54*** (0.03)	0.49*** (0.02)	0.52*** (0.04)
Continent Fixed Effects	No	Yes	Yes
Geography Controls	No	No	Yes
Observations	56	56	56
R^2	0.01	0.74	0.78

Robust standard errors clustered by most widely spoken language in all specifications. The GENDER ATTITUDES INDEX is constructed by taking the first principal component of the 8 World Values Survey questions relating to gender norms (described in Figure 3) at the individual level, and then calculating the average of this index within a country. Numbers closer to 1 indicate more support for gender equality. Additional controls are log GDP per capita (PPP-adjusted 2011 US dollars); current population; and the percentage of land area in the tropics or subtropics. GDP data is missing from the World Development Indicators database for Argentina, Myanmar, North Korea, Somalia, and Syria.

Table 4: Individual-Level Regressions of Women’s Labor Force Participation

<i>Dependent variable:</i>	IN LABOR FORCE		
<i>Specification:</i>	OLS	OLS	OLS
	(1)	(2)	(3)
Native language is gender language	-0.24*** (0.05)	-0.2*** (0.04)	-0.18*** (0.04)
Constant	0.67*** (0.02)	0.58*** (0.02)	0.27*** (0.09)
Country-Wave Fixed Effects	No	Yes	Yes
Individual Controls	No	No	Yes
Observations	13137	13137	13137
R^2	0.04	0.07	0.1

Robust standard errors clustered at the language level. Data is from Afrobarometer Rounds 2 through 5. The analysis includes data from Kenya, Niger, Nigeria, Uganda, and Tanzania; Niger was only added to the Afrobarometer in Round 5, while the other countries appear in all four rounds. Individual controls are age and age squared; indicators for being identifying as Muslim, Catholic, Protestant, or another religion; and indicators for having completed primary and secondary school.

Table 5: Individual-Level Regressions of Gender Differences in Labor Force Participation

<i>Dependent variable:</i>	IN LABOR FORCE		
<i>Specification:</i>	OLS	OLS	OLS
	(1)	(2)	(3)
Female \times gender language	-0.17*** (0.05)	-0.16*** (0.05)	-0.17*** (0.05)
Native language is gender language	-0.08*** (0.02)	-0.04* (0.02)	-0.04 (0.03)
Female	-0.1*** (0.01)	-0.1*** (0.01)	-0.11*** (0.01)
Constant	0.77*** (0.01)	0.7*** (0.02)	0.33*** (0.08)
Country-Wave Fixed Effects	No	Yes	Yes
Individual Controls	No	No	Yes
Observations	26287	26287	26287
R^2	0.04	0.07	0.11

Robust standard errors clustered at the language level. Data is from Afrobarometer Rounds 2 through 5. The analysis includes data from Kenya, Niger, Nigeria, Uganda, and Tanzania; Niger was only added to the Afrobarometer in Round 5, while the other countries appear in all four rounds. Individual controls are age and age squared; indicators for being identifying as Muslim, Catholic, Protestant, or another religion; and indicators for having completed primary and secondary school.

Table 6: Individual-Level OLS Regressions of Women’s Educational Attainment

<i>Dependent variable:</i>	PRIMARY SCHOOL			SECONDARY SCHOOL		
<i>Specification:</i>	OLS (1)	OLS (2)	OLS (3)	OLS (4)	OLS (5)	OLS (6)
Native language is gender language	-0.31*** (0.04)	-0.3*** (0.06)	-0.22*** (0.05)	-0.19*** (0.04)	-0.23*** (0.06)	-0.16*** (0.04)
Constant	0.7*** (0.03)	0.67*** (0.03)	0.93*** (0.04)	0.35*** (0.04)	0.33*** (0.03)	0.49*** (0.05)
Country-Wave Fixed Effects	No	Yes	Yes	No	Yes	Yes
Individual Controls	No	No	Yes	No	No	Yes
Observations	13125	13125	13125	13125	13125	13125
R^2	0.06	0.12	0.21	0.02	0.1	0.15

Robust standard errors clustered by language in all specifications. Data is from Afrobarometer Rounds 2 through 5. The analysis includes data from Kenya, Niger, Nigeria, Uganda, and Tanzania; Niger was only added to the Afrobarometer in Round 5, while the other countries appear in all four rounds. Individual controls are age, age squared and indicators for being identifying as Muslim, Catholic, Protestant, or another religion.

Table 7: Individual-Level OLS Regressions of Gender Differences in Educational Attainment

<i>Dependent variable:</i>	PRIMARY SCHOOL			SECONDARY SCHOOL		
<i>Specification:</i>	OLS (1)	OLS (2)	OLS (3)	OLS (4)	OLS (5)	OLS (6)
Female \times gender language	-0.12*** (0.01)	-0.11*** (0.01)	-0.12*** (0.01)	-0.06*** (0.01)	-0.06*** (0.01)	-0.06*** (0.02)
Female	-0.08*** (0.009)	-0.08*** (0.009)	-0.12*** (0.01)	-0.08*** (0.009)	-0.08*** (0.008)	-0.1*** (0.009)
Native language is gender language	-0.19*** (0.04)	-0.17*** (0.05)	-0.09** (0.04)	-0.13*** (0.04)	-0.17*** (0.05)	-0.1*** (0.03)
Constant	0.78*** (0.02)	0.76*** (0.03)	0.98*** (0.04)	0.43*** (0.04)	0.41*** (0.03)	0.5*** (0.05)
Country-Wave Fixed Effects	No	Yes	Yes	No	Yes	Yes
Individual Controls	No	No	Yes	No	No	Yes
Observations	26253	26253	26253	26253	26253	26253
R^2	0.06	0.12	0.21	0.03	0.11	0.15

Robust standard errors clustered by language in all specifications. Data is from Afrobarometer Rounds 2 through 5. The analysis includes data from Kenya, Niger, Nigeria, Uganda, and Tanzania; Niger was only added to the Afrobarometer in Round 5, while the other countries appear in all four rounds. Individual controls are age, age squared and indicators for being identifying as Muslim, Catholic, Protestant, or another religion.

Table 8: Coefficient Stability

	$\hat{\beta}$	$\tilde{\beta}$	$\tilde{\beta}^*(R_{max}, 1)$	δ
<i>Panel A. Cross-Country Regressions</i>				
Female labor force participation	-16.91	-13.99	-5.99	1.19
Gender difference in labor force participation	-12.17	-15.34	-19.26	2.60
Gender attitude index	-0.03	-0.12	-0.21	$\delta < 0$
<i>Panel B. Individual-Level Regressions</i>				
In labor force	-0.24	-0.18	-0.13	2.11
Female \times in labor force	-0.17	-0.17	-0.17	1719.14
Completed primary school	-0.25	-0.15	-0.10	1.98
Female \times completed primary school	-0.12	-0.12	-0.12	$\delta < 0$
Completed secondary school	-0.16	-0.13	-0.11	3.48
Female \times completed secondary school	-0.06	-0.06	-0.06	$\delta < 0$