

Discussion Paper Series

IZA DP No. 18955

September 2026

Configurations of Ancestral Norms and Women's Empowerment

Ulugbek Aminjonov

Trier University
and Bordeaux School of Economics

Olivier Bargain

Bordeaux School of Economics,
IUF and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

Configurations of Ancestral Norms and Women's Empowerment*

Abstract

This paper examines how combinations of ancestral cultural norms are associated with women's empowerment across ethnic groups worldwide. Linking Demographic and Health Surveys to ancestral traits from the Ethnographic Atlas, we develop a hybrid supervised–unsupervised approach combining Partial Least Squares with k-means clustering to identify multidimensional normative configurations. We identify five probabilistic regimes spanning the empowerment gradient. Severe disempowerment is associated with an accumulation of restrictive norms, while religious norms and women's role in agriculture distinguish higher-empowerment environments within otherwise broadly patriarchal configurations. Further up the gradient, contrasts concern first hierarchical social institutions, then kinship structures and the loosening of patriarchal arrangements, before less traditional, more gender-neutral configurations become widespread at the highest levels of empowerment. The salience of norms varies with the configurations in which they are embedded, revealing strongly non-monotonic patterns. These findings remain robust to accounting for contemporary socioeconomic conditions and to a range of sensitivity checks. Overall, the analysis provides a global mapping of ancestral normative environments associated with women's agency and helps identify the configurations most closely related to severe intra-household inequality.

JEL classification

J16, O15, Z13, C38, D13

Keywords

women's empowerment, ancestral norms, cultural persistence, kinship systems, machine learning, gender inequality, DHS, Ethnographic Atlas

Corresponding author

Olivier Bargain

olivier.bargain.eco@gmail.com

* The authors used ChatGPT solely for language editing and polishing. This study has received financial support from the French State in the framework of the Investments for the Future program IdEx Université de Bordeaux / GPR HOPE. We are grateful to colleagues at Namur University, TU Munich, the London School of Economics and the Paris School of Economics for comments. The usual disclaimer applies.

1. Introduction

Recent years have seen rapid growth in the economics of culture, reflecting broad recognition that deeply rooted traditions continue to shape contemporary social and economic outcomes (Nunn, 2020; Baland et al., 2020). Gender roles and women's empowerment are areas where these legacies are especially salient (Lowes, 2020; Jayachandran, 2015, 2021; Bau and Fernández, 2023; Giuliano, 2020). A substantial body of work examines ancestral institutions, such as patrilineal descent (La Ferrara, 2007; La Ferrara and Milazzo, 2017; Lowes, 2022), patrilocal residence (Bargain et al., 2022; D'Agostini and Tiberti, 2025; Aminjonov et al., 2024), marriage payments (Corno et al., 2020; Calvi & Keskar, 2021), polygyny (Fenske, 2015), and historical agricultural practices (Alesina et al., 2013, 2021; Becker, 2025).

This literature documents persistent associations with women's agency, labor market participation, and access to resources, while also shedding light on underlying mechanisms. Ancestral norms may shape women's cooperative and bargaining behavior, intra-household sharing and agency (Lowes, 2020; Baland et al., 2024), as well as their control over resources, assets, and income-generating activities (Calvi and Keskar, 2021; La Ferrara, 2007; Alesina et al., 2013). Yet several features are worth highlighting.

First, individual norms do not map unambiguously into women's empowerment. The implications of particular practices may depend on the rights and resources they confer on women and on the broader normative environment in which they operate. Dowry, for instance, is often viewed as part of a restrictive marriage system but can also increase women's bargaining power when transferred assets remain under their control (Calvi and Keskar, 2021).² This ambiguity motivates looking beyond individual practices to the broader normative environments in which they operate.

Second, and relatedly, most studies examine one ancestral norm at a time. This approach is well suited to identifying the association or causal effect of a particular practice, but is less informative about which practices tend to coexist and how they combine into the normative environments actually observed across societies. Some studies exploit the richness of specific local contexts to examine interactions among multiple norms (e.g., Guirking and

² Similar ambiguity surrounds kinship and post-marital residence. Beyond descriptive anthropological evidence, relatively few economic studies establish an empowering role for matrilineal arrangements (Aminjonov et al., 2024). Patrilineal descent likewise encompasses considerable variation in women's authority and autonomy. Contemporary Western societies, for instance, retain patrilineal conventions such as paternal surname transmission despite otherwise largely bilateral kinship systems (Goody, 1983).

Platteau, 2020), but remain geographically circumscribed. Others provide evidence at a broad regional or global scale, but generally focus on particular cultural practices or a limited set of normative dimensions (see [Table A1](#)).³ What remains largely unexplored is the global organization of ancestral norms into multidimensional configurations associated with women’s empowerment.

Third, such a configurational perspective requires broadening the set of practices considered jointly. Existing empirical work has naturally focused on relatively specific and measurable economic mechanisms, such as women’s labor supply, control over assets, or intra-household bargaining. Yet ancestral cultural environments also encompass dimensions less readily summarized by these channels, including social hierarchy, community organization, religious beliefs, and symbolic practices. Incorporating these dimensions allows us to characterize the wider normative environment within which the better-understood economic mechanisms operate.

Against this background, we seek to identify which configurations of ancestral norms are associated with women’s empowerment and disempowerment, with particular attention to the most severe forms of intra-household inequality. Rather than assigning a universal role to individual practices, we characterize the probabilistic normative environments in which they co-occur across ethnic groups worldwide. We further examine whether, within broadly similar normative environments, particular norms emerge as pivotal margins, distinguishing configurations associated with markedly different levels of women’s empowerment.

We implement this approach by combining contemporary measures of women’s decision-making with a broad set of ancestral cultural traits at the ethnic-group level. Our empirical strategy links 188 country-year waves of the Demographic and Health Surveys (DHS), covering around two million women, to historical cultural traits from the Ethnographic Atlas. We analyze these data using a hybrid supervised–unsupervised learning approach that combines Partial Least Squares (PLS) with k-means clustering. This procedure identifies distinct normative environments associated with different levels of women’s empowerment. Given the diversity of norm combinations across ethnic groups, these environments are best

³ [Table A1](#) reviews Ethnographic Atlas–based research on the link between women’s outcomes and ethnic norms. In particular, it highlights the few studies examining interactions among norms, or relatively broad normative mixes, including complementarities between kinship and residence rules ([Dessy et al., 2023](#); [D’Agostini and Tiberti, 2025](#)) and combinations of marriage, residence, and women’s productive roles ([Alesina et al., 2021](#)).

understood as probabilistic regimes, characterized by a higher prevalence of particular sets of ancestral practices rather than by fixed cultural configurations.⁴

The resulting framework provides a structured global mapping of how inherited normative environments are associated with contemporary intra-household gender dynamics. Our analysis identifies five probabilistic mixtures of ancestral norms associated with markedly different average levels of women's empowerment. Many norms emerge as salient, but at different points along the empowerment gradient. The most severe forms of women's disempowerment are associated with tightly bundled environments where patriarchal family arrangements co-occur with broader unequal social structures. Contrasts between the lowest-empowerment regimes are particularly informative: broadly similar patriarchal kinship and marital arrangements coexist with different empowerment levels, distinguished by the prevalence of male-dominated agriculture and moralizing gods. Further up the gradient, the main contrasts concern hierarchical social institutions, followed by kinship structures, with matrilineal and matrilineal arrangements becoming more prevalent before more gender-neutral settings emerge at the highest empowerment levels. Traditional practices such as polygyny and bride price nonetheless remain common. These contrasts are not uniformly monotonic: plough agriculture, for example, recurs in otherwise distinct configurations at both low and high empowerment levels. Their salience thus depends on the broader configurations in which they are embedded, underscoring the importance of examining combinations of practices rather than isolated traits.

Our approach also yields insights into geographic diversity, cultural persistence, and the robustness of the results. Each world region contains multiple clusters, revealing substantial within-region heterogeneity that defies simple regional or religious stereotypes. We also show that ancestral norm mixtures retain substantial explanatory power even after controlling for contemporary factors, including country wealth, inequality, women's labor force participation, broad regions, and religions, which indicates that the clusters capture historical structures not reducible to present-day circumstances. As an alternative to

⁴ This representation is consistent with work in institutional economics, cultural transmission, and cultural evolution emphasizing complementarities among norms and the persistence of mutually reinforcing institutional arrangements (Greif, 1994; Boyd and Richerson, 1985; Bisin and Verdier, 2000, 2001; Aoki, 2001; Henrich, 2001). Note also that our approach remains deliberately agnostic with respect to detailed historical and anthropological narratives. As such, it is consistent with both structuralist traditions, which highlight how coherent bundles of norms shape gender relations, and more recent anthropological perspectives emphasizing hybridity, historical contingency, and practice-based dynamics in the evolution of social systems. Its interpretation could nonetheless be enriched by anthropological insights into how kinship, descent, and marriage systems form interdependent structures that shape gendered authority.

controlling ex post for these factors, we also residualize empowerment ex ante and obtain cluster patterns similar to the baseline.

We perform additional sensitivity checks. The main conclusions are robust to potential matching bias, alternative clustering algorithms, and the pre-grouping of individual norms into broader domains. The precise composition of the probabilistic regimes naturally varies with the number of clusters, but the central patterns remain. We find that differences across regimes extend to other dimensions of women's outcomes, including attitudes toward domestic violence. Conversely, we use attitudes towards domestic violence as an alternative supervised outcome in constructing the regimes and recover broadly similar patterns. Finally, we use continuum regression to vary the degree of supervision in constructing the cultural space. Although the partitions change appreciably, the principal religious and agricultural contrasts at the lower end remain visible, suggesting that their prominence is not solely attributable to the outcome-oriented nature of PLS.

The contributions of this paper are threefold. *First*, we provide the first global mapping of cultural regimes relevant to women's empowerment. While most existing studies examine specific norms or subsets of norms (see [Table A1](#)), our approach characterizes probabilistic normative mixtures derived from the full set of relevant ancestral traits and reveals the configurations most closely related to women's outcomes. This allows us to identify which traits play pivotal roles within these configurations, situating well-studied practices such as matrilineality and plough agriculture within broader cultural contexts and bringing to light the importance of other factors such as religious norms.

Second, our methodological design, a hybrid supervised–unsupervised learning framework, can be replicated to characterize multidimensional cultural environments in relation to many other outcomes. In the present context, it allows us to extract multidimensional ethnic-level normative configurations aligned with women's empowerment. Unlike PCA, which identifies components that maximize variation in cultural norms alone, PLS constructs components that also maximize their covariance with the outcome; k-means then identifies recurrent configurations within the resulting empowerment-relevant cultural space. Alternative methods, such as linear regression, hierarchical Lasso, or random forests, can estimate the predictive contributions of individual norms and their interactions, but they do not directly produce a low-dimensional space of co-occurring norms that can be partitioned into interpretable probabilistic normative regimes. Our approach provides a scalable and portable method to study multidimensional cultural environments in relation to any outcome for which comparable data exist.

Finally, the approach is not only analytically meaningful but also potentially policy-relevant, since policy design is often constrained by limited information ([Brown et al., 2019](#)). Reaching deprived individuals, rather than only poor households, requires observable indicators that go beyond conventional poverty-targeting criteria, such as those used in proxy means tests, and include characteristics associated with intra-household inequality. In multi-ethnic societies especially, culturally informed indicators could complement socioeconomic targeting criteria by helping to identify groups facing greater risks of intra-household disadvantage and the normative configurations most closely associated with women’s relative position.

The remainder of the paper is organized as follows. Section 2 describes the data, Section 3 the empirical strategy, Section 4 the results and sensitivity analyses, and Section 5 concludes.

2. Data

2.1. Women’s Empowerment

Demographic and Health Surveys (DHS). For a large-scale, cross-country characterization of women’s empowerment, it is essential to rely on homogeneous and comparable measures of women’s outcomes. We therefore use data from the Demographic and Health Surveys (DHS), which provide nationally representative and harmonized information on women’s socio-economic status, health, gender-related attitudes, and intra-household decision-making across low- and middle-income countries (LMICs). At the time of analysis, the DHS program had implemented 358 country-wave surveys across 84 countries, with multiple waves for many countries and roughly 85% of surveys conducted after 2000. We pool all available waves to maximize spatial and temporal coverage, expand sample size and increase precision.

Measuring Empowerment. The DHS provide several subjective and self-reported indicators relevant to women’s agency, including measures of intra-household decision-making, attitudes toward domestic violence, well-being, and health. We primarily focus on decision-making variables, which capture cultural variation in women’s authority within the household and are therefore well suited to examining associations with ancestral norms.⁵

⁵ By contrast, distal outcomes such as nutrition, health, and broader well-being indices depend more on household resources and contemporaneous constraints, making them less suitable for examining the expression of traditional gender norms. Later, we nonetheless examine the sensitivity of our normative clusters using alternative outcomes, notably women’s attitudes towards domestic violence. We also considered estimating women’s resource shares as a proxy for bargaining power using recent approaches in the estimation of collective household models. [Aminjonov et al. \(2026\)](#) show that this is feasible across multiple countries, but

Specifically, we use DHS “final say” questions, which indicate who holds authority over relevant domains and are widely employed as proxies for women’s empowerment due to their comparability across surveys and cultural contexts. These questions may sometimes capture delegation of power rather than genuine agency (Baland et al., 2026). Thus, we focus on domains that more clearly reflect strategic life choices, those critical for a woman to “live the life she values” (Kabeer, 1999). We retain (i) decision-making over access to health care, (ii) freedom of movement (proxied by the ability to visit family or relatives), and (iii) participation in major household purchase decisions.⁶ These three domains are widely used in the literature to construct women’s agency indices based on DHS data (e.g., Ewerling et al., 2017; Miedema et al., 2018). Each indicator is coded as follows: for (i) and (ii), it takes the value 1 if the woman reports making the decision independently, and 0 otherwise; for (iii), it takes the value 1 if the woman decides either alone or jointly, and 0 otherwise. We construct each woman’s empowerment index as the mean of the three decision-making indicators.

Selection. We restrict the sample to country-waves in which all three empowerment items are available; since these items are typically missing jointly, this effectively amounts to requiring availability of the empowerment module. The resulting sample includes 2,344,672 ever-married women aged 15–49 from 71 LMICs, representing 85% of the initial global sample. **Figure A1** illustrates the geographical coverage in this sample (in blue). The dataset spans diverse cultural and regional contexts, covering approximately 80% of LMICs in Africa, 50% in Asia and the Caucasus, with additional representation from Latin America and the Middle East. This broad cross-country coverage enables us to capture substantial variation in both cultural norms and women’s empowerment outcomes across settings. **Figure A2** maps the average empowerment index defined above. The figure reveals substantial variation both across and within global regions, underscoring the relevance of contextual factors associated with women’s autonomy. These descriptive patterns also motivate our focus on how mixtures of cultural norms may help explain observed differences in empowerment outcomes.

in practice only for a more limited set of settings and with a relatively parsimonious sharing rule, as richer specifications can lead to substantial estimator instability.

⁶ We do not include control over own earnings because it is only defined for women who have income, making it structurally missing where female labor force participation is low and therefore potentially confounding empowerment with employment, or in contexts of self-consumption.

2.2. Cultural Norms

To examine how historical cultural norms relate to contemporary patterns of women's empowerment, we draw on the Ethnographic Atlas (EA), which documents the pre-industrial social, economic, and kinship structures of 1,267 ethnic societies worldwide (Murdock, 1967). The EA is widely used to study the long-term influence of ancestral cultural traits on contemporary outcomes, often in combination with DHS data (e.g., Alesina et al., 2013; Michalopoulos et al., 2019).⁷

The EA provides systematic information on ancestral practices central to family organization and gender relations, including the following *normative domains*: post-marital residence (e.g., patrilocality), kinship and descent systems (e.g., patrilineality), marriage forms (e.g., polygyny) and marriage payments (bride price, dowry), as well as women's economic role (e.g., their participation in agriculture, relative to men's) and subsistence and agricultural practices (e.g., pastoralism and plough use). These domains have direct conceptual or well-established empirical links to gender roles and women's position.

A second set of EA traits captures broader social and cultural institutions, such as community organization and governance (e.g., settlement patterns and community segmentation), systems of social stratification (e.g., slavery), and ritual or religious practices. These are not intrinsically gender-specific but may nevertheless affect women's position. For instance, slavery traditions may legitimize unequal social status, sometimes with particularly adverse consequences for women (Corno et al., 2020), while male circumcision or initiation practices may confer symbolic status and authority disproportionately on men, thereby reinforcing gender hierarchies (Silverman, 2004).

Finally, the EA contains other dimensions, such as jurisdictional hierarchy, for which there is no clear conceptual or established empirical connection to gender relations or women's empowerment. These are not retained in our baseline analysis.

3. Empirical Approach

Our empirical strategy proceeds in three steps. First, we link DHS households to their corresponding ancestral societies in the EA using information on contemporary ethnicity,

⁷ **Table A1** summarizes applications focusing on women's outcomes and related gender dimensions. Alesina and Giuliano (2015) discuss the role of EA kinship, descent, and agricultural variables in the economics of culture; Giuliano (2020) reviews the use of historical datasets including the EA; Lowes (2020) discusses the incorporation of ethnographic data into historical economics; and Bahrami-Rad et al. (2022) combine several EA variables to construct synthetic indices of kinship structure, most notably the Kinship Intensity Index.

language, and geographic location. This allows us to assign historically rooted cultural norms to present-day individuals. Second, we select and harmonize a set of gender-relevant ancestral norms based on both theoretical relevance and empirical associations with women’s empowerment, ensuring comparability across contexts. Third, we employ a hybrid supervised–unsupervised machine-learning approach to identify probabilistic mixtures of ancestral norms that are systematically associated with varying levels of women’s empowerment.

3.1. Matching DHS Households to Ancestral Cultural Norms

We match DHS households to their ancestral norms in the EA using a hierarchical matching procedure, combining ethnic names, linguistic classifications, and geographic information. Our approach draws on and integrates the matching strategies adopted in previous studies linking DHS data and the EA (Nunn & Wantchekon, 2011; Giuliano & Nunn, 2018; Michalopoulos et al., 2019; Ashraf et al., 2020; Alesina et al., 2021). An extensive description of this procedure is provided in Appendix A2.

The matched sample includes 669 distinct ethnicities for which both DHS-based empowerment measures and EA-based ancestral norms are observed, covering nearly all countries in the initial DHS selection (70 out of 71; see [Table A2](#)). [Table A3](#) summarizes the matching procedure: 49% of matches rely on individual-level DHS information (ethnic-group names and/or language), while 50% rely on geographic matching (GPS coordinates or subnational region).⁸

To reduce sampling noise in ethnic-group empowerment measures, we restrict the sample to ethnic groups with at least 50 women observed in the microdata (excluding 70 ethnic groups) and non-missing empowerment outcomes (excluding 54 ethnicities). The final sample comprises 545 ethnic groups, of which 284 (52%) are observed in one country, 163 (30%) in two countries, and 98 (18%) in three or more countries. Despite this cross-border coverage, ethnic and national contexts remain closely intertwined in our sample.

3.2. Pre-Selection and Categorization of Cultural Norms

Following the discussion in Section 2.2, we retain ancestral traits that are either directly related to gender roles or belong to broader social institutions with plausible implications for women’s position, while excluding EA dimensions with no clear conceptual connection

⁸ We provide sensitivity analyses of the clustering approach below.

to gender relations (e.g., jurisdictional hierarchy). This *pre-selection* is intended primarily to reduce complexity and improve the interpretability of the resulting cultural profiles.

We then *harmonize* the level of detail across normative domains. The original EA classifications vary substantially in granularity, with some domains containing many closely related categories. Feeding these categories directly into the analysis would result in an uneven representation of normative domains and make the resulting profiles more difficult to interpret. We therefore collapse closely related categories within each domain into a small number of substantively meaningful groups. **Table 1** presents, for each domain, the mapping from the original EA categories to the broader groups used in the clustering analysis. We retain two or three groups per domain, with a reference category shown in grey, which encompasses gender-neutral arrangements or the absence of a practice.

These groupings are primarily guided by the substantive meaning of the EA categories and the existing literature. For example, in the domain of post-marital residence, we collapse detailed categories into patrilocal (including virilocal), matrilineal (including uxorilocal), and a more gender-neutral category (combining ambilineal and neolocal residence). These groupings are consistent with patterns established in the literature, including higher women’s empowerment under ambilineal/neolocal than under traditional arrangements, and under matrilineal relative to patrilocal systems ([Aminjonov et al., 2024](#)).

Where the appropriate grouping is ambiguous, we additionally use the empirical relationship with empowerment to guide the aggregation. Specifically, within each domain, we regress the empowerment measure on the detailed EA categories and group categories displaying similar associations with empowerment. For instance, within the “high gods” domain, the category “supportive of human morality” is negatively associated with empowerment, while the other categories display positive associations, leading us to distinguish “moralizing gods” from “other or no gods”.⁹

While this grouping harmonizes the level of detail across normative domains, providing a parsimonious and interpretable simplification of the original EA categories, we later assess the sensitivity of our results to this choice by re-estimating the full procedure using the original unaggregated categories.¹⁰

⁹ In this case, the negative association may reflect the fact that moralizing religions are often embedded in institutional structures and behavioral prescriptions that constrain women’s autonomy.

¹⁰ The detailed procedure, code, and intermediate results underlying these groupings are available in the project [repository](#).

Table 1. Pre-categorization of Ancestral Norms

Domain	Resulting Categories	Original EA Categories
Kinship system	Patrilineal Matrilineal Neutral or others	Patrilineal Matrilineal Duolateral; quasi-lineal; ambilineal; bilateral; mixed
Post-marital residence	Patrilocal Matrilocal Neutral or others	Patrilocal; virilocal; optionally patrilocal; avunculocal Matrilocal; uxorilocal; optionally uxorilocal Ambilocal; neolocal; no common residence
Marriage payments	Bride price Dowry Reciprocal, none or others	Bride price or wealth transferred to the bride's family; bride service to the bride's family; token bride price; sister or female relative exchanged for the bride. Dowry Absence of consideration; Gift exchange, reciprocal
Domestic structure	Extended Non-extended or other	Polygynous: unusual co-wives; polygynous: usual co-wives; minimal extended families; small extended families; large extended families Independent nuclear family (monogamous); independent nuclear family (polygynous); independent polyandrous families
Marriage form	Polygyny, general or sororal Monogamy/Polyandrous or other	Independent nuclear, polygyny; preferentially sororal, same dwelling; preferentially sororal, separate dwelling; non-sororal, separate dwellings; non-sororal, same dwelling Independent nuclear, monogamous; independent polyandrous families
Male circumcision	Male circumcision No male circumcision	Within 2 months after birth; 2 months to 2 years; 2 to 5 years; 6 to 10 years; 11 to 15 years; 16 to 25 years; 25 to 50 years; after 50 years; normal age unclear Absent
High gods	Gods supportive of human morality No or other gods	Supportive of human morality Absent or not reported; not active in human affairs; active but not supportive
Community segmentation	Segmented community Non-segmented or other	Segmented communities; Segmented communities without local exoga Demes, not segregated into clan barrios; agamous communities; exogamous communities; clan communities or clan barrios
Slavery	Slavery No or near absent slavery	Incipient or nonhereditary; reported but type not identified; hereditary and socially significant Absence or near absence
Participation in agriculture	More male participation in agriculture More/equal female part. or no agriculture	Males only; males appreciably more; differentiated Equal participation; female appreciably more; females only; no agriculture (or unimportant activity)
Animal and plow cultivation	Animal and plow cultivation No or non-aboriginal plow	Aboriginal to prior contact Absent; not aboriginal
Settlement type	Non-sedentarity Sedentarity or other	Nomadic or fully migratory; seminomadic; semisedentary Compact but impermanent settlements; neighborhoods of dispersed family homes; separated hamlets forming a single community; compact and relatively permanent settlements; complex settlements

Note: Ethnographic Atlas initial categories within each normative domain are grouped into broad categories, starting with the set of more traditional traits (when there are three broad categories: the first one tends to be the more patriarchal set of traits and the second more likely in favor of women's empowerment) and ending with more modern, gender-neutral traits (or other, unclassified and possibly neutral traits). The latter category, in light grey, will be used as reference (omitted) category when presenting main results in Figure 1. When classification was not intuitive, we have used simple linear regressions to check which initial categories were more positively or negatively correlated with empowerment. When domains cover several dimensions that overlap (for instance in the case of "domestic structure" and "marriage form"), we have designed broad categories that focus on one dimension (e.g., "domestic organization" on extended vs. non-extended families, and "marriage form" on polygamous vs. non-polygamous).

3.3. Identifying Ancestral Norm Mixtures

This last step constitutes our main contribution and addresses our central objective: to identify combinations of ancestral norms that tend to co-occur across ethnic contexts with varying levels of women’s empowerment. The methodological challenge is to reduce a multidimensional space of correlated ancestral norms while preserving the dimensions most relevant to women’s empowerment. This challenge is compounded by the sparsity of the cultural space: even after harmonization, our 12 domains imply $3^3 \times 2^9 = 13,824$ possible norm configurations for only 545 ethnic groups. Dimension reduction is therefore particularly useful for identifying interpretable patterns of similarity across otherwise heterogeneous configurations.

To address this challenge, we combine supervised dimension reduction with unsupervised clustering: the former constructs a parsimonious cultural space oriented toward empowerment, while the latter identifies recurrent normative configurations within that space. We proceed in two stages. First, we construct the cultural space using *Partial Least Squares* (PLS), which extracts linear combinations of ancestral norms that maximize covariance with women’s empowerment. Second, we apply *k-means clustering* to the resulting PLS components to group ethnicities into recurrent cultural mixtures. This hybrid supervised–unsupervised design distinguishes our approach from existing empirical work, which typically relies on either single-norm regressions or unsupervised classifications unrelated to gender outcomes.

Step 1: Extracting Latent Cultural Components (Supervised Dimension Reduction).

PLS is particularly well suited to our setting because ancestral norms are multidimensional and highly correlated, while our objective is to identify combinations specifically related to women’s empowerment. Unlike PCA, which extracts directions that maximize variation in the predictors alone, PLS extracts latent components that maximize covariance between the norms and empowerment. It therefore provides a parsimonious, outcome-oriented representation of the cultural space while accommodating multicollinearity among norms. Appendix A2 provides further discussion and comparisons with alternative dimension-reduction approaches.

Formally, let y_g denote the average empowerment index in ethnic group g (based on health care, family visits, and major household purchases), and let X_g denote the vector of ancestral norms for group g , consisting of 15 binary indicators representing the non-reference categories across the 12 cultural domains described in [Table 1](#). PLS constructs latent components Z_g^c as linear combinations of the ancestral norms:

$$Z_g^c = X'_g w^c, \quad c = 1, 2, \dots, C,$$

where the vector w^c is chosen to maximize the covariance between the resulting components and women's empowerment:¹¹

$$w^c = \arg \max_{\|w\|=1} \text{Cov}^2(X'_g w, y_g).$$

Successive components are extracted from the residual variation after removing the information captured by the preceding components, yielding mutually distinct latent dimensions of the norm space. We determine the number of latent components C using cross-validation (see more details in Appendix A2). Panel (a) of Figure A3 presents the results: based on this criterion, we retain two components, Z_g^1 and Z_g^2 , which define the low-dimensional cultural space subsequently used for clustering.

A key design choice in our baseline specification is to train the PLS model on ethnic-level empowerment outcomes that capture variation at both ethnic and country levels. One might consider partialling out country-level variation to isolate within-country ethnic variation. However, deeply rooted ancestral norms may also be reflected in contemporary national institutions, legal frameworks, and macro-level gender environments, so removing all country-level variation could partial out some of the manifestations of the cultural phenomena we seek to characterize. We therefore retain this variation in the baseline. In Section 4.4, we assess the sensitivity of the results to adjusting for contemporary socioeconomic factors, either ex post (Check 1) or ex ante, by residualizing empowerment on the selected controls before estimating the PLS components (Check 2).

Step 2: Recovering Cultural Regimes (Unsupervised Clustering). While PLS identifies the dimensions of ancestral norms most closely associated with women's empowerment, *k-means* partitions this low-dimensional space into groups of ethnicities with similar positions along these dimensions. This second stage allows us to recover recurrent normative configurations and characterize them in terms of the prevalence and co-occurrence of the underlying ancestral norms. Intuitively, the algorithm assigns each ethnic group to the nearest cluster centroid and iteratively updates the centroids until assignments stabilize.

¹¹ In our baseline specification, each of the 545 ethnic groups is given equal weight, so that the extracted components capture global cross-cultural diversity rather than being mechanically dominated by a few demographically large ethnic groups.

Formally, let $Z_g = (Z_g^1, Z_g^2)$ denote the position of ethnic group g in the two-dimensional PLS space. For a given number of clusters K , *k-means* partitions the ethnic groups into clusters $S_1, \dots, S_K$ by minimizing within-cluster variation:

$$\min_{S_1, \dots, S_K} \sum_{k=1}^K \sum_{g \in S_k} \|Z_g - \mu_k\|^2,$$

where μ_k denotes the centroid of cluster k . As in the PLS step, each ethnic group receives equal weight. The number of clusters K is selected using several statistical criteria, discussed in detail in Appendix A2. [Panel \(b\) of Figure A3](#) presents the results, which motivate our baseline choice of $K = 5$ clusters. [Panel \(c\)](#) displays the resulting five-cluster partition of ethnic groups in the two-dimensional PLS space. We assess the sensitivity of this classification to alternative numbers of clusters and clustering procedures in Section 4.4.

4. Results

We first discuss the interpretation of the five clusters as probabilistic normative regimes, rather than discrete cultural types, then characterize their underlying norm configurations and key differences. We then assess robustness through an extensive set of checks and discuss the broader implications.

4.1. Preliminary Evidence and Cluster Interpretation

Empowerment Gradient. For ease of exposition, we order the five clusters identified by our hybrid procedure according to increasing levels of empowerment. Because the PLS components are constructed to maximize covariance with empowerment, the emergence of an empowerment gradient across clusters is expected by design. [Panel \(a\) of Figure A4](#) illustrates this gradient, with average ethnic-group empowerment increasing from Cluster 1 to Cluster 5. Our analysis therefore focuses on the configurations of ancestral norms that characterize different positions along this gradient, which we examine in Section 4.2.

Interpretations. The clusters do not sharply partition ethnic groups by empowerment. Ancestral norms account for only a limited share of cross-ethnic variation in women’s empowerment, so substantial overlap across clusters is to be expected despite differences in their average levels, as shown in [Panel \(b\)](#). Moreover, given the high-dimensional and sparse cultural space described in Section 3.3, clusters should be interpreted as *probabilistic normative regimes*: they bring together ethnicities that are similar in the reduced cultural space, rather than ethnicities sharing identical norm configurations. Each regime is therefore

characterized by the relative prevalence and co-occurrence of ancestral norms rather than by a fixed cultural bundle. Finally, the ordering of clusters is purely expositional and should not be interpreted as a dynamic process or developmental path toward greater intra-household equality.

The Explanatory Value of Probabilistic Normative Regimes. A first result is that the probabilistic cluster regimes provide a more informative representation of women's empowerment than either isolated norms or the complete set of norm indicators entered additively. In **Table 2**, we compare the cluster specification with separate models for each normative domain and with a kitchen-sink model containing all individual norm indicators. We focus on the most severe forms of women's disempowerment rather than merely capturing marginal variation along the continuous empowerment gradient. Accordingly, we define the outcome as equal to one when an ethnic group's average empowerment lies above the bottom quintile of the empowerment distribution and zero otherwise. This threshold is particularly appropriate because, as shown below, the prevalence of individual norms varies in markedly non-monotonic ways across the cluster-based empowerment gradient, making a single linear association over the full distribution potentially misleading.

Table 2 shows that individual norms display meaningful associations with empowerment but have limited explanatory power when considered separately. Patrilineality, for example, is associated with a 13.9-percentage-point lower probability of escaping the bottom empowerment quintile, yet the kinship model has an adjusted R^2 of only 0.024. The best-performing separate domain, gods supportive of human morality, has an adjusted R^2 of only 0.047. Statistical significance at the level of a particular norm therefore does not imply that this norm adequately represents the wider cultural environment associated with women's empowerment.

Within our sample, the cluster specification achieves a higher adjusted R^2 than the kitchen-sink model containing all individual norm indicators (0.146 versus 0.114). In the probit specifications, clusters achieve nearly the same likelihood-based fit as the full additive model, with a pseudo- R^2 of 0.139 versus 0.146, despite using only four cluster indicators. This more parsimonious representation is clearly preferred by the information criteria: the cluster specification has an Akaike Information Criterion (AIC) of 480 rather than 490 and a Bayesian Information Criterion (BIC) of 501 rather than 555.¹² These results show that the

¹² The AIC balances fit against model complexity by penalizing the number of free parameters, rewarding specifications that achieve good fit efficiently. The BIC applies a stricter penalty for model complexity that grows with sample size, thus favoring parsimony more strongly than AIC in our sample.

cluster regimes do more than reproduce associations attributable to isolated cultural traits. They summarize the joint structure of the norm space more effectively than separate domains and even than an additive specification containing all norms simultaneously.

Table 2: Explanatory Power and Parsimony of Clusters vs. Individual Norms

Model/domain	Norm	Coeff.	Robust SE	Adjusted R ²	Pseudo-R ²	AIC	BIC
		(1)	(2)	(3)	(4)	(5)	(6)
Clusters				0.146	0.139	480	501
All norm dummies (kitchen-sink)				0.114	0.146	490	555
Kinship system	Patrilineal	-0.139***	(0.036)	0.024	0.029	535	548
	Matrilineal	-0.007	(0.047)				
Post-marital residence	Patrilocal	-0.143***	(0.044)	0.016	0.025	538	551
	Matrilocal	0.028	(0.054)				
Marriage payments	Bride price	-0.205***	(0.027)	0.025	0.032	526	534
	Dowry	0.020	(0.019)				
Domestic structure & marriage form	Extended household	-0.103***	(0.037)	0.009	0.012	543	552
	Polygyny	-0.122***	(0.041)	0.008	0.012	543	552
Religion & Rituals	Gods supportive of human morality	-0.234***	(0.053)	0.047	0.043	526	535
	Male circumcision	-0.175***	(0.032)	0.044	0.050	522	531
Community segmentation & social hierarchy	Segmented community	-0.057	(0.043)	0.002	0.004	547	556
	Slavery	-0.097***	(0.034)	0.012	0.014	542	550
Settlement type	Non-sedentarity	-0.129**	(0.064)	0.008	0.008	545	553
Agricultural production	More male participation in agriculture	-0.099**	(0.046)	0.007	0.010	546	559
	Animal and plow cultivation	0.060	(0.055)				

Notes: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The outcome equals one if an ethnic group's average empowerment level falls above the bottom quintile of the empowerment distribution, and zero otherwise. Coefficients and heteroskedasticity-robust standard errors are obtained from separate linear probability models for each cultural-norm domain; gender-neutral or other categories constitute the omitted reference group. Adjusted R² values are from the corresponding linear probability models, whereas pseudo-R², AIC, and BIC are from probit models using the same outcome and regressors. Lower AIC and BIC indicate better fit after penalizing model complexity. Perfectly predicting categories are omitted from the probit estimation. Agricultural production combines participation in agriculture and animal/plow cultivation. The cluster and kitchen-sink rows report model-level fit statistics; coefficients are not displayed for these specifications. Clusters 1 and 2 contain 74.3% of ethnic groups in the bottom empowerment quintile, while accounting for 48.1% of the sample. Robust standard errors are reported in parentheses. *** (p<0.01), ** (p<0.05), * (p<0.10).

4.2. Cluster Profiling: Cultural Norm Configurations

We characterize each cluster by examining the prevalence of ancestral norms across cultural domains. **Figure 1a** presents a heatmap in which each column corresponds to a cluster and each cell reports the share of ethnic groups practicing a given norm. For instance, patrilineal descent is practiced by 86% of ethnic groups in Cluster 1. Darker shading indicates a higher

prevalence of a given norm. Norms are grouped by domain, with reference categories, corresponding to gender-neutral arrangements, omitted from the figure for clarity (see [Table 1](#) for the detailed classification, with reference categories in grey). Thus, an overall lighter profile indicates that ethnic groups are more frequently characterized by less restrictive arrangements. We now characterize the five regimes in turn.

Cluster 1 combines the highest prevalence of strongly patriarchal norms across several domains: patrilineal descent, patrilocal residence, bride price, polygyny, and male circumcision are all extremely common,¹³ alongside a high prevalence of moralizing gods and male-dominated agriculture, including plough-based cultivation. The cluster is therefore characterized by an accumulation of restrictive norms across multiple domains and corresponds to the lowest level of women's empowerment.

Cluster 2 maintains most of these patriarchal features but is distinguished by a markedly lower prevalence of male-dominated agriculture and moralizing gods. In particular, moralizing religious systems are less common; the cluster is also characterized by almost universal sedentarity (correspondingly, a lower prevalence of pastoralism has been associated with more restrictive gender norms, in Becker, 2025). Finally, it displays a lower prevalence of male-dominated agriculture and plough cultivation. This contrast in agricultural organization is consistent with [Boserup's \(1970\)](#) argument and echoes [Alesina et al. \(2013\)](#), who link ancestral plough use to a more male-dominated division of labor and persistent gender inequality.

Cluster 3 retains many of the same kinship, marriage, and agricultural arrangements as Cluster 2 but differs in broader social and ritual dimensions. Male circumcision, segmented communities, ancestral slavery practices, and extended households are all less prevalent, possibly pointing to more socially integrated and relatively more equal societies.¹⁴ This suggests that the higher empowerment observed in Cluster 3 is associated primarily with differences in broader social organization rather than in kinship, marital residence or agricultural arrangements.

¹³ Female economic and ritual status is often limited in such settings, where inheritance, lineage membership, and authority flow predominantly through male kinship networks ([Goody, 1976](#); [Kandiyoti, 1988](#)). The association between kinship and religious configurations is consistent with research linking kinship structures to broader moral and cooperative systems ([Enke, 2019](#)).


¹⁴ Extended-family arrangements may embed household decisions within wider obligations to kin. Evidence of social pressure to share income with relatives illustrates how such obligations can constrain individual economic choices ([Jakiela & Ozier, 2016](#)), although it does not directly establish their implications for women's empowerment.

Cluster 4 is distinguished by a very contrasted configuration of kinship structure. Patriarchal institutions are replaced by matrilineal and matrilocal systems to a large extent; dowry also appears more frequently, which may be consistent with greater female autonomy when transferred assets remain under women's control (Calvi & Keskar, 2021). Importantly, this increase in empowerment does not coincide with a general relaxation of norms across all domains: polygyny remains common, while male-dominated agriculture, plough cultivation, slavery, segmented communities, and moralizing religious systems are more prevalent than in Cluster 3. Plough cultivation thus reappears in a higher-empowerment configuration, illustrating why its presence alone does not identify a uniformly disadvantaged normative environment.

Finally, **Cluster 5** marks a more fundamental shift toward gender-neutral arrangements. Both patrilineal and matrilineal descent become uncommon, as do patrilocal and matrilocal residence, implying that the omitted gender-neutral categories become increasingly prevalent. Male dominance in agriculture and hierarchical practices also decline substantially. Yet, some of the traditional practices, bride price and polygyny, remain common.¹⁵


¹⁵ This cautions against a uniform interpretation of bride price, whose implications for women depend on the incentives surrounding marriage payments and marital stability (Gaspard & Platteau, 2010).

Figure 1a: Prevalence of Ancestral Cultural Norms by Cluster



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS and k-means estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Figure 1b: Difference in Cultural Norm Prevalence across Adjacent Clusters and Highlights of Pivotal Norms



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS and k-means estimates. Notes: Each cell reports the difference in cultural norm prevalence between adjacent clusters.

4.3. Additional Results

Geography, Religion, and the Limits of Regional Explanations. The regional distribution of clusters ([Figure A5](#)) reveals a clear geographical structure, with lower-empowerment configurations especially concentrated in the Middle East and North Africa and the highest-empowerment regime particularly prevalent in Latin America. Yet no region maps uniquely into a single cluster, and substantial heterogeneity remains within broad regions, especially in Sub-Saharan Africa and Asia. Where multiple ethnic groups are observed, the ethnic-level analysis can be extended within countries, further revealing meaningful variation and confirming that normative environments are related to geography but not reducible to it. In particular, some countries located in different world regions display remarkably similar cluster distributions, further indicating that the identified regimes capture normative configurations that cut across geographical boundaries.


Pivotal Margins between Neighboring Regimes at Low-Empowerment Levels. In **Figure 1b**, we complement our baseline comparison by reporting changes in norm prevalence across adjacent clusters, emphasizing the norms that are most pivotal in distinguishing neighboring regimes. Overall, the cluster comparisons reveal both broad configurational differences and more specific pivotal margins. Most notably, moralizing religious systems and the gender organization of agricultural production, characterized by male-dominated agriculture and greater reliance on the plough, sharply distinguish the lowest-empowerment regimes. Configurations in which these features are less prevalent are associated with higher empowerment even when patriarchal kinship structures and other forms of social inequality remain widespread. Moving up the gradient, differences first involve less hierarchical social organization, then more women-oriented kinship and family structures, and finally more gender-neutral cultures, as discussed in detail above.

To identify the norms that most clearly distinguish the two lowest-empowerment regimes, we move beyond comparisons of average norm prevalence and construct an optimal one-to-one matching between ethnic groups in Clusters 1 and 2. Each of the 80 ethnic groups in Cluster 1 is matched to a distinct ethnic group in Cluster 2 so as to minimize the total Euclidean distance between their configurations across the 12 ancestral norm domains, yielding 80 exclusive pairs. For each normative domain, we then calculate the share of matched pairs for which the two ethnic groups differ. Results are reported in **Figure 2**. Because norm profiles are discrete, several assignments can attain the same minimum total distance. We therefore repeat the optimal matching 1,000 times using random tie-breaking among equally distant matches. The bars thus report the average frequency of norm differences across these optimal assignments, while the capped intervals report the minimum and maximum frequencies obtained across the 1,000 solutions.¹⁶ The differences between the closest configurations are highly concentrated in a small number of domains. Moralizing gods differ in approximately 65% of matched pairs (range about 64–66%), male-dominated agriculture in 48% (43–55%), and animal/plough cultivation in 24% (24–26%). Non-sedentarity follows at about 18%, while each of the remaining domains differs in fewer than 12% of matched pairs. Thus, even after pairing ethnic groups so that their overall ancestral normative configurations are as similar as possible, the distinction between the two lowest-empowerment regimes is disproportionately concentrated in religious beliefs and the organization of agricultural production, with women’s role in agriculture and plough

¹⁶ These are therefore not conventional sampling confidence intervals, but measures of the sensitivity of the results to alternative equally optimal matchings.

use emerging as particularly important pivotal margins. As noted above, the exercise should not be interpreted as an observed transition from Cluster 1 to Cluster 2; it merely identifies the dimensions that most systematically distinguish the empirically closest configurations on either side of these neighboring regimes.

Figure 2: Which Norms Distinguish the Closest Pairs of Ethnicities across Clusters 1 and 2



Notes: Authors' calculations based on 545 ethnic groups from the pooled DHS-Ethnographic Atlas sample. Each of the 80 ethnic groups in Cluster 1 is matched to a distinct ethnic group in Cluster 2 so as to minimize the total Euclidean distance between their configurations across the 12 ancestral norm domains. For each domain, the bars report the mean percentage of matched pairs with different norms across 1,000 optimal assignments obtained through random tie-breaking. Capped intervals report the minimum and maximum percentages across these equally optimal assignments; they are not statistical confidence intervals.

4.4. Sensitivity Checks and Discussion

A comprehensive sensitivity analysis is presented in **Appendix A5** (Checks 1 to 7). Here, we summarize the main findings and discuss their implications.

Check 1: Cultural Persistence Beyond Contemporary Conditions. We assess whether the identified clusters capture variation in women's empowerment beyond contemporary socioeconomic conditions by controlling for income, inequality, women's education, and female labor force participation (**Table A4**). The empowerment gradient across clusters

remains largely stable. A particularly informative finding is that contemporary female labor force participation does not fully account for the differences associated with historically more balanced economic roles. This suggests that ancestral divisions of labor remain relevant to women's agency beyond the traces they have left on contemporary female employment, potentially through enduring norms governing women's roles and authority within the household. Adding religious and geographical controls mainly attenuates differences at the lower end of the gradient. Overall, the ancestral configurations remain informative about women's empowerment after accounting for these contemporary characteristics.

Check 2: Sensitivity to Country-Level Residualization. As an ex ante counterpart to Check 1, we assess sensitivity to adjusting empowerment for contemporary country-level socioeconomic characteristics, broad regions and broad religious groups before estimating the PLS components and applying k-means. The resulting norm profiles and contrasts between adjacent clusters remain broadly similar to the baseline ([Figure A6](#)), indicating that the main configurations are robust to accounting for income and broad regional differences before constructing the cultural space.

Check 3: Sensitivity to Matching Methods and Sample-Size Threshold. Part of our matching strategy relies on geographic proximity rather than direct ethnicity or language identifiers, potentially introducing measurement error. We assess sensitivity to this choice by applying geographic matching to individuals for whom direct matches are available and, separately, restricting the analysis to direct matches ([Figure A7](#)). Both exercises yield configurations broadly similar to the baseline. Some differences arise among intermediate clusters, but the principal contrasts associated with severe disempowerment remain stable. We also raise the minimum ethnic-group sample-size threshold from 50 to 100 women, retaining 496 of the 545 ethnic groups (91%). The resulting cluster profiles and empowerment gradient remain closely comparable to the baseline.

Check 4: Sensitivity to Clustering Choices. We assess the sensitivity to the clustering algorithm and the number of clusters ([Figure A8](#)). Ward's hierarchical clustering yields profiles closely similar to the baseline. Specifications with seven and nine clusters reveal finer configurations without altering the main conclusions. These partitions preserve the same type of non-monotonicity and, most importantly, the pivotal role of religious norms and agricultural organization, including women's participation and plough use, in distinguishing configurations at the lower end of the empowerment gradient.

Check 5: Alternative Outcomes (Attitudes towards Domestic Violence). To assess whether the identified clusters capture meaningful variation beyond the specific outcome used in estimation, we first hold baseline cluster assignments fixed and examine an index of women’s disapproval of wife-beating (**Figure A9**). Clusters associated with greater decision-making autonomy also exhibit greater disapproval, extending the observed gradient to a related but distinct outcome not used to construct the baseline cultural space. Alternatively, we re-estimate the full PLS and k-means procedure using disapproval of wife-beating as the sole supervised outcome and, separately, jointly with the decision-making index (**Figure A10**). Both specifications yield normative profiles broadly similar to the baseline. In particular, the principal religious and agricultural contrasts remain salient, showing that the main patterns are consistent when a substantively meaningful alternative outcome guides the construction of the cultural space.

Check 6: Robustness to Norm Pre-Grouping. Section 3.2 aggregates the original EA categories within each domain into a smaller number of groups prior to the PLS step, collapsing categories whose estimated relationship with women’s empowerment is statistically indistinguishable. Because this grouping uses information about the outcome before the PLS step, a natural concern is whether it drives - rather than merely simplifies - the patterns we identify. To address it, we re-estimate the PLS and k-means procedure using the original categories without prior grouping (**Figure A11**), aggregating them only afterward to facilitate comparison of cluster profiles. The resulting configurations closely resemble the baseline, indicating that the main patterns do not depend on the outcome-informed pre-grouping.

Check 7: Sensitivity to the Degree of Supervision. Because PLS uses empowerment to construct the cultural space, we assess whether the main configurations persist when less emphasis is placed on their relationship with the outcome. Continuum regression (see the methodological discussion in **Appendix A2-2**) provides a unified framework to vary the degree of supervision (parameter γ), encompassing OLS ($\gamma=0$), PLS ($\gamma=1$), and an increasingly PCA-like representation (as γ rises above 1). Using this approach, we progressively shift component extraction toward the variance-maximizing PCA benchmark while retaining two latent components and five clusters (**Figure A12**). Although cluster profiles change under reduced supervision, non-monotonic patterns persist. Most importantly, the principal religious and agricultural contrasts remain visible at the lower end of the empowerment gradient (**Panel (e)**), even as mean empowerment levels become less clearly separated. Kinship and residence arrangements continue to differentiate

configurations elsewhere. The persistence of these contrasts suggests that their prominence is not solely attributable to the outcome-oriented nature of PLS.

5. Conclusion

Communities may inherit multiple normative strands, some restrictive, others enabling, and the balance between them can profoundly influence how women respond to modern institutions, reforms, and policies. Our research builds precisely on this insight. We develop a novel framework to study how combinations of ancestral cultural norms are associated with women's empowerment across ethnic groups worldwide. Moving beyond single-norm approaches, we combine supervised dimension reduction (PLS) with unsupervised clustering (k-means) to identify empirically relevant mixtures of norms that are directly aligned with variation in women's agency. This hybrid strategy allows us to recover a low-dimensional cultural space structured around empowerment, and to characterize probabilistic cultural regimes that capture how norms co-occur in practice.

Our results identify five probabilistic configurations of ancestral norms spanning the empowerment gradient. The lowest levels of women's agency are characterized by an accumulation of restrictive patriarchal practices across several domains, reinforced by unequal social structures. At the lower end, religious beliefs and agricultural organization distinguish different empowerment levels within otherwise broadly patriarchal environments. Moving up the gradient, differences increasingly reflect broader social organization, while the highest levels of empowerment are associated with a broader relaxation of restrictive norms across domains, even though some traditional practices remain common. The recurrence of some norms, including plough cultivation, at different positions along the gradient highlights the non-monotonic character of these configurations. These patterns persist after accounting for contemporary socioeconomic conditions, geography, and religion, and display substantial within-region heterogeneity.

The findings point to the association between inherited normative environments and intra-household inequality, with potential relevance for policy design. They suggest that interventions aimed at improving women's outcomes could benefit from considering broader normative environments, rather than targeting single constraints in isolation. In particular, the identified configurations help flag contexts associated with greater vulnerability among women. In multi-ethnic settings, these configurations could inform

assessments of intra-household disadvantage and motivate evaluations of whether interventions work differently across normative environments.¹⁷

A complementary avenue for future research would be to examine whether the uneven contemporary salience of ancestral norms partly reflects differences in their persistence. Some norms may appear less consequential because the practices they represent have weakened, while others may remain salient because those practices have endured.¹⁸ Linking our configurations to evidence on cultural change and historical disruptions could help distinguish these trajectories, adding a dynamic perspective to the mapping developed here. That said, our characterization already shows that most norms play a role, although their salience varies across the empowerment gradient, with some (including plough agriculture) distinguishing configurations at both its lower and upper ends. This non-monotonicity is an important feature of our results and underscores the relevance of examining individual norms within the broader configurations in which they are embedded. Examining actual transitions between configurations would, however, entail a trade-off between the historical depth achievable within selected ethnic or geographical settings and the global scope of the characterization offered here.

A related extension would seek to identify the causal mechanisms through which particular norms operate within broader configurations, drawing on historical shocks or policy discontinuities. For example, [Alesina et al. \(2013\)](#) investigate the historical origins of contemporary gender roles using variation in ancestral plough use, instrumented by crop suitability. The same trade-off applies here: identifying these mechanisms would require more focused comparisons of specific practices or settings, complementing the global characterization offered by our approach.

More broadly, the framework offers a general way to organize multidimensional information about ethnic, social or economic groups into interpretable configurations relevant to contemporary outcomes. Its application need not be limited to ancestral norms and gender outcomes: it could characterize combinations of historical formal or informal institutions associated with regional development, inherited community characteristics related to

¹⁷ Such an extension is motivated by evidence that social organization shapes redistribution and the effects of transfer programs ([Moscona & Seck, 2024](#)).

¹⁸ Ancestral beliefs in moralizing gods are strongly associated with contemporary religious composition: the average Muslim share is 70.9% among ethnic groups historically characterized by such beliefs, compared with 21.1% among other groups. This overlap suggests that the contemporary salience of moralizing gods may partly reflect the persistence of religious traditions and associated prescriptions, although it does not by itself establish historical continuity.

children's education and health, or social and cultural environments associated with political participation and institutional trust, for instance. Across these applications, the common purpose is to move beyond isolated attributes and broad group labels to identify the configurations that structure differences in outcomes. The approach thus provides a portable framework for studying how multiple dimensions of social and historical heterogeneity combine to shape contemporary inequalities.

References

- Alesina, A., & Giuliano, P. (2015). Culture and institutions. *Journal of Economic Literature*, 53(4), 898–944.
- Alesina, A., Brioschi, B., & La Ferrara, E. (2021). Violence against women: A cross-cultural analysis for Africa. *Economica*, 88(349), 70–104.
- Alesina, A., Giuliano, P., & Nunn, N. (2013). On the origins of gender roles: Women and the plough. *Quarterly Journal of Economics*, 128(2), 469–530.
- Aminjonov, U., Bargain, O., Colacce, M., & Tiberti, L. (2024). Culture, intrahousehold distribution, and individual poverty. *Economic Development and Cultural Change*, 73(1), 127–165.
- Aminjonov, U., Bargain, O., Colacce, M. (2026). Global Evidence on Gender Gaps and Child Poverty in Consumption, forthcoming: *World Bank Economic Review*
- Aoki, M. (2001). *Toward a Comparative Institutional Analysis*. Cambridge, MA: MIT Press
- Ashraf, N., Bau, N., Nunn, N., & Voena, A. (2020). Bride price and female education. *Journal of Political Economy*, 128(2), 591–641.
- Bahrami-Rad, D., Becker, A., & Henrich, J. (2021). Tabulated nonsense? Testing the validity of the Ethnographic Atlas. *Economics Letters*, 204, 109880.
- Bahrami-Rad, D., Beauchamp, J., Schulz, J. F., & Henrich, J. (2022). Kin-based institutions and economic development. (Working paper)
- Baland, J.-M., Bourguignon, F., Platteau, J.-P., & Verdier, T. (Eds.). (2020). *The handbook of economic development and institutions*. Princeton University Press.
- Baland, J.-M., Bequet, L., Guirking, C., & Manuel, C. C. (2024). Sharing norm, household efficiency and female demand for agency in the Philippines. *World Development*, 174, 106463.
- Baland, J.-M., Boltz, M., Guirking, C., Jolivet, A., & Ziparo, R. (2026). Delegation in the Family. DeFiPP Working Paper 2605, University of Namur.

- Bargain, O., D. Boutin and H. Champeaux (2019). Women's political participation and intrahousehold empowerment: Evidence from the Egyptian Arab Spring, *Journal of Development Economics*, 141, 102379
- Bargain, O., Loper, J., & Ziparo, R. (2022). Traditional norms, access to divorce and women's empowerment. IZA Discussion Paper No. 15374.
- Bau, N., & Fernández, R. (2023). Culture and the family. In S. Lundberg & A. Voena (Eds.), *Handbook of the economics of the family* (pp. 1–48). North-Holland.
- Beaman, L., Duflo, E., Pande, R., & Topalova, P. (2012). Female leadership raises aspirations and educational attainment for girls: A policy experiment in India. *Science*, 335(6068), 582–586.
- Becker, A. (2025). On the economic origins of concerns over women's chastity. *Review of Economic Studies*, 92(4), 2303–2329.
- Bertocchi, G., & Bozzano, M. (2014). Family structure and the demographic transition. *European Economic Review*, 70, 197–220.
- Bisin, A., & Verdier, T. (2000). Beyond the melting pot: Cultural transmission, marriage, and the evolution of ethnic and religious traits. *Quarterly Journal of Economics*, 115(3), 955–988.
- Bisin, A., & Verdier, T. (2001). The economics of cultural transmission and the dynamics of preferences. *Journal of Economic Theory*, 97(2), 298–319
- Boserup, E. (1970). *Woman's role in economic development*. George Allen & Unwin.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Brown, C., Ravallion, M., & van de Walle, D. (2019). Most of Africa's nutritionally deprived women and children are not found in poor households. *Review of Economics and Statistics*, 101, 631–644.
- Calvi, R., & Keskar, A. (2021). Dowries, resource allocation, and poverty. *Journal of Economic Behavior & Organization*, 192, 268–303.
- Chattopadhyay, R., & Duflo, E. (2004). Women as policymakers: Evidence from a randomized policy experiment in India. *Econometrica*, 72(5), 1409–1443.
- Corno, L., Hildebrandt, N., & Voena, A. (2020). Age of marriage, weather shocks, and the direction of marriage payments. *Econometrica*, 88(3), 879–915.
- D'Agostini, C., & Tiberti, L. (2025). Culture, Kinship and Women's Labour Opportunities: Evidence from Malawi and Indonesia, *Journal of Development Studies*, 62/2, 199–217.
- Dessy, S., Tiberti, L., & Zoundi, D. (2023). The gender education gap in developing countries: Roles of income shocks and culture. *Journal of Comparative Economics*, 51(1), 160–180.

- Doepke, M., Tertilt, M., & Voena, A. (2012). The economics and politics of women's rights. *Annual Review of Economics*, 4, 339–372.
- Edlund, L., & Lagerlöf, N.-P. (2006). Individual versus parental consent in marriage: Implications for intra-household resource allocation and growth. *American Economic Review*, 96(2), 304–307.
- Enke, B. (2019). Kinship, cooperation, and the evolution of moral systems. *Quarterly Journal of Economics*, 134(2), 953–1019.
- Ewerling, F., Lynch, J. W., Victora, C. G., van Eerdewijk, A., Tyszler, M., & Barros, A. J. D. (2017). The SWPER index for women's empowerment in Africa: Development and validation of an index based on survey data. *The Lancet Global Health*, 5(9), e916–e923.
- Gaspart, F., & Platteau, J.-P. (2010). Strategic behavior and marriage payments: Theory and evidence from Senegal. *Economic Development and Cultural Change*, 59(1), 149–185.
- Giuliano, P. (2020). Gender and culture. *Oxford Review of Economic Policy*, 36(4), 944–961.
- Goody, J. (1976). *Production and reproduction: A comparative study of the domestic domain*. Cambridge University Press.
- Goody, J. (1983). *The development of the family and marriage in Europe*. Cambridge University Press.
- Giuliano, P., & Nunn, N. (2018). Ancestral characteristics of modern populations. *Economic History of Developing Regions*, 33(1), 1–17.
- Greif, A. (1994). Cultural beliefs and the organization of society. *Journal of Political Economy*, 102(5), 912–950.
- Guirking, C., & Platteau, J.-P. (2020). The dynamics of family systems: Lessons from past and present times. In J.-M. Baland, F. Bourguignon, J.-P. Platteau, & T. Verdier (Eds.), *The Handbook of Economic Development and Institutions* (pp. 449–519). Princeton University Press.
- Henrich, J. (2001). Cultural transmission and the diffusion of innovations. *American Anthropologist*, 103(4), 992–1013.
- Jayachandran, S. (2015). The roots of gender inequality in developing countries. *Annual Review of Economics*, 7, 63–88. <https://doi.org/10.1146/annurev-economics-080614-115404>
- Jayachandran, S. (2021). Social norms as a barrier to women's employment in developing countries. *IMF Economic Review*, 69(3), 576–595.
- Jakiela, P., & Ozier, O. (2016). Does Africa need a rotten kin theorem? Experimental evidence from village economies. *The Review of Economic Studies*, 83(1), 231–268.

- Kabeer, N. (1999). Resources, agency, achievements: Reflections on the measurement of women's empowerment. *Development and Change*, 30, 435–464.
- Kandiyoti, D. (1988). Bargaining with patriarchy. *Gender & Society*, 2(3), 274–290.
- La Ferrara, E. (2007). Descent rules and strategic substitutability in social structures. *Journal of Economic Theory*, 137(1), 561–583.
- La Ferrara, E., & Milazzo, A. (2017). Customary norms, inheritance, and human capital: Evidence from a reform of the matrilineal system in Ghana. *American Economic Journal: Applied Economics*, 9(4), 166–185.
- Lowes, S. (2020). Kinship structure & women: Evidence from economics. *Daedalus*, 149(1), 119–133.
- Lowes, S. (2022). *Kinship structure and the family: Evidence from the matrilineal belt* (NBER Working Paper No. 30509). <https://doi.org/10.3386/w30509>
- Lowes, S., & Nunn, N. (2018). Bride price and the well-being of women. In S. Anderson, L. Beaman, & J.-P. Platteau (Eds.), *Towards gender equity in development* (pp. 117–138). Oxford University Press.
- MacQueen, J. B. (1967). Some methods for classification and analysis of multivariate observations. In L. M. Le Cam & J. Neyman (Eds.), *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, pp. 281–297. University of California Press.
- Miedema, S. S., Haardörfer, R., Girard, A. W., & Yount, K. M. (2018). Women's empowerment in East Africa: Development of a cross-country comparable measure. *World Development*, 110, 453–464.
- Michalopoulos, S., Putterman, L., & Weil, D. N. (2019). The influence of ancestral lifeways on individual economic outcomes in Sub-Saharan Africa. *Journal of the European Economic Association*, 17(4), 1186–1231.
- Michalopoulos, S., & Papaioannou, E. (2013). Pre-colonial ethnic institutions and contemporary African development. *Econometrica*, 81(1), 113–152.
- Moscona, J., & Seck, A. A. (2024). Age set versus kin: Culture and financial ties in East Africa. *American Economic Review*, 114(9), 2748–2791.
- Murdock, G. P. (1967). Ethnographic atlas: A summary. *Ethnology*, 6(2), 109–236. <https://doi.org/10.2307/3772751>
- Nunn, N., & Wantchekon, L. (2011). The slave trade and the origins of mistrust in Africa. *American Economic Review*, 101(7), 3221–3252.
- Nunn, N. (2008). The long-term effects of Africa's slave trades. *Quarterly Journal of Economics*, 123(1), 139–176.

Nunn, N. (2014). Gender and missionary influence in colonial Africa. In E. Akyeampong, R. Bates, & J. A. Robinson (Eds.), *African development in historical perspective* (pp. 489–512). Cambridge University Press.

Nunn, N. (2020). The historical roots of economic development. *Science*, 367(6485), eaaz9986.

Silverman, E. K. (2004). Anthropology and circumcision. *Annual Review of Anthropology*, 33, 419–445.

Stone, M., & Brooks, R. J. (1990). Continuum regression: Cross-validated sequentially constructed prediction embracing ordinary least squares, partial least squares and principal components regression. *Journal of the Royal Statistical Society: Series B (Methodological)*, 52(2), 237–258.

Appendix A1: EA-references and Descriptive Statistics

Table A1: Women's Outcomes Varying with Ethnic Customs from the Ethnographic Atlas

Authors (Year)	Country/Region	Cultural Variable (from Ethnographic Atlas)	Outcome Studied
<i>Agricultural practice & women's productive role</i>			
Alesina, Giuliano & Nunn (2013)	Global*	Plough use vs. Hoe; women's participation in agriculture	Female labor force participation, sector of employment, and gender-role attitudes today
Alesina, Giuliano & Nunn (2021)	Global*	Traditional plough agriculture	Male-biased sex ratios, demographic imbalance
Hansen, Jensen & Skovsgaard (2015)	Global*	Traditional plough agriculture; historical gender division of labour	Gender roles, female labour force participation, gender-equality indices
Fredriksson & Gupta (2023)	Global*	Ancestral irrigation potential; plough use	Female labour force participation, gender-role attitudes, women's rights
<i>Marriage payments</i>			
Ashraf, Bau, Nunn & Voena (2020)	Indonesia & Zambia	Bride price	Female education, age at marriage, women's bargaining power
Corno, Hildebrandt, & Voena (2020)	Sub-Saharan Africa & India	Marriage market institutions	Variation in child marriage and daughter's human capital between bride-price and dowry societies
<i>Other marriage customs (kinship, residence, polygyny, etc.)</i>			
Aminjonov, Bargain, Colacce & Tiberti (2023)	Ghana & Malawi	Post-marital residence (+ kinship structure)	Intra-household resource shares
Bargain, Loper & Ziparo (2022)	Indonesia	Post-marital residence (+ kinship structure)	Ethnic heterogeneity in response to legal reforms using various women's and children's outcomes
Dessy, Tiberti & Zoundi (2023)	Malawi & Indonesia	Post-marital residence (+ kinship structure)	Gender gap in education in response to climate income shock
Fenske (2015)	Africa	Polygyny prevalence	Economic correlates of polygyny, women's outcomes, inequality
Lowes (2022)	Sub-Saharan Africa (Matrilineal Belt)	Matrilineal vs patrilineal kinship (+post-marital residence)	Women's autonomy, domestic violence, children's education & health
<i>Slavery</i>			
Corno, La Ferrara, & Voena (2020)	Africa	Slave trade	Persistence of female genital cutting
<i>Normative mix</i>			
Alesina, Brioschi & Nunn (2021)	Global (with focus on Africa, Asia)	Marriage customs (bride price, polygyny, cousin marriage)	Fertility transition and family size
Alesina, Brioschi & La Ferrara (2021)	Sub-Saharan Africa	Norms about marriage patterns, living arrangements and the productive role of women	Intimate partner violence, attitudes toward wife-beating
D'Agostini & Tiberti (2026)	Malawi & Indonesia	Matrilineal vs patrilineal, matrilineal vs patrilineal, extended family arrangements	Women's labor opportunities
Miotto (2021)	Sub-Saharan Africa	Pre-colonial cash crop intensity, but controlling for polygyny, bride price, women's role in agriculture	Women's labor outcomes and female economic empowerment
Giuliano (2018)	Global*	Family structure variables (patrilineal, matrilineal, polygyny, etc.)	Historical roots of gender differences in labor, politics, health

* Global indicates that the study is based on global data on historical ethnic groups worldwide.

Table A2: List of Selected Countries and Statistics on “Final Say”

Country	Number of DHS surveys	Sample of women	Average share of women "with final say"		Number of DHS surveys	Sample of women	Average share of women "with final say"
	(1)	(2)	(3)		(1)	(2)	(3)
Afghanistan	1	28,603	0.176	Malawi	4	56,300	0.261
Albania	2	16,225	0.360	Maldives	2	12,172	0.424
Angola	1	8,033	0.444	Mali	4	44,490	0.153
Armenia	4	18,226	0.371	Mauritania	1	9,759	0.335
Azerbaijan	1	5,249	0.264	Moldova	1	4,880	0.563
Bangladesh	6	93,831	0.298	Morocco	1	16,790	0.200
Benin	4	47,051	0.275	Mozambique	4	34,367	0.364
Bolivia	2	27,825	0.493	Myanmar	1	7,868	0.519
Burkina Faso	3	38,731	0.175	Namibia	2	7,331	0.496
Burundi	2	14,820	0.307	Nepal	5	50,141	0.315
Cambodia	4	41,118	0.510	Nicaragua	1	13,043	0.371
Cameroon	3	28,471	0.279	Niger	2	18,725	0.141
Chad	1	13,386	0.236	Nigeria	4	87,616	0.195
Colombia	4	145,160	0.527	Pakistan	2	27,494	0.206
Comoros	1	3,270	0.358	Papua New Guinea	1	9,847	0.415
Congo	1	6,745	0.353	Peru	6	169,375	0.491
Congo Dem. Rep.	2	22,413	0.291	Philippines	5	63,019	0.531
Cote d'Ivoire	2	16,063	0.217	Rwanda	5	42,688	0.373
Dominican Republic	2	23,009	0.477	Sao Tome and Principe	1	1,746	0.382
Egypt	4	69,510	0.307	Senegal	8	72,016	0.124
Ethiopia	3	28,633	0.352	Sierra Leone	3	25,840	0.244
Gabon	2	9,469	0.434	South Africa	1	2,830	0.517
Gambia	2	14,978	0.274	Swaziland	1	2,026	0.403
Ghana	4	22,903	0.384	Tajikistan	2	13,922	0.248
Guatemala	1	14,979	0.427	Tanzania	4	33,977	0.276
Guinea	3	22,542	0.231	Timor-Leste	2	15,499	0.448
Guyana	1	2,984	0.529	Togo	1	6,357	0.253
Haiti	4	31,879	0.415	Turkey	1	5,479	0.345
Honduras	2	24,864	0.415	Uganda	4	29,326	0.377
India	3	251,580	0.321	Ukraine	1	4,160	0.600
Indonesia	4	129,418	0.475	Uzbekistan	1	5,408	0.223
Jordan	5	53,252	0.425	Yemen	1	15,618	0.276
Kenya	4	40,534	0.423	Zambia	4	29,204	0.379
Kyrgyz Republic	1	5,460	0.419	Zimbabwe	4	22,888	0.447
Lesotho	3	14,828	0.431				
Liberia	3	14,999	0.392	Average			0.366
Madagascar	3	31,430	0.462	Total	188	2,344,672	

Note: authors' calculation using the Demographic and Health Surveys (DHS).

* Women making decisions independently on healthcare and family visits, and independently or jointly for large household purchases.

**Table A3: Distribution of DHS Observations
by Type of Matching to the Ethnographic Atlas**

Matching strategy	Percent of women in pooled DHS data	Cumulative percent of women in pooled DHS data
	(1)	(2)
Based on reported ethnicity	41.57	41.57
Based on reported native language	7.47	49.04
Based on GPS coordinates	32.68	81.72
Based on subnational region	17.67	99.39
Unmatched	0.61	100.00

Notes: Percentages refer to the share of women in the pooled DHS sample (with non-missing values for one of the three decision-making variables) that are successfully linked to an ancestral group through each matching strategy. "Reported ethnicity" and "reported native language" represent direct matches based on self-identified ethnic group or language information in the DHS data. "GPS coordinates" and "subnational region" correspond to indirect geographic matching used when no direct ethnic or linguistic link is available. "Unmatched" indicates observations for which no reliable match to the Ethnographic Atlas could be established. Only one country (Maldives) drops from the sample because no reliable link could be established between DHS and ancestral groups in the Ethnographic Atlas.

Figure A1: Geographical Coverage of the DHS Surveys in the Final Sample





Figure A2: Women's Empowerment across Countries



Source: DHS data for 71 countries for which at least one of the selected decision-making variables are non-missing. Note: The figure plots the average share of women making decisions independently on healthcare and family visits, and independently or jointly on large household purchases.

Appendix A2: Matching & the Supervised-Unsupervised Approach

A2-1. Detailed EA-DHS Matching Approach

To match DHS households (from contemporaneous ethnicities) to ancestral ethnicities and the corresponding norms, we implement the following sequence.

1. We first attempt a **direct match** between *ethnic group names* reported in the DHS and ethnic group identifiers in the EA. When this is not possible, we match DHS ethnicities to the ethnic classifications in [Murdock \(1959\)](#)'s map and then use the *recent concordance* between Murdock's map and the Atlas developed by [Kincaide et al. \(2025\)](#) to link those groups to the EA. In addition, we match DHS ethnic groups to ethnic classifications in the Afrobarometer and then draw on the concordance provided by [Nunn and Wantchekon \(2011\)](#) to map these to the EA. To further improve the direct match, we also check alternate names of unmatched DHS ethnic groups or the names corresponding to their native languages using Joshua Project, Ethnologue, and the database of Linking Ethnic Data from Africa (LEDA). For the languages spoken by ethnic groups in the DHS, we rely on the *language-ethnicity concordance* constructed by [Giuliano and Nunn \(2018\)](#) to link these groups to the EA.
2. For ethnic groups that could not be matched directly, we proceed with **linguistic matching** using the aforementioned information on ethnic language. We rely on the native languages reported in the DHS to match those to the Ethnologue or the language family trees in the LEDA database. To link matched languages to the EA, we again use Giuliano and Nunn's language-ethnicity concordance.
3. In cases where direct ethnic or linguistic matches were unsuccessful, we rely on **geographic matching**, as in the related literature ([Alesina et al., 2013](#); [Cao et al., 2021](#); [Giuliano and Nunn, 2021](#)). We start by matching DHS household GPS coordinates to Murdock's Map to infer the most likely ancestral group inhabiting the area. To expand matching, we also locate DHS clusters in the language map of [Giuliano and Nunn \(2018\)](#) to identify which language is spoken in the location where households reside. When household GPS information is not available, we further rely on other geographic identifiers (district or another administrative division) in the DHS data to identify which language is mostly spoken in a given area and use the language-ethnicity concordance. When several ethnic groups/languages overlap geographically, we use the dominant language/ethnic group in that area.

A2-2. Additional Information on the PLS and Clustering Procedure

Partial Least Squares. PLS is widely used in high-dimensional fields such as chemometrics, bioinformatics, genetics, and psychometrics (Höskuldsson, 1988; Wold et al., 2001; Boulesteix & Strimmer, 2007; Abdi, 2010). It became particularly prominent in chemometrics, where researchers often seek to relate an outcome, such as the concentration or physical property of a chemical compound, to a large number of highly correlated measurements, such as spectral intensities. Rather than estimating the contribution of each measurement separately, PLS extracts a small number of latent components formed from combinations of predictors that maximize covariance with the outcome. This makes it possible to summarize complex, correlated signatures while retaining the dimensions most relevant to the outcome of interest.

Our setting has a similar statistical structure. Ancestral norms are multidimensional, interdependent, and often strongly correlated, and our interest lies in combinations of norms rather than in estimating a separate effect for each cultural trait. PLS allows us to summarize this joint structure while orienting the resulting cultural space toward women's empowerment. The analogy with chemometrics is useful in this respect: just as a chemical property may be associated with a combination of correlated spectral measurements rather than with a single wavelength, women's empowerment may be associated with configurations of mutually related ancestral traits rather than with any norm considered in isolation.

K-means Clustering. The second stage complements the supervised construction of the cultural space with an unsupervised classification of ethnic groups. K-means is one of the most widely used clustering methods in machine learning and is particularly appropriate when observations are represented in a low-dimensional continuous feature space and the objective is to obtain a transparent partition into relatively compact groups (Suyal and Sharma, 2024).

This division of labor between the two stages is important. PLS determines the dimensions along which ethnic groups are compared, using their relationship with empowerment, whereas k-means subsequently groups ethnicities according to proximity in this latent space without using empowerment directly in the clustering criterion. We then return to the original ancestral traits to characterize each cluster in terms of the prevalence and co-occurrence of norms. The resulting clusters are therefore interpreted as probabilistic normative regimes rather than as deterministic cultural types.

K-means also has the advantage of keeping the second stage transparent. More elaborate alternatives, such as supervised latent-class or finite-mixture models, could estimate latent regimes jointly with their relationship to empowerment. Such approaches are closer to a fully model-based representation of probabilistic regimes, but require stronger distributional and identification assumptions and make the construction of the outcome-relevant space and the formation of groups part of the same estimation problem. Our two-stage procedure instead keeps these operations separate: PLS identifies an empowerment-relevant cultural space, and k-means provides a simple unsupervised partition of that space.

Alternative Approaches. Alternative methods address only some of our objectives. **PCA** provides an efficient reduction of correlated predictors, but extracts directions that maximize variation in the norms themselves, irrespective of their relevance for empowerment. Regression-based approaches instead focus on explaining or predicting empowerment. **OLS** can include multiple norms simultaneously, but strong collinearity makes their separate associations difficult to disentangle, while allowing for complementarities requires specifying a potentially large number of interaction terms. Penalized methods such as **Lasso** or elastic net can accommodate larger sets of predictors and interactions, while **tree-based methods**, including random forests, offer greater flexibility in capturing nonlinearities and complex interactions. These approaches nevertheless produce regression coefficients or prediction rules rather than a low-dimensional cultural space that can naturally be partitioned into recurrent, interpretable normative regimes.

PLS therefore occupies a useful middle ground for our purpose: it reduces the dimensionality of a correlated norm space, as PCA does, while using information from the outcome to identify directions relevant to empowerment. This supervision is deliberate. Because Y enters the construction of the PLS components, the subsequent clusters are expected to differ along the empowerment dimension. Our objective is not to interpret this gradient itself as an independently discovered relationship, but to examine which configurations of ancestral norms characterize different positions along it.

Finally, **continuum regression** provides a natural framework for assessing sensitivity to the degree of supervision used to construct the cultural space, connecting OLS, PLS and principal components regression (Stone and Brooks, 1990). Applying the same clustering procedure while moving from PLS toward the PCA end of this continuum reveals whether the main normative configurations, and especially the norms identified as pivotal, persist when

covariance with empowerment receives less emphasis relative to variation in the norms themselves. This is investigated in **Appendix A5, Check 7**.

A2-3. Optimal Number of Latent Components: Cross-validation

We determine the optimal number of latent components using cross-validation, selecting the number that minimizes the out-of-sample mean squared prediction error (MSE). Based on this criterion, we retain two latent components, Z_g^1 and Z_g^2 . The first PLS component, Z_g^1 , captures the direction in X_g that maximizes covariance with y_g . The second component, Z_g^2 , captures the next most predictive direction in the residual variation of X_g , after removing the variation explained by Z_g^1 , and so on for subsequent components. **Panel (a)** of **Figure A3** illustrates the cross-validation procedure and the resulting choice of two components.

A2-4. Optimal Number of Clusters


We determine the number of clusters using three complementary criteria: the elbow method, the Calinski–Harabasz index, and the gap statistic.¹⁹ As shown in **Panel (b)** of **Figure A3**, these criteria do not point to a unique solution. The Calinski–Harabasz index and gap statistic favor relatively coarse partitions, with their global maxima occurring at low values of K . However, the elbow criterion indicates a marked flattening of the within-cluster sum of squares around $K=5$, while both the Calinski–Harabasz index and the gap statistic display a local improvement at five clusters relative to the immediately preceding partition. We therefore retain $K=5$ as our baseline, providing a sufficiently rich partition to distinguish alternative normative configurations while remaining parsimonious and interpretable. We assess the sensitivity of the results to alternative numbers of clusters in Section 4.4. **Panel (c)** displays the resulting five-cluster partition of ethnic groups in the two-dimensional PLS

¹⁹ The *elbow method* provides a simple heuristic by plotting within-cluster dispersion (WCSS) against the number of clusters K and identifying the point at which marginal gains sharply diminish, i.e. the “elbow” (Thorndike, 1953). Formally, it corresponds to the number of clusters beyond which adding an additional cluster yields only limited improvements in within-cluster fit. This approach remains widely used and is discussed in modern clustering surveys such as Suyal and Sharma (2024), which highlight its practicality. The *Calinski–Harabasz (CH) index* measures the ratio of between-cluster to within-cluster dispersion, with higher values indicating clusters that are both internally compact and well separated (Calinski and Harabasz, 1974). The *gap statistic* (Tibshirani et al., 2001) compares the observed within-cluster dispersion with that expected under a reference distribution with no clustering structure; larger values indicate stronger evidence of clustering. Its formal selection rule chooses the smallest number of clusters such that $\text{Gap}(K) \geq \text{Gap}(K + 1) - s_{K+1}$, where s_{K+1} denotes the simulation-adjusted standard error of the gap statistic at $K + 1$. Thus, an additional cluster is retained only when the improvement is sufficiently large relative to sampling uncertainty.

space, illustrating how k-means groups ethnicities according to their proximity along the empowerment-relevant cultural dimensions.


Figure A3: Selection of the Number of PLS Components and the Number of Clusters for k-means

Panel (a): Selection of the Number of PLS Components




Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. Notes: The figure plots the cross-validated mean squared error (MSE) for PLS models estimated with different numbers of latent components. The MSE indicates how well each model predicts the empowerment outcome in unseen data. The optimal number of components corresponds to the point where the MSE is minimized, beyond which adding additional components yields no improvement in predictive accuracy.

Panel (b): Selection of the Number of Clusters for k-means



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. Notes: Graph (i) displays the total within-cluster sum of squares (WCSS) for different values of k (the number of clusters). WCSS measures cluster compactness—the lower the WCSS, the more similar the observations within each cluster. Graph (ii) displays the Calinski-Harabasz (CH) index for different values of k . The CH index measures the ratio of between-cluster variance to within-cluster variance — higher values indicate clusters that are simultaneously compact internally and well-separated from each other. Graph (iii) displays the gap statistic (Tibshirani et al., 2001) with bootstrap standard error bands for different values of k . The gap statistic measures the extent to which clustering in the data exceeds what would be expected under a null reference distribution with no cluster structure - larger values indicate stronger evidence of clustering.

Panel (c): Dispersion of Ethnic Groups in the PLS Component Space and Cluster Visualization




Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS and k-means estimates. Notes: The figure illustrates the clustering of ethnic groups by PLS components.

Appendix A3: Additional Results on the Cluster Gradient of Empowerment


Figure A4: Clusters and Empowerment Distribution

Panel (a): Mean Women's Empowerment by Cluster



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS and k-means estimates. Notes: The figure illustrates cluster-level means of the women's empowerment outcome. Capped spikes indicate 95% confidence intervals around the means. p-values indicate the significance of the difference between adjacent clusters.

Panel (b): Distribution of Women's Empowerment within Clusters



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS and k-means estimates. Notes: The figure illustrates the (kernel) distribution of the women's empowerment outcome across clusters.

Appendix A4: Distribution of Empowerment Clusters by Geographical or Religious Groups

We examine how the identified cultural–empowerment clusters are distributed across regions. **Figure A5 (Panel a)** highlights clear geographical contrasts, with a gradient from clusters associated with lower empowerment, predominantly found in the Middle East and North Africa, to more egalitarian configurations prevailing in Latin America and Europe. However, no region maps one-to-one into a single cluster. Overall, geography strongly structures normative configurations but does not determine them.

The high concentration of Cluster 1 in the Middle East and North Africa is consistent with the central role of kinship and patriarchal norms in these contexts, although this cluster is also present in other regions. Sub-Saharan Africa (SSA) exhibits substantial diversity, but a predominance of Clusters 2 and 3, reflecting patriarchal structures combined with more gender-balanced economic roles. Asia displays a relatively even distribution across clusters, illustrating substantial within-region heterogeneity and the coexistence of distinct normative regimes. The most gender-neutral configurations are concentrated in Latin America and parts of Europe; yet even in these regions, more restrictive configurations remain present. This pattern underscores that normative environments are related to geography but not reducible to it.

Panel (b) illustrates how our ethnic-level approach can be used to zoom in on particular regions and countries, here focusing on Sub-Saharan Africa. This finer resolution reveals substantial heterogeneity not only across countries but also within them, with several countries, such as Namibia, encompassing ethnic groups assigned to markedly different normative regimes. Some systematic patterns nevertheless emerge. Part of the within-region variation appears correlated with religious context, although substantial heterogeneity remains, as illustrated by the broad representation of clusters in Guinea. Naturally, within-country analysis is less informative where there is little or no ethnic variation in our sample, as in Yemen. Cluster 4, in particular, largely corresponds to ethnic groups belonging to the African Matrilineal Belt, especially in Malawi, Zambia, Mozambique, and Tanzania, consistent with its distinctive prevalence of matrilineal and matrilineal arrangements. These patterns illustrate the flexibility of our framework: where ethnic variation exists, it can uncover substantial geographical variation in normative configurations below the national level.


Panel (c) presents selected pairs of geographically distant countries with similar normative-regime distributions, where similarity is measured by the Euclidean distance between their

respective vectors of ethnic-group shares across the five clusters. Some geographically distant countries display remarkably similar cluster distributions. This is the case, for example, for The Gambia and Pakistan, Mozambique and Nepal, Zambia and India, and the Philippines and Peru. These comparisons reinforce the interpretation that the regimes capture configurations of norms that can arise in very different geographical settings.

Finally, as seen above and further described in **Panel (d)**, religion is associated with the distribution of regimes, but substantial variation remains within each major religious group. Muslim ethnic groups are disproportionately represented in the two lowest-empowerment clusters, whereas Buddhist groups are concentrated in Clusters 4 and 5. Christian and Hindu groups display more dispersed distributions, with substantial representation across intermediate and higher-empowerment regimes. Thus, religion helps structure ancestral normative configurations but does not provide a one-to-one explanation of cluster membership.


Figure A5: Geographical & Religious Distribution of Clusters

Panel (a): Broad Regions




Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS and k-means estimates. Notes: This panel illustrates the share of ethnic groups, within each region, that belong to one of the clusters.

Panel (b): Zoom on Africa and MENA




Notes: This panel illustrates the share of ethnic groups, within each African/MENA country, that belong to one of the clusters.

Panel (c): Pairs of Similar Countries



Notes: This panel illustrates the share of ethnic groups within countries, focusing on pairs of "similar" countries. Similarity is measured by the Euclidean distance between their respective vectors of ethnic-group shares across the five clusters.

Panel (d): Religions



Notes: This panel illustrates the share of ethnic groups, for each religion, that belong to one of the clusters.

Appendix A5: Detailed Sensitivity Analyses

Check 1: Cultural Persistence Beyond Contemporary Conditions

A natural question is whether the cultural–empowerment clusters we identify capture variation beyond contemporary socioeconomic conditions. While women’s agency today is shaped by current levels of development, education, and labor market opportunities, it may also reflect persistent cultural influences. To assess this, we regress women’s empowerment on cluster dummies while progressively controlling for present-day factors.

Table A4 shows that the cluster gradient remains remarkably stable across specifications. Even after controlling for country-level GDP per capita, inequality, women’s educational attainment, and female labor force participation, higher-numbered clusters continue to be associated with significantly higher empowerment levels. Notably, contemporary female employment does not subsume the role of clusters linked to historically more balanced economic roles (such as Cluster 2, where male-dominated agriculture is less prevalent despite broadly patriarchal family arrangements), suggesting that current labor market outcomes do not fully capture deeper cultural influences. These findings are consistent with the literature on cultural persistence (e.g., [Giuliano and Nunn, 2021](#)) and indicate that clusters reflect enduring normative structures rather than simply differences in current development.

We further examine the role of geography and religion. Adding religious controls (model 6) and regional fixed effects (model 7) attenuates coefficient magnitudes, and the difference between Clusters 1 and 2 becomes insignificant once religion is included. This is consistent with the partial overlap between certain normative configurations and religious environments discussed above. However, differences across higher clusters remain large and statistically significant, confirming that the combinations of ancestral norms captured by our clusters explain variation in women’s empowerment beyond what is accounted for by geography and religion.

**Table A4: Cluster Differences in Women's Empowerment
After Adjusting for Contemporaneous Socio-Economic Controls**

Variables	Dep. var.: Average share of women making decision independently for health and visits, and independently or jointly for large purchases						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Cluster # = 2	0.057*** (0.014)	0.061*** (0.013)	0.070*** (0.013)	0.051*** (0.013)	0.038*** (0.014)	0.008 (0.014)	0.012 (0.014)
Cluster # = 3	0.107*** (0.014)	0.110*** (0.013)	0.117*** (0.013)	0.101*** (0.013)	0.091*** (0.013)	0.058*** (0.014)	0.051*** (0.014)
Cluster # = 4	0.137*** (0.018)	0.134*** (0.018)	0.134*** (0.017)	0.108*** (0.017)	0.102*** (0.017)	0.073*** (0.017)	0.052*** (0.017)
Cluster # = 5	0.208*** (0.017)	0.182*** (0.017)	0.189*** (0.017)	0.158*** (0.016)	0.150*** (0.017)	0.107*** (0.017)	0.075*** (0.017)
Log of country's GDP per capita		0.043*** (0.006)	0.047*** (0.006)	0.017** (0.007)	0.020*** (0.007)	0.023*** (0.007)	0.020*** (0.007)
Share of the country's richest 40%			0.102*** (0.022)	0.045** (0.022)	0.045** (0.022)	0.036* (0.021)	0.066*** (0.021)
Share of women with at least sec. education				0.161*** (0.022)	0.170*** (0.022)	0.142*** (0.022)	0.098*** (0.023)
Share of currently employed women					0.057** (0.024)	-0.006 (0.024)	0.019 (0.025)
Muslim (=1 if >50% Muslim)						-0.016 (0.018)	-0.015 (0.018)
Christian (=1 if >50% Christian)						0.058*** (0.016)	0.069*** (0.016)
Hindu (=1 if >50% Hindu)						-0.030 (0.024)	-0.065*** (0.024)
Buddhist (=1 if >50% Buddhist)						0.184*** (0.033)	0.142*** (0.033)
MENA (=1 if >50% from MENA)							-0.043 (0.028)
Sub-Saharan Africa (=1 if >50% from SSA)							-0.061*** (0.019)
Asia (=1 if >50% from Asia)							0.020 (0.019)
Europe (=1 if >50% from Europe)							-0.015 (0.031)
Observations	545	545	545	545	545	545	545
Adjusted R-squared	0.245	0.303	0.329	0.388	0.393	0.482	0.511
Joint significance of cluster identifiers:							
F-stat	45.07	37.48	40.35	31.08	29.58	20.32	8.70
p-value	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Notes: Authors' estimations based on the pooled DHS data for 70 countries and Ethnographic Atlas. Regressions are implemented at the ethnic group level, over 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS-k-means estimates. Religion variables are binary indicators of whether a given religion is practiced by more than 50% of women in each ethnic group. Regional variables indicate whether more than 50% of women in each ethnic group is from a give region. Latin and America and Caribbean (LAC) is dropped as the reference region. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1


Check 2: Sensitivity to Country-Level Residualization

Whereas Check 1 introduces contemporary controls ex post, this exercise applies a related adjustment ex ante, before constructing the cultural space. Our baseline trains the PLS model on the pooled, unresidualized ethnic-group-level empowerment outcome. Here, we first regress empowerment on the country-specific socioeconomic characteristics and broad regional and religious controls used in Check 1. We then use the residualized empowerment outcome to train the PLS model and apply k-means to the resulting components. This allows us to assess whether similar normative configurations emerge when the supervised stage targets variation in empowerment net of these observed characteristics.

Figure A6 reports the resulting cluster profiles in **Panel (a)** and differences in norm prevalence between adjacent clusters in **Panel (b)**. The main configurations and the principal contrasts associated with severe disempowerment remain broadly similar to those in **Figures 1a and 1b**. These findings complement Check 1 by showing that the central patterns persist when income and broad regional differences are accounted for before constructing the cultural space.


Figure A6: Residualization Based on Contemporaneous Socio-Economic Controls

Panel (a): Prevalence of Ancestral Cultural Norms by Cluster



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS based on residualized empowerment (using contemporary socio-economic characteristics, broad regions and broad religious groups) and k-means estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Panel (b): Difference in Cultural Norm Prevalence
across Adjacent Clusters and Highlights of Pivotal Norms



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS based on residualized empowerment and k-means estimates. Notes: Each cell reports the difference in cultural norm prevalence between adjacent clusters.

Check 3: Sensitivity to Matching Methods and Sample-Size Threshold

Some links between DHS respondents and ancestral groups rely on geographic proximity when direct ethnicity- or language-based matching is unavailable. Although geographic linkage expands sample coverage and is used in the broader literature on ancestral institutions,²⁰ it may introduce measurement error: current residents need not descend from the group historically associated with their location, and ancestral groups may have overlapped geographically.

We assess sensitivity to matching methods through two exercises (**Figure A7**). In **Panel (a)**, we apply geographic matching to respondents for whom direct ethnicity- or language-based matches are also available and re-estimate the PLS and k-means procedure. In **Panel (b)**, we restrict the analysis to directly matched respondents. Both exercises yield normative configurations broadly similar to the baseline. Some differences arise among intermediate clusters, but the principal contrasts associated with severe disempowerment remain stable.

As an additional check, we calculate each ethnic group's shares of DHS observations linked through the different matching methods and include these shares as controls in the regression of empowerment on cluster dummies. The cluster coefficients remain comparable in magnitude, statistically significant and ordered as in the baseline, indicating that differences in matching composition do not account for the empowerment gradient.²¹


A separate concern is that ethnic-group averages of empowerment differ in precision. Groups close to the minimum threshold of 50 sampled women may have noisier estimates than larger groups, although both receive equal weight in PLS and clustering. We therefore re-estimate the full procedure after raising this threshold to 100 women. This retains 496 of the 545 ethnic groups (91%) and yields cluster profiles and an empowerment gradient closely comparable to the baseline (**Panel (c)**). Taken together, these exercises show that the main findings are robust to alternative matching methods and a stricter sample-size requirement.

²⁰ For example, [Michalopoulos and Papaioannou \(2013\)](#) link historical ethnic homelands and ancestral institutions to contemporary regional economic outcomes. Their analysis illustrates the use of geographic linkage, rather than direct identification of individual ancestry.

²¹ Relative to Cluster 1, baseline coefficients for Clusters 2–5 are 0.057, 0.107, 0.137 and 0.208. After controlling for matching-method shares, the corresponding coefficients are 0.058, 0.100, 0.121 and 0.200; all remain significant at the 1% level.


**Figure A7: Prevalence of Ancestral Cultural Norms by Cluster,
Sensitivity to Matching Methods and Sample-Size Threshold**

Panel (a): Using Geographic Matching when Direct Individual Matching is Available



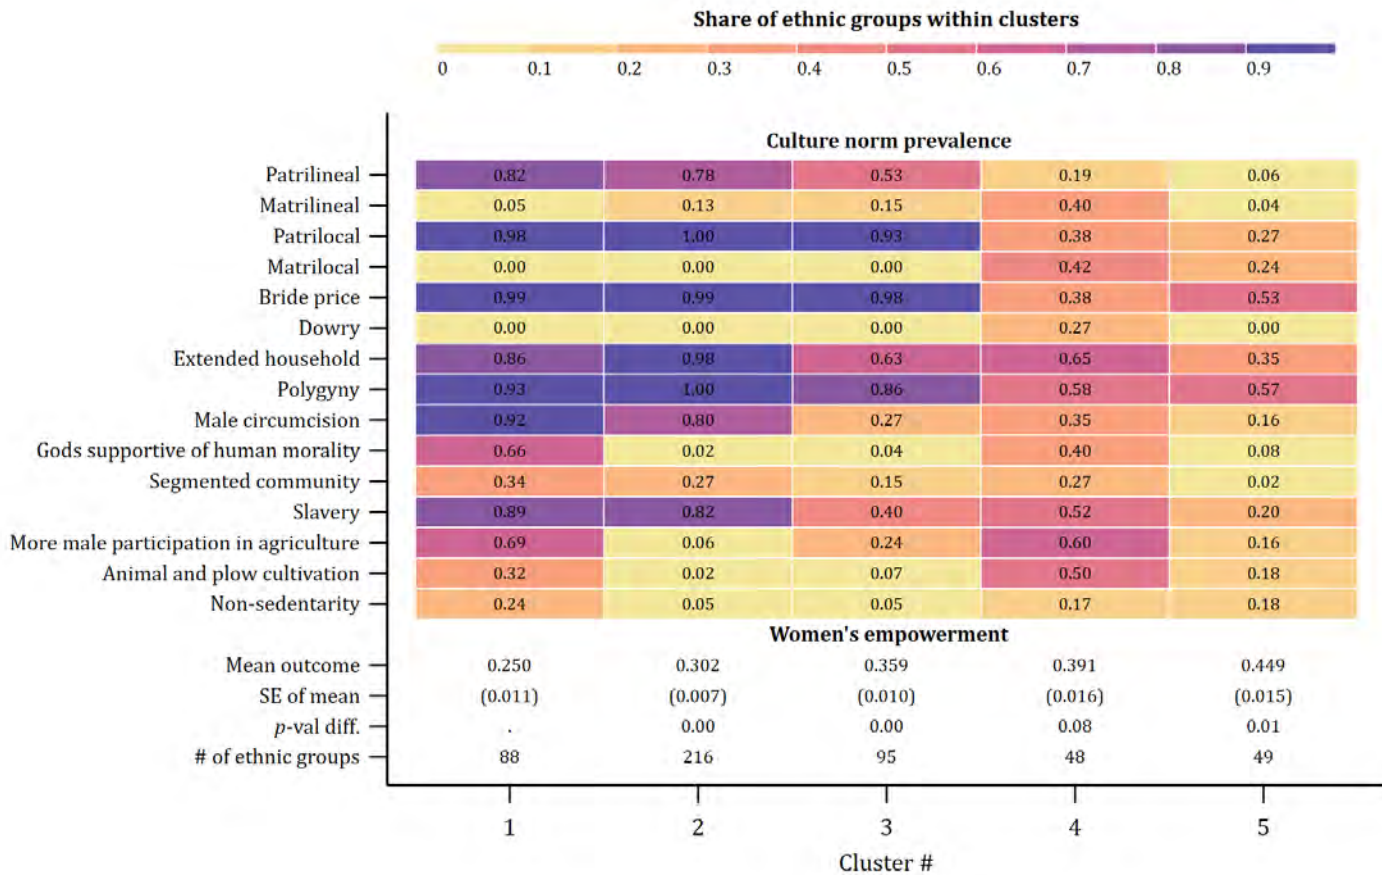
Source: Authors' estimations based on the pooled DHS data for 70 countries and 569 ethnic groups from the Ethnographic Atlas (matched based on geographic information even when direct matching is available). The cluster information comes from the PLS and k-means estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Panel (b): Restricting the Analysis to Ethnic Groups Matched Directly



Source: Authors' estimations based on the pooled DHS data for 70 countries and 317 ethnic groups from the Ethnographic Atlas (matched based on individual-level information only). The cluster information comes from the PLS and k-means estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Panel (c): Restricting the Analysis to Ethnic Groups with at Least 100 Women



Source: Authors' estimations based on the pooled DHS data for 70 countries and 496 ethnic groups from the Ethnographic Atlas. Unlike Figure 1, this panel restricts the sample to ethnic groups with at least 100 sampled women, hence retaining 496 of the 545 baseline groups (91%), and re-estimates the full PLS and k-means procedure. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Check 4: Sensitivity to Clustering Method and Number of Clusters


We examine whether the identified configurations depend on the clustering algorithm or the number of clusters. First, we replace k-means with Ward's agglomerative hierarchical clustering, applied to the same PLS components (Ward, 1963). This method progressively merges clusters to minimize the increase in within-cluster variation at each step. Although both methods favour compact groups, they construct partitions through different procedures. As shown in **Figure A8, Panel (a)**, the resulting normative profiles closely resemble the baseline configurations.

We next increase the number of clusters from five to seven and nine, in **Panels (b) and (c)** respectively. These specifications reveal finer distinctions among normative configurations, while preserving the principal contrasts of the baseline. With seven clusters, additional intermediate configurations emerge, including distinct matrilineal profiles associated with different empowerment levels. With nine clusters, broadly patriarchal configurations at the lower end of the distribution are further differentiated, while matrilineal arrangements again characterize several regimes at intermediate and higher levels of empowerment.

Although empowerment differences between some adjacent clusters become less pronounced as the partition becomes finer, the same type of non-monotonicity remains evident. Most importantly, moralizing gods and agricultural organization, particularly women's participation and plough use, retain their pivotal role at the lower end of the empowerment gradient.


**Figure A8: Prevalence of Ancestral Cultural Norms by Cluster,
Alternative Clustering Approaches**

Panel (a): Alternative Clustering Algorithm



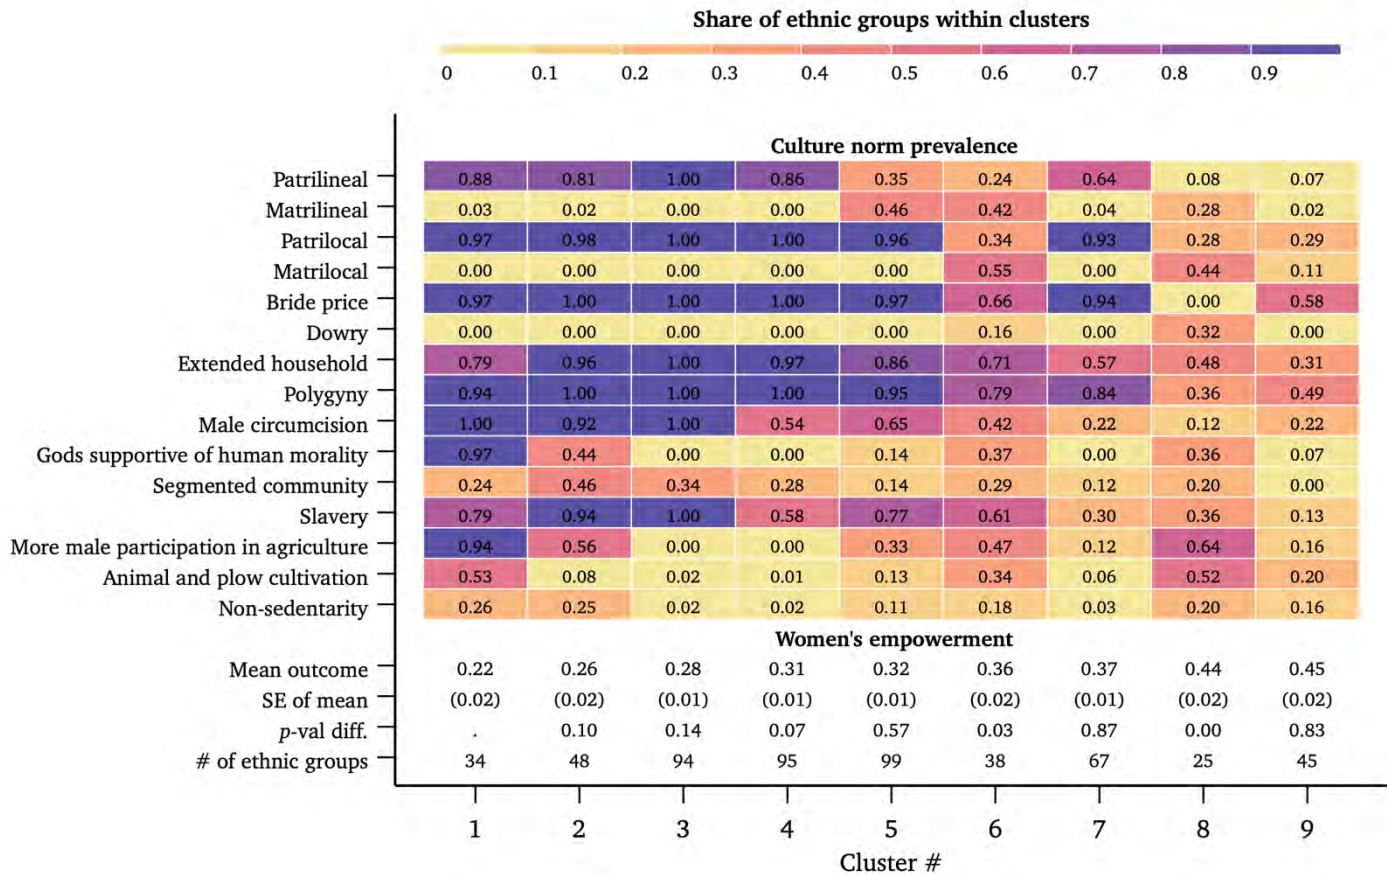
Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS and Ward hierarchical estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Panel (b): Cluster set with k=7



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS and k-means estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Panel (c): Cluster set with k=9



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS and k-means estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Check 5: Alternative Empowerment Outcome – Attitudes towards Domestic Violence

Our baseline outcome captures women’s participation in household decision-making. As a robustness check, we examine women’s attitudes toward wife-beating, which capture a related but distinct dimension: acceptance of gendered authority and the perceived legitimacy of violence within marriage.²²


(a) Attitudes toward Wife-Beating across Baseline Clusters

We first examine whether the configurations identified using decision-making also differentiate attitudes toward wife-beating. Holding baseline cluster assignments fixed, we compare women’s rejection of five commonly surveyed justifications: going out without informing the husband, neglecting children, arguing with the husband, burning food, and refusing sexual relations. The patterns are similar across the five items; for concision, we report an index averaging disapproval across them.

Figure A9 shows that clusters associated with greater decision-making autonomy also exhibit greater disapproval of wife-beating. This correspondence extends the empowerment gradient to an outcome not used to construct the baseline PLS components, supporting the relevance of the identified normative configurations beyond the particular decision-making measure used in the main analysis.

²² We focus on attitudes toward wife-beating because they directly concern the legitimacy of gendered authority within marriage. Outcomes such as nutrition, health and broader well-being also depend heavily on household resources and service access, making their relationship with intra-household agency less direct. Reported attitudes nevertheless remain distinct from experienced violence and may reflect contemporary circumstances (Díaz & Saldarriaga, 2023) as well as inherited norms.

Figure A9: Disapproval of Wife-Beating across Baseline Clusters



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the baseline PLS and k-means estimates based on the empowerment index. The figure illustrates cluster-level means of disapproval of wife-beating as an alternative outcome. Capped spikes indicate 95% confidence intervals around the means. p-values indicate the significance of the difference between adjacent clusters.


(b) Sensitivity to Using Attitudes toward Wife-Beating as a Supervised Outcome

We re-estimate the full PLS and k-means procedure under two alternative supervision specifications: (i) women's disapproval of wife-beating alone, and (ii) disapproval of wife-beating jointly with the baseline decision-making index. The second specification uses PLS's capacity to accommodate multiple outcomes, assessing whether constructing the cultural space around both dimensions yields configurations comparable to the baseline.

Figure A10 shows broadly similar normative profiles under both specifications compared with **Figure 1a**. Configurations corresponding to baseline Clusters 1, 2 and 5 remain particularly recognizable, with some differences in the prevalence of individual norms, such as community segmentation. Most importantly, the principal religious and agricultural contrasts identified in the baseline appear pivotal again: configurations associated with greater disapproval of wife-beating exhibit a lower prevalence of moralizing gods and male-dominated agriculture. These results show that the main normative contrasts remain broadly consistent when using a substantively meaningful alternative outcome to guide the construction of the cultural space.

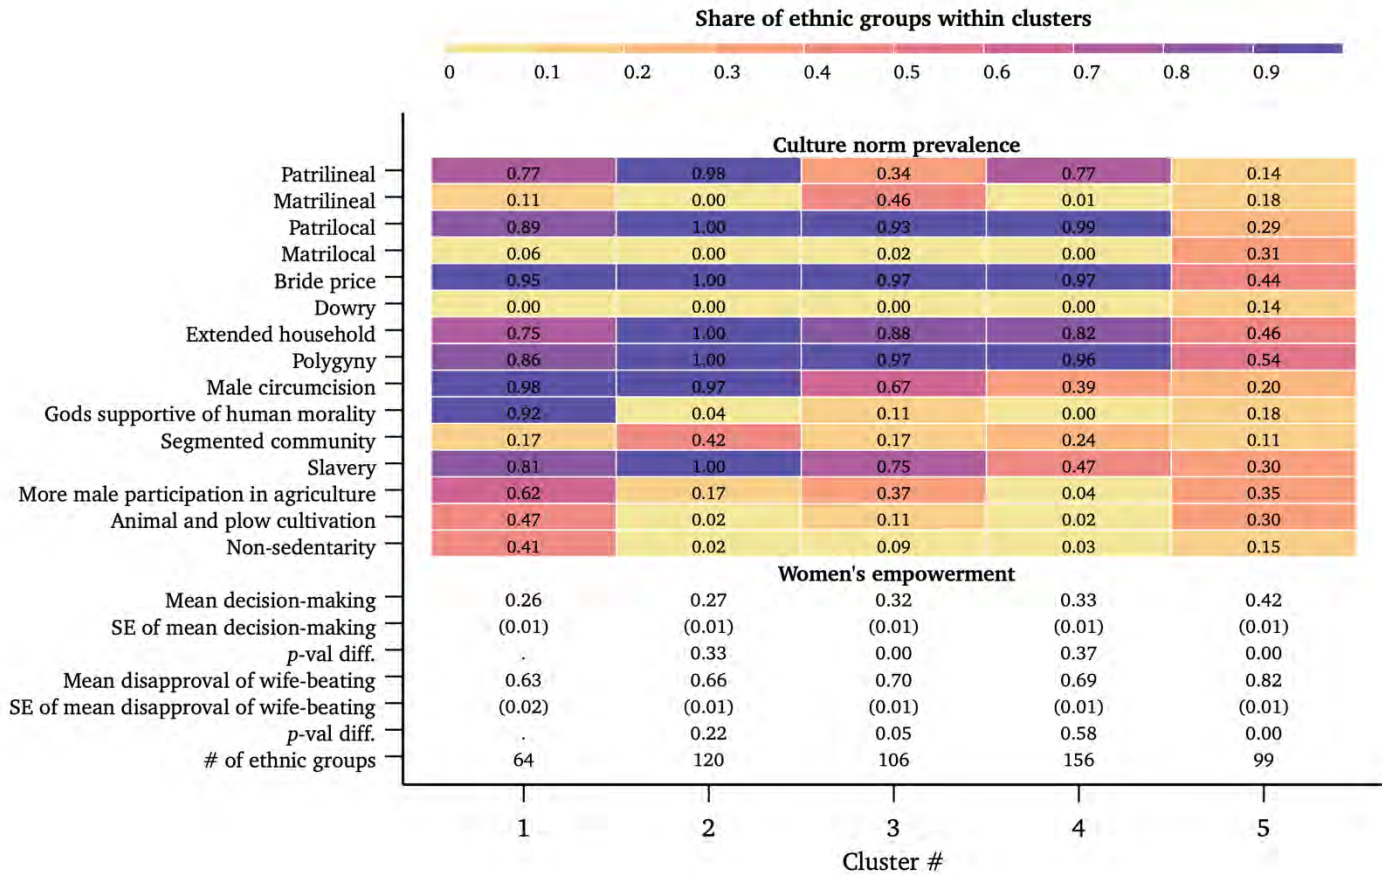
Figure A10: Prevalence of Ancestral Cultural Norms by Cluster – Alternative Measures of Women’s Empowerment as Supervised Outcome

Panel (a) Women’s disapproval of wife-beating



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS and k-means estimates, with women’s disapproval of wife-beating used as a supervised outcome. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Panel (b) Women's decision-making autonomy + disapproval of wife-beating



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS and k-means estimates, with women's decision-making and disapproval of wife-beating used as multiple supervised outcomes. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral cultural norm. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

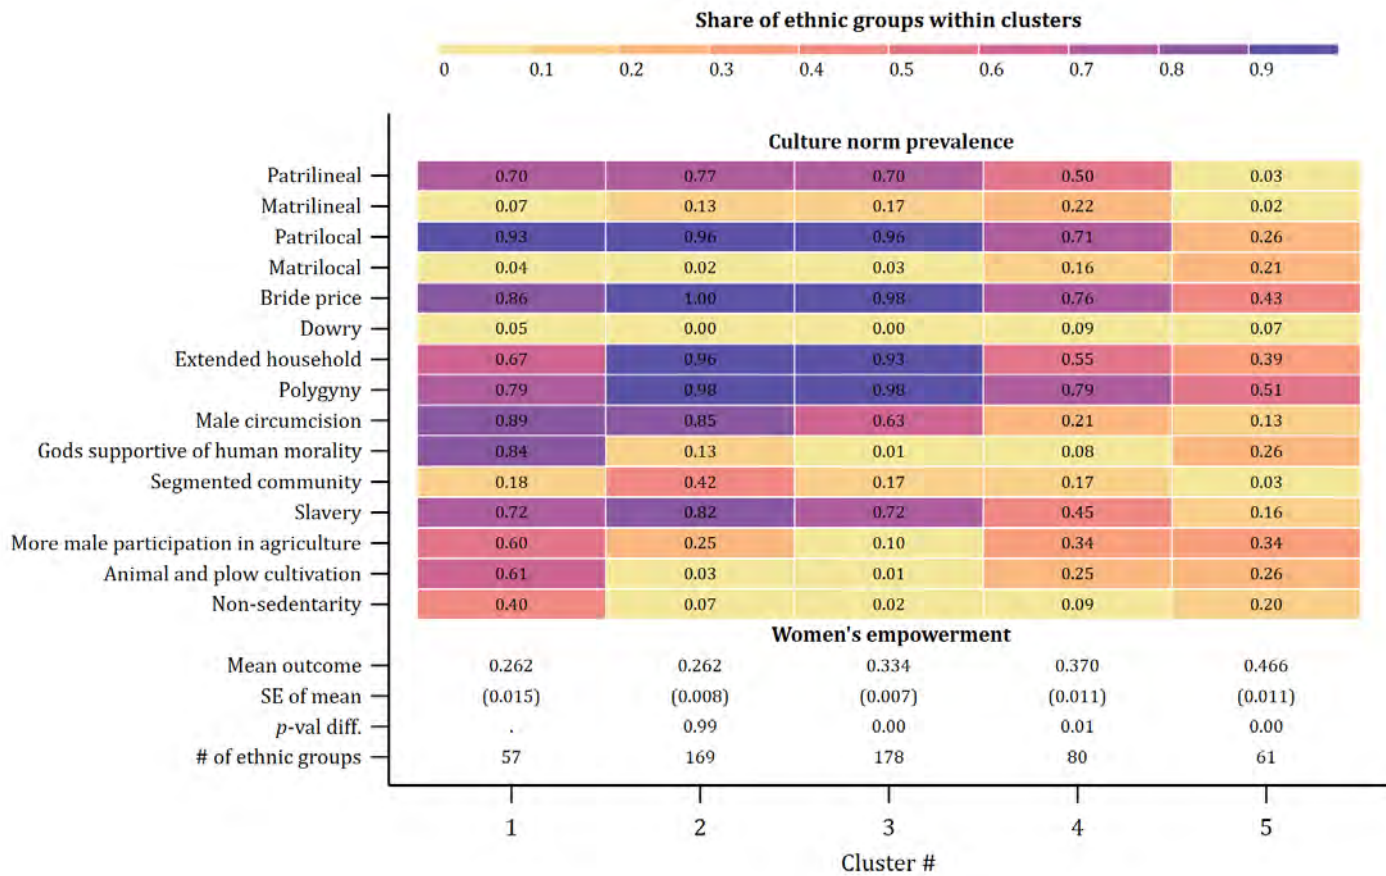
Check 6: Robustness to Norm Pre-Grouping

Our baseline specification combines original EA categories within each normative domain when their estimated associations with empowerment are statistically indistinguishable (Section 3.2, Table 1). This simplifies the representation of ancestral norms before the PLS step. Because the grouping itself uses outcome information, we assess whether the identified configurations depend on these preliminary restrictions.

We re-estimate the PLS and k-means procedure using disaggregated EA categories, with domain-specific reference categories selected to align with the baseline coding. PLS is estimated using these category indicators, and k-means is applied to the resulting components. For comparison with **Figure 1a**, we then display the prevalence of norms within the newly estimated clusters using the baseline groupings in **Table 1**. This regrouping is performed solely for presentation and does not affect component extraction or cluster assignments in this exercise.

The resulting profiles (**Figure A11**) closely resemble the baseline configurations and preserve their principal normative contrasts. These findings indicate that the main patterns do not depend on the outcome-informed pre-grouping: a similar characterization emerges when the finer distinctions among EA categories are retained during estimation.

Figure A11: Prevalence of Ancestral Cultural Norms by Cluster – PLS and Clustering on Original, Un-Aggregated Norm Categories



Source: Authors' estimations based on the pooled DHS data for 70 countries and 545 ethnic groups from the Ethnographic Atlas. The cluster information comes from the PLS and k-means estimates. Notes: Each cell reports the share of ethnic groups within a given cluster that practice an ancestral norm, aggregated post hoc into our baseline groupings (Table 1) for display purposes only. Darker shading indicates higher prevalence of a norm within the cluster. Norms are grouped vertically by cultural domain (e.g., kinship system, post-marital residence, marriage payments, etc.), allowing visual comparison of how combinations of norms vary across clusters.

Check 7: Sensitivity to the Degree of Supervision

Our baseline PLS specification deliberately uses women's empowerment to identify the most relevant directions in the cultural norm space. A possible concern is therefore that the resulting regimes and pivotal margins reflect the particular degree of supervision imposed by PLS. We assess this sensitivity using *continuum regression* (Stone and Brooks, 1990), which varies the relative importance attached to covariance with the outcome and to variation in the predictors. The parameter γ acts as a supervision parameter: $\gamma=0$ corresponds to OLS, $\gamma=1$ to PLS, and the solution approaches PCA as γ increases beyond 1. At each value, we sequentially extract two components with orthogonal scores and apply the same five-cluster k-means procedure.

Panel (a) of **Figure A12** reports, for each component, the share of the maximum predictor variance attained by the corresponding PCA component. As supervision decreases, the continuum components progressively approach the PCA benchmark. For the first component, this share rises from 90.0% under PLS to 98.8% at $\gamma=5$ and 99.9% at $\gamma=20$. For the second component, it increases more markedly, from 57.0% under PLS to 89.2% at $\gamma=5$ and 98.5% at $\gamma=20$. Thus, the larger parameter values yield nearly variance-maximizing components, providing a substantive sensitivity test beyond small departures from PLS.


Panels (b)–(d) present cluster profiles for $\gamma = 2, 5$ and 20 , respectively; **Panel (e)** compares differences in norm prevalence between adjacent clusters. The partitions change appreciably as supervision decreases. Components increasingly prioritize differences that account for variation in the norms themselves, even when those differences are less closely associated with empowerment. Consequently, culturally distinct configurations need not be as clearly separated on the outcome. This is particularly apparent among the higher-ranked clusters, which exhibit different normative profiles despite similar average empowerment.

Most importantly, the principal religious and agricultural contrasts at the bottom of the empowerment distribution survive this reorganization. Across the three specifications, the lowest-ranked cluster exhibits substantially higher prevalence of moralizing gods, male-dominated agriculture and plough cultivation than the next cluster (**Panel (e)**), while both retain broadly patriarchal kinship and marital arrangements. Although the difference in mean empowerment between these two clusters is no longer statistically significant at conventional levels at $\gamma = 5$ and $\gamma = 20$ ($p = 0.13$ and $p = 0.22$, respectively), their contrast in normative composition remains pronounced. Kinship and post-marital residence continue to differentiate configurations further up, and non-monotonic patterns remain evident, although the specific configurations and positions at which norms reappear change. The


strongest robustness finding is therefore the continued prominence of religious and agricultural distinctions at the lower end of the distribution, rather than the stability of the entire partition.

Figure A12: Sensitivity to the Degree of Supervision using Continuum Regression


Panel (a): Predictor Variance Captured Relative to PCA



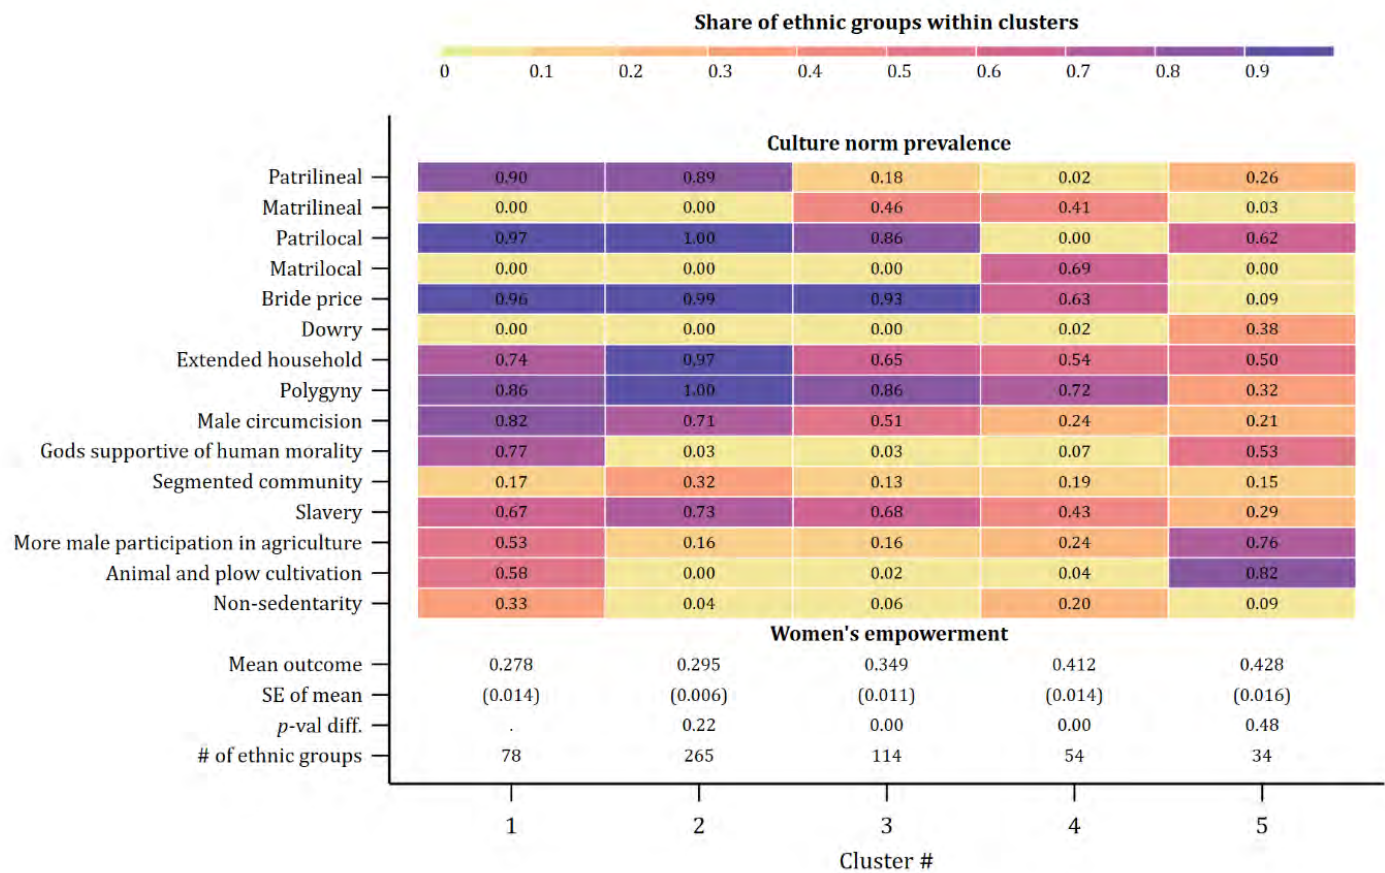
Panel (b): Prevalence of Ancestral Cultural Norms by Cluster, Reduced Supervision ($\gamma=2$)




Panel (c): Prevalence of Ancestral Cultural Norms by Cluster, Reduced Supervision ($\gamma=5$)



Panel (d): Prevalence of Ancestral Cultural Norms by Cluster, Reduced Supervision ($\gamma=20$)



Panel (e): Differences in Norm Prevalence across Adjacent Clusters under Reduced Supervision



Additional References for the Appendix

- Abdi, H. (2010). Partial least squares regression and projection on latent structure regression (PLS-Regression). *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(1), 97–106.
- Alesina, A., Michalopoulos, S., & Papaioannou, E. (2016). Ethnic inequality. *Journal of Political Economy*, 124(2), 428–488.
- Boulesteix, A.-L., & Strimmer, K. (2007). Partial least squares: A versatile tool for the analysis of high-dimensional genomic data. *Briefings in Bioinformatics*, 8(1), 32–44.
- Caliński, T., & Harabasz, J. (1974). A dendrite method for cluster analysis. *Communications in Statistics*, 3(1), 1–27.
- Cao, Y., Enke, B., Falk, A., Giuliano, P., & Nunn, N. (2021). *Herding, armed conflict, and a culture of honor: Global evidence* (NBER Working Paper No. 29250). National Bureau of Economic Research.
- Corno, L., La Ferrara, E., & Voena, A. (2020). Female genital cutting and the slave trade. CEPR Discussion Paper No. DP15577.
- Díaz, J. J., & Saldarriaga, V. (2023). A drop of love? Rainfall shocks and spousal abuse: Evidence from rural Peru. *Journal of Health Economics*, 89, 102739.
- Giuliano, P. (2018). Gender: A Historical Perspective, in Susan L. Averett, Laura M. Argys, and Saul D. Hoffman (eds), *The Oxford Handbook of Women and the Economy*, Oxford Handbooks.
- Giuliano, P., & Nunn, N. (2021). Understanding cultural persistence and change. *Review of Economic Studies*, 88(4), 1541–1581.
- Fenske, J. (2015). African polygamy: Past and present. *Journal of Development Economics*, 117, 58–73.
- Fredriksson, A., & Gupta, B. (2023). Irrigation, gender roles, and female labor supply. *Journal of Development Economics*, 161, 103032.
- Hansen, C. W., Jensen, P. S., & Skovsgaard, C. V. (2015). Modern gender roles and agricultural history: The Neolithic inheritance. *Journal of Economic Growth*, 20(3), 365–404.
- Höskuldsson, A. (1988). PLS regression methods. *Journal of Chemometrics*, 2(3), 211–228.
- Kincaide, L.-T., McGuirk, E. F., & Nunn, N. (2025). *A comprehensive concordance between the Ethnographic Atlas and Murdock's map of ethnic groups of the African continent*.
- Miotto, M. (2023). Colonialism, cash crops and women in Africa, CERGE-EI Working Paper No. 750.
- Murdock, G. P. (1959). *Africa: Its peoples and their culture history*. McGraw-Hill.

- Suyal, M., & Sharma, S. (2024). A review on analysis of K-means clustering machine learning algorithm based on unsupervised learning. *Journal of Artificial Intelligence and Systems*, 6, 85–95.
- Thorndike, R. L. (1953). Who belongs in the family? *Psychometrika*, 18(4), 267–276.
- Tibshirani, R., Walther, G., & Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2), 411–423.
- Ward, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301), 236–244.
- Wold, H. (1975). Soft modelling by latent variables: The non-linear iterative partial least squares (NIPALS) approach. In J. Gani (Ed.), *Perspectives in Probability and Statistics* (pp. 520–540). Academic Press.
- Wold, S., Sjöström, M., & Eriksson, L. (2001). PLS-regression: A basic tool of chemometrics. *Chemometrics and Intelligent Laboratory Systems*, 58(2), 109–130.