Gendered Language

Pamela Jakiela
CGD, UMD, BREAD, & IZA

Owen Ozier
World Bank, BREAD, & IZA

October 2018

Motivation

Language structures thought:

> *"Languages differ not only in how they build their sentences but also in how they break down nature to secure elements to put in those sentences."*
>
> – Benjamin Lee Whorf (1941)

Motivation

Language structures thought:

> *"Languages differ not only in how they build their sentences but also in how they break down nature to secure elements to put in those sentences."*
>
> – Benjamin Lee Whorf (1941)

Sapir-Whorf Hypothesis: linguistic determinism

- Our native language limits the scope of our thinking
- Example [now debunked]: the Inuit have 7 different words for snow

Motivation

Language structures thought:

> *"Languages differ not only in how they build their sentences but also in how they break down nature to secure elements to put in those sentences."*
>
> – Benjamin Lee Whorf (1941)

Sapir-Whorf Hypothesis: linguistic determinism

- Our native language limits the scope of our thinking
- Example [now debunked]: the Inuit have 7 different words for snow

Nonetheless, there is mounting evidence that the languages we speak influence our thoughts and actions in subtle, subconscious ways

Motivation: Language Structures Thought

Example: agentive language impacts perceptions of responsibility

Agentive language is the norm in English:

> *"Non-agentive language sounds evasive in English, the province of guilt-shirking children and politicians."*

Motivation: Language Structures Thought

Example: agentive language impacts perceptions of responsibility

Agentive language is the norm in English:

> *"Non-agentive language sounds evasive in English, the province of guilt-shirking children and politicians."*

Agentive language less common in other languages (e.g. Spanish)

Motivation: Language Structures Thought

Fausey and Boroditsky (2011) show Spanish, English monolinguals videos depicting "**intentional**" and "**accidental**" versions of the same event

Action	**Intentional**	**Accidental**
Knocks box	Faces table, knocks box off table	Knocks box off table while gesturing
Breaks pencil	Sits at table, breaks pencil in half	Breaks pencil in half while writing

Motivation: Language Structures Thought

Fausey and Boroditsky (2011) show Spanish, English monolinguals videos depicting "**intentional**" and "**accidental**" versions of the same event

Action	**Intentional**	**Accidental**
Knocks box	Faces table, knocks box off table	Knocks box off table while gesturing
Breaks pencil	Sits at table, breaks pencil in half	Breaks pencil in half while writing

Experiment 1: subjects describe what happened

- Spanish-speakers less likely to use agentive language to describe accidental events, no differences observed for intentional acts

Motivation: Language Structures Thought

Fausey and Boroditsky (2011) show Spanish, English monolinguals videos depicting "**intentional**" and "**accidental**" versions of the same event

Action	Intentional	Accidental
Knocks box	Faces table, knocks box off table	Knocks box off table while gesturing
Breaks pencil	Sits at table, breaks pencil in half	Breaks pencil in half while writing

Experiment 1: subjects describe what happened

- Spanish-speakers less likely to use agentive language to describe accidental events, no differences observed for intentional acts

Experiment 2: (other) subjects try to remember who did what

- Both groups equally likely to remember intentional actors; Spanish-speakers less likely to remember who caused accidents

Motivation: Language Structures Thought

"Languages are far from impartial 'containers' for the packaging of underlying thoughts, but rather are active players in the construction of those thoughts."

– Ogunnaike *et al*. (2010)

Language shapes our thoughts and actions in subtle ways:

- Using agentive language makes actors, responsibility more salient (Fausey *et al*. 2010, Fausey and Boroditsky 2011)
- Language of instructions, stimuli impacts implicit prejudices of bilinguals (Danzinger and Ward 2009, Ogunnaike *et al*. 2010)
- Speakers of languages that treat the future as a separate tense save less than those that treat the future like the present (Chen 2013)

Motivation: Gender Norms

Percent agreeing: **when jobs are scarce, men should have more of a right to a job than women**

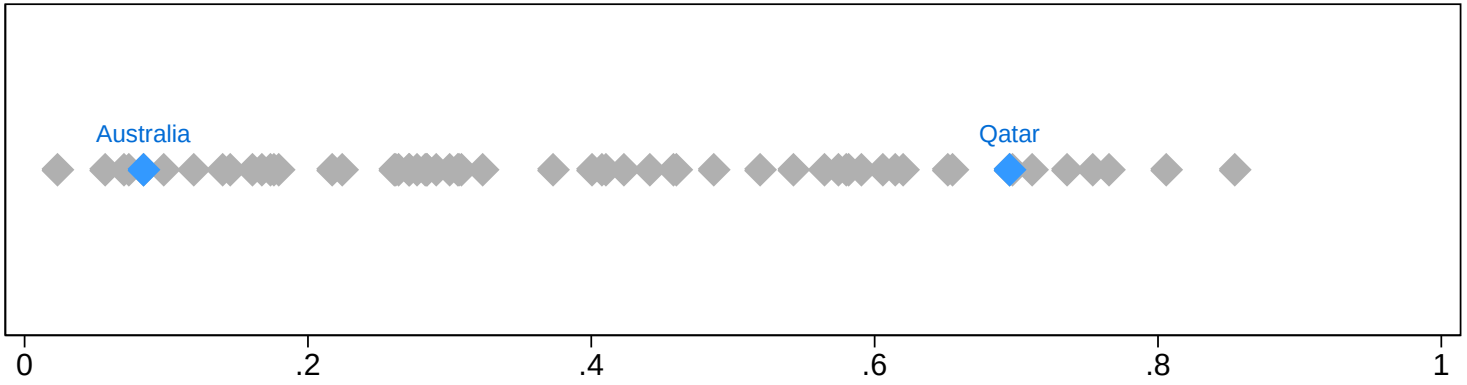

Percent agreeing: **when a woman works, the children suffer**

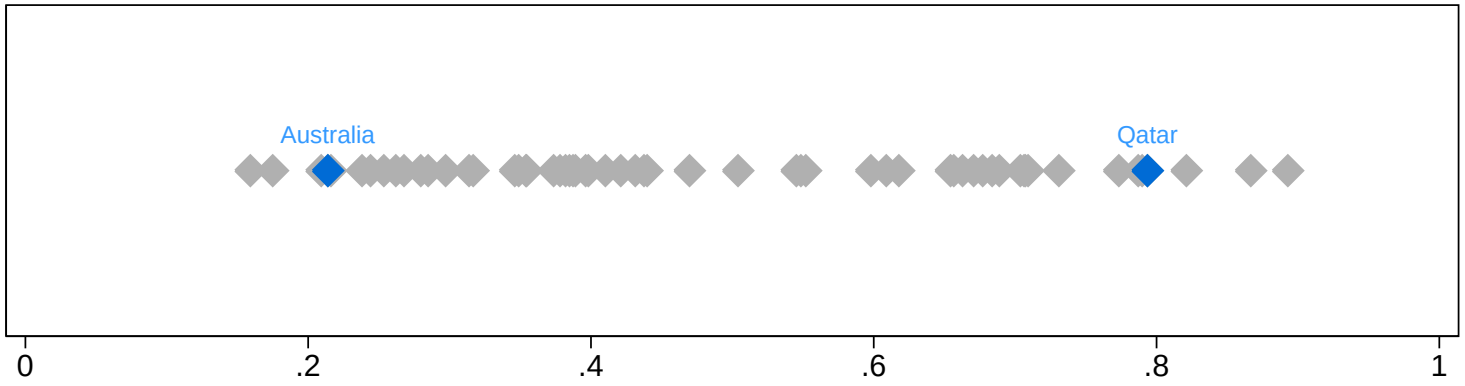

Motivation: Gender and Language

Linguistic gender distinctions:

- Pronominal distinctions between men and women
- Nominal classification systems (grammatical gender)

Motivation: Gender and Language

Linguistic gender distinctions:

- Pronominal distinctions between men and women
- Nominal classification systems (grammatical gender)

Do linguistic gender distinctions impact gender norms?

Grammatical gender creates "a habitual consciousness of two sex classes as a standard classifacatory fact in our thought-world."

– Benjamin Lee Whorf (1936)

Builds on arguments advanced by Durkheim and Mauss (1903)

Motivation: Gender and Language

Suggestive evidence of a link between grammar and gender norms:

- Givati and Troiano (2012) show that countries with gendered pronouns have shorter government-mandated maternity leaves
- Perez and Tavits (2018) show that grammatical gender impacts gender attitudes among Estonian/Russian bilinguals
- Santacreu-Vasut *et al.* (2013) and Shoham and Lee (2017) use World Atlas of Language Structures to estimate cross-country relationship between grammatical gender, gender-related outcomes
- Hicks *et al.* (2015) use WALS data to look at US immigrants

Our Contribution

1. Characterize the grammatical gender structure of 4,334 languages which together account for 99 percent of the world's population
   - India: 6 languages coded in WALS, we code 281
   - Kenya: 3 languages coded in WALS, we code (all) 51
2. Estimate the proportion of each country's population whose native language uses a grammatical gender system to classify nouns
   - Estimate the cross-country relationship between grammatical gender and women's labor force participation and educational attainment, and the relationship with gender attitudes among men and women
3. Use individual-level data from countries where both gender and non-gender languages are indigenous and widely spoken
   - Replicate cross-country results within countries

Outline of the Talk

1. What is grammatical gender?
2. Identifying gender languages
3. Cross-country analysis
   - 3.1 Labor force participation
   - 3.2 Educational attainment
   - 3.3 Gender attitudes
4. Within-country analysis
   - 4.1 Labor force participation
   - 4.2 Educational attainment
5. Discussion, policy implications, and conclusion

Non-Preview of Main Results

Grammatical Gender

Motivation: Gender and Language

Languages differ in their treatment of gender:

- Pronominal distinctions between men and women
- Nominal classification systems (grammatical gender)

Motivation: Gender and Language

Languages differ in their treatment of gender:

- Pronominal distinctions between men and women
- Nominal classification systems (grammatical gender)

Example: Swahili does not make pronominal gender distinctions

$$\left.\begin{array}{r}\text{she goes to school}\\ \text{he goes to school}\end{array}\right\}\ \text{[yeye] anaenda shuleni}$$

Motivation: Gender and Language

Languages differ in their treatment of gender:

- Pronominal distinctions between men and women
- Nominal classification systems (grammatical gender)

Example: Swahili does not make pronominal gender distinctions

$$\left.\begin{array}{r}\text{she goes to school}\\ \text{he goes to school}\end{array}\right\}\ \text{[yeye] anaenda shuleni}$$

There are different words for males and females (e.g. "boy" vs. "girl"), but genders are treated identically from a grammatical perspective

Motivation: Gender and Language

Example: Spanish uses different pronouns for males and females

she goes to school	=	ella va a la escuela
he goes to school	=	él va a la escuela

Motivation: Gender and Language

Example: Spanish uses different pronouns for males and females

she goes to school	=	ella va a la escuela
he goes to school	=	él va a la escuela

Spanish uses a system of **grammatical gender** to classify nouns

- All Spanish nouns are either masculine or feminine
- Grammatical gender determines agreement (e.g. with adjectives)

Nominal classification

Most languages have a system for categorizing nouns (Aikhenvald 2003)

- Many languages partition nouns into **noun classes** or genders

Nominal classification

Most languages have a system for categorizing nouns (Aikhenvald 2003)

- Many languages partition nouns into **noun classes** or genders

Elements of a noun class often often share morphological properties:

- **Spanish:**
  - Masculine words end in O
  - Feminine words end in A
- **Swahili:** class prefixes are used as class names
  - small items belong in ki-/vi- class, humans in m-/wa- class

Nominal classification

Most languages have a system for categorizing nouns (Aikhenvald 2003)

- Many languages partition nouns into **noun classes** or genders

Elements of a noun class often often share morphological properties:

- **Spanish:**
  - Masculine words end in O
  - Feminine words end in A
- **Swahili:** class prefixes are used as class names
  - small items belong in ki-/vi- class, humans in m-/wa- class

Typical noun class system = semantic core + many exceptions

Nominal classification

Noun classes are defined by **agreement** — eg. nouns with adjectives

Example: Swahili has nine distinct noun classes, each characterized by a set of prefixes for verbs, adjectives, demonstratives, possessives, etc.

		[noun]	**new**	**these**
these new chairs	=	viti	vipya	hivi
these new teachers	=	walimu	wapya	hawa

Example: agreement depends on gender (masc./fem.) in Spanish

		the	**[noun]**	**white**
the white shirt	=	la	camisa	blanca
the white hat	=	el	sombrero	blanco

Grammatical Gender

A **grammatical gender** system is a system of noun classification that:

- Includes masculine and feminine as two of the classes
- Characterizes (some) inanimate objects as masculine or feminine
  - English is not a gender language* (though it uses gender pronouns)

Grammatical Gender

A **grammatical gender** system is a system of noun classification that:

- Includes masculine and feminine as two of the classes
- Characterizes (some) inanimate objects as masculine or feminine
  - English is not a gender language* (though it uses gender pronouns)

Languages that use grammatical gender — a.k.a. **gender languages** — differ in grammatical gender intensity along several dimensions

- Do the masculine and feminine classes partition the noun space?
  - Many languages have a neuter class (eg. German, Russian)
- How many parts of speech must change to reflect agreement?
  - Example: verbs agree with gender in Russian, but not in Spanish

Does Grammatical Gender Matter?

Conventional wisdom is that grammatical gender is arbitrary:

> *"In German, a young lady has no sex, while a turnip has."*
>
> – Mark Twain

Does Grammatical Gender Matter?

Conventional wisdom is that grammatical gender is arbitrary:

> *"In German, a young lady has no sex, while a turnip has."*
>
> – Mark Twain

Some linguists have questioned this assumption (cf. Lakoff 1987), arguing that gender categories have a certain... cultural intelligibility

- In Dyirbal, women are grouped with fire and "dangerous things"
- In Ket, one linguist suggested that certain small mammals are feminine "because they are of no importance to the Kets"
- Assignment of inanimate objects to grammatical gender categories often reflects stereotypes about male vs. female body types

Does Grammatical Gender Matter?

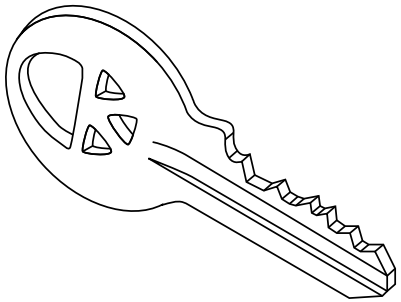

Does Grammatical Gender Matter?

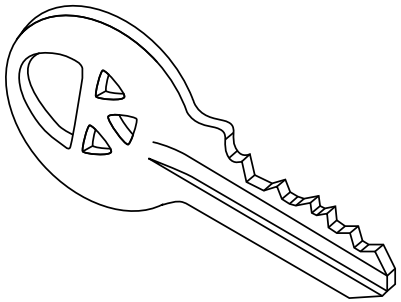

Native **German** speakers said:

hard
heavy
jagged
metal
serrated

Native **Spanish** speakers said:

golden
intricate
little
lovely
shiny

Does Grammatical Gender Matter?

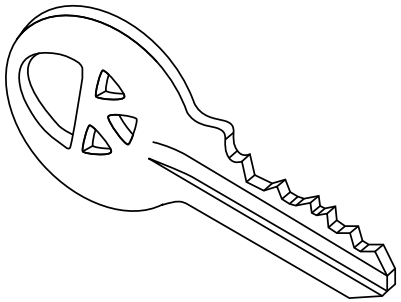

Native **German** speakers said:	Native **Spanish** speakers said:
hard	golden
heavy	intricate
jagged	little
metal	lovely
serrated	shiny
der Schlüssel (masculine)	**la llave** (feminine)

Source: Boroditsky *et al.* (2002)

Does Grammatical Gender Matter?

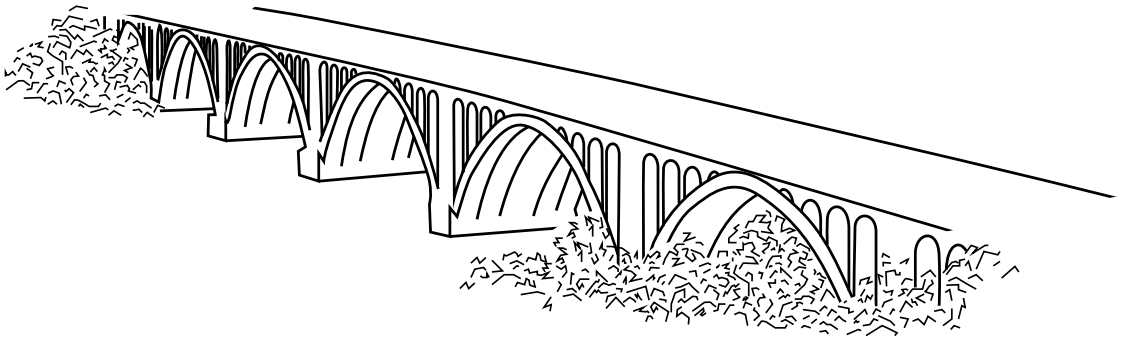

Native **German** speakers said:	Native **Spanish** speakers said:
beautiful	big
elegant	dangerous
fragile	long
peaceful	strong
pretty	sturdy

Does Grammatical Gender Matter?

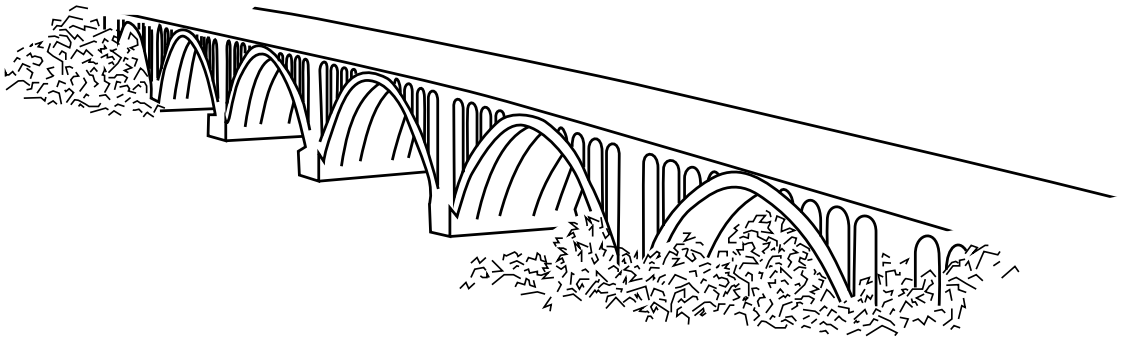

Native **German** speakers said:	Native **Spanish** speakers said:
beautiful	big
elegant	dangerous
fragile	long
peaceful	strong
pretty	sturdy
die Brücke (feminine)	**el puente** (masculine)

Source: Boroditsky *et al.* (2002)

Does Grammatical Gender Matter?

Evidence that grammatical gender matters:

- Santacreu-Vasut *et al.* (2013): political quotas for women are more common in countries where the national language is non-gender
- Hicks *et al.* (2015): immigrants are more likely to divide household tasks along gender lines if they grew up speaking a gender language
- Perez and Tavits (2018): Estonian/Russian bilinguals show greater support for gender equality when (randomly) interviewed in Estonian

Does Grammatical Gender Matter?

Evidence that grammatical gender matters:

- Santacreu-Vasut *et al.* (2013): political quotas for women are more common in countries where the national language is non-gender
- Hicks *et al.* (2015): immigrants are more likely to divide household tasks along gender lines if they grew up speaking a gender language
- Perez and Tavits (2018): Estonian/Russian bilinguals show greater support for gender equality when (randomly) interviewed in Estonian

Studies suggest grammatical gender associated with women's exclusion from public life, labor market, etc; specialization in domestic sphere

- Existing work hampered by data limitations

Identifying Gender Languages

The World's Languages

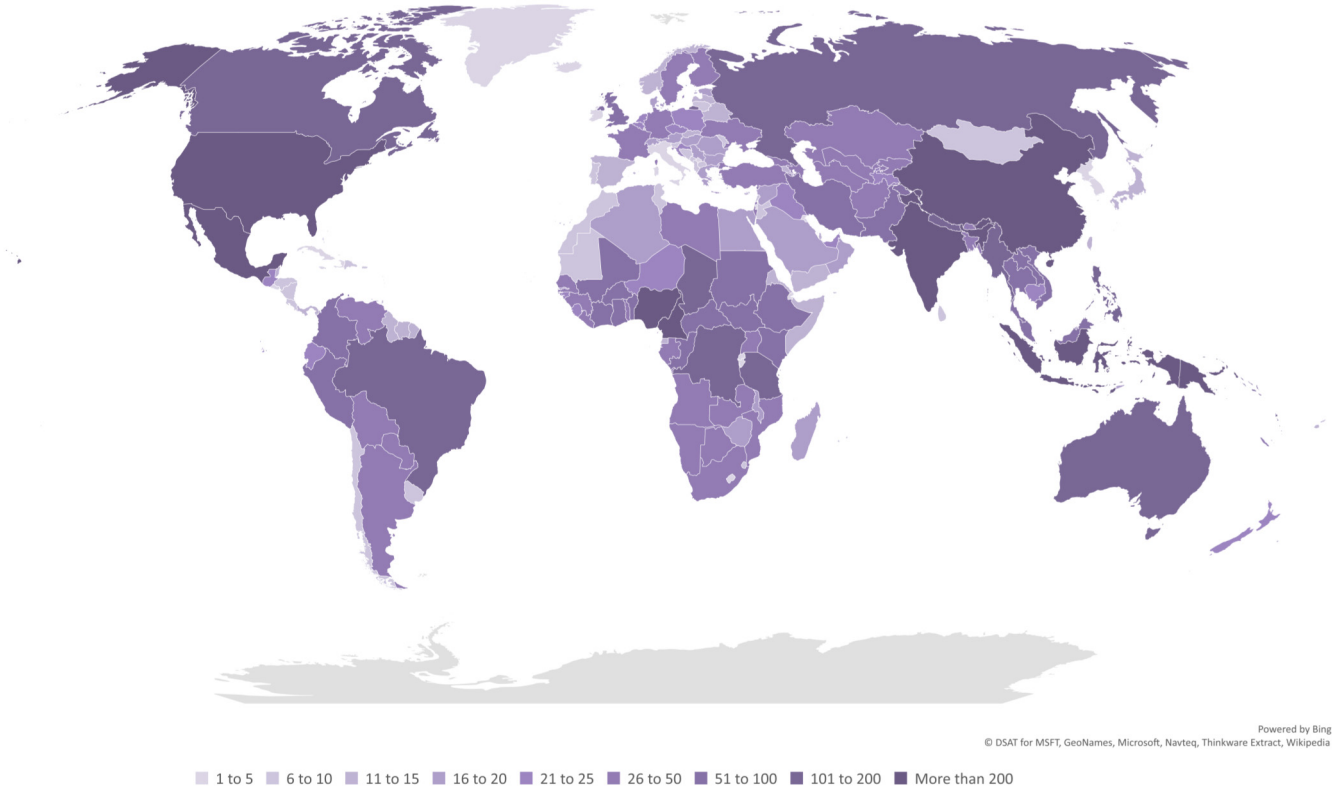

The **Ethnologue** is the most comprehensive database of languages

- Includes over 7,000; 6,190 of them living oral native languages

The World's Languages

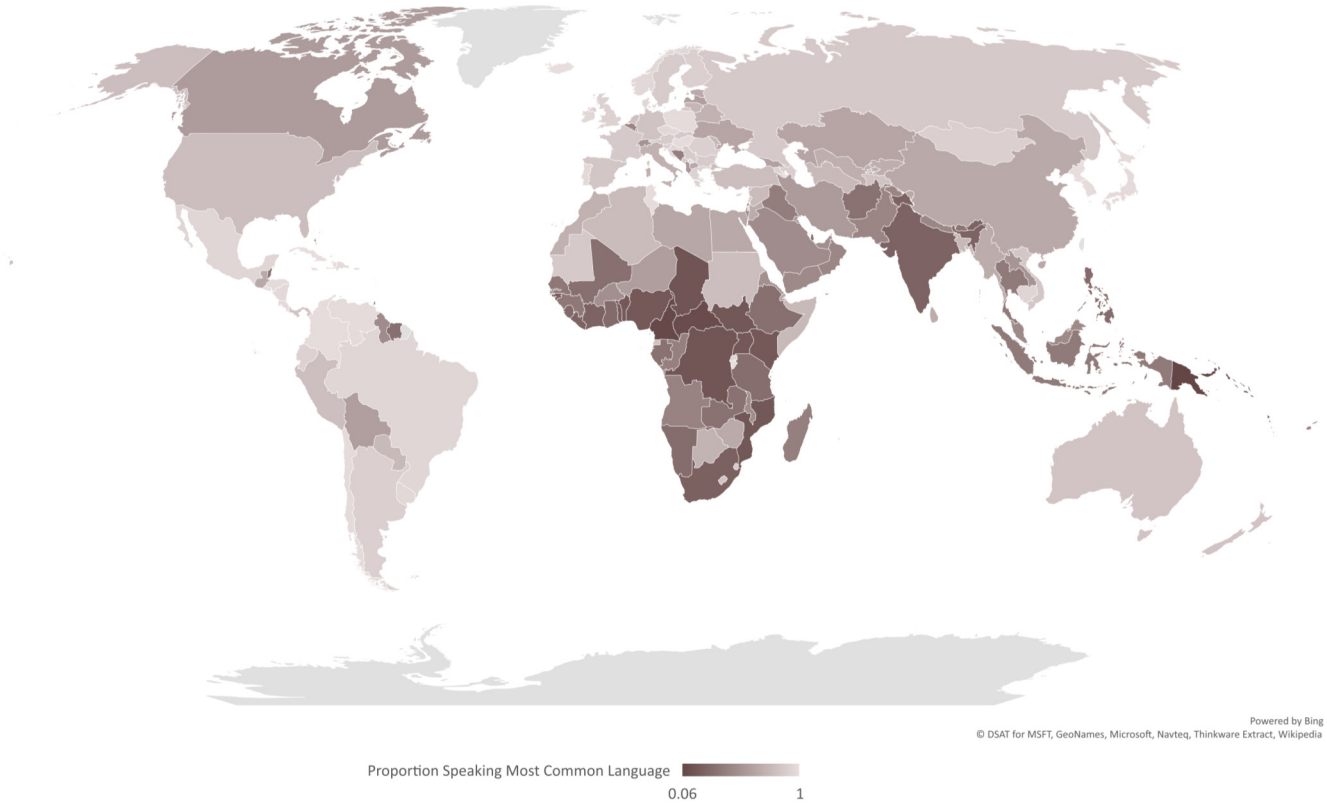

In many (LIC/LMIC) countries, the most widely spoken native language accounts for a small fraction of the population (e.g. 0.18 in Nigeria)

Classifying Gender Structures

We compile data on grammatical structures from a range of sources:

- World Atlas of Language Structures


WHORF
LANGUAGE, THOUGHT, & REALITY
Corbett Gender
Cambridge
Aikhenvald Classifiers
OXFORD
SiHoLS 121
Marcin Kilarski
Nominal Classification
IMPACT 9
Hellinger/Bußmann Gender Across Languages
IMPACT 11
Hellinger/Bußmann Gender Across Languages
Heine and Nurse
A Linguistic Geography of Africa
McWhorter
Defining Creole
Grammatical Gender - Ibrahim
Cinque
Kayne
THE OXFORD HANDBOOK OF
COMPARATIVE SYNTAX
COMPENDIUM OF THE
WORLD'S LANGUAGES
SECOND EDITION
GEORGE L. CAMPBELL
VOLUME I:
ABAZA TO KURDISH
Trends in Linguistics – S. a. M. 124
Barbara Unterbeck and Matti Rissanen
(Editors)
Gender in Grammar and Cognition
THE
EASTERN
SUDANIC
LANGUAGES
TUCKER
VOLUME I
M. Blackings
N. Fabb
Ma'di
Mouton
de Gruyter
DAWSONS

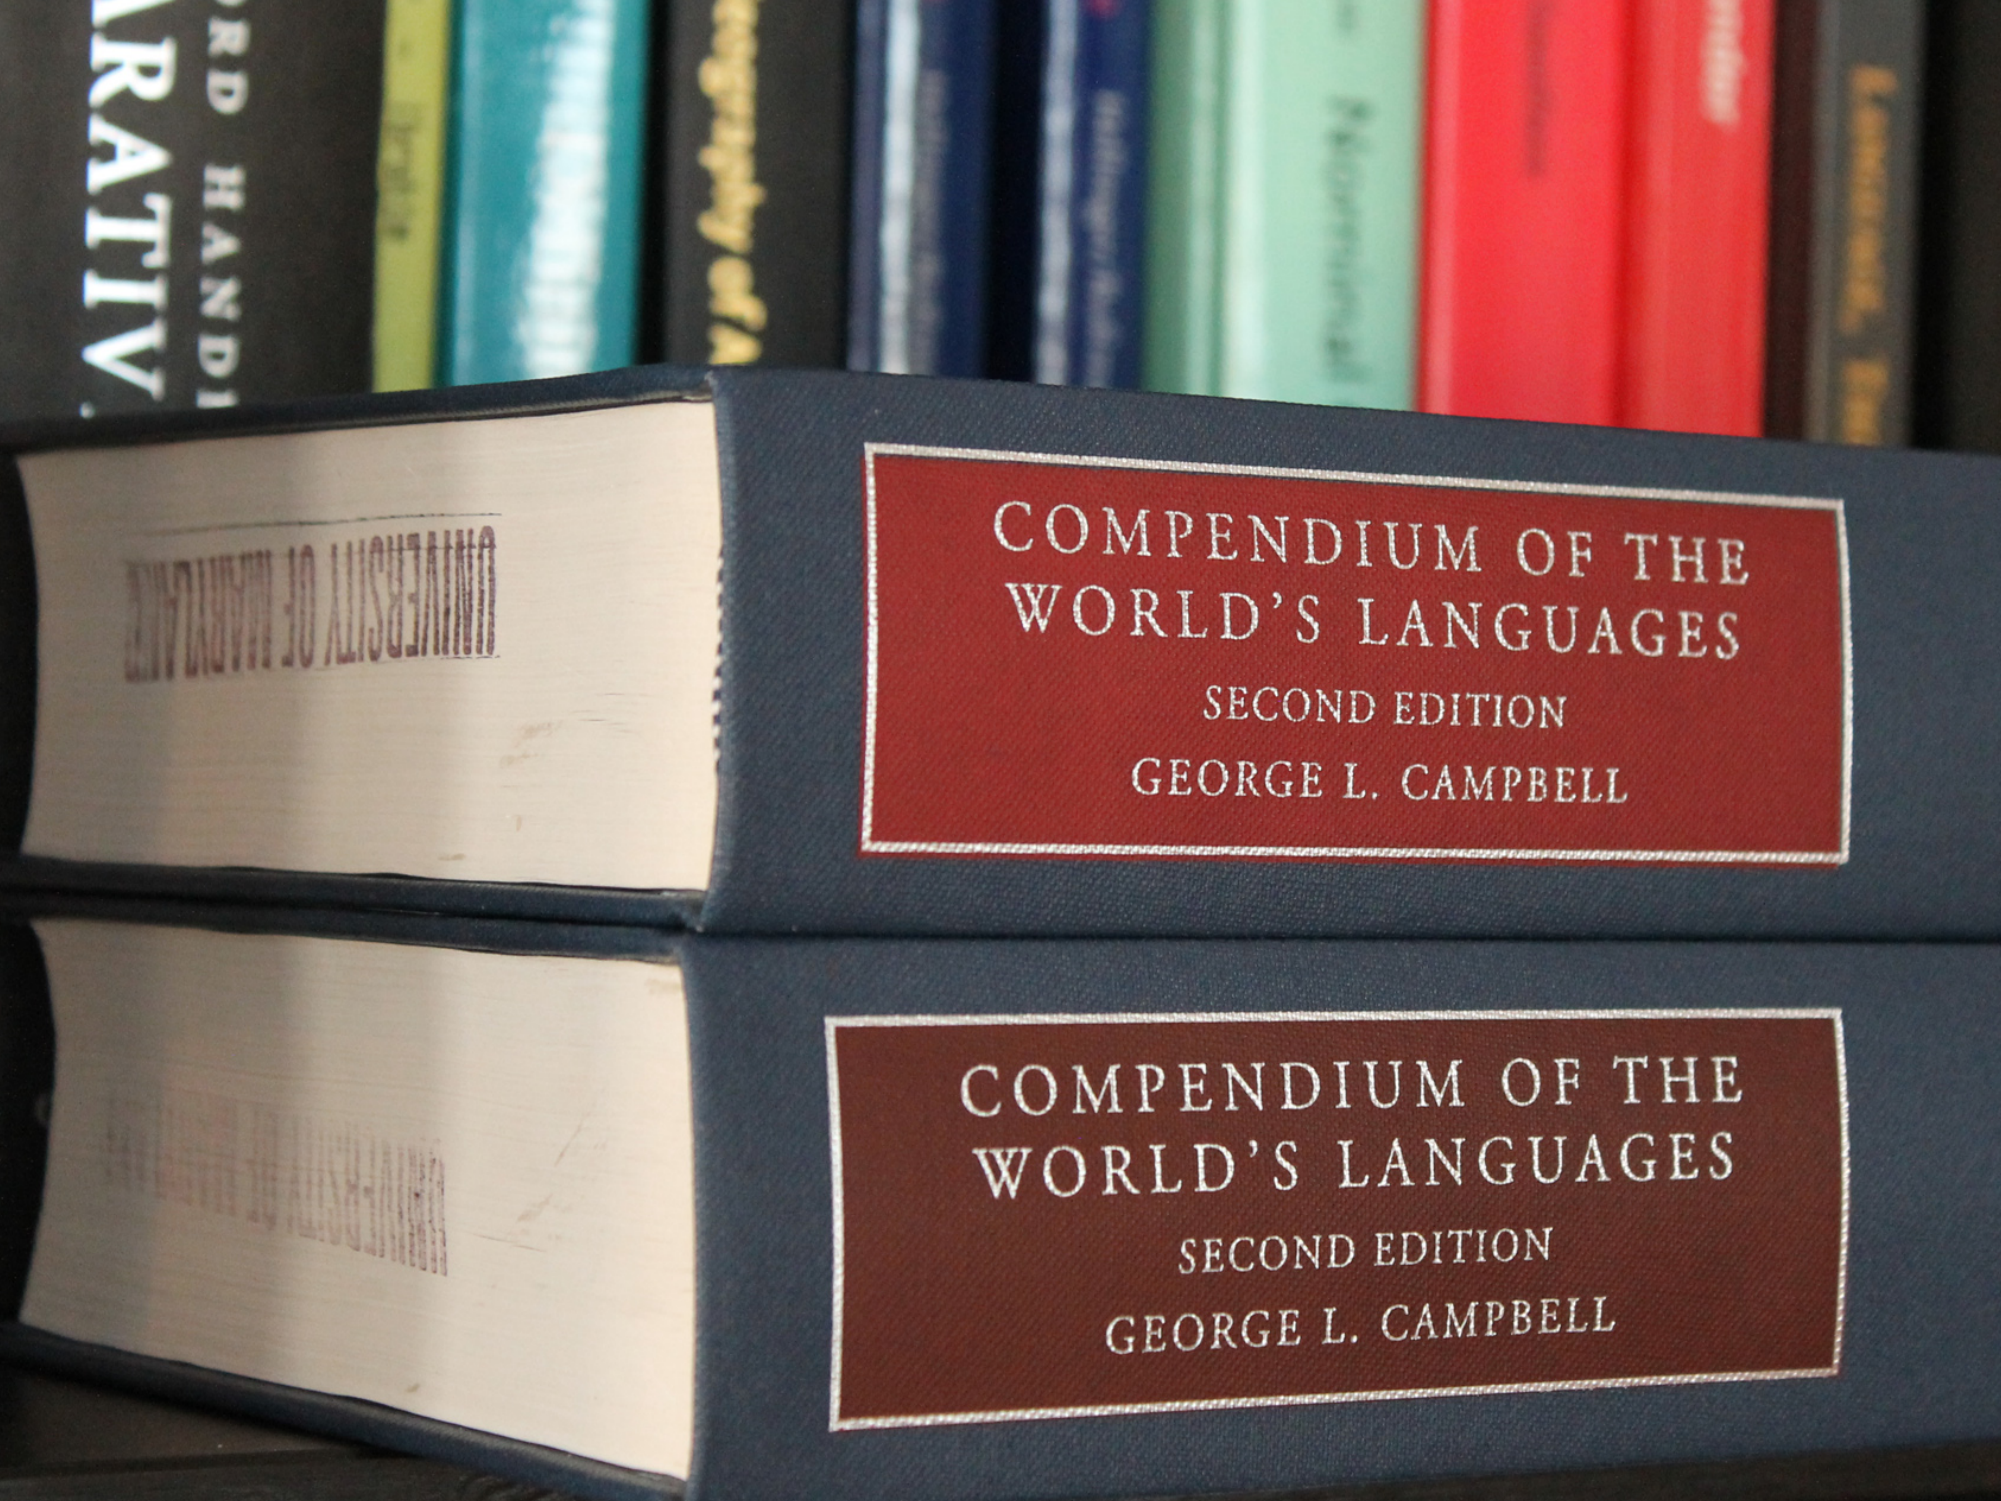
COMPENDIUM OF THE
WORLD'S LANGUAGES
SECOND EDITION
GEORGE L. CAMPBELL
COMPENDIUM OF THE
WORLD'S LANGUAGES
SECOND EDITION
GEORGE L. CAMPBELL

Classifying Gender Structures

We compile data on grammatical structures from a range of sources:

- World Atlas of Language Structures
- *Linguistic Survey of India*
  - Compiled by George A. Grierson between 1891 and 1928
- George L. Campbell's *Compendium of the World's Languages*
- Language-specific data sources:
  - Grammatical monographs
  - Language textbooks and online learning materials
  - Academic work by (modern) linguists
  - Interviews with native speakers and translators

Classifying Gender Structures

For each language, we attempt to code two variables:

- A indicator for using any system of grammatical gender
- A indicator for using a dichotomous system of grammatical gender
  - All nouns must be either masculine or feminine

Classifying Gender Structures

For each language, we attempt to code two variables:

- A indicator for using any system of grammatical gender
- A indicator for using a dichotomous system of grammatical gender
  - All nouns must be either masculine or feminine

We do not attempt to determine:

- The number of genders/classes, if there are more than two
- The intensity of the agreement system (i.e. what must agree)
- The presence of gendered personal pronouns (for humans)

Classifying Gender Structures

Languages positively identified as **gender** languages in two ways:

1. Explicit statement about grammatical gender structure

Serbian: "Three grammatical genders (masculine, feminine, and neuter) and two numbers (singular and plural) are also distinguished."

Tigrinya: "Tigrinya nouns are either masculine or feminine and are inflected for number. Gender is not marked on the noun, but on nominal dependents like articles and adjectives. Verbs agree with their subjects and objects in person, number, and gender."

Classifying Gender Structures

Languages positively identified as **gender** languages in two ways:

1. Explicit statement about grammatical gender structure

Serbian: "Three grammatical genders (masculine, feminine, and neuter) and two numbers (singular and plural) are also distinguished."

Tigrinya: "Tigrinya nouns are either masculine or feminine and are inflected for number. Gender is not marked on the noun, but on nominal dependents like articles and adjectives. Verbs agree with their subjects and objects in person, number, and gender."

2. A textbook or language-specific grammar indicates that:

   - There are masculine and feminine noun classes (genders), at least one of which includes nouns other than male/female animates
   - Adjectives or another part of speech must agree in gender

Classifying Gender Structures

Languages identified as **non-gender** languages in the same ways:

1. Explicit statement about grammatical gender structure

Gamo: "The use of gender is governed by non-linguistic factors — i.e. by the actual sex of the referent."

Maithili: "Modern Maithili, however, has no grammatical gender. In other words, in modern Maithili, distinctions of gender are determined soley by the sex of the animate noun."

Nuosu: "There is no grammatical gender, and such words as do not denote animate beings have no gender at all."

Classifying Gender Structures

Languages identified as **non-gender** languages in the same ways:

1. Explicit statement about grammatical gender structure

Gamo: "The use of gender is governed by non-linguistic factors — i.e. by the actual sex of the referent."

Maithili: "Modern Maithili, however, has no grammatical gender. In other words, in modern Maithili, distinctions of gender are determined soley by the sex of the animate noun."

Nuosu: "There is no grammatical gender, and such words as do not denote animate beings have no gender at all."

2. A textbook or language-specific grammar describes nouns or nominals without mentioning any noun class system, or describes a system of classes that do not include either masculine or feminine

Classifying Gender Structures

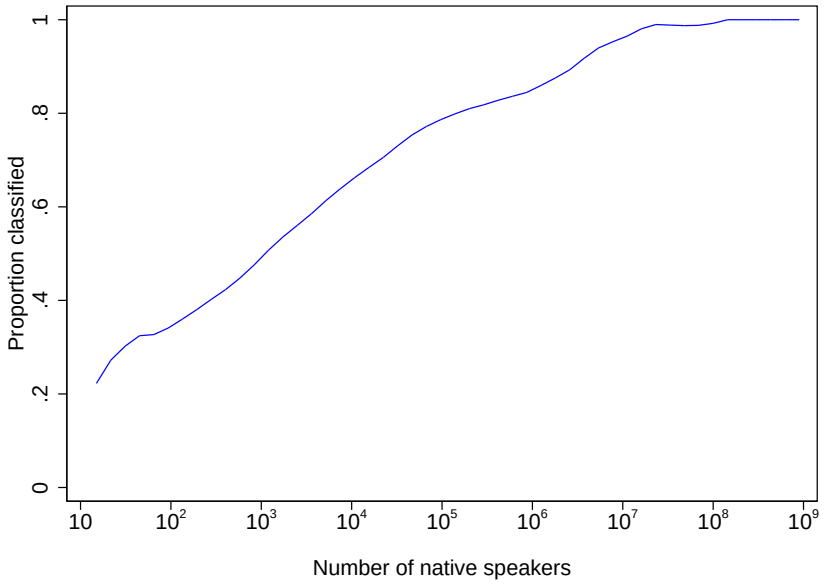

We identify the grammatical structure of 4,334 of 6,190 languages

- All but four of the languages with more than one million speakers
- Verify gender structure w/ two sources whenever possible

Classifying Gender Structures

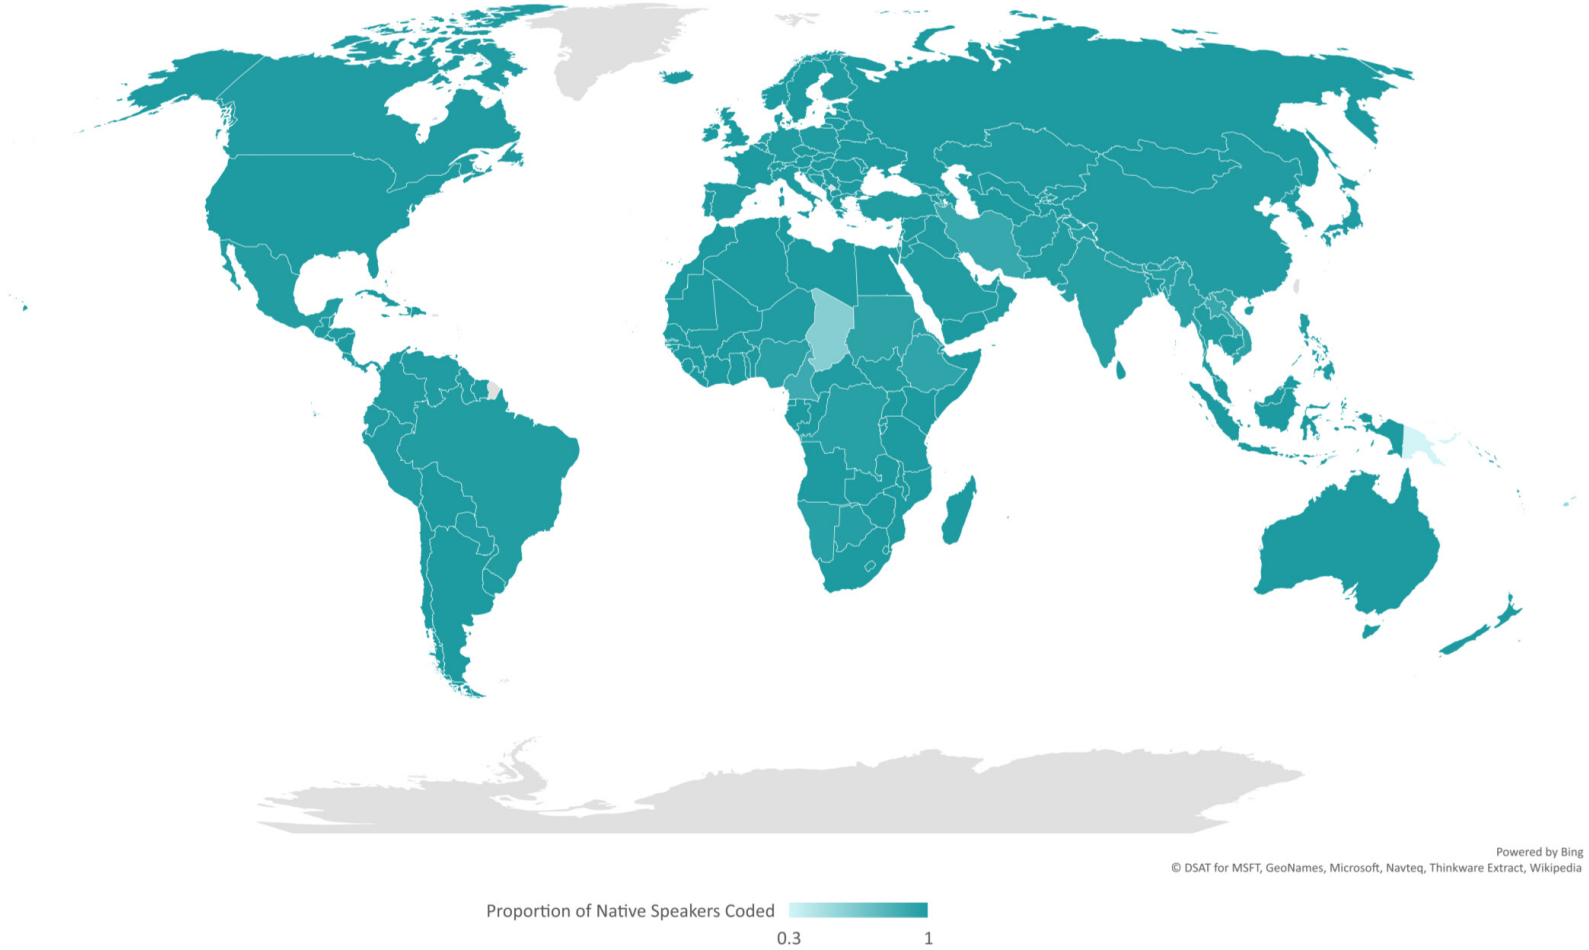

We classify more than 95 percent of population in all but eight countries

The Distribution of Gender Languages

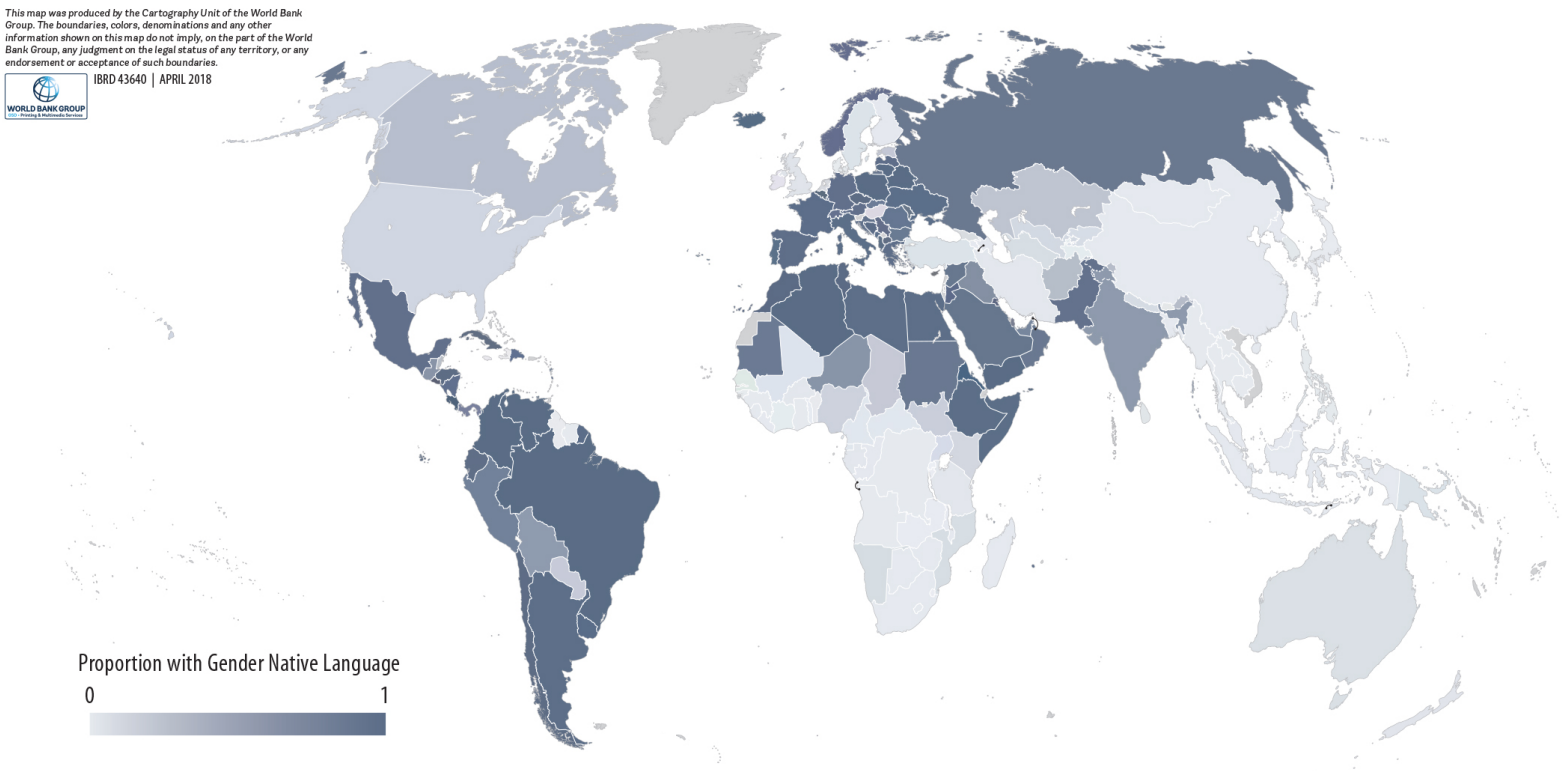

Native speakers of gender languages: 38 percent of world's population

→ [Comparison with WALS]

Comparing country-level WALS data to full data

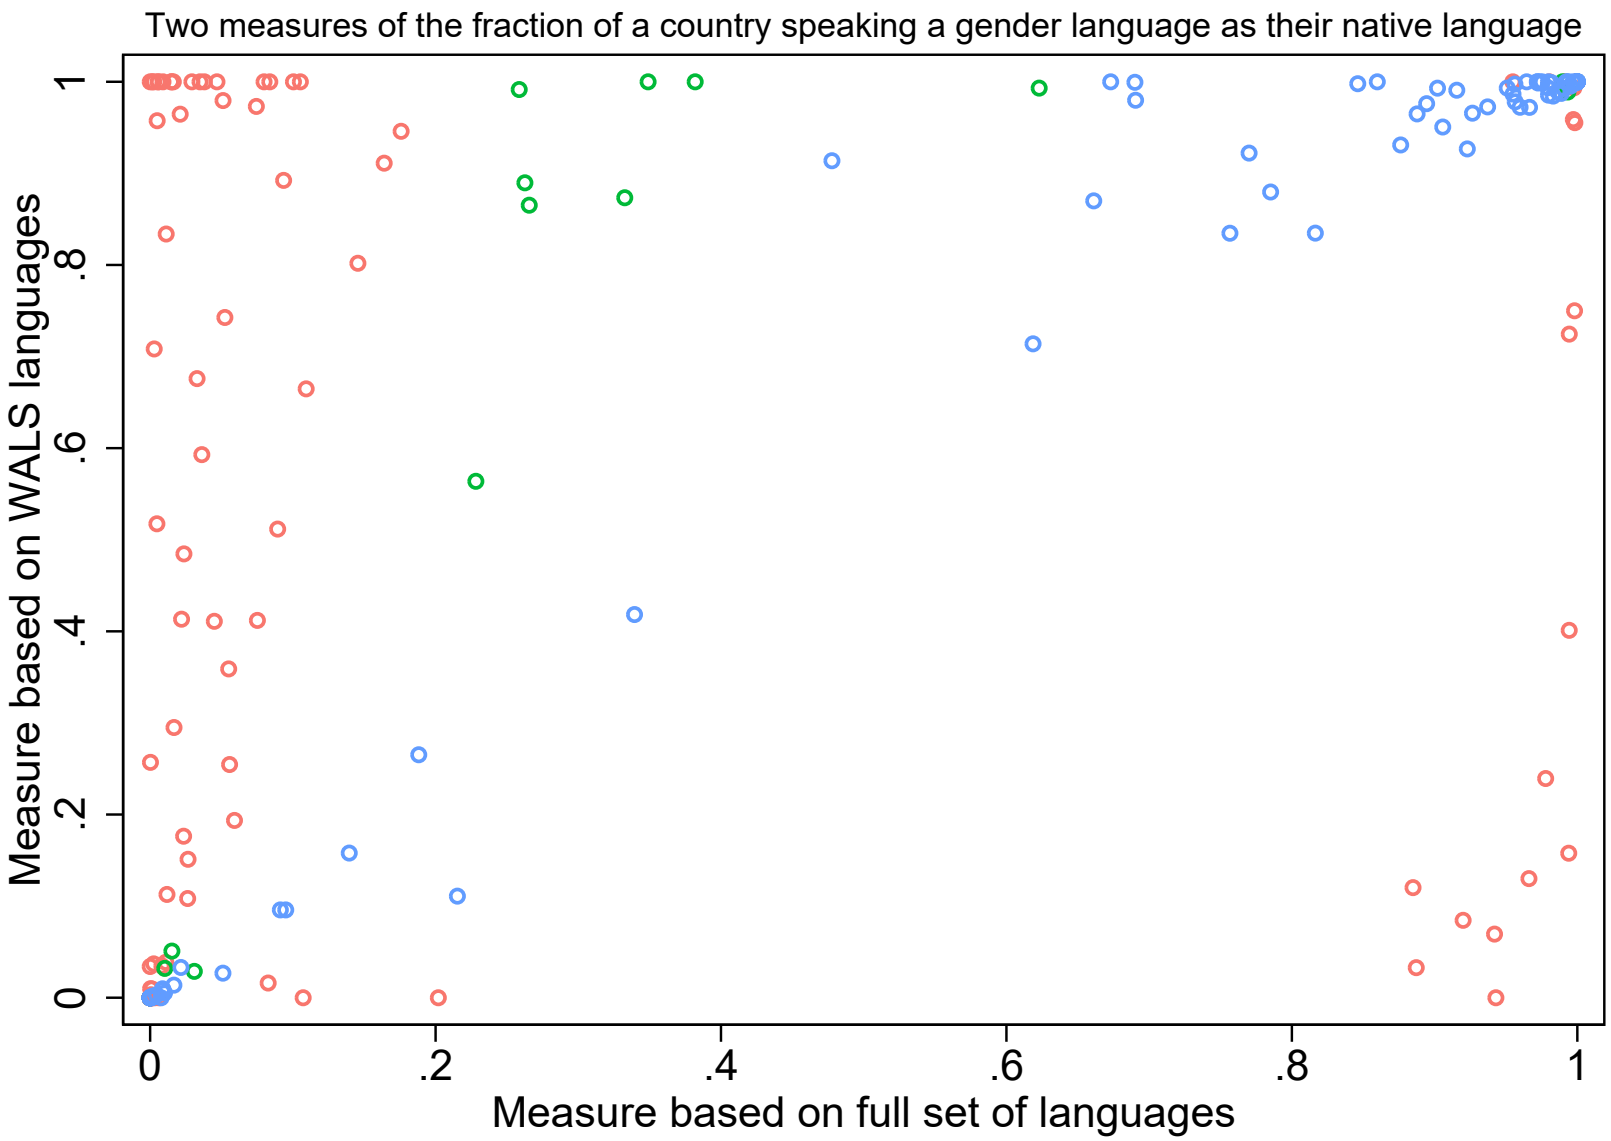

Cross-Country Analysis

Cross-Country Analysis: Data

1. Labor force participation
   - World Development Indicators
   - Available for 177 countries

Cross-Country Analysis: Data

1. Labor force participation
   - World Development Indicators
   - Available for 177 countries
2. Educational attainment (primary and secondary school completion)
   - Barro-Lee Educational Attainment Data
   - Available for 142 countries

Cross-Country Analysis: Data

1. Labor force participation
   - World Development Indicators
   - Available for 177 countries
2. Educational attainment (primary and secondary school completion)
   - Barro-Lee Educational Attainment Data
   - Available for 142 countries
3. Gender attitudes
   - World Values Survey, Round 6
   - Available for 56 countries

Cross-Country Analysis: Empirical Specifications

We estimate OLS regressions of the form:

$$Y_c = \alpha + \beta \mathit{Gender}_c + \delta_{continent} + \lambda X_c + \varepsilon_c$$

where:

- Gender_c is the proportion of population speaking gender language
- $\delta_{continent}$ is a vector of continent fixed effects
- X_c is a vector of country-level geographic controls:
  - Average rainfall, average temperature, proportion tropical, indicator for being landlocked, suitability for the plough
- ε_c is a mean-zero error term

Cross-Country Analysis: Robust Inference

1. Measurement error in country-level prevalence of gender languages
   - Bounding exercise following Imbens and Manski (2004)

Cross-Country Analysis: Robust Inference

1. Measurement error in country-level prevalence of gender languages
   - Bounding exercise following Imbens and Manski (2004)
2. Non-independence of languages with families
   - Permutation test based on structure of the language tree

Cross-Country Analysis: Assessing Causality

1. Examine within-country gender differences, where applicable
   - Applies to LFP and education, not gender attitudes

Cross-Country Analysis: Assessing Causality

1. Examine within-country gender differences, where applicable
   - Applies to LFP and education, not gender attitudes
2. Examine coefficient stability, robustness to observable controls
   - Follow Altonji *et al.* (2005), Oster (forthcoming)

Cross-Country Analysis: Assessing Causality

1. Examine within-country gender differences, where applicable
   - Applies to LFP and education, not gender attitudes
2. Examine coefficient stability, robustness to observable controls
   - Follow Altonji *et al.* (2005), Oster (forthcoming)
3. Replicate cross-country results using within-country variation

Cross-Country Analysis: Female LFP

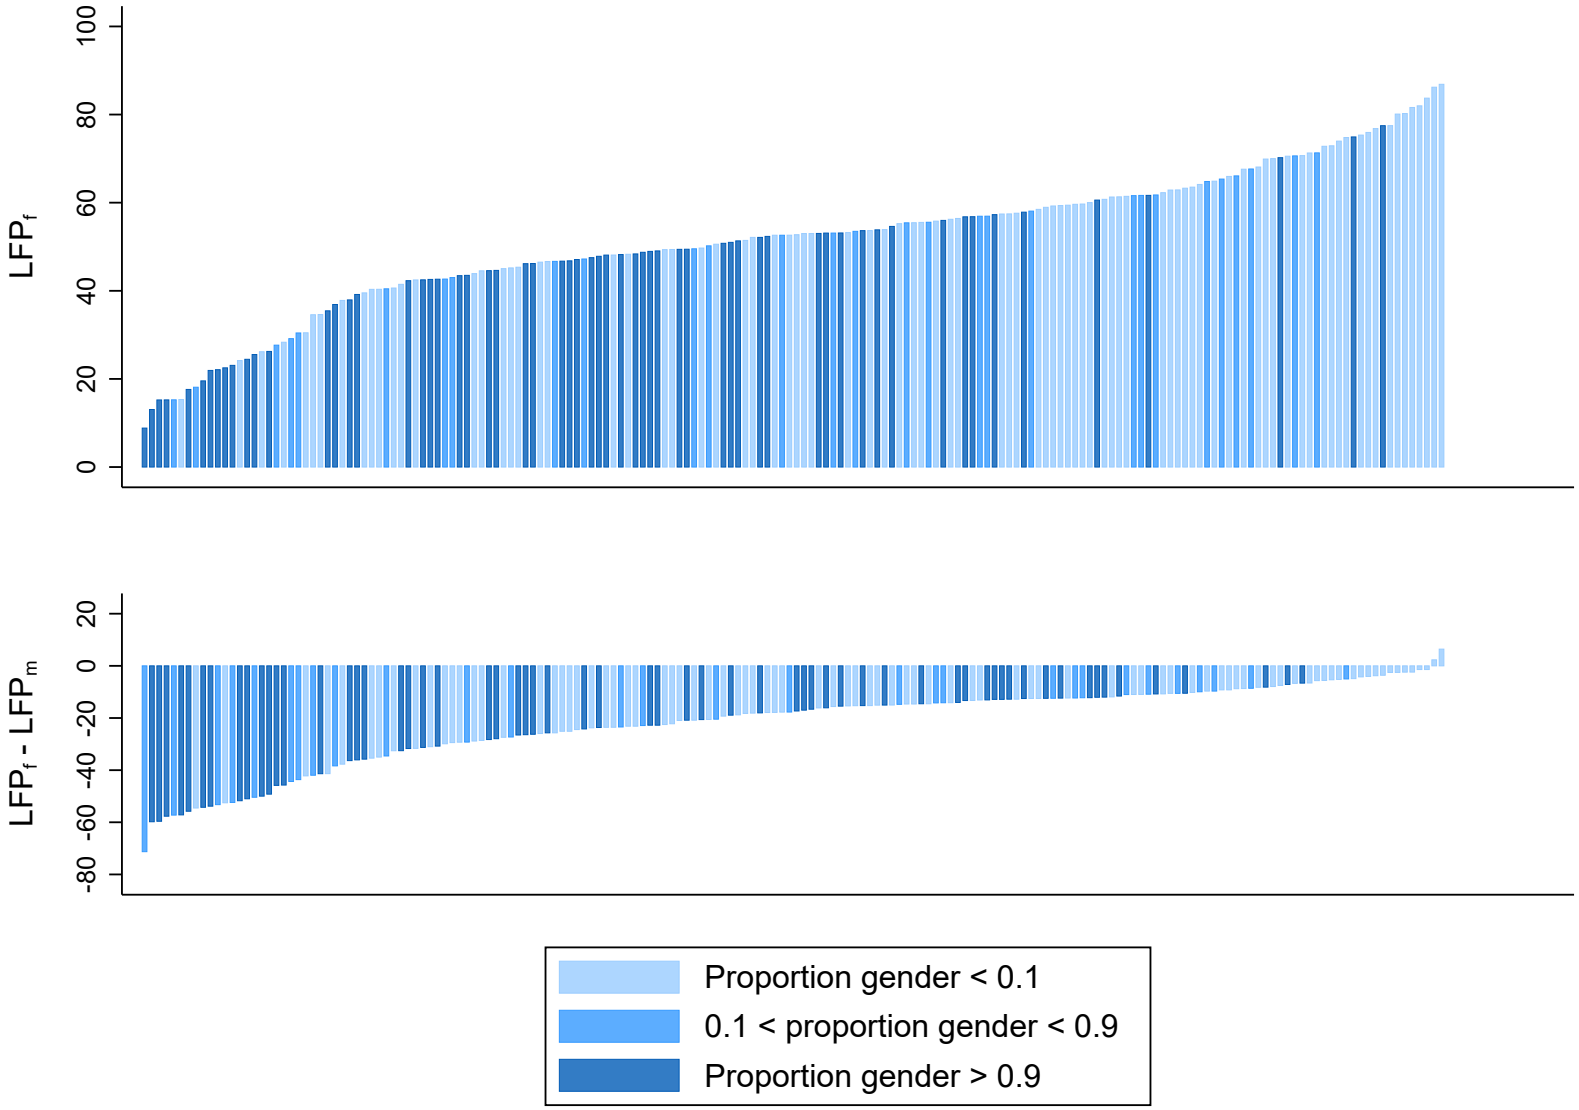

Cross-Country Analysis: Female LFP

Dependent variable:	$\mathbf{LFP}_f$		$\mathbf{LFP}_f$ - $\mathbf{LFP}_m$	
Specification:	**OLS** (1)	**OLS** (2)	**OLS** (3)	**OLS** (4)
Proportion gender	-13.83	-11.92	-11.61	-14.66
	(2.80)	(3.34)	(2.47)	(3.25)
	[$p < 0.001$]	[$p < 0.001$]	[$p < 0.001$]	[$p < 0.001$]
Continent Fixed Effects	No	Yes	No	Yes
Country-Level Geography Controls	No	Yes	No	Yes
Observations	178	178	178	178
R^2	0.15	0.33	0.12	0.47

Robust standard errors are clustered by the most widely spoken language in all specifications; they are reported in parentheses. P-values are reported in square brackets. LFP_f is the percentage of women in the labor force, measured in 2011. LFP_f - LFP_m is the gender difference in labor force participation — i.e. the difference between female and male labor force participation, again measured in 2011. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough.

Cross-Country Analysis: Female LFP

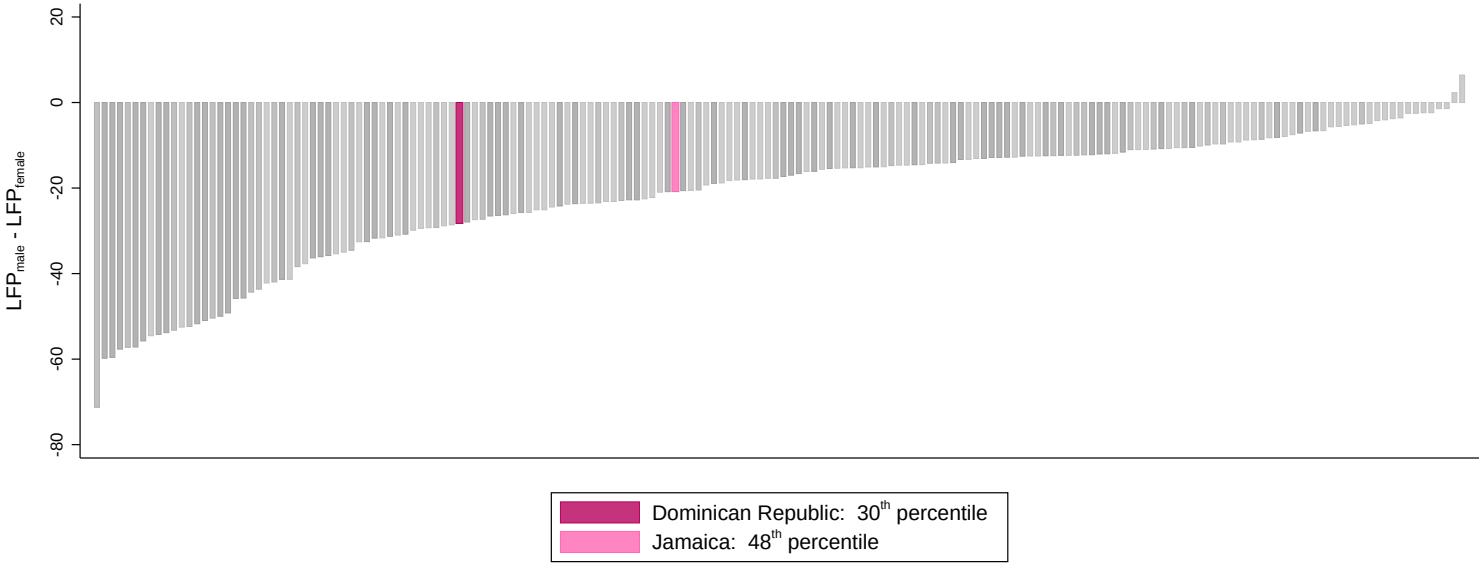

Estimated coefficients are economically significant:

- Grammatical gender could fully explain the disparity in female labor force participation between Jamaica and the Dominican Republic
- Grammatical gender keeps 125 million women out of work force

Cross-Country Analysis: Female LFP

Robustness checks:

- Marginal impact of stronger grammatical gender systems
- Including "bad" controls
- Omitting major world languages

Cross-Country Analysis: Educational Attainment

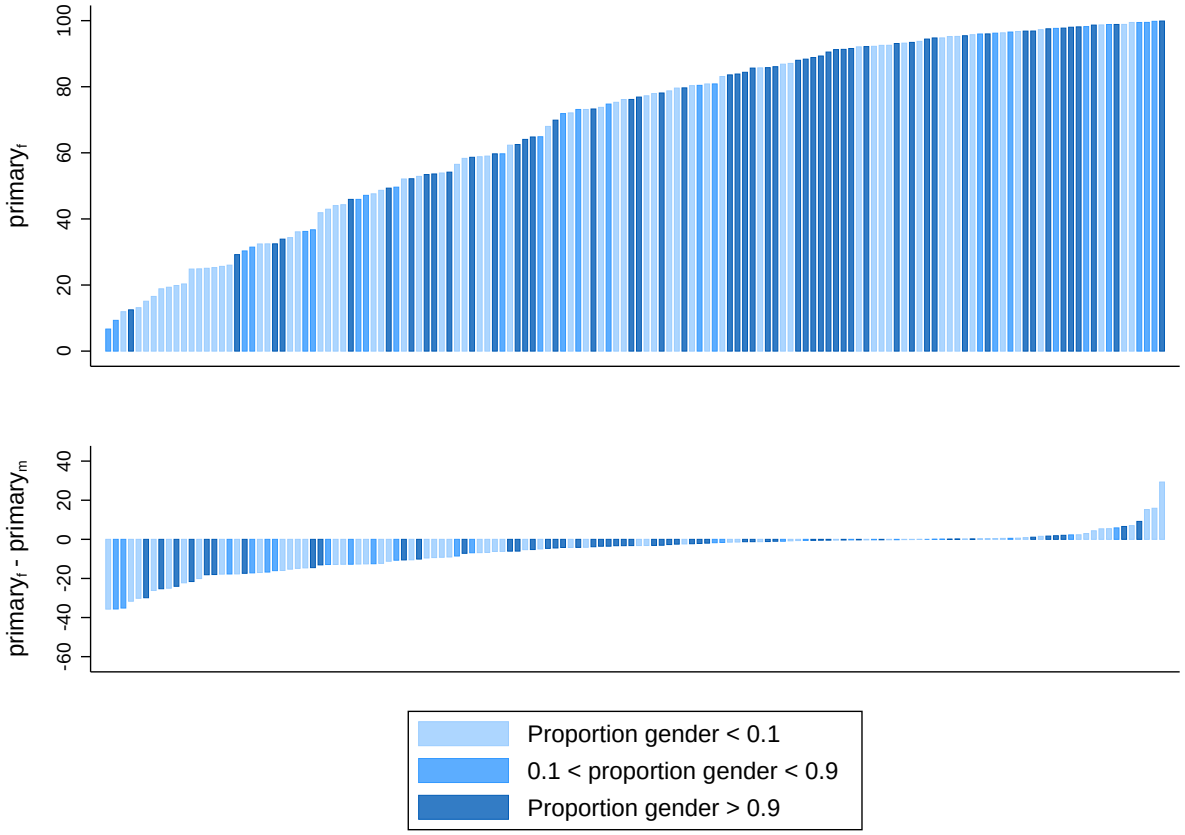

→ Primary education by continent

Cross-Country Analysis: Educational Attainment

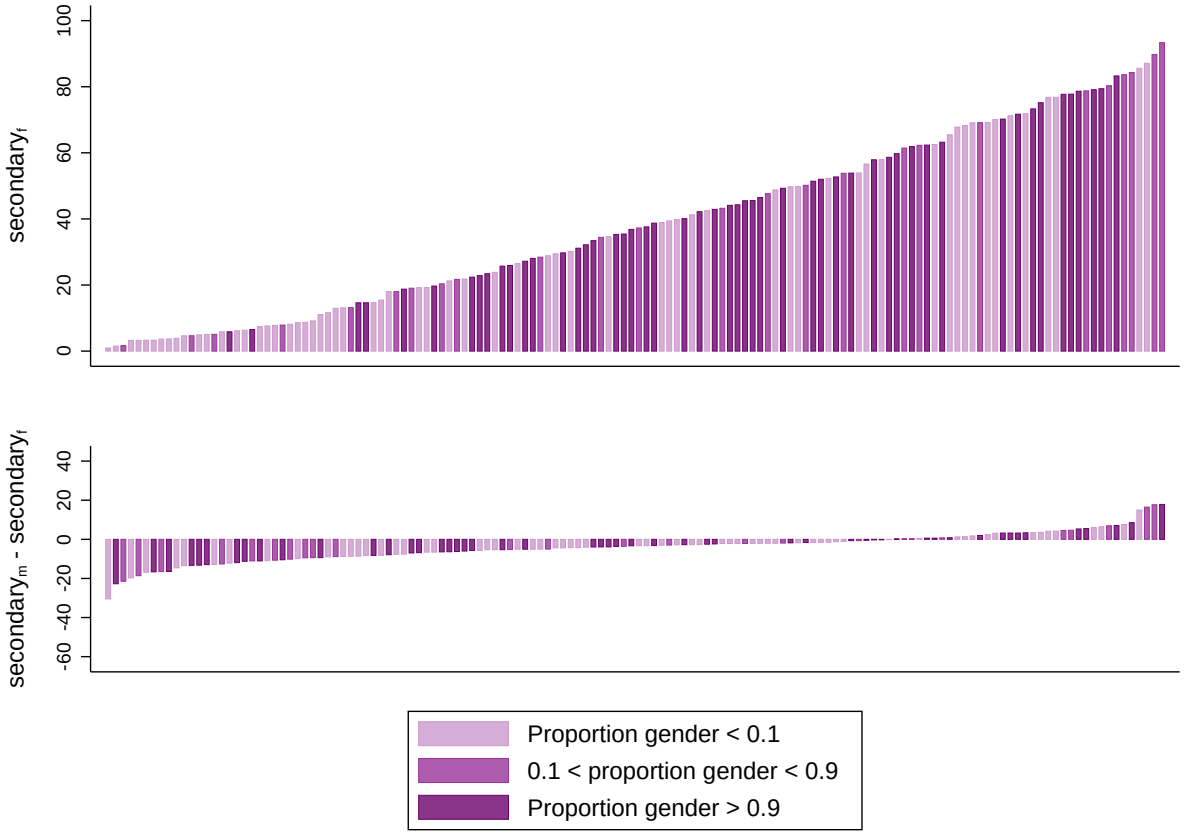

→ Secondary education by continent

Cross-Country Analysis: Educational Attainment

Dependent variable:	$\mathbf{PRI}_f$		$\mathbf{PRI}_f$ - $\mathbf{PRI}_m$	
Specification:	**OLS** (1)	**OLS** (2)	**OLS** (3)	**OLS** (4)
Proportion gender	14.79	-6.71	1.21	-3.72
	(5.83)	(4.40)	(2.14)	(2.16)
	[0.013]	[0.130]	[0.573]	[0.088]
Continent Fixed Effects	No	Yes	No	Yes
Country-Level Geography Controls	No	Yes	No	Yes
Observations	142	142	142	142
R^2	0.06	0.61	0.00	0.20

Robust standard errors are clustered by the most widely spoken language in all specifications; they are reported in parentheses. P-values are reported in square brackets. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough.

Cross-Country Analysis: Educational Attainment

Dependent variable:	$\mathbf{SEC}_f$		$\mathbf{SEC}_f$ - $\mathbf{SEC}_m$	
Specification:	**OLS** (1)	**OLS** (2)	**OLS** (3)	**OLS** (4)
Proportion gender	14.52	0.43	0.48	-0.86
	(5.77)	(3.70)	(1.93)	(2.35)
	[0.013]	[0.907]	[0.802]	[0.716]
Continent Fixed Effects	No	Yes	No	Yes
Country-Level Geography Controls	No	Yes	No	Yes
Observations	142	142	142	142
R^2	0.06	0.67	0.00	0.10

Robust standard errors are clustered by the most widely spoken language in all specifications; they are reported in parentheses. P-values are reported in square brackets. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough.

Cross-Country Analysis: Gender Attitudes

World Values Survey includes 8 questions on gender attitudes:

- When a mother works for pay, the children suffer [1]
- When jobs are scarce, men should have more right to a job than women [1]
- On the whole, men make better political leaders than women do [1]
- On the whole, men make better business executives than women do [1]
- Being a housewife is just as fulfilling as working for pay [1]
- If a woman earns more money than her husband, it's almost certain to cause problems [1]
- A university education is more important for a boy than for a girl [1]
- Having a job is the best way for a woman to be an independent person [0]

Cross-Country Analysis: Gender Attitudes

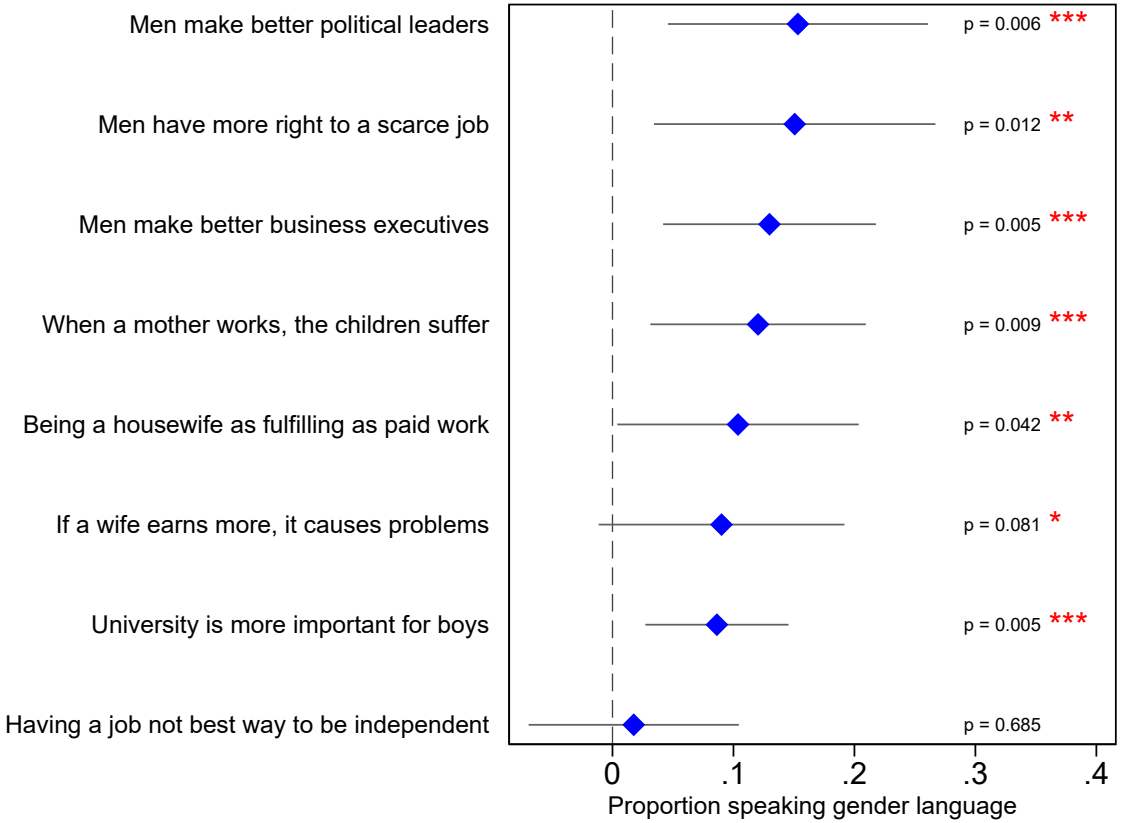

Cross-Country Analysis: Gender Attitudes

Dependent variable:	**Gender Attitude Index**	
Specification:	**OLS (1)**	**OLS (2)**
Proportion gender	-0.03	-0.12
	(0.05)	(0.04)
	[0.576]	[0.002]
Continent Fixed Effects	No	Yes
Country-Level Geography Controls	No	Yes
Observations	56	56
R^2	0.01	0.78

Robust standard errors clustered by most widely spoken language in all specifications. The **Gender Attitude Index** is the first principal component of responses to the eight questions on gender attitudes included in the World Values Survey. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough.

Cross-Country Analysis: Gender Attitudes

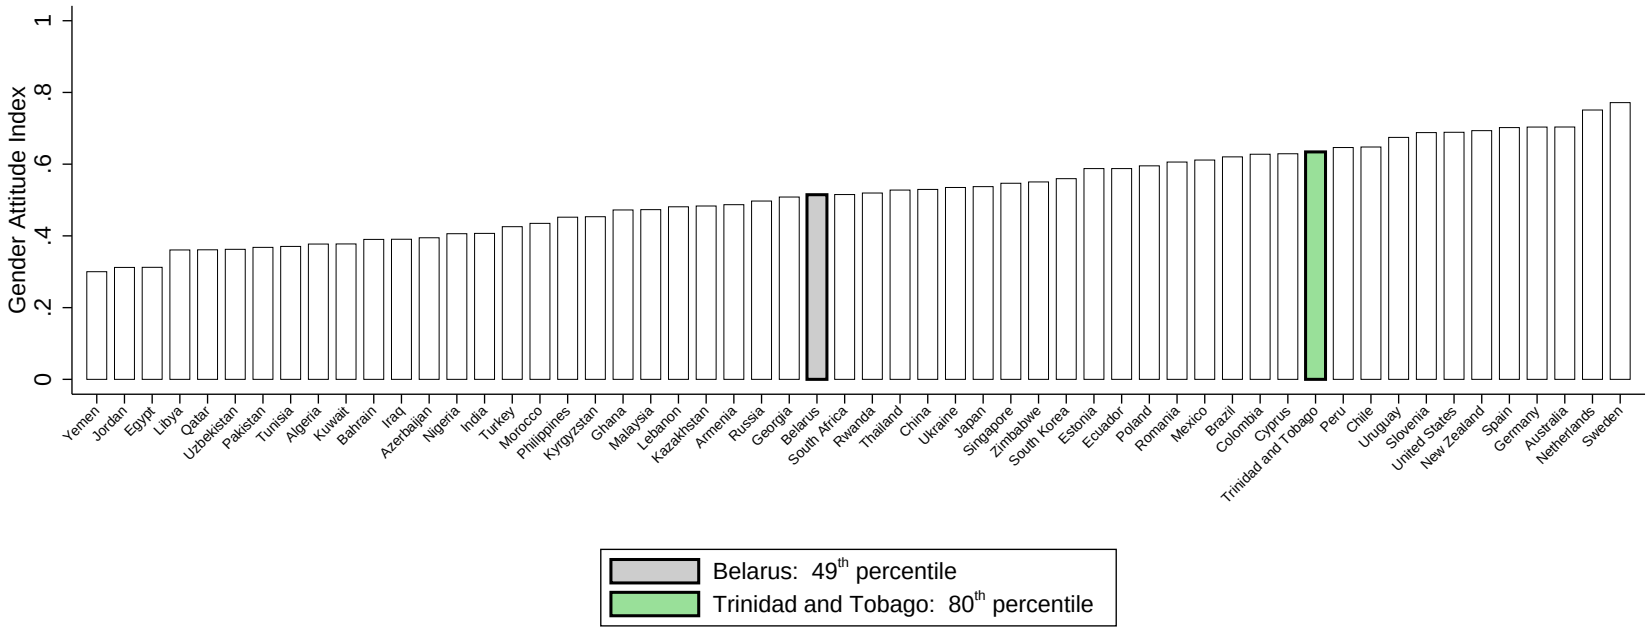

Cross-Country Analysis: Gender Attitudes

Attitudes among Women:

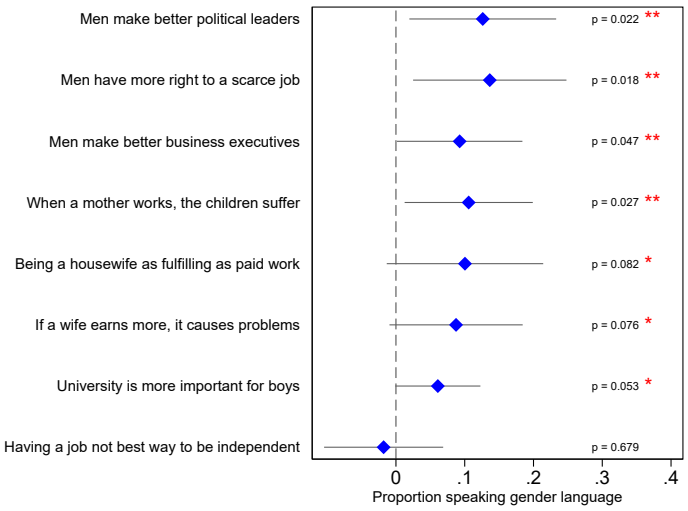

Attitudes among Men:

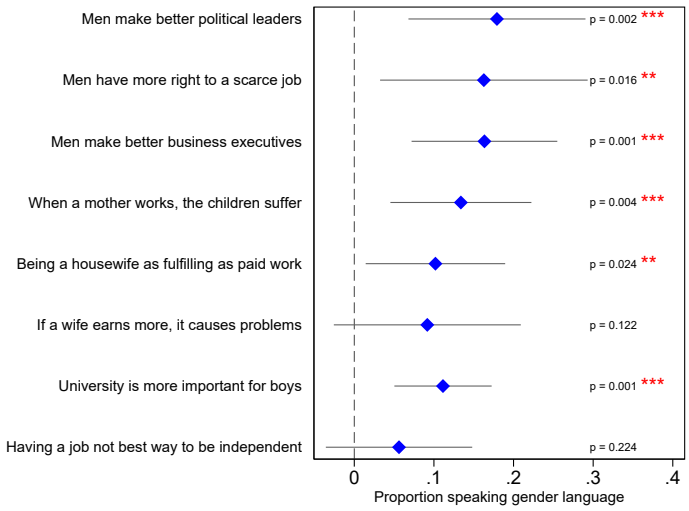

Cross-Country Analysis: Gender Attitudes

Sample:	**Attitude Index: Women**		**Attitude Index: Men**	
Specification:	**OLS (1)**	**OLS (2)**	**OLS (3)**	**OLS (4)**
Proportion gender	-0.02	-0.10	-0.04	-0.14
	(0.05)	(0.04)	(0.06)	(0.04)
	[0.714]	[0.012]	[0.508]	[$p < 0.001$]
Continent Fixed Effects	No	Yes	No	Yes
Geography Controls	No	Yes	No	Yes
Observations	56	56	56	56
R^2	0.00	0.73	0.02	0.78

Robust standard errors clustered by most widely spoken language in all specifications. The **Gender Attitude Index** is the first principal component of responses to the eight questions on gender attitudes included in the World Values Survey. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough.

Cross-Country Analysis: Measurement Error

The problem: RHS variable is an interval for 85 of 193 countries

- Analysis thus far assumes missingness is ignorable
- Measurement error is not classical, could bias estimates

Cross-Country Analysis: Measurement Error

The problem: RHS variable is an interval for 85 of 193 countries

- Analysis thus far assumes missingness is ignorable
- Measurement error is not classical, could bias estimates

Our approach: calculate bounds following Imbens and Manski (2004)

1. Identify highest and lowest coefficient estimates numerically
2. Calculate associated naïve confidence intervals, take the union
3. Symmetrically tighten the confidence interval for correct coverage

Cross-Country Analysis: Measurement Error

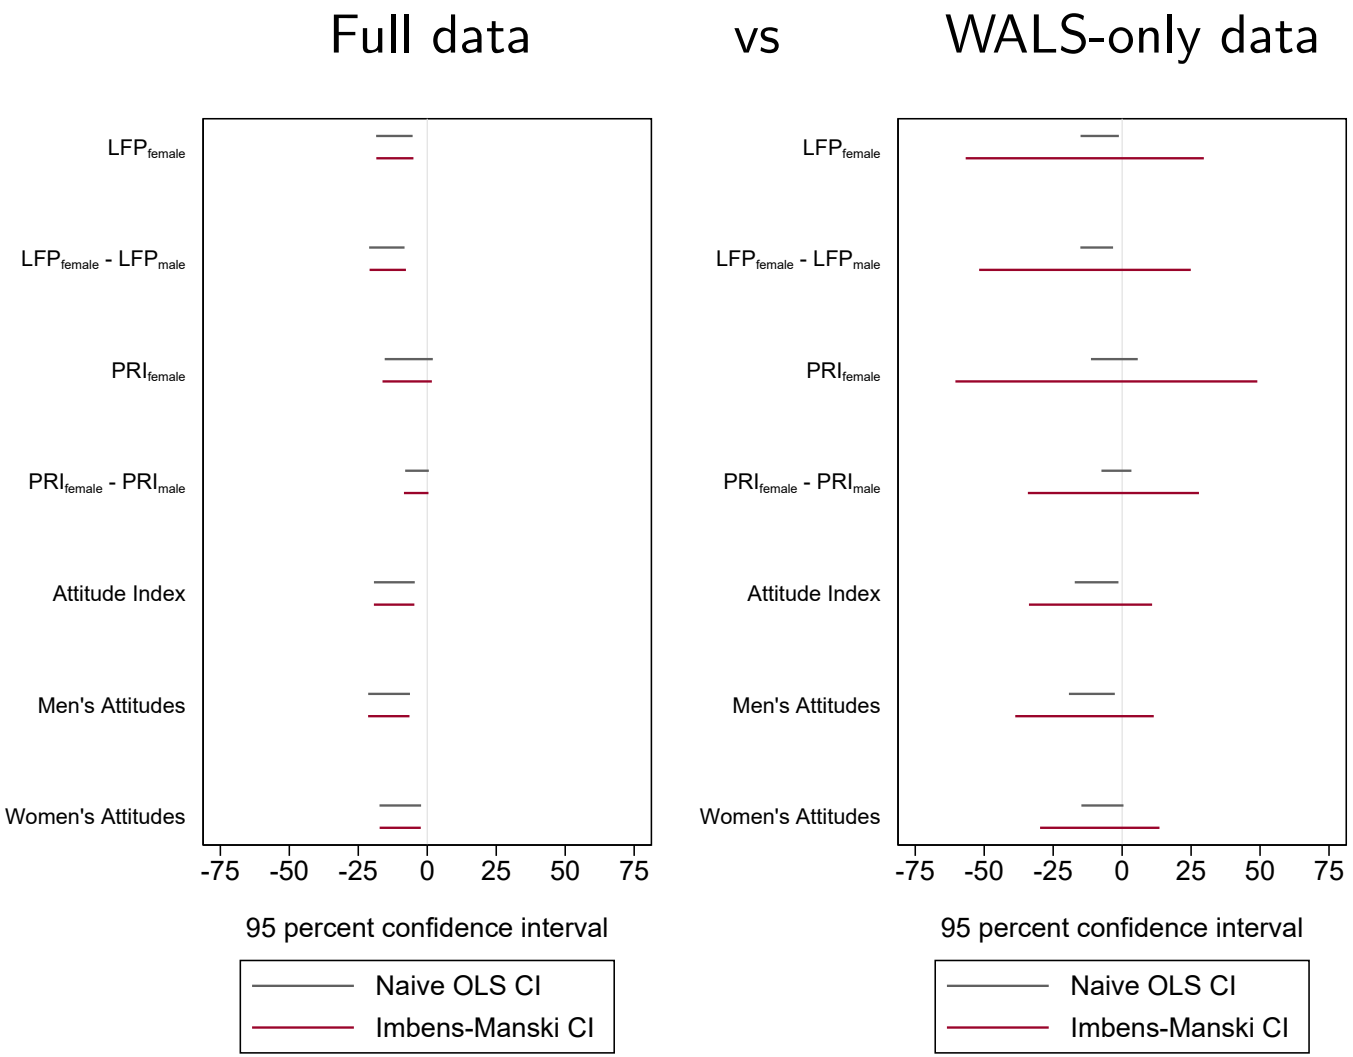

→ [Manski table]

Cross-Country Analysis: Independence

The problem: languages are not independent (Roberts *et al.* 2015)

- Useful variation in grammatical structure within and between families
- Intuitively, this is a clustering problem, but countries not nested

Cross-Country Analysis: Independence

The problem: languages are not independent (Roberts *et al.* 2015)

- Useful variation in grammatical structure within and between families
- Intuitively, this is a clustering problem, but countries not nested

Our approach: permutation tests based on the language tree

1. Assign languages to largest possible homogeneous clusters
2. Randomly permute treatment (grammatical gender) across clusters
3. Replicate cross-country analysis for each hypothetical treatment

   $\Rightarrow$ Allows us to calculate permutation-test p-values

Cross-Country Analysis: Permutation Tests

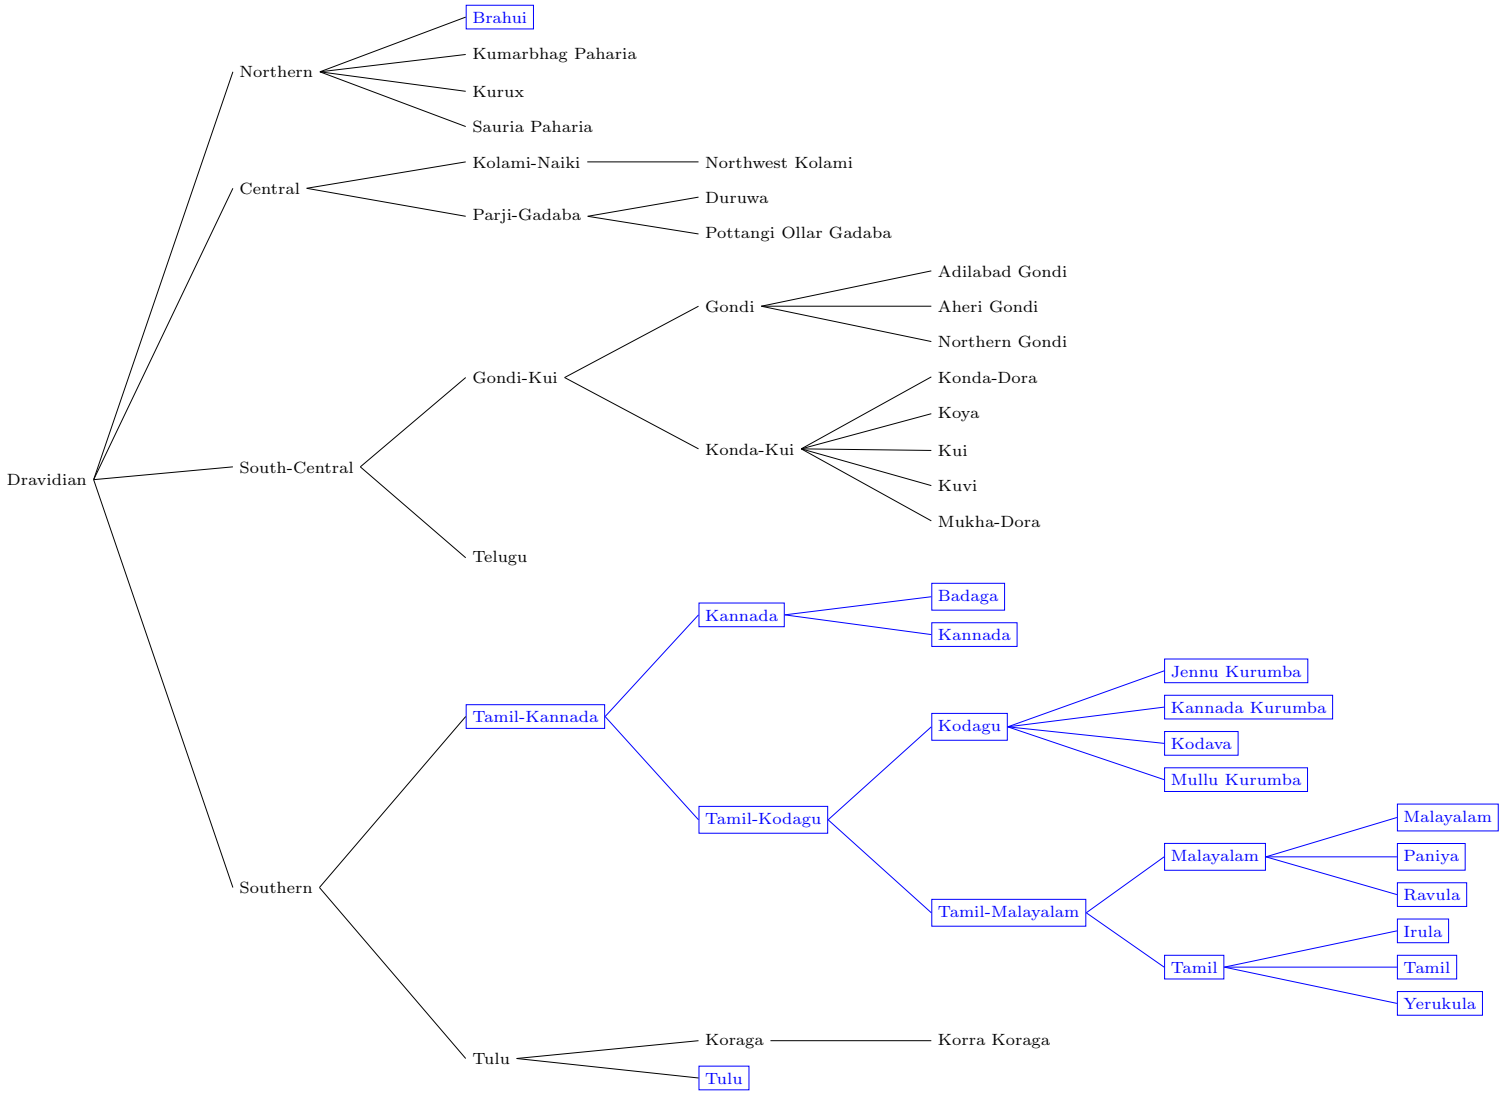

Cross-Country Analysis: Permutation Tests

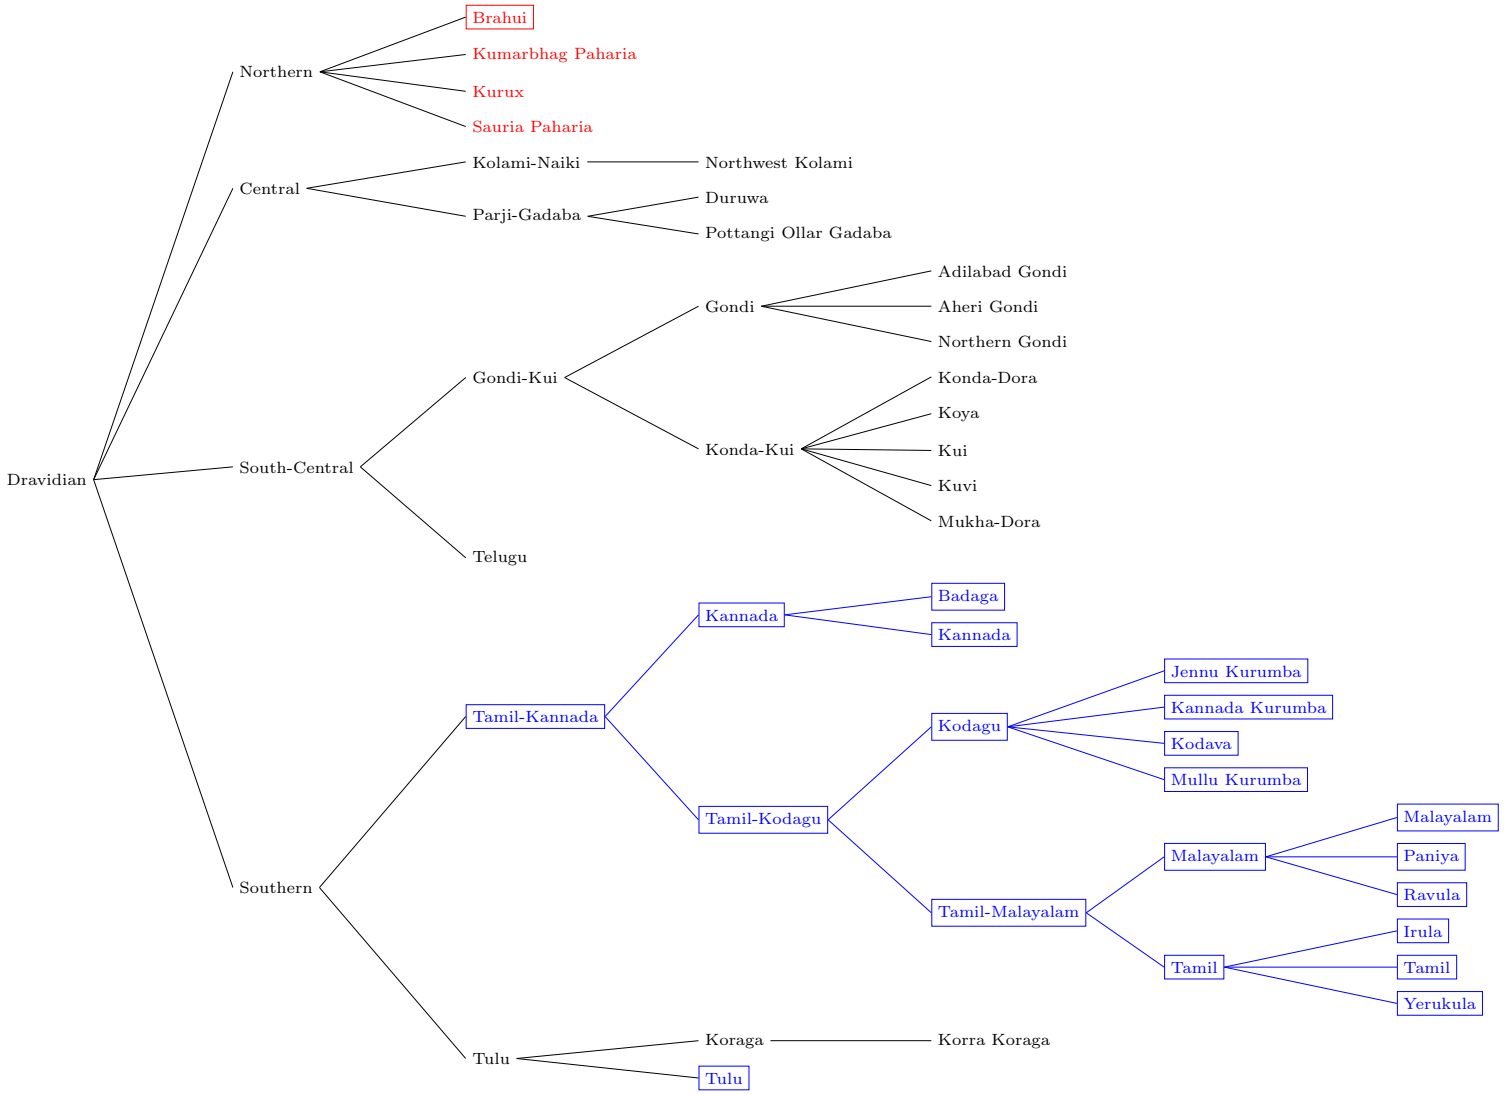

Cross-Country Analysis: Permutation Tests

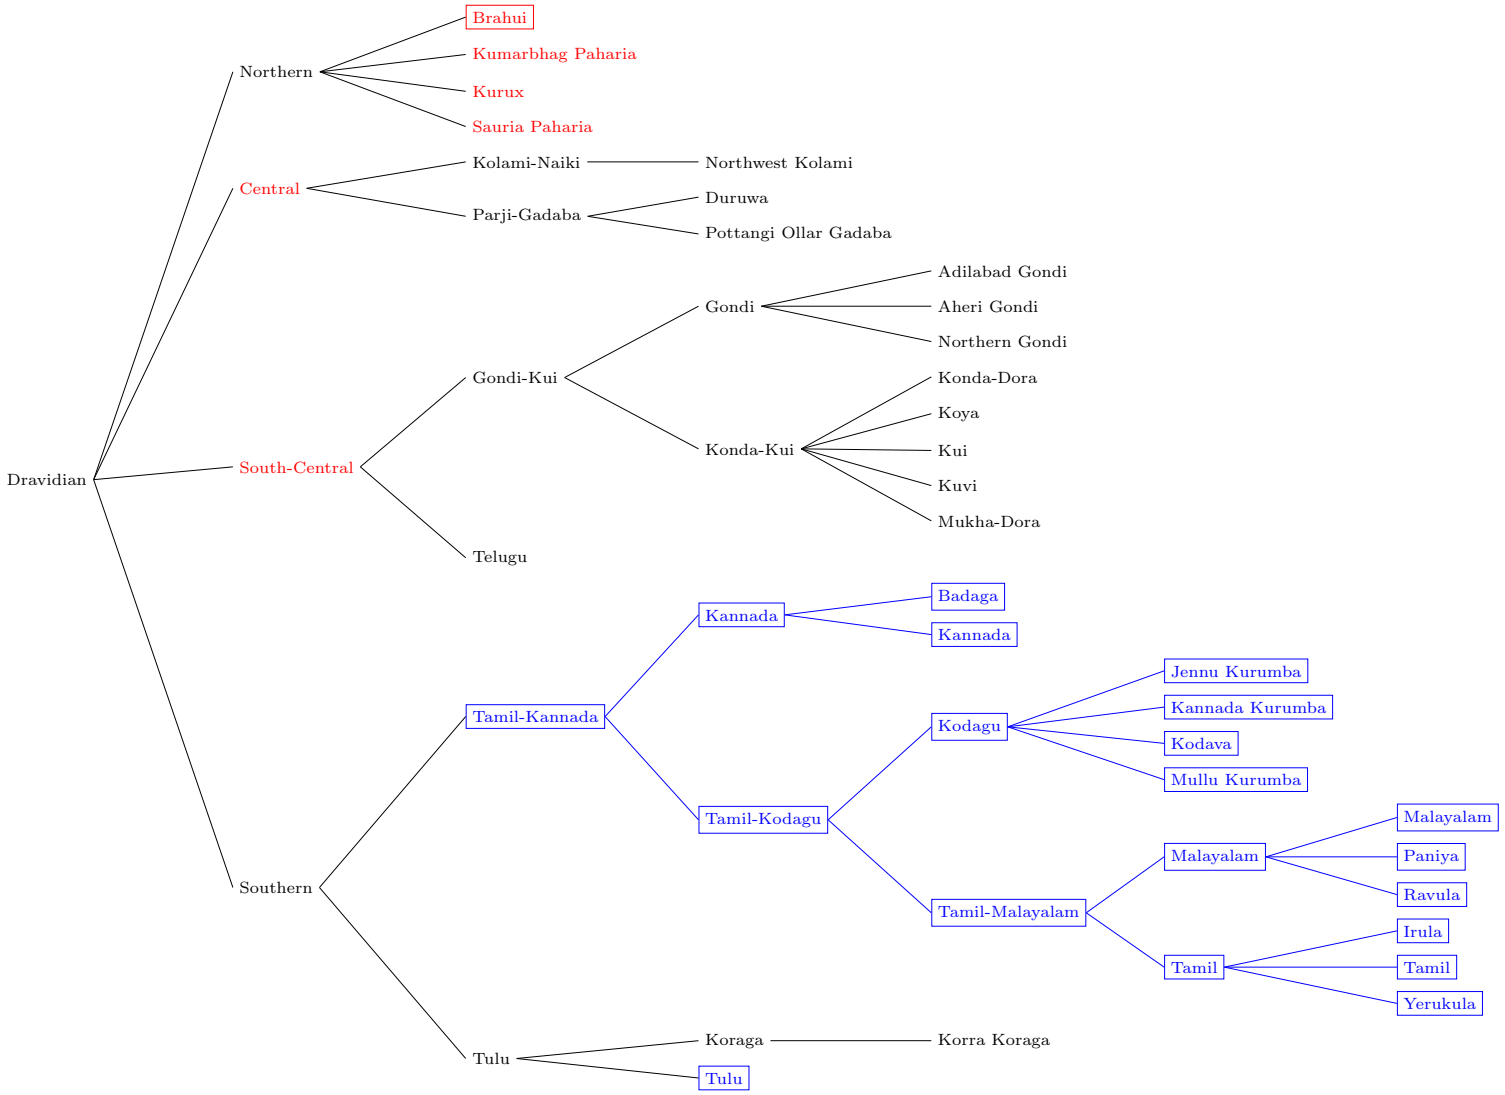

Cross-Country Analysis: Permutation Tests

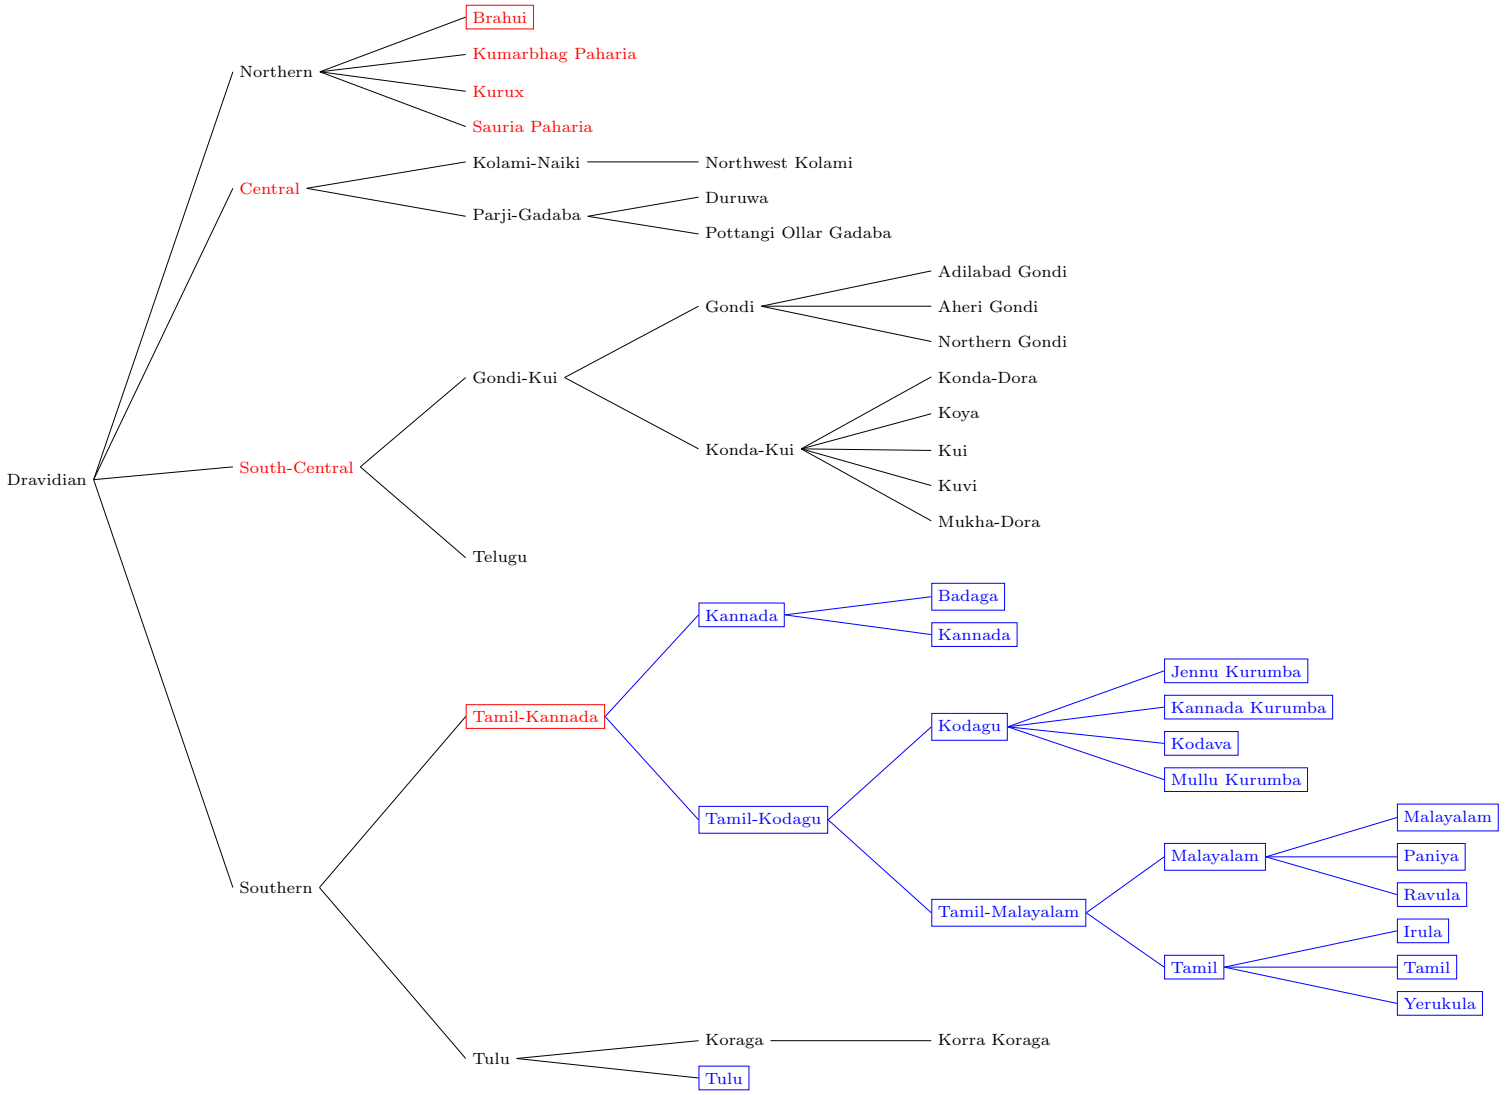

Cross-Country Analysis: Permutation Tests

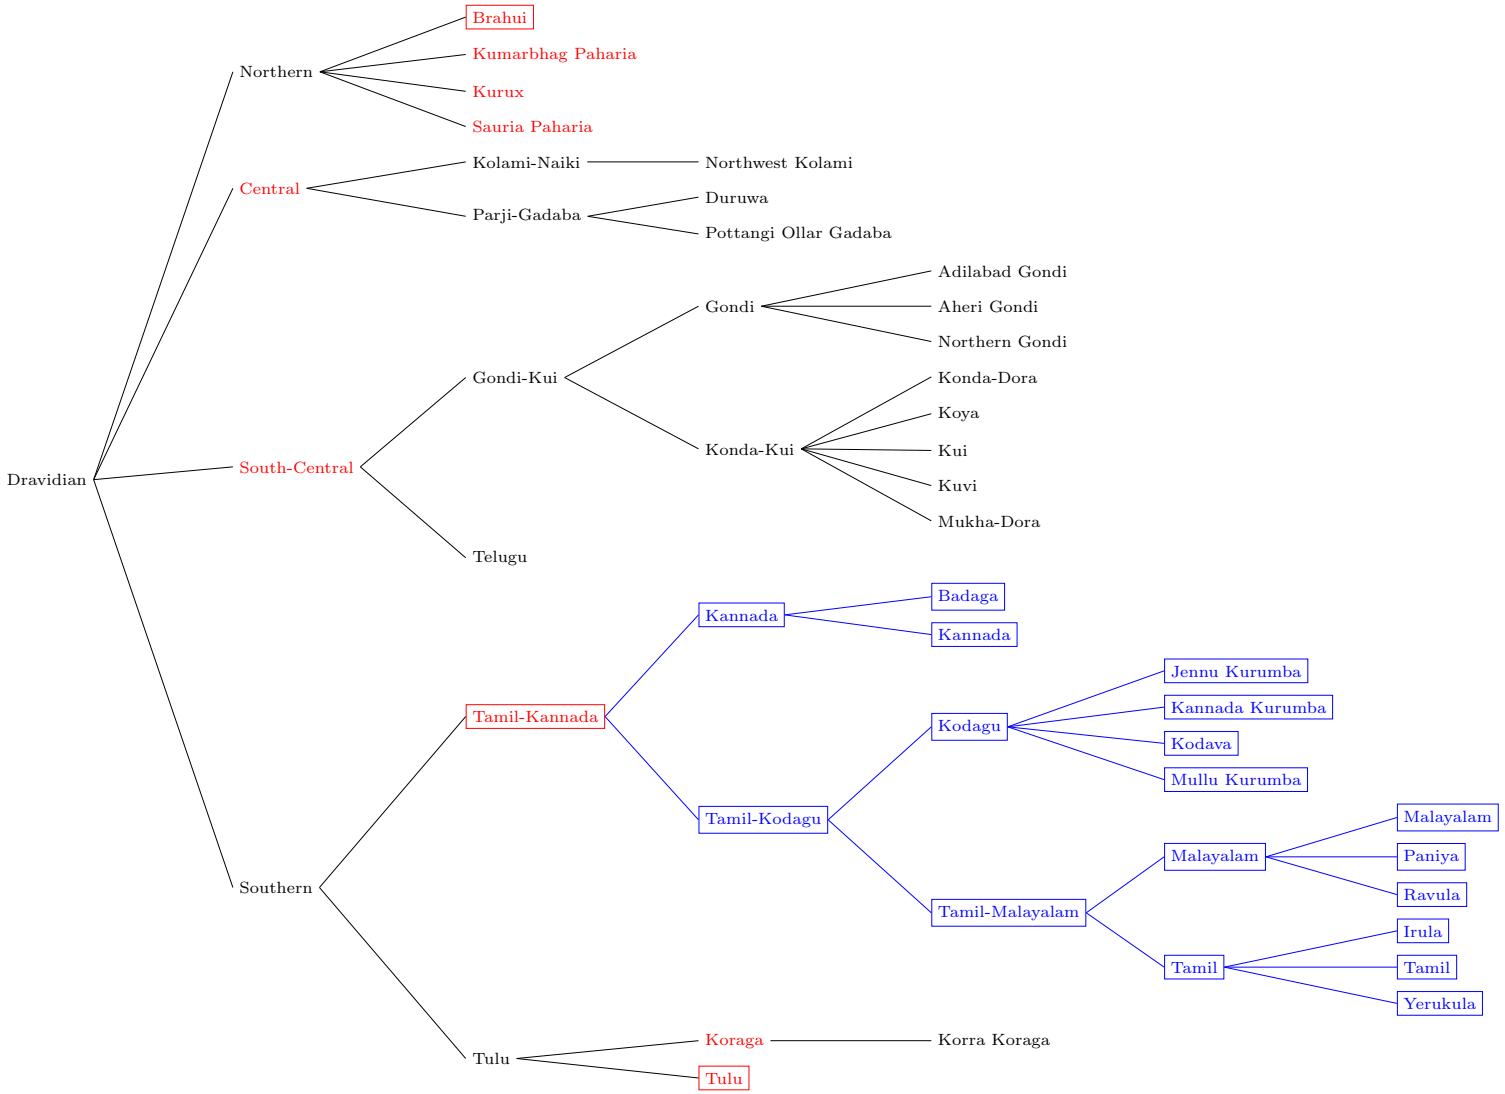

Cross-Country Analysis: Permutation Tests

Female LFP:

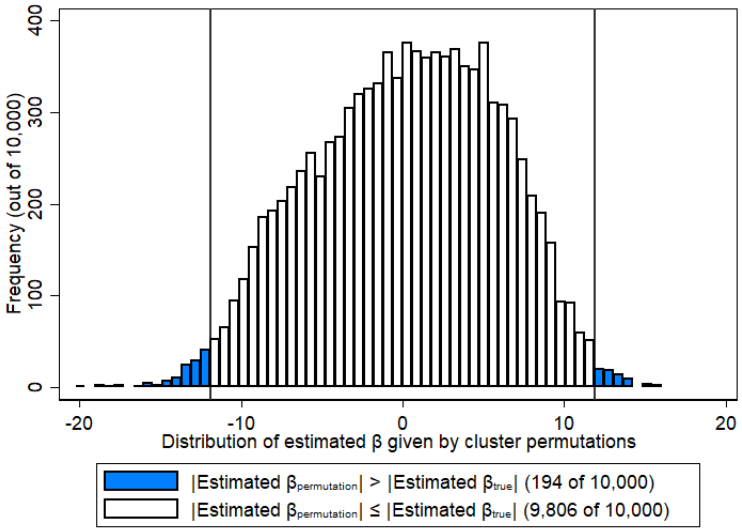

Gender Difference in LFP:

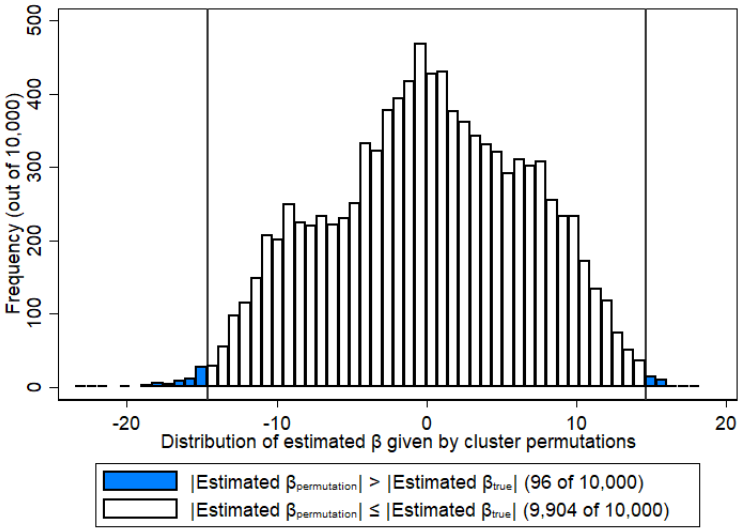

Cross-Country Analysis: Permutation Tests

	Naïve OLS p-values	Permutation-based p-values
Female labor force participation	0.00050	0.01520
Gender difference in labor force participation	0.00001	0.00810
Female primary school completion	0.13012	0.16920
Gender difference in primary school completion	0.08773	0.08820
Female secondary school completion	0.90692	0.92410
Gender difference in secondary school completion	0.71638	0.73140
Gender attitudes index	0.00225	0.05030
Gender attitudes index among women	0.01223	0.09620
Gender attitudes index among men	0.00063	0.03040

P-values estimated using 10,000 permutations. For each outcome, the naïve p-value comes from the associated regression in a previous table. The permutation-based p-value is the fraction of permutations in which the magnitude of the estimated coefficient (from a hypothetical permutation of the gender indicator that respects the cluster structure of the language tree) exceeds the magnitude of the estimated coefficient in the true (non-permuted) data set.

Cross-Country Analysis: Coefficient Stability

Altonji *et al.* (2005) and Oster (2017) propose using robustness to observable controls to assess the magnitude of omitted variable bias

- Bias from unobservables is proportional to coefficient movements
- Coefficient movements must be scaled by changes in R^2

Cross-Country Analysis: Coefficient Stability

Altonji *et al.* (2005) and Oster (2017) propose using robustness to observable controls to assess the magnitude of omitted variable bias

- Bias from unobservables is proportional to coefficient movements
- Coefficient movements must be scaled by changes in R^2

Consider a true model:

$$Y = \alpha + \beta X + \eta W_{observable} + \gamma W_{unobservable} + \varepsilon$$

Cross-Country Analysis: Coefficient Stability

Altonji *et al.* (2005) and Oster (2017) propose using robustness to observable controls to assess the magnitude of omitted variable bias

- Bias from unobservables is proportional to coefficient movements
- Coefficient movements must be scaled by changes in R^2

Consider a true model:

$$Y = \alpha + \beta X + \eta W_{observable} + \gamma W_{unobservable} + \varepsilon$$

Data on Y, X, and $W_{observable}$ tells us:

- How much does $\hat{\beta}$ change when $W_{observable}$ is included?
- How much of the residual variation in Y is explained by $W_{observable}$?

Cross-Country Analysis: Coefficient Stability

In this framework, δ is a **proportional selection coefficient**:

> δ denotes ratio of (i) covariance between treatment and unobserved controls to (ii) covariance between treatment and observed controls

Cross-Country Analysis: Coefficient Stability

In this framework, δ is a **proportional selection coefficient**:

> δ denotes ratio of (i) covariance between treatment and unobserved controls to (ii) covariance between treatment and observed controls

Regression results with and without controls allow us to calculate:

- True causal β^* under the assumption that $\delta = 1$
- Value of δ^* that would be required for omitted variable bias from unobservables to fully explain observed association between X and Y
  - Altonji *et al.* (2005) suggest results are robust if $\delta > 1$

Cross-Country Analysis: Coefficient Stability

	OLS Coefficients			
	$\mathring{\beta}$	$\tilde{\beta}$	$\beta^*(R_{max}, 1)$	δ^*
Female LFP	-13.83	-11.92	-8.35	1.44
Gender difference in LFP	-11.61	-14.66	-17.87	3.24
Female primary completion	14.79	-6.71	-19.40	$\delta < 0$
Gender difference in primary	1.21	-3.72	-6.27	$\delta < 0$
Female secondary completion	14.52	0.43	-9.69	0.05
Gender difference in secondary	0.48	-0.86	-1.77	$\delta < 0$
Gender attitude index	-0.03	-0.12	-0.20	$\delta < 0$
Gender attitudes: women	-0.02	-0.10	-0.18	$\delta < 0$
Gender attitudes: men	-0.04	-0.14	-0.23	$\delta < 0$

Where:

- β^* = implied causal impact of X on Y if $\delta = 1$
- δ^* = implied proportional selection coefficient under null

Within-Country Analysis

Within-Country Analysis: Afrobarometer Data

Gender languages account for between 10 and 90 percent of population of Chad, **Kenya**, Mauritania, **Niger**, **Nigeria**, S. Sudan, **Uganda**

Within-Country Analysis: Afrobarometer Data

We pool Afrobarometer data from Kenya, Niger, Nigeria, Uganda:

Survey Round	**Kenya**	**Niger**	**Nigeria**	**Uganda**	**Total**
Round 2: 2002–2003	2,353	0	2,116	2,238	6,707
Round 3: 2005	1,261	0	2,120	2,345	5,726
Round 4: 2008	1,092	0	2,291	2,420	5,803
Round 5: 2011–2013	2,373	1,192	2,366	2,379	8,310
Total	7,079	1,192	8,893	9,382	26,546

Identify grammatical gender structure for 99 percent of respondents

- Respondents speak 167 different African languages

Within-Country Analysis: IHDS Data

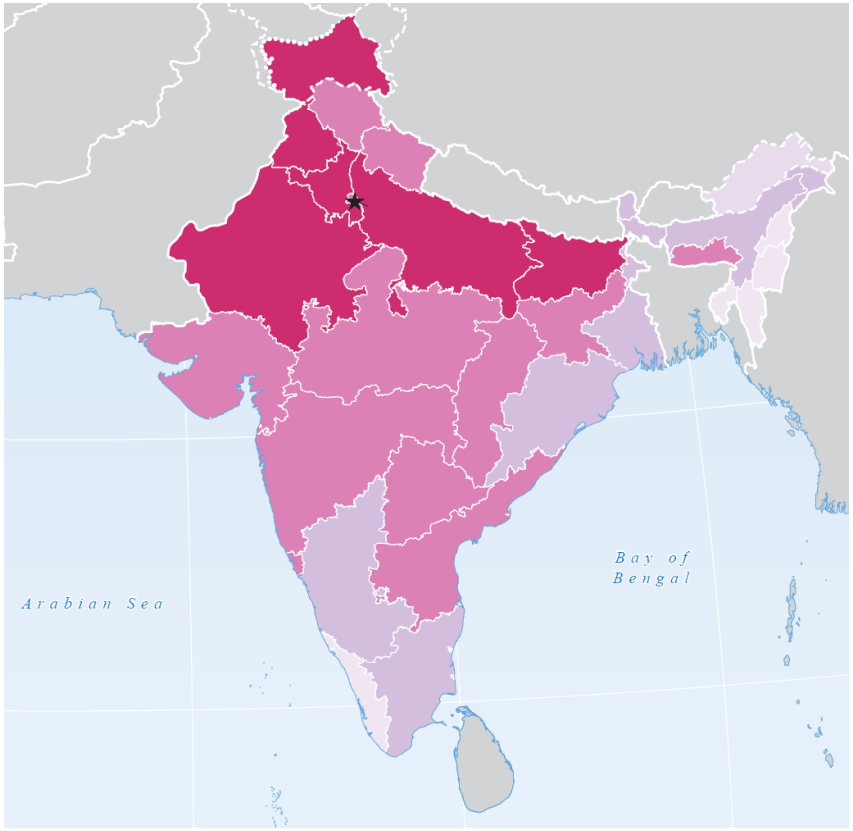

62 percent of the Indian population speaks a gender native language

- India Human Development Survey (IHDS) includes data on 75,966 household heads and spouses who speak 57 different languages

Within-Country Analysis: Empirical Specifications

When we restrict the sample to women:

$$Y_i = \alpha + \beta \mathit{Gender}_i + \zeta Z_i + \epsilon_i$$

where:

- Gender_i is an indicator for having a gender native language
- X_i is a vector of individual-level controls
  - Age, age^2, religion indicators
- Regressions of Afrobarometer data also include country$\times$round FEs
- ϵ_i is a mean-zero error term

Within-Country Analysis: Empirical Specifications

When we include data on both women and men:

$$Y_i = \alpha + \beta Gender_i + \eta Female_i + \theta Gender \times Female_i + \gamma_{country \times round} + \zeta Z_i + \epsilon_i$$

where:

- $Gender_i$ is an indicator for having a gender native language
- $Female_i$ is an indicator for being female
- $Gender \times Female_i$ is a $Gender_i \times Female_i$ interaction
- X_i is a vector of individual-level controls (age, religion, interactions)
- Regressions of Afrobarometer data also include country$\times$round FEs
- ϵ_i is a mean-zero error term

Within-Country Analysis: Sets of eight coefficients

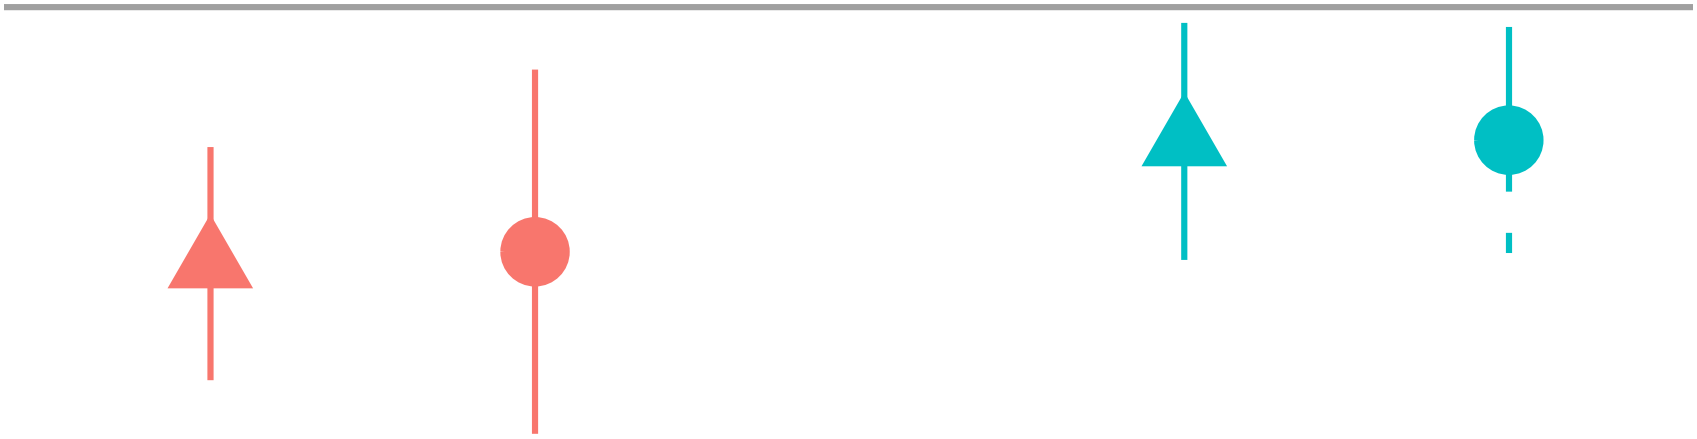

Within-Country Analysis: LFP

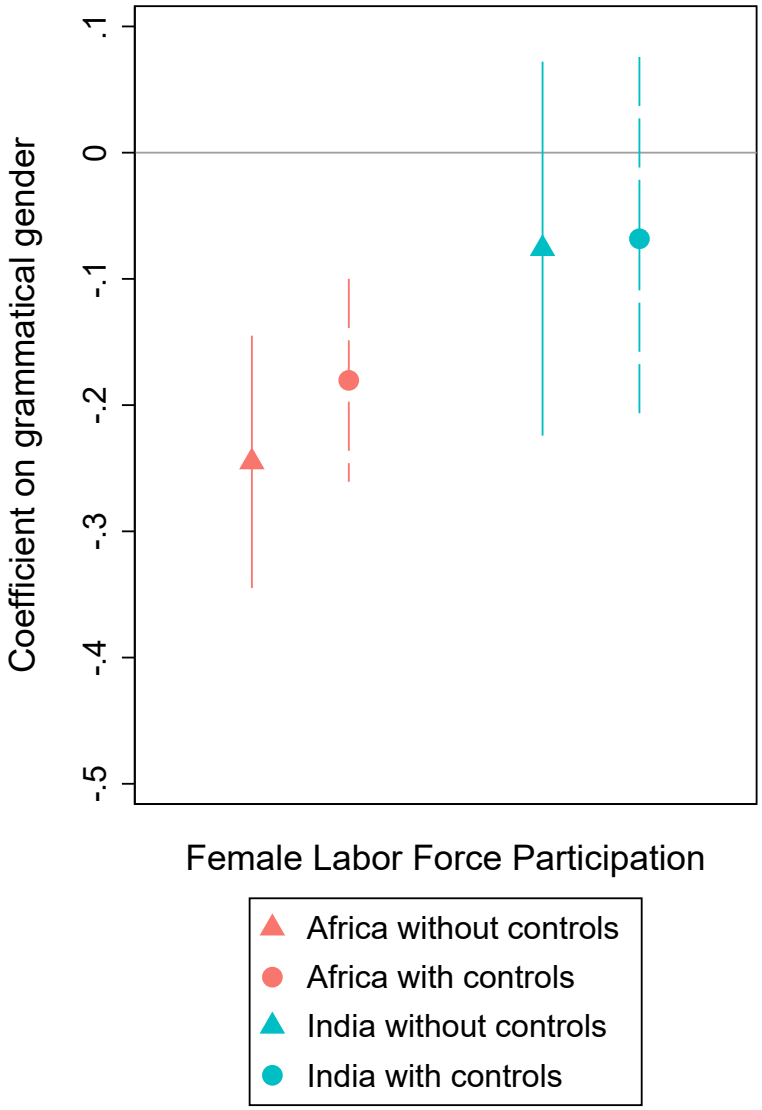

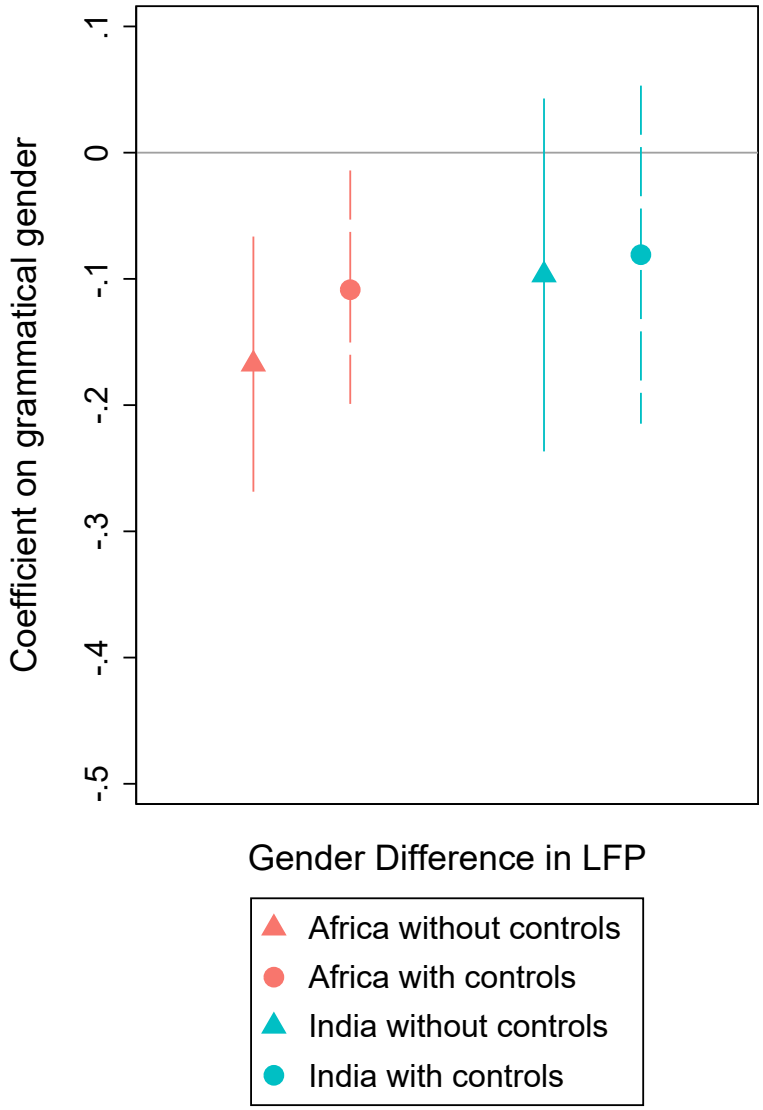

Within-Country Analysis: Primary Schooling

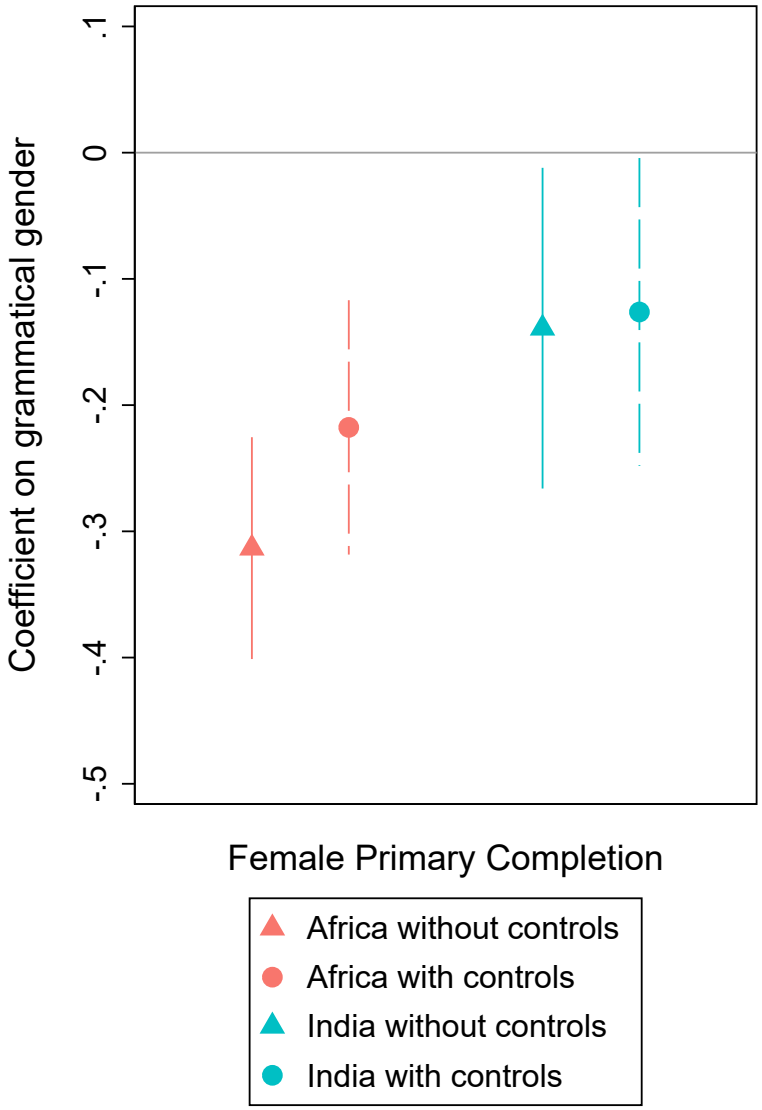

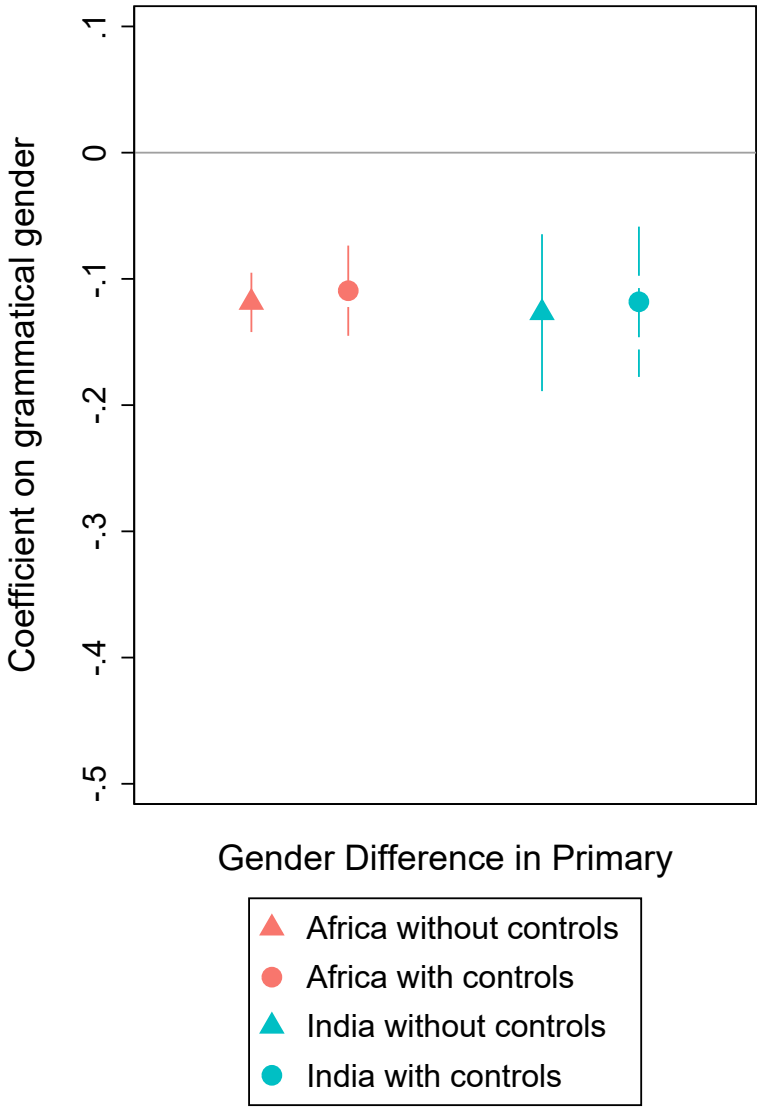

Within-Country Analysis: Secondary Schooling

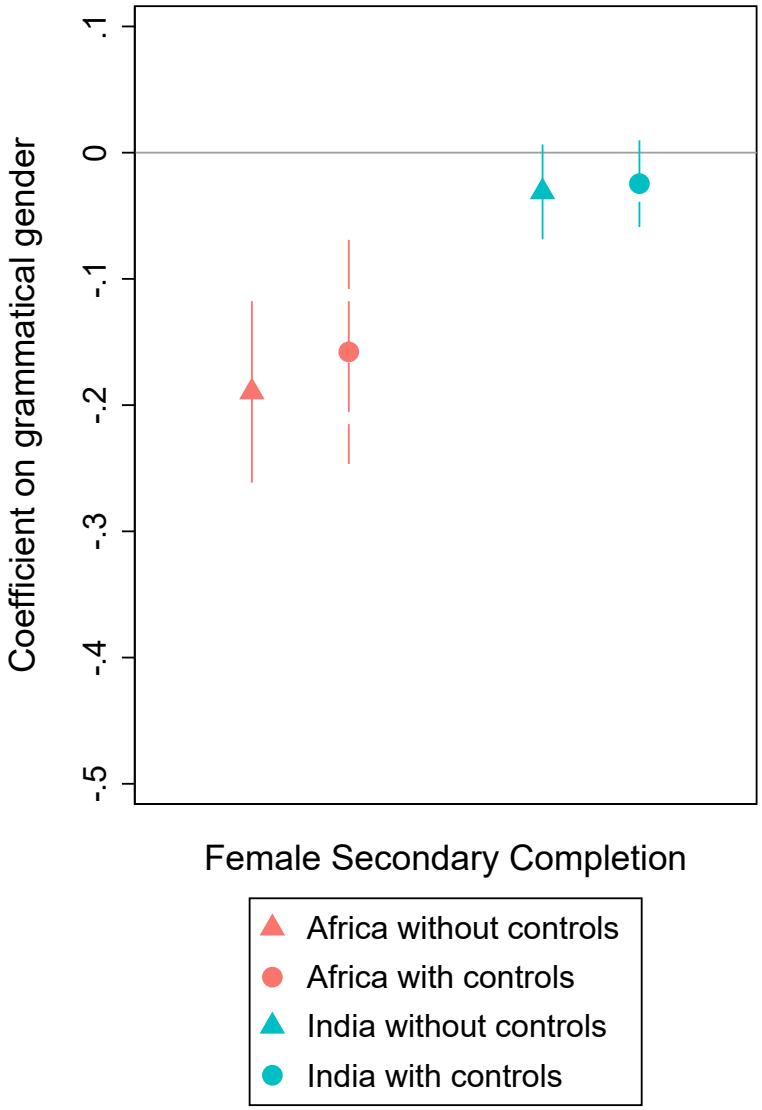

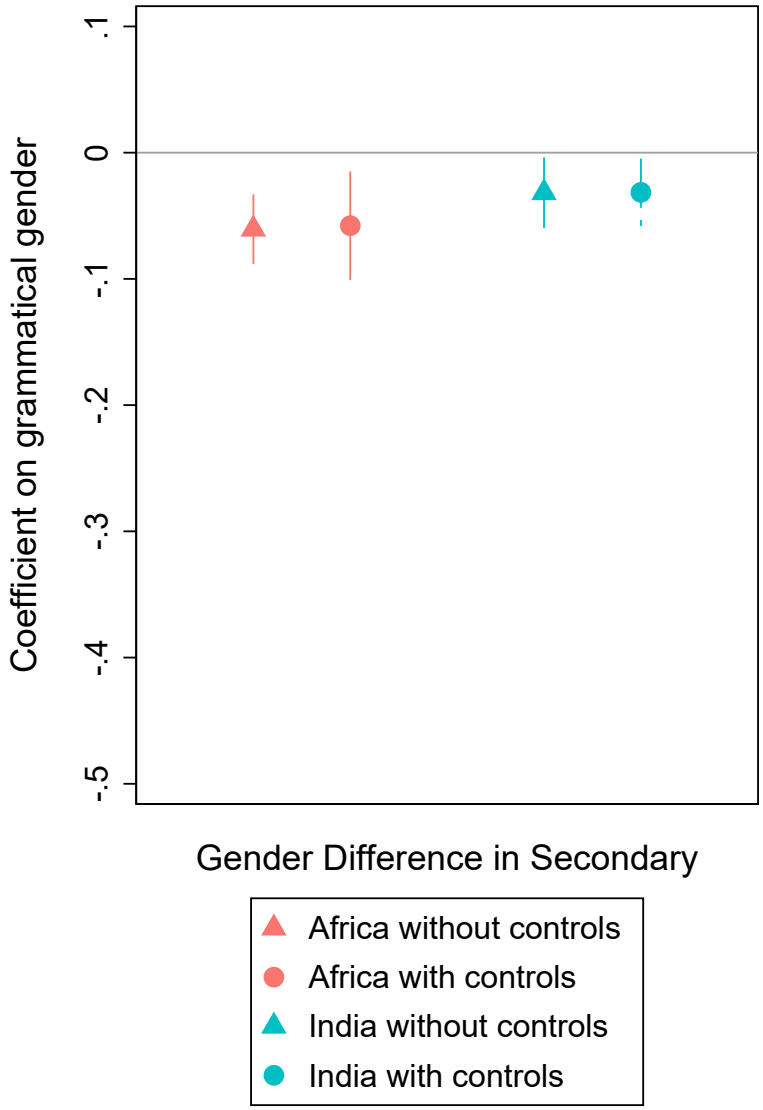

Within-Country Analysis: Results

Labor Force Participation

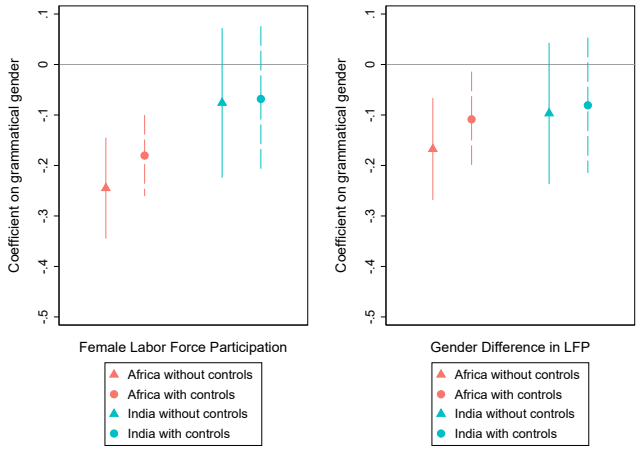

Primary Completion

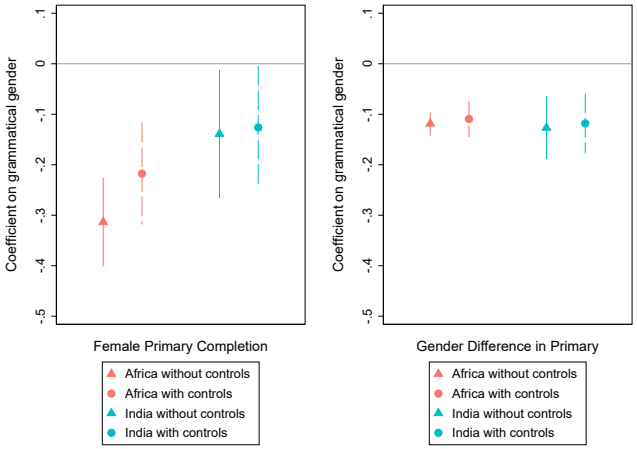

Secondary Completion

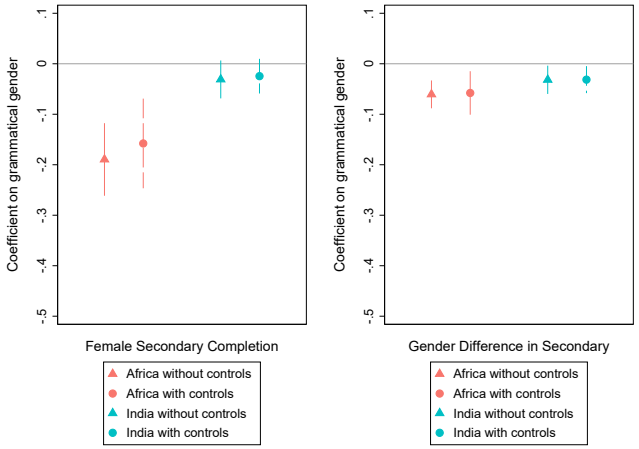

Within-Country: Coefficient Stability

	OLS Coefficients			
	$\mathring{\beta}$	$\tilde{\beta}$	$\beta^*(R_{max}, 1)$	δ^*
Panel A. Afrobarometer Data from Kenya, Niger, Nigeria, and Uganda				
In labor force (women only)	-0.24	-0.18	-0.13	2.11
Female × in labor force (pooled)	-0.17	-0.11	-0.06	1.86
Completed primary (pooled)	-0.31	-0.22	-0.15	2.18
Female × primary (Table A8)	-0.12	-0.11	-0.10	4.64
Completed secondary (pooled)	-0.19	-0.16	-0.14	3.47
Female × secondary (Table A8)	-0.06	-0.06	-0.06	6.01
Panel B. India Human Development Survey III (IHDS) Data				
In labor force (women only)	-0.08	-0.07	-0.07	11.70
Female × in labor force (pooled)	-0.10	-0.08	-0.04	1.90
Completed primary (women only)	-0.14	-0.13	-0.12	12.14
Female × primary (pooled)	-0.13	-0.12	-0.11	13.19
Completed secondary (women only)	-0.03	-0.02	-0.02	7.20
Female × secondary (pooled)	-0.03	-0.03	-0.03	25.89

Policy ramifications

Policy Implications: More Than Words

Gender matters when it shouldn't.

Policy Implications: More Than Words

Gender matters when it shouldn't.

- Bohren, Imas, and Rosenberg (2018a,b) show experimentally that randomizing the gender of the account name:

Policy Implications: More Than Words

Gender matters when it shouldn't.

- Bohren, Imas, and Rosenberg (2018a,b) show experimentally that randomizing the gender of the account name:
  - Elicits differently-toned responses (using more opinion words) when the account posing the question is female-named;

Policy Implications: More Than Words

Gender matters when it shouldn't.

- Bohren, Imas, and Rosenberg (2018a,b) show experimentally that randomizing the gender of the account name:
  - Elicits differently-toned responses (using more opinion words) when the account posing the question is female-named;
  - and elicits a lower subjective rating — fewer "upvotes" — when new users posting questions a female-named (though the pattern reverses with more established accounts).

Policy Implications: More Than Words

Gender matters when it shouldn't.

- Bohren, Imas, and Rosenberg (2018a,b) show experimentally that randomizing the gender of the account name:
  - Elicits differently-toned responses (using more opinion words) when the account posing the question is female-named;
  - and elicits a lower subjective rating — fewer "upvotes" — when new users posting questions a female-named (though the pattern reverses with more established accounts).
- Boring, Ottoboni, and Stark (2016) show that students give higher evaluation ratings to instructors whom they perceive to be male.

Policy Implications: More Than Words

Gender matters when it shouldn't.

- Bohren, Imas, and Rosenberg (2018a,b) show experimentally that randomizing the gender of the account name:
  - Elicits differently-toned responses (using more opinion words) when the account posing the question is female-named;
  - and elicits a lower subjective rating — fewer "upvotes" — when new users posting questions a female-named (though the pattern reverses with more established accounts).
- Boring, Ottoboni, and Stark (2016) show that students give higher evaluation ratings to instructors whom they perceive to be male.
  - True even when (a) the instructors are actually female (online randomization by MacNell, Driscoll and Hunt 2014)

Policy Implications: More Than Words

Gender matters when it shouldn't.

- Bohren, Imas, and Rosenberg (2018a,b) show experimentally that randomizing the gender of the account name:
  - Elicits differently-toned responses (using more opinion words) when the account posing the question is female-named;
  - and elicits a lower subjective rating — fewer "upvotes" — when new users posting questions a female-named (though the pattern reverses with more established accounts).
- Boring, Ottoboni, and Stark (2016) show that students give higher evaluation ratings to instructors whom they perceive to be male.
  - True even when (a) the instructors are actually female (online randomization by MacNell, Driscoll and Hunt 2014)
  - and when (b) the male instructors produce worse learning outcomes (natural experiment analyzed by Boring 2017).

Policy Implications: More Than Words

Interventions can leverage the salience of gender.

Policy Implications: More Than Words

Interventions can leverage the salience of gender.

- Porter and Serra (2018) show that having a female role model visit a Principles of Economics class makes female students more likely to take Intermediate Micro, and to consider majoring in economics (no impact on male students).

Policy Implications: More Than Words

Phiona Mutesa

Policy Implications: More Than Words

Interventions can leverage the salience of gender.

- Porter and Serra (2018) show that having a female role model visit a Principles of Economics class makes female students more likely to take Intermediate Micro, and to consider majoring in economics (no impact on male students).
- Riley (2018) shows that watching *Queen of Katwe* causes Ugandan secondary school students to perform better on a mathematics examination, with largest effects for female and lower-ability students.

Policy Implications: More Than Words

Interventions can leverage the salience of gender.

- Porter and Serra (2018) show that having a female role model visit a Principles of Economics class makes female students more likely to take Intermediate Micro, and to consider majoring in economics (no impact on male students).
- Riley (2018) shows that watching *Queen of Katwe* causes Ugandan secondary school students to perform better on a mathematics examination, with largest effects for female and lower-ability students.
- sadietannerconference.org
  "You can't be what you can't see" - Dr. Joycelyn Elders

Policy Implications: More Than Words

Interventions can address misperceptions about common beliefs.

- Bursztyn, González, and Yanagizawa-Drott (2018) show that

Policy Implications: More Than Words

Interventions can address misperceptions about common beliefs.

- Bursztyn, González, and Yanagizawa-Drott (2018) show that
  - Men in Saudi Arabia *believe* that other men are less supportive of female labor force participation than they really are;

Policy Implications: More Than Words

Interventions can address misperceptions about common beliefs.

- Bursztyn, González, and Yanagizawa-Drott (2018) show that
  - Men in Saudi Arabia *believe* that other men are less supportive of female labor force participation than they really are;
  - Correcting beliefs experimently makes men more willing to forgo income so that wives can participate in online job-matching;

Policy Implications: More Than Words

Interventions can address misperceptions about common beliefs.

- Bursztyn, González, and Yanagizawa-Drott (2018) show that
  - Men in Saudi Arabia *believe* that other men are less supportive of female labor force participation than they really are;
  - Correcting beliefs experimently makes men more willing to forgo income so that wives can participate in online job-matching;
  - Months later, this increases the likelihood that women have participated in a job interview outside the home.

Policy Implications: More Than Words

Interventions can address misperceptions about common beliefs.

- Bursztyn, González, and Yanagizawa-Drott (2018) show that
  - Men in Saudi Arabia *believe* that other men are less supportive of female labor force participation than they really are;
  - Correcting beliefs experimently makes men more willing to forgo income so that wives can participate in online job-matching;
  - Months later, this increases the likelihood that women have participated in a job interview outside the home.
- Patnaik (2018) shows that that the Quebec Parental Insurance Program's "daddy-only" label for some parts of parental leave

Policy Implications: More Than Words

Interventions can address misperceptions about common beliefs.

- Bursztyn, González, and Yanagizawa-Drott (2018) show that
  - Men in Saudi Arabia *believe* that other men are less supportive of female labor force participation than they really are;
  - Correcting beliefs experimently makes men more willing to forgo income so that wives can participate in online job-matching;
  - Months later, this increases the likelihood that women have participated in a job interview outside the home.
- Patnaik (2018) shows that that the Quebec Parental Insurance Program's "daddy-only" label for some parts of parental leave
  - increased fathers' use of parental leave (53 percentage pts, 200 pct);

Policy Implications: More Than Words

Interventions can address misperceptions about common beliefs.

- Bursztyn, González, and Yanagizawa-Drott (2018) show that
  - Men in Saudi Arabia *believe* that other men are less supportive of female labor force participation than they really are;
  - Correcting beliefs experimently makes men more willing to forgo income so that wives can participate in online job-matching;
  - Months later, this increases the likelihood that women have participated in a job interview outside the home.
- Patnaik (2018) shows that that the Quebec Parental Insurance Program's "daddy-only" label for some parts of parental leave
  - increased fathers' use of parental leave (53 percentage pts, 200 pct);
  - also increased fathers' long-term share of household and child-rearing responsibilities, increasing mothers' labor supply as well.

Portrait of the Acting CE as a Young[er] Man

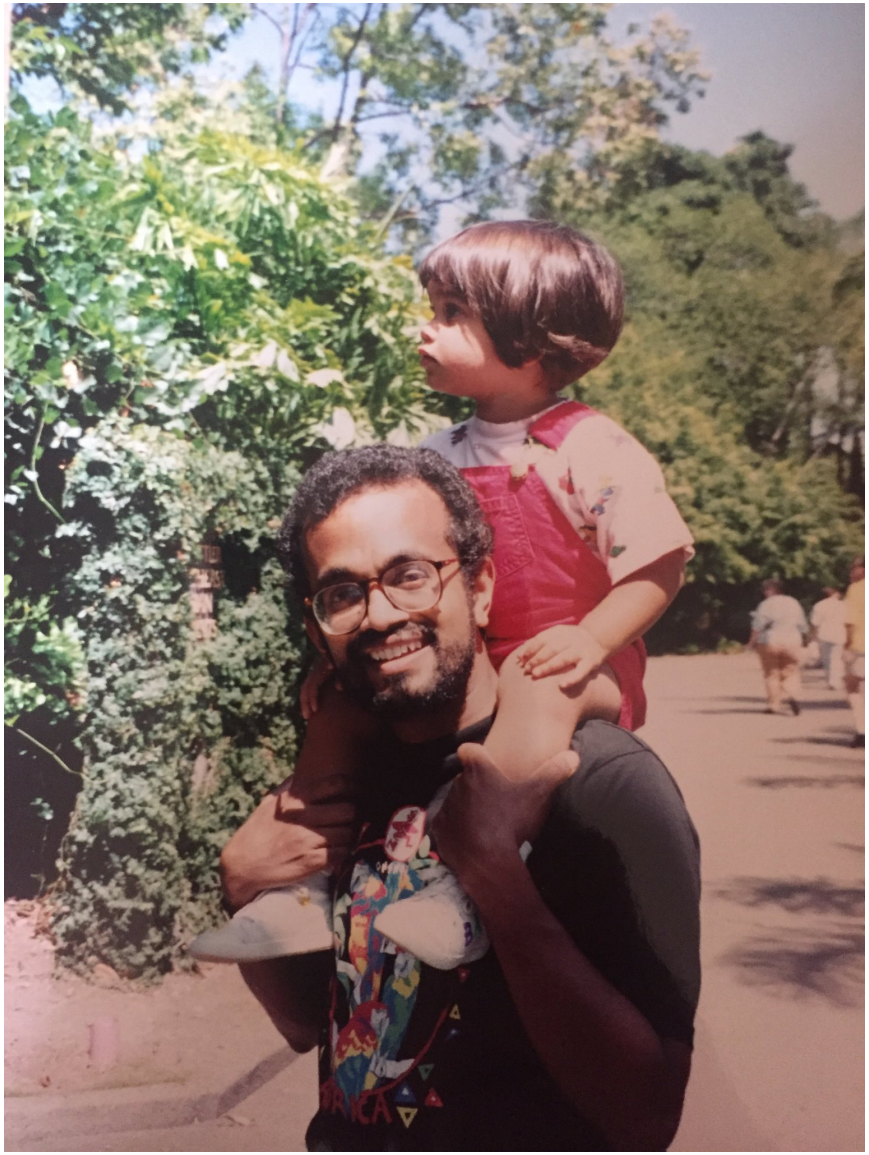

Source: @Shanta_WB

Conclusions

We characterize the grammatical gender structure of most of the world's living languages, accounting for 99 percent of the population

We present cross-country evidence that gender languages:

- Predict lower female labor force participation, gender attitudes

We present within-country evidence that gender languages:

- Predict lower female labor force participation, schooling

Languages have inherent cultural value, but they change over time; some changes result from policy choices (e.g. Académie Francaise)

- Our results suggest that linguistic choices - and many other nudges - should be seen as policy

Thank you!

Additional Slides

Marginal Impact of Dichotomous Gender Categories

Dependent variable:	$\mathbf{LFP}_f$		$\mathbf{LFP}_f$ - $\mathbf{LFP}_m$	
Specification:	**OLS (1)**	**OLS (2)**	**OLS (3)**	**OLS (4)**
Proportion (any) gender	-6.66	-7.19	4.29	-5.77
	(2.54)	(3.91)	(1.65)	(4.34)
	[0.010]	[0.068]	[0.010]	[0.185]
Proportion dichotomous gender	-10.58	-6.57	-23.44	-12.35
	(4.78)	(4.16)	(3.54)	(4.53)
	[0.029]	[0.116]	[$p < 0.001$]	[0.007]
Continent Fixed Effects	No	Yes	No	Yes
Geography Controls	No	Yes	No	Yes
Observations	178	178	178	178
R^2	0.19	0.33	0.30	0.50

Robust standard errors clustered by most widely spoken language in all specifications. $\mathbf{LFP}_f$ is the percentage of women in the labor force, measured in 2011. $\mathbf{LFP}_m$ - $\mathbf{LFP}_f$ is the difference between male and female labor force participation in 2011. Strong gender languages are those that partition the space of nouns into two gender categories, masculine and feminine. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough.

→ Return to robustness checks

Robustness to Potentially Endogenous Controls

Dependent variable:	**LFP$_f$**	**LFP$_f$ - LFP$_m$**
Specification:	**OLS**	**OLS**
	(1)	**(2)**
Proportion speaking gender language	-6.66	-10.42
	(2.80)	(2.84)
	[$p < 0.001$]	[$p < 0.001$]
Continent Fixed Effects	Yes	Yes
Country-Level Geography Controls	Yes	Yes
Observations	176	176
R^2	0.57	0.68

Robust standard errors clustered by most widely spoken language in all specifications. **LFP$_f$** is the percentage of women in the labor force, measured in 2011. **LFP$_m$ - LFP$_f$** is the difference between male and female labor force participation in 2011. Strong gender languages are those that partition the space of nouns into two gender categories, masculine and feminine. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough. Bad controls are log GDP per capita (in 2011), log population (in 2011), and the percentage Catholic, Protestant, other Christian, Muslim, and Hindu (taken from Alesina *et al.* 2013), and an indicator for former communist countries.

→ Return to robustness checks

Omitting Major World Languages

Dependent variable:	$\mathbf{LFP}_f$			$\mathbf{LFP}_f - \mathbf{LFP}_m$		
Omitted Language:	**Arabic**	**English**	**Spanish**	**Arabic**	**English**	**Spanish**
Specification:	**OLS** (1)	**OLS** **(2)**	**OLS** **(3)**	**OLS** **(4)**	**OLS** **(5)**	OLS **(6)**
Proportion speaking gender language	-6.18	-12.33	-10.10	-9.09	-15.31	-11.31
	(3.56)	(3.84)	(3.87)	(3.52)	(3.59)	(3.39)
	[0.085]	[0.002]	[0.010]	[0.011]	[$p < 0.001$]	[0.001]
Continent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Country-Level Geography Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	159	167	160	159	167	160
R^2	0.21	0.34	0.37	0.31	0.49	0.51

Robust standard errors are clustered by the most widely spoken language in all specifications; they are reported in parentheses. P-values are reported in square brackets. $\mathbf{LFP}_f$ is the percentage of women in the labor force, measured in 2011. $\mathbf{LFP}_f - \mathbf{LFP}_m$ is the difference between male and female labor force participation in 2011. Geography controls are the percentage of land area in the tropics or subtropics, average yearly precipitation, average temperature, an indicator for being landlocked, and the Alesina *et al.* (2013) measure of suitability for the plough.

→ Return to robustness checks

Primary School Completion by Continent

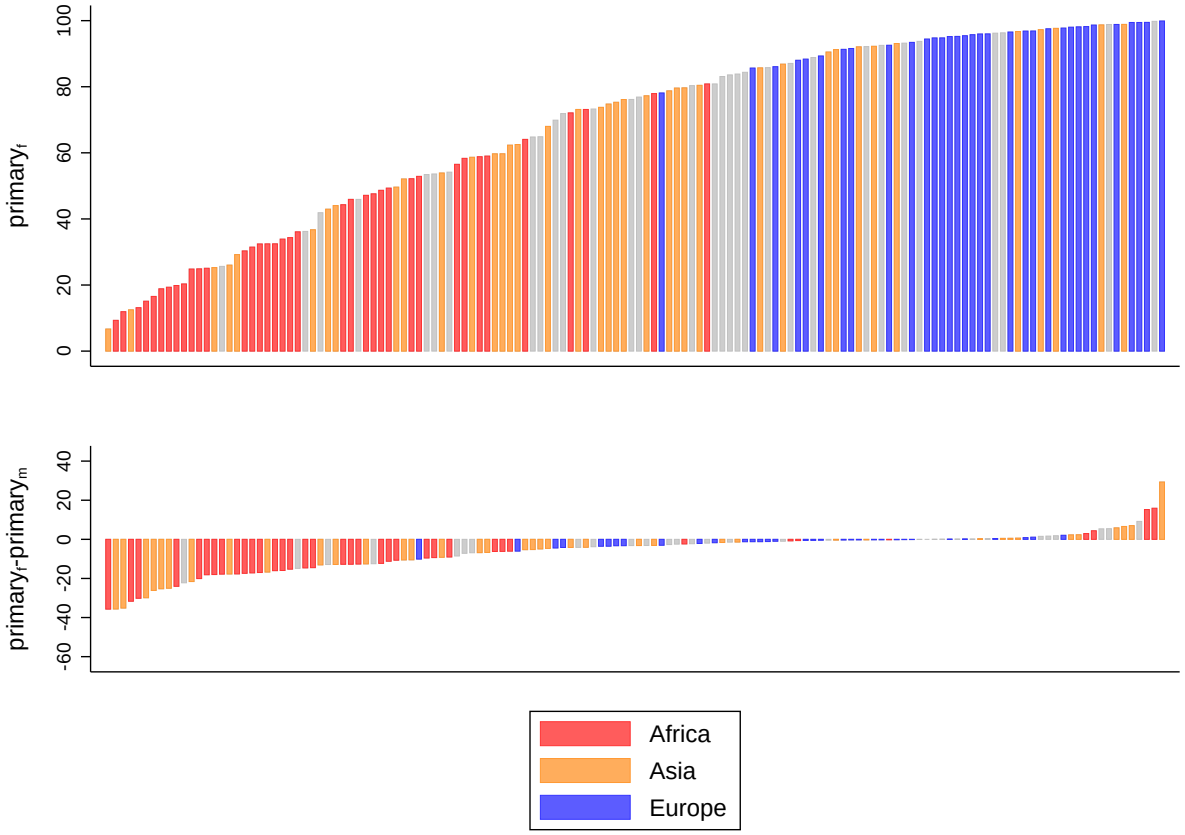

→ Return to secondary education figure

Secondary School Completion by Continent

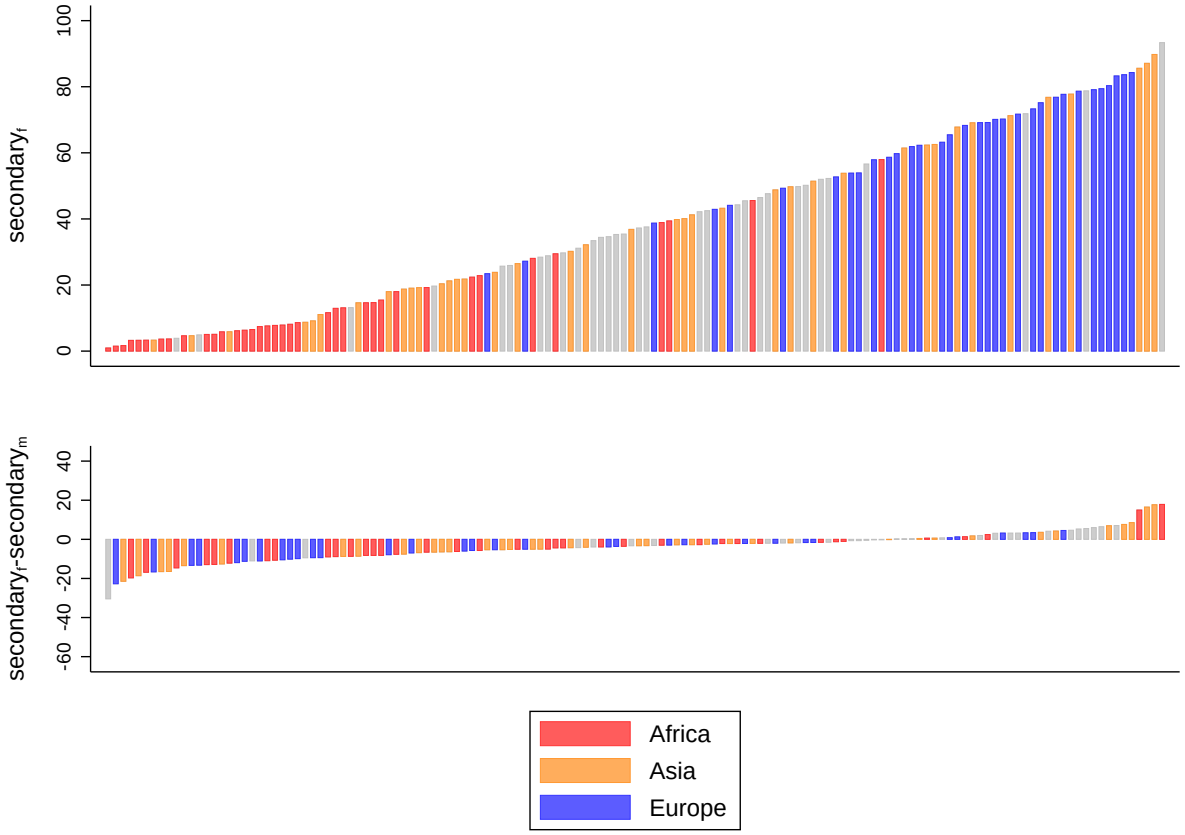

→ Return to secondary education figure

Comparing country-level WALS data to full data

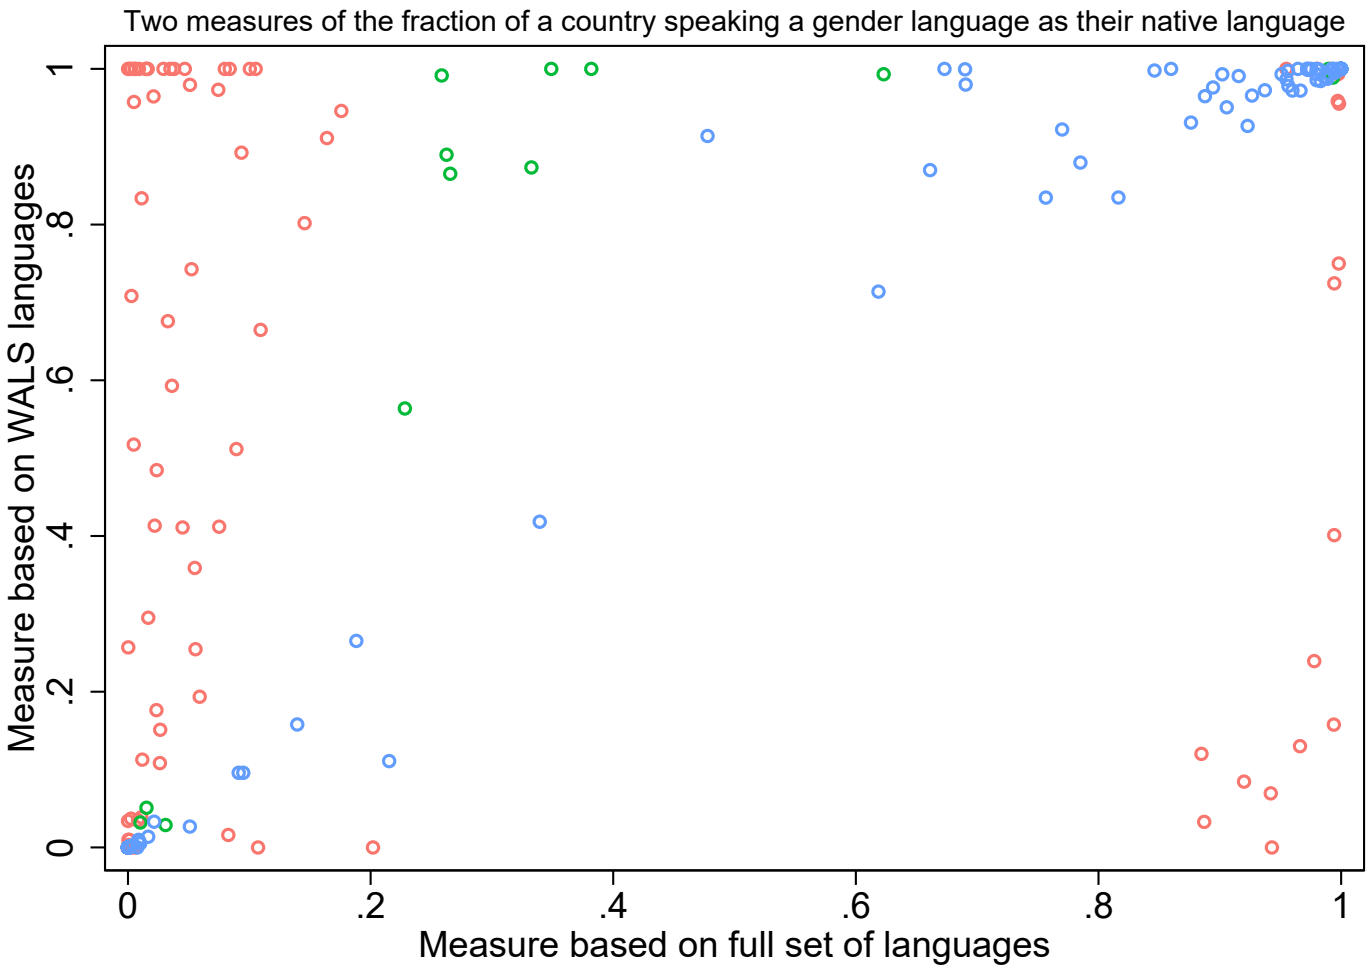

→ Return to distribution of gender languages

Cross-Country Analysis: Measurement Error

	Naïve OLS CI	Imbens-Manski CI
Female labor force participation	[−18.533, −5.305]	[−18.467, −5.013]
Gender difference in labor force participation	[−21.077, −8.233]	[−20.916, −7.741]
Female primary school completion	[−15.431, 2.010]	[−16.221, 1.673]
Gender difference in primary school completion	[−8.003, 0.559]	[−8.446, 0.432]
Female secondary school completion	[−6.901, 7.769]	[−8.261, 7.327]
Gender difference in secondary school completion	[−5.510, 3.799]	[−5.401, 3.746]
Gender attitudes index	[−0.193, −0.045]	[−0.194, −0.047]
Gender attitudes index among women	[−0.173, −0.022]	[−0.173, −0.023]
Gender attitudes index among men	[−0.214, −0.063]	[−0.215, −0.064]

footnotesizeFor each outcome, the naïve confidence interval comes from the associated regression in a previous table. The Imbens-Manski worst-case confidence interval is calculated by finding the minimum and maximum possible point estimates of the relevant coefficient based on the interval nature of the dataset (without complete data on the grammatical structure of all languages, the right-hand-side variable–the fraction of a country's population speaking a gender language–is only observed up to an interval in some cases), then by tightening the confidence interval for correct coverage following Imbens and Manski (2004).

→ Return to Manski graph