FOLKLORE

Stelios Michalopoulos
Melanie Meng Xue

We are extremely grateful to Yuri Berezkin for generously sharing his lifetime work on folklore classification and for clarifying many of its aspects. The Editor and 3 referees offered excellent suggestions. We would like to also thank Alberto Alesina, Elias Papaioannou, Benjamin Enke, Rafael La Porta, Ernesto Dal Bo, Jesse Shapiro, Roland Bénabou, David Weil, and Max Winkler for their feedback. Seminar and workshop participants at numerous venues offered many useful suggestions. Kush Bavaria, Vafa Behnam, Rohit Chaparala, Elizabeth Dimen, Adrien Foutelet, Masahiro Kubo, Xueyun Luo, Zhihan Wang, and Xiaochen Yan provided superlative research assistance. Stelios Michalopoulos acknowledges financial support from the PSTC at Brown University; Corresponding author: Stylianos_Michalopoulos@brown.edu. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

ABSTRACT

Folklore is the collection of traditional beliefs, customs, and stories of a community passed through the generations by word of mouth. We introduce to economics a unique catalogue of oral traditions spanning approximately 1,000 societies. After validating the catalogue's content by showing that the groups' motifs reflect known geographic and social attributes, we present two sets of applications. First, we illustrate how to fill in the gaps and expand upon a group's ethnographic record, focusing on political complexity, high gods, and trade. Second, we discuss how machine learning and human-classification methods can help shed light on cultural traits, using gender roles, attitudes towards risk, and trust as examples. Societies with tales portraying men as dominant and women as submissive tend to relegate their women to subordinate positions in their communities, both historically and today. More risk-averse and less entrepreneurial people grew up listening to stories where competitions and challenges are more likely to be harmful than beneficial. Communities with low tolerance towards antisocial behavior, captured by the prevalence of tricksters getting punished, are more trusting and prosperous today. These patterns hold across groups, countries, and second- generation immigrants. Overall, the results highlight the significance of folklore in cultural economics, calling for additional applications.

Stelios Michalopoulos
Brown University
Department of Economics
64 Waterman Street
Providence, RI 02912
and CEPR
and also NBER
smichalo@brown.edu

Melanie Meng Xue
New York University - Abu Dhabi
Social Science Division
Abu Dhabi
melanie.xue@nyu.edu

1 Introduction

Two broad observations motivate our study. First, narratives are central building blocks of our societies. We think in stories and explain the world by telling stories. Harari (2015), for example, identifies in the myths present in the peoples' collective imagination the roots of their successes and failures. Despite their central role in connecting actions to values and needs, economists have only recently turned to the study of narratives, (e.g., Akerlof and Snower [2016]; Shiller [2017]).[1]

Second, during the last two decades, a burgeoning body of work exploring the cultural, historical, and institutional roots of comparative development highlights the significance of ethnic, linguistic, and religious groups (Michalopoulos and Papaioannou [2017]; Nunn [2020]). Much of this research, however, relies on valuable but incomplete ethnographic sources including the widely-used *Ethnographic Atlas* (EA). Furthermore, the absence of proxies of historical norms renders inquiries into how attitudes change and why they persist intractable.[2] Other weaknesses of the EA, the celebrated compilation of George Peter Murdock (1967), regard the uneven coverage of groups and attributes and measurement error.

In this study, we leverage a group's oral tradition to shed light on its cultural heritage and past social and economic structures. According to the Oxford Dictionary, folklore consists of the traditional beliefs, customs, and stories of a community, passed through the generations by word of mouth. This corpus is the subject of the discipline of folklore. We do three things to reveal the potential of integrating traditional narratives in the toolset of economists and political scientists interested in the origins of comparative development, gender norms, morality, psychology, and culture.

First, we introduce to economics a catalogue of folklore that codes the distribution of thousands of motifs across 958 world societies. This database is the lifetime work of the eminent anthropologist and folklorist Yuri Berezkin. A motif, according to the author, is an episode or an image found in the set of narratives recorded within an ethnolinguistic community. We validate the catalogue's content by establishing that images and episodes in a group's oral tradition reflect salient features of its physical environment. For example, groups closer to earthquake-prone regions have a higher incidence of earthquake-related motifs, groups on fertile land have more crop-related images, and groups living close to rivers (or in colder climates) have more

[1] See Benabou, Falk, and Tirole (2018) for a theoretical investigation of the interplay between narratives and moral reasoning. Rodrik (2017) highlights the parallels between the structure of fables and that of economic models.

[2] Attempts to address these issues have been made in the context of specific traits, regions, and periods, but a comprehensive answer is lacking. Algan and Cahuc (2010), for example, use immigrants in the US over time to uncover trust values for most of the 20th century and Voigtlander and Voth (2012) document the persistence of anti-Semitism in Germany. Nunn and Qian (2020) show that historical climate shocks weaken the transmission of cultural traits today. Finally, see Chen (2013) and Galor, Ozak, and Sarid (2016) for attempts to link linguistic features regarding, for example, the structure of the future tense to contemporary cultural attributes.

episodes reflecting their respective landscapes. Then, we link the groups in Berezkin's collection to the EA and show that the folklore-based measures of political complexity, family structure, and subsistence mode robustly correspond to EA's analogous traits.

Second, we illustrate how to use a group's oral tradition to fill in gaps in the ethnographic record, focusing on the degree of political complexity and the presence of high gods. In addition, we show how one can use folklore to quantify the extent of the pre-industrial market economy, a key economic aspect that the EA does not cover.

Third, we present a method to uncover a group's cultural heritage that involves the reading and classification of motifs by multiple individuals.[3] We focus on trust, risk-taking, and gender norms to illustrate our approach. To capture trust, we look at how tricksters (a common archetype in oral traditions) are depicted in the motifs, distinguishing between instances where their deceiving behavior is successful or punished. Regarding risk-taking, we look at how challenges and competitions are portrayed, differentiating between tragedies and victories. To measure gender norms, we classify the various stereotypical roles males and females play in the motifs.

These folklore-based measures of historical attitudes are robust predictors of contemporary values and economic choices. Folks that grew up listening to stories where tricksters often fail to deceive their victims are more trusting and prosperous today. Groups with oral traditions rich with heroes who successfully tackle challenging situations tend to display more appetite for risk and appear more entrepreneurial. Societies whose folklore portrays females as less dominant, more submissive, and more likely to engage in domestic affairs than males tend to relegate their women to inferior roles in their communities, both historically and today. These patterns hold across countries, second-generation immigrants, and ethnic groups, suggesting that folklore may be one of the vehicles via which norms are intergenerationally transmitted.

We organize the paper as follows. In Section 2, we introduce and describe the work of folklorist Yuri Berezkin. In Section 3, we validate the catalogue's content using various checks. In Section 4, we use folklore to improve upon the ethnographic record. In Section 5, we show how to quantify ancestral beliefs and attitudes from a group's oral tradition. In Section 6, we offer some thoughts on future work.

[3] In the Appendix, we discuss the promise and pitfalls of applying machine learning techniques to extracting values from folklore. We apply shrinkage methods over thousands of keywords in the folklore catalogue to uncover concepts that predict contemporary values as reflected in the Global Preference Survey (GPS) dataset of Falk et al. (2018).

2 Berezkin's Folklore and Mythology Catalogue

For the most part, folklorists limit their studies to small, well-defined collections with a regional focus (see El-Shamy [2004], for example, for a classification of folktales in the Arab world). Yuri Berezkin's Folklore and Mythology Catalogue is an exception, representing a global comparative perspective of oral traditions. In Appendix Section 1, we offer a brief overview of folklore studies. Berezkin has constructed a unique mythology and folklore database for 958 groups worldwide, initially cataloguing motifs of indigenous societies in the Americas and later extending his classification to groups in the Old World. From the 958 world groups he has categorized 2,564 motifs. A motif reflects a combination of images, episodes, or structural elements found in two or more texts, including sacred and profane ones (Berezkin [2015a]). See Figure *I* for the spatial distribution of groups in Berezkin's catalogue.[4]

Berezkin is interested in understanding the historical spread of motifs across societies. His research explores the areal distribution of specific (sets of) motifs in relation to large-scale population movements, migrations, cultural contacts and interactions in history and prehistory, see Berezkin and Duvakin (2016), and Berezkin (2015b). To encode motifs, he has consulted a list of 6,239 books (67%) and journal articles (33%) from 4,041 authors, edited by 4,932 publishing houses in 32 different languages.[5] For the average society, Berezkin consulted 15 publications from 12 authors, written in 3.5 different languages from 14 publishers. The bulk of materials appeared during the 20th century with the earliest and usually the key publication for the median society dating back to 1904. The date of the earliest publication consulted per oral tradition is in Appendix Figures 1. The most common languages are English, Russian, Spanish, German, French, and Portuguese (Appendix Table 1 Panel *A*).

The median group in Berezkin has 62 motifs. See Figure *II* for a continent-specific histogram of the number of motifs across groups. Russians have the largest number of motifs, and one group, the Yeyi in northwestern Botswana, has only a single recorded motif. In Appendix Table 1, Panel *B*, we report the top 10 societies in terms of motifs. There are two broad classifications of motifs, those describing adventures and tricks (roughly 50%) and those related to cosmology and etiology.

The building blocks of our analysis are the motifs in Berezkin's catalogue. Some motifs appear to be universal, whereas others are only found in a handful of folklore traditions. The

[4] Since introducing his catalogue in 2015, Berezkin updates the motifs and groups sporadically. We use the October 2019 version. In personal correspondence, he indicated that the updating of oral traditions for groups in the Americas and Europe is complete. For groups in Africa, Asia, and Oceania, the coverage is above 85%. The updates mostly affect small, isolated groups.

[5] Examples of folklore journals include: The Folk-Lore Journal, Journal of American Folklore, Journal of Latin American Lore, Latin American Indian Literature Journal, The Journal of the Polynesian Society, Journal de la Societe des Americanistes, The Journal of the Anthropological Institute of Great Britain and Ireland, The Journal of Indo-European Studies, Journal de la Societe Finno-Ougrienne, Oral Traditions, etc.

median motif spans 18 oral traditions. Appendix Figure 2 shows the histogram of the incidence of motifs across oral traditions. The most widespread motif, $k27n$, is titled "Tasks of the in-laws" and reads, "Father or other kinsmen of hero's wife or bride try to kill or test him and/or suggest to him difficult tasks." This theme can be found in 355 out of the 958 oral traditions spanning all continents, see Appendix Figure 3a. Eleven motifs appear in exactly two oral traditions. For example, motif $b107$ found among the Heiltsuk in today's British Columbia and the Flathead in Western Montana, reads: "During a (world) fire, the oyster burns, so the oyster shells are black," see Appendix Figure 3b. In Appendix Section 2, we trace some of the original tales used to construct a common motif and describe the underlying sources of the Telugu oral tradition.

Each motif in Berezkin's catalogue is accompanied by a title and a short description of an image or an episode in the group's oral tradition. Therefore, the underlying data is text.[6] We employ two approaches to quantify the aspects of folklore that we are interested in. The first one is a dictionary-based method that breaks down the text into words and connects word counts to the attributes we wish to measure. It is the simplest and most frequently used method.[7] We use the related words from the word embeddings in ConceptNet for any keywords of interest. The second approach relies on the human classification of a motif's content.

2.1 Caveats

There are several aspects of this database worth discussing. One natural question regards when the stories appeared in a given oral tradition. Folklorists are interested in collecting stories untouched by modernization and collectors perceived narratives as having a precise origin in space and time. This idea has been criticized by Jason (1970) and Goldberg (1984) due to the inherent uncertainty in developing convincing estimates. Nevertheless, certain types of motifs likely predate others. For example, cosmological ones are considered significantly older than other types of motifs. In a seminal work, Silva and Jamshid (2016) integrate insights from phylogenetic analysis of languages to establish that the tales of the Beauty and the Beast and Rumpelstiltskin originated sometime between $2,500$ and $6,000$ years ago. So, it is useful to keep in mind that what appears in folklore images is potentially from a period long before the one captured by the EA.

Second, in his catalogue, Berezkin (as all works in comparative mythology) does not code images found in a single oral tradition, but focuses on motifs shared by at least two groups. In

[6] Over the last few years, text analysis has seen advances thanks to the abundance of digital text and the development of computing techniques. See Grimmer and Stewart (2013) and Evans and Aceves (2016) for reviews of studies in text analysis in political science and sociology, respectively. Gentzkow, Kelly, and Taddy (2019) provide an excellent entry into the weaknesses and advantages of the available text-analysis methods.

[7] Gentzkow and Shapiro (2010) count the number of newspaper articles containing partisan phrases, whereas Saiz and Simonsohn (2013) enter search queries into Google to obtain document-frequency measures of corruption. Tetlock (2007), Baker, Bloom, and Davis (2016) and Enke (2020), also use a pre-specified dictionary of terms.

Appendix Section 3, we shed some light on the magnitude of motif borrowing by looking at how the incidence of specific concepts in an oral tradition influences the frequency of such images in a neighboring group's folklore.

Third, this catalogue does not account for the popularity of a motif, nor the number of alternative versions of a particular tale. Instead, our folklore images simply reflect whether or not the motif exists in a society. How deeply a given oral tradition has been studied will naturally influence the number of variants recorded per tale. So, focusing on whether a given motif is present may help mitigate concerns regarding differential sampling across societies.[8] Similarly, we do not know the number of tales and legends in a given oral tradition since there is no one-to-one mapping between myths and motifs. One tale may map into multiple motifs and vice versa.

Fourth, the definition of a group in Berezkin's catalogue is usually along linguistic lines. Sometimes, he groups related societies when he does not have enough information.

Fifth, the number of stories recorded for a group is influenced by the degree to which that group was studied by folklorists. To address this issue, we often control for country-specific constants in our empirical analysis. By comparing oral traditions within the same country, we mitigate concerns related to the unbalanced coverage across countries in terms of the quality and breadth of the recorded oral traditions.[9] Taking a step further, we control for the number of publications in a group's folklore.

Sixth, there is substantial heterogeneity in folklore collection over time. Early folklorists were often missionaries or curious intellectuals, whereas later generations were frequently academics, including anthropologists conducting fieldwork in small-scale societies. Differences in the timing of those being surveyed may reflect differences in the tastes, biases, and interests of the story collectors. To partially account for this, we control for the year of the earliest publication per oral tradition.[10]

Finally, Berezkin's goal is to produce a catalogue of images and episodes that appear in a given oral tradition. He *does not* attempt to infer the moral or reconstruct the hidden meaning of

[8] For instance, Ross, Greenhill, and Atkinson (2013) study the diffusion of 700 variants of the story "The Kind and the Unkind Girls" in Europe. In Berezkin's work, it is a *single* motif, *k*56 (present in 169 oral traditions), that subsumes these numerous variants. This example illustrates the usefulness of a motif as an aggregator of similar themes.

[9] A way to gauge the explanatory power of country-specific features on the coverage of oral traditions is to look at the R^2 value from regressing features of Berezkin's references per group on country fixed effects. Country-specific constants explain around 30% of the variation in the number of titles per group and year of the first publication across groups.

[10] Here is a publication on the Inuit oral tradition: "Stories of the Black River People. From stories told by Robert (Nasruk) Cleveland of Shungnak (Alaska). Tape-recorded by Don Charles Foote. Transcribed by Minnies (Aliitchak) Gray. Translated and proofread by Ruth (Taqaviñ) Ramoth-Sampson and Angeline (Ipiilik). Newlin Anchorage: University of Alaska (1980)". This example highlights that even in the production of a *single* publication there are several individuals involved.

the stories. Symbols, metaphors, analogies, and indirect representations of abstract notions are the motifs' building blocks. As a result, abstract keywords, like honesty, patience, and courage, rarely appear in the catalogue itself. This approach has the clear benefit of minimizing the subjective interpretation of folklore content.

3 Empirical Analysis - Validation

In this section, we assess whether the motifs in a society's oral tradition describe (i) natural features in the group's landscape and (ii) societal aspects known from its ethnographic record.

3.1 ConceptNet

We classify motifs into different concepts using ConceptNet, a knowledge representation project, designed to reflect the general human knowledge and its expression in natural language. It originated from the MIT Media Lab and we use the Web API version: api.conceptnet.io.[11] To construct a ConceptNet-based list of related terms we retrieve the top-50 list for each seed word.

To find the intensity of a given keyword within an oral tradition, we first break down each motif into individual words. Berezkin lists the title and the description of a motif both in Russian and English. To maximize the information extracted, we translate the Russian version of the motifs using Google's API translation services, and use *both* English translations. Then, we look up all the related words for every concept and tag the motif accordingly. Hence, we assign each motif to multiple concepts. To arrive at our society-specific estimate of a particular keyword, we add up all motifs tagged by at least one word related to this concept and divide that by the total number of motifs within a group's oral tradition.

3.2 Folklore and the Natural Environment

Folklorists generally believe that elements appearing in a group's oral tradition reflect aspects of its migration history and its interactions with other traditions. This implies that images can be preserved even if the landscapes, climates, and social configurations change. Moreover, stories travel due to the interactions among members of different traditions. Eventually, the borrowed stories are integrated into the stock of indigenous folklore (Boas [1916]). These migration and cross-pollination processes suggest that the link between a society's landscape at the time of initial observation (in the late 19th and early 20th centuries) and the geographic attributes mentioned in its oral tradition is not straightforward. Hence, tracing a group's contemporary ecology in its

[11] The embeddings are a version of ConceptNet Numberbatch. In an ACL 2017 SemEval study, Speer and Lowry-Duda (2017) show that these precomputed vectors are the state of the in terms of vector representations.

narratives will increase our confidence in using the motifs to learn about other unknown aspects of these societies.

We examine the following traits we can reliably measure in the group's folklore and physical environment; (i) proximity to earthquake zones (ii) intensity of lightning strikes (iii) annual temperature (iv) caloric potential from agriculture and (v) proximity to rivers. In Appendix Table 2-Panel A, we report the corresponding summary statistics across the 958 societies in Berezkin's catalogue.[12] Below we ask whether these natural phenomena are salient enough to manifest themselves in a society's folklore.

We estimate the following OLS specifications:[13]

$$\ln(0.01 + \frac{\text{\# Concept-Specific Motifs}_{i,c}}{\text{\# Motifs}_{i,c}}) = \mathbf{a}_c + \beta Geo_i + \boldsymbol{\gamma} Baseline\ Controls_i + \varepsilon_i. \quad (1)$$

The dependent variable is the share of motifs in a group's oral tradition, i, mentioning a given keyword (along with related words); a_c is a vector of continental or country-specific constants; and Geo_i is a vector of geographic traits. We use the group's centroid (recorded by Berezkin) to compute the distance terms and a radius of 200 kilometers to construct the values of the respective geographic attributes. The baseline controls are (i) the log number of publications Berezkin has consulted for a group and (ii) the log year of the earliest publication cited for group i. We cluster the standard errors at the language family level as recorded by Berezkin. There are a total of 75 language families. These include the Indo-European, Algic, Sino-Tibetan, Altaic, Uto-Aztecan, Uralic, Nilo-Saharan, and Niger-Congo families.

Images and Episodes in Folklore Reflecting the Physical Environment

In Table I, we present the results. In the first two columns, the dependent variable is the share of motifs tagged by ConceptNet with words related to earthquakes. There are six motifs that typically explain why earthquakes occur. Motif $i119$, for example, titled: "The dead shake the earth" has the following description "The earthquakes are produced by the dead who are in the underworld or during the earthquakes the inhabitants of the lower world try to come out." Another motif, $i28a$, is titled "Animals produce earthquakes," and the description reads: "Big game animals disappear under the earth and produce earthquakes." We construct the distance from the centroid of each group to the nearest high-intensity earthquake region and follow Bentzen

[12] To avoid our results being driven by societies where little is known of their oral tradition, we restrict our attention to groups with at least five motifs. This number coincides with the bottom 1% of the number of motifs across groups and reduces the sample to 943 societies. The number of observations in the individual regressions may be smaller because of missing geographic information for groups on small remote islands. The patterns are stronger when we weight by the number of motifs per group, effectively allowing the coefficient estimates to be more sensitive to groups with better-sampled oral traditions (results not shown).

[13] We take logs after adding a small number to account for the skewed nature of concept intensity across oral traditions. Results are similar if we use the shares instead of the logarithmic transformation or focus on the extensive margin; i.e., an indicator reflecting whether a concept is present.

(2019) to define the latter as those in zones 3 and 4. An average group in Berezkin's dataset has 0.16% of earthquake-related episodes. However, those located within 100 kilometers of an earthquake zone have 0.3% of such motifs. This pattern is robust to accounting for continental and country fixed effects.

In columns (3) and (4), we use the share of motifs related to "storms" as a dependent variable. A total of 10 motifs mention the words storm or thunderstorm. Motif, $i2$, present in 56 societies, reads: "Lightning emerges from eyes or mouth of the being who is thought to produce thunderstorms." Societies located in regions experiencing intense thunder strikes are weakly more likely to feature storm-related episodes in their oral tradition.[14] The link becomes statistically significant at conventional levels when we compare groups within the same country.

In columns (5) and (6), we ask whether societies in colder climates feature folklore images typical of their environment. To answer this question, we input the word "frozen" in ConceptNet. 39 motifs are tagged by the following terms: cold, ice, frosty, and freeze. Motif $c27$ titled "A horn in the ice" is present in 14 oral traditions, and its description reads: "A horned monster breaks ice on river or lake. Usually, people walk across a frozen body of water, get to see a horn protruding from the ice, and try to cut it off. The monster breaks the ice, many people drown". To explore how climate shapes the frequency of frozen-related motifs, we construct for each society the mean annual temperature between 1900 and 2000. In regions with a mean annual temperature below 10 Celsius, groups' oral traditions have, on average, 1.4% of frozen-tagged motifs, twice as large compared to groups located in warmer places.

The environment exerts a significant influence on a group's mode of production. For example, groups residing on more fertile lands are, on average, more likely to depend on agriculture. In columns (7) & (8), we ask whether this relationship is also evident in the oral tradition. We use the caloric potential of agriculture in a given region for crops available before AD 1500 from Galor and Ozak (2015).[15] To get a proxy for the importance of agricultural activities in a society's folklore, we combine three keywords that reflect the most commonly cultivated crops worldwide, namely "rice," "wheat," and "corn." Doing so, we tag 35 motifs. The tagged words are: cereal, grain, cob, corn, maize, crop, wheat, flour, and rice. Motif $h34g$, for instance, can be found in 35 oral traditions and is titled: "One-grain porridge." Its description reads: "One cereal grain (cob, etc.) is enough to prepare a meal." Among the 510 groups with 0 crop-related motifs, the median regional caloric potential for agriculture is $5,323$. This number jumps to $7,235$ for groups with

[14] We use the mean lightning flash rates observed by the spaceborne Optical Transient Detector (OTD) and Lightning Imaging Sensor (LIS) instruments from 1995 to 2010, see Cecil, Buechler, and Blakeslee (2014).

[15] To arrive at the caloric potential the authors calculate for each cell of size $5\prime \times 5\prime$ the calories per hectare per year that one would attain if they plant the crop that maximizes the caloric yield given the climatic characteristics of the particular area. Therefore, this is not a suitability measure for any specific crop, but it is a measure of a cell's potential crop yield if planted with the crop that delivers the most calories.

at least one farming-related motif. The pattern is the same when we exploit within-continent variation and within-country variation. Because rice, wheat, and corn share many related words, the results are similar if we use them separately. See Appendix Table 3, Panel A. In the last two columns of Table I, we estimate the relationship between the distance to a river and the incidence of river-related concepts. These include words like river, brook, watercourse, stream, creek, and dam. The median group is 56 kilometers from a river.[16] Those further away are less likely to have motifs describing riverine environments.

Appendix Figures $4a-4e$ depict the binned scatterplots of the within-country associations. In Appendix Section 3 we discuss additional robustness checks.

3.3 Folklore and the Ethnographic Record

We matched an oral tradition in Berezkin's database to $1,245$ out of the $1,265$ groups in the EA. From the 958 groups with oral traditions, 619 mapped into a society present in the EA. Hence, for 339 groups, we have information on their folklore but no ethnographic records. After establishing the empirical relationship between folklore images and ethnographic traits, one can use the former to reconstruct the ethnographic background missing for these ethnicities. We run OLS specifications of the following type:

$$\ln(0.01 + \frac{\text{\# Concept-Specific Motifs}_{i,c}}{\text{\# Motifs}_{i,c}}) = \mathbf{a}_c + \beta EA\ Trait_i + \boldsymbol{\gamma} Baseline\ Controls_i + \varepsilon_i \quad (2)$$

where the $EA\ Trait_i$ is the attribute of interest from the EA for group i located in continent/country c. See Appendix Table 2, Panel B, for the summary statistics in this sample.

3.3.1 Mode of Subsistence

For each type of subsistence, we construct the relevant concept as follows. For hunting activities, we use the words related to "hunt." The related terms, found in 144 motifs, are: hunt, chase, deer, scavenger, hunter, pursuit, search, and quest. The following hunt-related motif, $h19$, titled: "Raven tries to starve people," is found in 27 societies, and its description reads: "Raven hides or scares game animals preventing hunters to kill them. People outwit him." To capture fishing activities, we use "fish," as our seed word. This results in 102 motifs. See Appendix Figures $5a$ and $5b$ for the distribution of fish and hunt-related motif frequencies across oral traditions. We use "pastoralism" as our seed in ConceptNet for animal husbandry. This results in 25 motifs including words like: cattle, agriculture, graze, herder, farm, herdsman, livestock, and pasture. Here is the description of a pastoral motif, $k136$: "A lad becomes a master and a leader of a

[16] The data we use come from Natural Earth, a public domain map dataset that includes a global set of single line drainages of sizeable rivers. The original source is the CIA World DataBank II, a collection of world map data created by the U.S. government in the 1980s.

great amount of cattle (cows or buffaloes) and meets a princess." Finally, for farming, we use the words related to the crops explained above.

In the odd-numbered columns of Table II, we include continental fixed effects, and compare groups in the same country in the even-numbered columns. In columns (1) to (4) in Panel A, the dependent variable is the intensity of crop-related motifs. In the first two columns, the main predictor is the share of subsistence from agriculture (with the omitted category being subsistence needs from all other sources). In the next two columns, instead, we add each of the other three subsistence shares with the omitted category now being that of farming. The measures range from 0 to 9, roughly mapping into the deciles of subsistence needs covered by the corresponding activity. Farming groups "talk" more about crops in their folklore than other societies. Columns (3) and (4) show that this difference is more pronounced when one compares fishing and foraging to farming.

In columns $(5) - (8)$ in Panel A, the outcome variable is the share of motifs related to pastoralism, whereas in columns $(1 - 4)$ and $(5 - 8)$ in Panel B, the dependent variable is the share of motifs related to fish and hunt, respectively. The beta coefficient on the share of subsistence from herding is 0.27 (column 6), suggesting a strong influence of the latter on folklore images related to pastoralism. Predominantly fishing-based societies have a median of 5% fishing motifs in their oral tradition, whereas the corresponding statistic for non-fishing groups is 2.3%. Finally, hunt-specific motifs are robust features of hunting and gathering societies. The former systematically decline as the shares of subsistence from fishing and farming increase.

A complementary way to gauge the predictive power of the mode of subsistence on related images in the group's folklore is to look at the R^2 in specifications without continental or country constants (not reported). The magnitudes range from 0.23 for images on pastoralism to 0.09 for the share of motifs on hunting. In Appendix Figures $6a - 6d$, we present binned scatterplots visualizing the relationships established in columns (1) and (5) of Panels A and B of Table II.

3.3.2 Family Structure

We explore whether groups organized along extended family lines are more likely to feature family members in their oral traditions. Unlike nuclear families, socialization is more likely to occur within the family's network in extended-family societies. If this is the case, one would expect to see family members often being the protagonists in the stories and tales passed from one generation to the next. We use the words related to "mother," according to ConceptNet (using other family members as seeds, we obtain similar patterns). These are mother, father, grandmother, daughter, stepmother, and stepfather. We use variable $v8$ from the EA that codes the prevailing form of domestic organization and we assign the value 1 to those groups organized as independent nuclear

families (with monogamy or occasional polygyny) and 0 otherwise. For 1,218 groups in the EA with such information, 356 are classified as independent nuclear families, while the remaining 862 are polygynous and extended families. In the latter, mother-related motifs show up in 7.9% of the total images and episodes recorded in the oral tradition. The corresponding statistic is 6.7% for societies with a nuclear structure. This difference is statistically significant at the 1% level. In Appendix Table 4, we show that this pattern is robust to exploiting within-continent or within-country variation, respectively. Appendix Figure 7*a* displays the binned scatterplot of the relationship estimated in column (1).

3.3.3 Political Centralization

In columns (3) and (4) in Appendix Table 4, we ask whether the strength of political complexity in the EA predicts folklore-based measures of political hierarchy. Using the keyword "king" in ConceptNet, the following words are tagged: king, kingdom, kingship, queen, prince, and princess. Appendix Figure 5*c* illustrates the global distribution of the intensity of king-related motifs across oral traditions. The median society in the EA has no motifs related to a king. This number jumps to 2.7% for centralized societies. Hierarchical groups during the pre-colonial era are systematically more likely to have king-related motifs. Appendix Figure 7*b* shows the binned scatterplot corresponding to column (3).

4 Filling in the Gaps in the Ethnographic Record

The fact that folklore-based measures of the economy and the polity are in line with the Ethnographic Atlas suggests that we can complement the ethnographic record by (i) filling in the gaps for groups with missing information and (ii) reconstructing traits that have not been recorded. The two applications below show how to improve the coverage of two extensively-used EA variables.

4.1 Political Centralization

Out of the 1,264 societies in the EA, there is missing information on the layers of jurisdictional hierarchy beyond the local community, variable $v33$, for 135 groups. We use their folklore to fill in these gaps. After regressing political complexity on the intensity of king-related motifs, we use the estimated coefficient of .66 to predict the jurisdictional hierarchy levels across all groups in the EA. In Appendix Figure 7*c*, we put together two density plots of predicted centralization –one for those groups with EA-recorded levels of hierarchy and the other for those with missing information. On average, societies without ethnographic records on political centralization are predicted to be slightly more complex although this difference is not statistically significant.

Naturally, one may wonder about the quality of the prediction. One way to go about it is to first look at commonly-used variables in the EA that correlate with the variable of interest. Then, for the sample with missing values explore whether the imputed measure displays a correlation structure similar to that present in the non-missing sample.

Table *III* Panel *A* has the following structure. We use eight variables from the EA and report two bivariate correlations between each of these variables and (i) the EA-recorded measure of centralization and (ii) the folklore-imputed measure in the sample of missing observations. Columns $(1) - (4)$ reflect subsistence shares from hunting and gathering, fishing, herding, and farming, respectively. Column (5) is an indicator equal to one if the group engages in intensive agriculture. The variables in columns $(6) - (8)$ are measures of social complexity, including the degree of class stratification, the role of high gods, and the layers of jurisdictional hierarchy above the local level. Overall, the folklore-imputed measure of centralization behaves in a qualitatively similar manner to that recorded by ethnographers. The only variable for which these two measures diverge is the dependence on farming. In the non-imputed sample, there is a significant correlation between reliance on agriculture and political complexity. In the imputed sample, there is no correlation between the two. Nevertheless, when we look at whether a group follows intensive agricultural practices, the correlation is the same in both samples.[17]

4.2 High Gods

The presence of moralizing gods has attracted the attention of scholars in economics, psychology, and beyond. There is a lively debate, for example, on whether complex societies gave birth to big gods or the other way around (Whitehouse et al. [2019]). However, the corresponding variable in the EA, $v34$, is missing for more than 40% of the groups, limiting the empirical investigations.

According to anthropologists, the crucial attribute of big god societies is the presence of supernatural entities that award and punish humans for their behavior (Schaller et al. [2010]). Establishing that such images are present in their folklore would first add large-scale empirical support to this foundational hypothesis. Second, we could use the intensity of such folklore episodes to obtain a measure of moralizing gods where this information is missing.

A *single* concept cannot capture complex plots of supernatural forces punishing and rewarding. In Table *IV*, we show how to uncover the intensity of these complex images by combining individual concepts. We focus on three keywords that reflect the themes we are after. In columns (1), (2), and (3), the dependent variable is the share of motifs tagged by "punishment," "award," and "supernatural," concepts, respectively.[18] High-gods societies are significantly more likely

[17] This happens because in the subset of groups with missing values on political complexity, those that rely more on farming are *less* likely to follow intensive agricultural practices. These two variables correlate positively and strongly in the sample of groups with known political complexity.

[18] The words tagged by "punishment" are: punishment, punish, and retribution. For "award" the terms are:

to display punishment in their oral tradition, and episodes featuring rewards and supernatural entities are also more common. We conduct a principal component analysis to derive an index that jointly captures these three concepts. The first component explains approximately 62% of the total variance. The second explains around 25%, with the third explaining the balance. Importantly, all three concepts load positively on the first principal component. Its eigenvalue is greater than one, while the eigenvalues of the other two are less than one (with one being the rule of thumb). We thus focus on the first principal component to construct an index that reflects the joint intensity of punishment, award, and supernatural entities in a group's oral tradition. We call this "punishing and awarding supernaturals." Columns (4) and (5) reveal that within continents and countries, groups with big gods, according to ethnographers, are systematically more likely to have an oral tradition that prominently features images of supernatural entities, punishments, and awards. Appendix Figure 8*a* shows the binned scatterplot.

In the last part, we fill in the values for groups without a record on $v34$. Specifically, we regress high gods on the folklore intensity of "punishing and awarding supernaturals" and obtain a coefficient of 0.327. In Appendix Figure 8*b*, we put together the density plots of the folklore-imputed moralizing gods for the two different samples. One reflects the distribution in the groups where such information is present, and the other corresponds to the groups with missing data. The two distributions largely overlap. In Panel B of Table III, we assess the quality of the prediction by comparing the correlations between the predicted measure and other variables from the EA in the imputed sample to the correlations between the EA-recorded high gods and the same variables in the non-imputed sample. Overall, the two sets of correlations are in broad agreement.

These two applications provide a guidance to using a group's folklore to complete the gaps across the many entries of the EA and the Standard Cross Cultural Sample (SCCS).

4.3 Using Folklore to Encode Non-extant, Group-level Characteristics

The Case of Trade Markets are at the heart of any economic analysis. However, the importance of the market economy is entirely missing from the EA. Here we use a society's folklore to measure this phenomenon. Using "trade" as the seed the following words are tagged: exchange, trade, swap, barter, trader, and merchant. Motif $m171a$ is present in 79 oral traditions in different parts of the Old World. Its title is: "The profitable exchange: gets a bride" and describes a person or an animal that gets to exchange less valuable goods for more valuable ones and ultimately secures a bride. The average group in Berezkin's database has 0.9% of trade motifs (Figure $IIIa$ portrays its global distribution). We do two things to explore how

award, reward, and winner. Finally, for "supernatural" the related words are: magical, superhuman, supernatural, supernaturally, vampire, and witchcraft. Using the keyword "reward" instead of "award" produces similar results.

informative this measure is. First, we look at the intensity of trade-related motifs as a function of a groups' proximity to pre-industrial trade routes. We used data from Michalopoulos, Naghavi, and Prarolo (2018) that compiled a comprehensive set of routes for the Old World from pre-600AD to 1700 AD. Using these two sets of routes, we constructed the average distance from each group in Berezkin's database. The summary statistics are telling of a robust, broad pattern. Among the 477 societies in Berezkin's catalogue located in the Old World (Africa, Asia, Europe), those within 200 kilometers of ancient trade routes have an average of 1.4% of trade-related motifs, twice as large compared to those located farther away. Columns (1) and (2) in Appendix Table 5 show that this pattern is not driven by cross-continental or cross-country differences.

Second, in columns (3) & (4) of Appendix Table 5, we test whether groups spanning different ecosystems are more likely to have an oral tradition rich in trade motifs. This is motivated by Bates' view that historical trade was more intense along ecological divides where products from different zones would be exchanged. We follow Fenske (2014) and construct the degree of ecological polarization within a 200km radius of each group using the dominant ecosystem classes as recorded in the FAO-GAEZ project (we obtain the same pattern if we use ecological diversity). Societies located along more ecologically polarized zones have a larger share of episodes in their folklore describing an exchange taking place. Figures *IIIb* and *IIIc* show the binned residual scatterplot from columns (1) and (3). Columns (5) & (6) show that both the distance to pre-industrial trade routes and ecological polarization are significant predictors of trade images in groups' oral traditions across and within countries.

5 Folklore and Contemporary Beliefs and Attitudes

Constructing a Country-Level Catalogue of Motifs To arrive at a catalogue of folklore across countries, we first match the oral traditions listed in Berezkin to the ethnic groups in the 1964 Atlas Narodov Mira (ANM). Second, we use a group's fraction in the country's population in 2000 (combining the ANM polygons with information on cell-level population density from the Gridded Population of the World within country borders as of 2000) as a weight in our aggregation of motifs at the country level. This procedure is accurate for countries where the indigenous population as of 1500 constitutes a significant fraction of the 1960 population (when the ANM dataset was built). For the countries, however, that the ANM lists groups that appeared after 1500 AD, like the Brazilians, Australians, US Americans, etc., we use the population percentages from the countries of origin as reflected in the Putterman and Weil (2010) global migration matrix.

5.1 From Machine Learning to Learning from Humans

In Appendix Section 4, we show that shrinkage methods can offer a glimpse into the keywords predictive of attitudes across countries today as measured in the GPS dataset. Among the *thousands* of features, the handful selected are sometimes intuitively linked to the outcome of interest. However, the fact that across specifications, the set of selected features varies is reminiscent of Mullainathan and Spiess (2017) observation about the inherent instability of the models chosen by machine learning methods, particularly in the presence of correlated predictors.

Below we show how one can improve upon shrinkage methods based on the assessments of multiple individuals. We focus on trust, risk-taking, and gender norms. Of course, this list is by no means exhaustive. We hope that the methodology developed below can offer some guidance in using folklore to reconstruct a proxy for any cultural value.

5.1.1 Trust

We start our analysis looking at trust. Besides its obvious economic significance, we focus on this trait because as discussed in Appendix Section 4 the keywords selected by Lasso change depending on whether trust is measured by the GPS or the World Value Surveys (WVS) dataset. Below, we show that how tricksters are treated in a group's folklore consistently predicts trust levels today irrespective of how the latter is measured.

Tricksters in Folklore Oral traditions worldwide abound in folktales where tricksters engage in antisocial behavior, including cheating, stealing, and generally deceiving others. A distinguishing feature of these stories is that sometimes the deceiving character seems to get away with her behavior, whereas in other tales, he/she gets punished due to his actions or is simply unsuccessful. To categorize these complex narratives, we hired workers from Amazon Mechanical Turk (MTurk).

First, we identified motifs that mention words related to "cheat," "deceive," and "trick," according to ConceptNet. There are a total of 281 motifs that contain at least one word related to these concepts.[19] On average, 15 MTurks read and classified each motif into four mutually exclusive categories: whether the character portrayed in the motif engages in antisocial (absurd, obscene, or deceiving) behavior, and whether he/she is successful. In Appendix Figure 9, we provide an example of a motif that MTurks classified along with the instructions provided. For more details on the MTurk platform and the classifications, see Appendix Section 5.

Here are two examples of motifs describing antisocial behavior. Motif $k122$, found in 26 societies, mainly in Europe and the Middle East, reads: "Hero gets to the powerful woman

[19] The related words tagged in the motifs are: betray, cheat, deceive, deception, trick, cheater, deceiver, deceit, deceitful, dupe, falsify, hack, hoodwink, unfaithful, cajole, confuse, fool, hoax, stratagem, ruse, and wile.

who lives in another world, then returns to the world of humans. Imposter claims the hero's deeds for himself. The powerful woman comes and finds the real hero, punishes (rejects) the deceiver". All MTurks classified this as an instance of antisocial behavior where the character is unsuccessful/punished. On the contrary, the motif depicted in Appendix Figure 9, mainly found in the Americas, is classified as "antisocial and successful" by 78% of the MTurks with 22% choosing the last option, i.e., "antisocial but outcome is not described." For each motif, we use the most common classification.[20]

Out of the 281 motifs, 57 are uniquely classified as "not antisocial" and hence are not considered. In 46% of the remaining 224, the trickster is depicted as successful, 36% as unsuccessful/punished, and for another 46 motifs, MTurks thought that the outcome of the antisocial behavior could not be determined. We aggregate the intensity of motifs at the group and at the country level. The key variable of interest is how often tricksters fail compared to instances where the character engaging in antisocial behavior is either successful or not explicitly punished.

Below, we explore whether a group's level of generalized trust is influenced by how cheating is portrayed in its oral tradition. Perhaps it is instructive to pause here and ask whether a society with a relative abundance of motifs where tricksters get punished reflects a high or low trust group. To the extent that narratives in a group's folklore encapsulate the equilibrium a community has converged to, one would expect oral traditions that typically highlight the punishment of tricksters to reflect an environment of low tolerance towards antisocial behavior.

Across Countries To arrive at the folklore measure of trust, we subtracted the country-level frequency of motifs where antisocial behavior is punished from the instances where the antisocial behavior was not (explicitly) punished. We normalize these frequencies by the total number of motifs in the country. The cross-country mean equals -2.5%, i.e., on average, tricksters are more likely to get away with their actions. On the one end, there are countries like Bahrain (with the maximum score of 4.2%), Sri Lanka, Denmark, and the Netherlands where antisocial behavior is disproportionately punished compared to countries like Togo (with the minimum score of -11.3%), Angola, Haiti, Burkina Faso, Bosnia and Herzegovina, and Myanmar where tricksters go unpunished much more often than not. Figure IVa illustrates the intensity of the trickster's relative punishment in folktales across countries, and Figure IVb shows the histogram.

Table V consists of three panels that show how the rate at which antisocial behavior is punished maps into current levels of trust and economic activity. In Panel A, the unit of analysis is a country. In the first two columns, the dependent variable is the average response to the commonly-used generalized trust question in the WVS/EVS. Column (1) reveals a striking

[20] For 22 out of 281 motifs, the mode was not unique, i.e., MTurks were evenly split between two classifications. In these cases, we treated the motifs as indicating *both* categories.

positive link between trust today and narratives of low tolerance of antisocial behavior across countries. The variation in the latter explains roughly 20% of the overall variation in self-reported trust levels. This correlation is also quantitatively meaningful. A one-standard-deviation increase in the rate at which tricksters are unsuccessful is associated with an increase in trust of roughly 0.35 standard deviations. See Figure IVc for the binned residual scatterplot. Column (2) shows that this association also holds within continents. The pattern is quantitatively similar in columns (3) and (4), where we use the GPS measure of trust. In the last two columns, we replace trust with log GDP per capita in 2010. Across 161 states, those with an oral tradition that strongly disapproves of antisocial behavior are more prosperous today, see Figure IVd for the residual scatterplot.

Across Second-Generation Immigrants Given the cross-country association between contemporary trust attitudes and tricksters' punishment, the natural question becomes whether the parental folklore of 2nd-generation immigrants may help explain differences in trust levels today. Finding such a pattern would be consistent with the intergenerational transmission of values via traditional narratives.

To answer this question, we turn to the European Social Survey rounds 1 to 8 conducted between 2002 and 2016. We focus on second-generation immigrants whose parents have the same country of origin (Fernandez [2011]).

There are a total of $6,554$ second-generation adult immigrants in our sample whose parents were born in 117 countries worldwide. In the first two columns of Table V Panel B, the dependent variable reflects the answer to the question "Most people can be trusted or you can't be too careful?" In columns (3) & (4), the dependent variable is the decile in the income distribution of the respondent's income in the specific country-round.

The empirical specification takes the following form:

$$\text{Outcome}_{i,c,o,r} = a_{c,r} + \beta \text{Folklore Measure}_{o} + \gamma \text{Baseline Controls}_{o} + \boldsymbol{\zeta} \mathbf{X}_{i,c,o,r} + \varepsilon_{i,c,o,r}, \tag{3}$$

where $\text{Outcome}_{i,c,o,r}$ is that of individual i, residing/born in country c, interviewed in round r, whose parents were both born in country o. $\mathbf{X}_{i,c,o,r}$ is a vector of personal characteristics, including age decile fixed effects, gender, self-reported health status fixed effects, and religion fixed effects.[21] We cluster the standard errors at the oral tradition level (country of origin of the

[21] The religion variable we use is labeled *rlddnm* and indicates "Religion or Denomination Belonging to at Present." It includes the following eight denominations: Roman Catholic, Protestant, Eastern Orthodox, Other Christian denomination, Jewish, Islamic, Eastern religions, Other non-Christian religions. We complement this classification with an indicator reflecting those that answer "No" in the question of whether they belong to a particular religion or denomination. This is the case for roughly 30% of the respondents.

parents) and $a_{c,r}$ reflect round-country-of-interview specific constants. The baseline controls are the same as those discussed above. All columns include the baseline controls and country-round fixed effects. In the even-numbered columns, we add on top the set of individual-level controls.

The pattern uncovered among 2nd generation immigrants within European countries echoes the cross-country picture. Those who trace their ancestry to oral traditions where tricksters are often unsuccessful are systematically more trusting and enjoy within the same country higher incomes than those with a heritage of an oral tradition where antisocial behavior is not prominently punished. The possibility that today's trust levels can be traced in how tricksters are portrayed in the narratives that children grow up listening to suggests the long-lasting influences of early life environments.

Across Ethnic Groups Trust attitudes across ethnic groups at a global scale are not available. This is why the literature on culture has often focused on the immigrants' country of origin exploiting within-country variation in cultural traits. Nevertheless, thanks to Berezkin's catalogue we may reconstruct values and beliefs at the ethnic group level. Figure Va portrays the intensity of trickster punishment across the 958 ethnic groups. There is substantial variation both across and within countries in how tricksters fare.

Take, for example, the three groups that Berezkin lists for Afghanistan. These are the Pashto, one of the main language groups, the Kafirs, Nuristani people, and the Wahki/Ishkashimi/Munji groups, who all speak Pamir languages. There are significant differences in the fate of tricksters across these three groups. Among Kafirs, tricksters are usually punished. Tales where the deceiving character fails, exceed those where he is successful by two percentage points. In contrast, in the Wahki/Ishkashimi/Munji folklore, the frequency at which tricksters succeed exceeds that of failing by almost five percentage points. The Pashto are somewhere in between, with the punishing motifs falling behind the motifs where tricksters are successful by roughly one percentage point.

In Panel C of Table V, we show that differences in tricksters' outcomes can explain part of the variation in economic performance across groups both historically and today. In columns (1) & (2), we use the average luminosity in 2008 per square kilometer in a radius of 200 kilometers of each group's centroid. Both within and across countries, we find that groups whose oral tradition sanctions tricksters' behavior often are more developed today. Figure Vb displays the binned scatterplot corresponding to column (1). Going back to the example of the groups in Afghanistan, the Wahki/Ishkashimi/Munji are at the bottom of the within-Afghanistan luminosity distribution, with Kafirs being on the top.

In the last four columns of Table V-Panel C, we show that historical development across ethnic groups also correlates strongly with peoples' values as deduced from folklore. To capture

historical development in a Malthusian, pre-industrial setting, we use local communities' mean size from the EA, v31. Groups are classified into eight bins in ascending order ranging from societies with less than 50 people to groups with more than $50,000$ members. Historically more populous groups depict tricksters failing in their attempts to deceive. See Figure Vc for the relationship estimated in column (3). This pattern is not driven by comparing farming societies to others. In columns (5) & (6), we focus on groups that derive most of their subsistence from agriculture. Among farmers, those with stronger norms against cheating are more developed historically. The association between economic performance and the folklore-based measure of trust is quantitatively stable over time. The beta coefficients estimated in columns (2) and (4) of Table V in Panel C are both around 0.08.

This example highlights the usefulness of folklore in shedding light on historical beliefs and attitudes across ethnic group worldwide and their economic consequences in the past and the present.

5.1.2 Risk-Taking

In this section we focus on risk attitudes. The latter play a pivotal role in shaping entrepreneurial activity, and the propensity to accept risk varies considerably across societies. Tales of challenges and competitions are recurrent themes across oral traditions and the machine learning estimates presented in Appendix Table 6 reveal that the keywords related to competitions predict higher risk-taking across countries today. Naturally, one may wonder whether tales where the character(s) wins a contest, are more likely to encourage risk-taking than narratives where such dangerous interactions are detrimental to the hero's well-being.

To get at this distinction, we had each of the 87 motifs mentioning a word related to challenges and competitions, classified by 9 MTurks on average.[22] The question we asked reads: "The character(s) in this motif finds him in a challenging/difficult situation. Is the character(s) successful in getting out of it?" We use the modal answer allowing for multiple classifications in case of more than one mode. On the one hand, 48% of the challenge-tagged motifs portray a character(s) who is successful. Take, for example, motif $k38f$ titled "Dragon winner" found in 115 oral traditions. Its description reads: "A reptile monster demands humans (usually virgins) as a sacrifice or abducts a girl or closes sources of water. Hero kills him." Several motifs describe challenges but without specifying the outcome. For example, motif $k27e$ titled "Eating or drinking contest" describes a person or animal that must consume enormous quantities of poisonous food.

On the other hand, in 19% of these competition-related themes, the character is harmed.

[22] The words related to challenges and competitions according to ConceptNet found in the motifs are: contest, competition, winner, challenge, challenger, compete, tournament, contender, competitor, confront, dare, difficult, and difficulty.

Motif $c33$, for instance, present in 12 oral traditions is titled "Prometheus (the chained strong man)" and describes a strong man who ventured to confront God and is for eternity chained to a mountain; see Appendix Figure 10a. Finally, 7 out of the 87 motifs are considered by the MTurks as not describing a challenging situation. Below, we will show that these different classifications have a qualitatively different relationship with attitudes towards risk-taking.

Across Countries In Appendix Table 7, Panels A and B, we present the summary statistics and the correlation structure among the outcomes of contests and competitions across the oral traditions of 199 countries. The average country has roughly 6% of such themes in its folklore. In one out of every six of these episodes (1% of all motifs), the character fails. The rest of the time, he is unharmed by the challenge.

In the first two columns of Table VI, we exploit across and within-continent variation. Countries with a larger share of challenge and competition-related motifs in their folklore are more risk-tolerant according to the GPS today. This pattern is expected, given that the Lasso exercise already selected these keywords. In column (3), we go a step further and distinguish between the different outcomes. The intensity of images where the character is successful maps onto an increased appetite for risk. Episodes describing competitions without specifying an outcome are also positive and significant predictors of risk-taking (although in terms of magnitudes, the estimated beta coefficient of 0.17 is roughly 60 percent smaller than the corresponding coefficient for the intensity of successful challenges). Notably, motifs that portray competitions as detrimental to the character's well-being, i.e., lead to tragedies, are inversely related to risk-taking. Finally, contest-related motifs that the MTurks classified as describing no such interactions are unrelated to risk attitudes. To put the numbers in context, consider the following example. According to the GPS dataset within Europe, the most risk-tolerant country is the Netherlands, whereas Russia is third from the bottom and thus one of the most risk-averse. The share of tales depicting a hero overcoming challenges is twice more common in the Dutch folklore than in the Russian one.

Motivated by these patterns and to get at a single measure of the outcome of challenges and competitions, we subtract from the share of motifs where challenges lead to failures the episodes where characters are not (explicitly) harmed. In column 4, we show that the relative frequency of harmful contests significantly lowers attitudes towards risk today. Appendix Figure 10b depicts the residual scatterplot.

The remaining columns in Table VI illustrate that these tales also predict entrepreneurial activity today. We capture the latter with two proxies. In columns (5) & (6), we use the number of patents filed by residents per $100,000$ people, and in (7) & (8), we use the number of new business

registrations (limited liability corporations) per $1,000$ people ages between 15 and 64.[23] Both these measures are in logs and reflect the $2006 - 2018$ averages. Across and within continents, countries whose folklore depicts challenges as opportunities instead of tragedies record higher entrepreneurial rates. In Figures VIa and VIb, we depict the variation in risk attitudes as proxied by the relative intensity of harmful contests across countries and ethnic groups, respectively.

Across Second-Generation Immigrants Comparing entrepreneurial activity across countries is not easy since patenting activity and incorporation rates, for example, are very much influenced by a country's institutional framework, among other things. Hence, the literature has converged to using self-employment rates across individuals within the same state as a more meaningful proxy of entrepreneurial attitudes. In the last two columns of Table VI, we focus on second-generation immigrants and ask how parental folklore influences the probability of self-employment.[24]

The empirical specification follows that of equation (3). There are a total of $6,063$ second-generation immigrants in our sample with a probability of self-employment of 12.7%. Their parents were born in 106 countries. The difference between columns (9) and (10) is that we add individual controls in the latter. In both specifications, we compare individuals born and currently living in the same country and hence face similar institutional arrangements. Second-generation immigrants whose parental folklore depicts contests and competitions unfavorably are systematically less likely to be self-employed. The estimated coefficient in column (9) suggests that a one-percentage-point increase in such images decreases the probability of self-employment by approximately two percentage points (see Appendix Figure 10c for the binned scatterplot).

The following example is indicative of the overall pattern. Within Israel, two large groups of second-generation immigrants are those whose parents came from Morocco and Romania. Individuals of Romanian ancestry have a 21% likelihood to be self-employed, whereas those whose parents hail from Morocco have a 12% probability. In Moroccan folklore, the likelihood that fictional characters get out of difficult and challenging situations instead of being hurt is a mere 1.5%. This number is three times as large in the Romanian folklore (4.8%).

5.1.3 Gender Roles in Folklore

In this section, we focus on extracting gender norms from folklore images. The importance of gender attitudes transcends the social sciences and humanities. Economists, for example,

[23] Data on patents come from the World Intellectual Property Organization Patent Report. The statistics on incorporation rates come from the World Bank's Entrepreneurship Survey and database.

[24] The variable *emplrel* in the ESS reflects each respondent's employment relation in their current job, or their last one, if currently out of work. Self-employment takes the value of 0 if the respondent answers being an employee and one if she answers that she is self-employed or working for her own family business.

motivated by famous works in anthropology have empirically linked specific agricultural traits to the role of women, ranging from the presence of the plow, Alesina, Giuliano, and Nunn (2013), to tea or cotton cultivation, Qian (2008) and Xue (2016), and the extent of the pastoral economy, Becker (2019). This line of thinking links the (geographically-determined) relative returns of female labor force participation to women's status in the past and via a process of intergenerational transmission to their role in the present. Tracing gender stereotypes in the oral tradition of a given society will allow us first to understand women's relative status in our ancestors' narratives. Second, assess whether the way women are depicted in a group's folklore is predictive of their position today.

Although pinpointing which motifs mention a male and/or female character is straightforward, what exactly should one classify is not. For example, asking whether women (men) engage in good or bad actions would deliver answers that would hardly be comparable as these assessments are value-system dependent. To circumvent this delicate issue, we ask whether the character is portrayed stereotypically drawing insights from research on media. Towbin et al. (2004), for example, look at how women and men are depicted in the Disney movies and Lauzen, Dozier, and Horan (2008) study the social roles enacted by female and male characters in prime-time television programs in the US.

The literature predominantly associates men with being violent/dominant/strong, (physically) active, angry, arrogant, and sexual. The women's categories include being emotional, beautiful, and domestic/dependent/submissive. We explore many of these stereotypes in our analysis. Appendix Figure 11 shows an example of a motif that MTurks classified along with the instructions.

MTurks classified a total of $1,073$ male and/or female motifs.[25] In 586, both a male and a female character are present, whereas in 296 (191) motifs, only a male (female) character is detected. We allow the MTurks to choose more than one classification as the characters may display a variety of qualities in the same motif. To classify a given motif, we choose the most common answer. For 29 (81) motifs out of the 777 (882) tagged as having a female (male) character, the unique mode chosen by MTurks was "no gendered entity present." We do not consider these motifs in the analysis. Again in case of a tie, we allow for multiple classifications.

We start by exploring the gender stereotypes across motifs in Berezkin's catalogue in Appendix Table 8. Being violent, dominant, and arrogant is the most common representation of male characters with a 33% frequency. The corresponding statistic for females suggests that for every image of a violent and dominant male, there is half of a motif where a woman displays such

[25] The words used to identify female characters are: woman, girl, female, wife, lady, mother, daughter, stepmother, grandmother, sister, niece, girlfriend, queen, princess, and for men: man, boy, male, husband, father, son, stepfather, grandfather, brother, nephew, boyfriend, king, and prince.

qualities. On the contrary, the most common theme among female characters is that of being submissive and dependent, with about 30% of female motifs having such representation. This number is two and a half times smaller for males. A similar pattern emerges when we look at the incidence of other stereotypes. Compared to men, women are almost twice as likely to appear engaging in domestic affairs and half as likely to be physically active. However, relatively more images depict women as sexual, and on average, females are more intelligent and less naive than men. MTurks classified the overwhelming majority of male and female motifs (80%) into the stereotypes provided. For the remaining 20%, the category "other" was chosen. This suggests that besides the attributes listed, there are other aspects that the current classification does not capture. Future work should shed light on the topics that comprise the residual category.

The catalogue of oral traditions reveals stark gender disparities.[26] Before showing how gender roles aggregate at the country and group-level it is vital to keep in mind prominent folklorist Mary Ellen Brown's. She highlights how the folkloristic procedures may affect and skew the collection and analysis of expressive culture. These include the selection of informants, the dimensions of traditional culture that become accessible to outsiders, and how the gender, age, class, or race, and the collector's biases influence the recorded narratives. This is why establishing a link between gender bias in folklore and women's current status is vital. It would suggest that, notwithstanding these considerations, variation in gender roles across oral traditions reflects cross-cultural differences in historical gender norms.

Across Countries In Appendix Table 9*a*, we present the summary statistics on gender representation across the folklore of 199 countries. We define as male bias the difference between male and female motifs for a given stereotype and normalize it by the total number of motifs in a country's oral tradition. For example, the cross-country mean of male bias in submissive motifs is -5.8%, i.e., in the oral tradition of the average country, women compared to men are roughly 6% more often depicted as submissive. So, negative numbers indicate a lower prevalence of the specific stereotype among men compared to women. On average, males appear as more dominant, less submissive, more intelligent, more physically active, and less engaged in domestic affairs than females. Overall, there are no gaps in gender representation in terms of sexuality and naiveté.

To assess whether these stereotypes appear in bundles or are distributed evenly across countries, in Appendix Table 9*b*, we look at the correlation structure of the different dimensions of gender bias across countries. Oral traditions portraying males as violent and arrogant feature

[26] This pattern is not unique to Berezkin's work. Karsdorp et al. (2015) note that male and female characters in Thompson's Motif Index also appear in different themes. Men are prominent in (i) mythological, (ii) magic, (iii) marvels, and (iv) wise and the foolish tales, whereas women dominate the categories (i) deception, and (ii) sex.

women in a submissive status and also depict males as more physically active and less engaged in domestic affairs. Male bias in sexuality, intelligence, and naiveté are weakly correlated with male bias in terms of violence. Informed by these patterns, we create two variables. The first captures the degree to which an oral tradition describes males relative to females as more violent, less dependent, less engaged in domestic affairs, and more physically active. We call this *male dominance bias*. The other variable reflects the degree to which males are portrayed as more intelligent and less naive than women; we call this *male mental capacity bias*. Figure $VIIa$ illustrates the spatial distribution of male dominance bias across the globe.

A few examples may help set the stage for the empirical analysis that follows. Philippines is a country with negligible bias in the its folklore against women ranking 195th out the 199 countries in the sample. Compared to other parts of Southeast Asia, women's legal and family property rights have not been questioned. On the contrary, in Afghanistan the extent of male dominance bias in the folklore images is almost twice that of the average country. Another striking comparison is between Lesotho and Swaziland (now officially known as the Kingdom of Eswatini). Both are small, ethnically homogeneous, landlocked countries in Southern Africa. Despite these similarities, the status of women differs dramatically between the two. Swazi women lack fundamental legal rights, including the right to own assets, and are subject to widespread discrimination and violence. According to the Gender Inequality Index created by the UNDP, Swaziland ranked 150th out of 188 countries in the world in 2012. At the level of traditional narratives, Swazi's folklore has the dubious distinction of being close to the top of the global distribution in terms of displaying males as physically active and disengaged from domestic affairs compared to women. On the contrary, Basotho women are historically known as agents of economic empowerment, wielding significant decision power within the household (Eldredge [1991]). As of 2012, women's literacy rate was 95% compared to 83% for men. Lesotho's oral tradition, mainly consisting of narratives from the Sotho group, is, in fact, the only one in the world where men appear less violent than women. Female labor force participation is about 41% in Swaziland versus 60% in Lesotho. Similar contrasts can be found in other continents.

A commonly-employed measure of women's status today is female labor force participation. This proxy is available for a large number of countries, and it is relatively precisely estimated. In Panel A of Table VII, we ask whether variation in male bias across oral traditions maps into variation in how much women are embedded in the labor market as of 2019. In columns (1) to (3), we include the baseline controls, and we add the three different measures of male bias sequentially. In column (1), we introduce the male dominance bias. It enters with a negative sign, and it is precisely estimated. In societies that feature relatively more images of dominant and physically active males and dependent and home-bound females, women are systematically

less integrated into the labor market. A one-standard-deviation increase in male dominance bias decreases female labor force participation by 0.39 standard deviations reflecting a high sensitivity of women's status in the labor force to the oral tradition they are exposed to. Variation in male dominance bias across the globe accounts for roughly 14% of the heterogeneity in female labor force participation. See Figure $VIIb$ for the residual scatterplot. In columns (2) and (3), we introduce the male mental capacity bias and the male sexual bias. Neither of these two aspects seems to be systematically related to female labor force participation. Columns (3) to (6) replicate the first three adding continental fixed effects. Exploiting within-continent variation delivers a similar pattern. One may wonder whether the aggregation of the dimensions of male bias masks some interesting heterogeneity. In Appendix Table 10, we introduce one dimension of male bias at a time. The results echo those of Table VII-Panel A.[27]

Across 2nd Generation Immigrants Having established a robust inverse association between male dominance themes in a country's oral tradition and contemporary women's status across the globe, we turn to 2nd-generation immigrants.

The empirical specification follows equation (3). In all columns of Panel B in Table VII, we include country-round fixed effects and the baseline controls. In the even-numbered specifications, we add individual level controls. We look at three outcomes. The first reflects an economic choice, whereas the other two capture attitudes towards gender equality. Specifically, in the first two columns, we focus on female adult respondents, and the dependent variable is an indicator that takes the value of 1 if the individual answers that her main activity in the last seven days before the survey is "housework, looking after children, and others." The variation in oral traditions comes from 102 countries of parental origin. In columns (3) to (6), the dependent variable encodes how much a respondent agrees with the following two statements. In columns (3) & (4), the statement reads: "Men should have more right to a job than women when jobs are scarce," and in columns (5) & (6), the statement is: "Women should be prepared to cut down on paid work for sake of the family." The possible answers range from 1 to 5 with 1, indicating that the respondent disagrees strongly and a value of 5 that he/she strongly agrees.

The message is consistent across specifications. Within European countries among 2nd-generation immigrants, females whose parents hail from a society with a strong male dominance bias in its oral tradition are more likely to be housewives. Moreover, male and female respondents from male-biased oral traditions are more comfortable sacrificing women's labor force participa-

[27] Our findings echo Towbin et al. (2004). The authors investigate gender representation in fiction stories from the Wattpad online writing community and find male over-representation and traditional gender stereotypes (e.g., dominant men and submissive women) are common throughout nearly every genre in 1.8 billion words corpus. Hence, contemporary fiction bears strong parallels to folklore in terms of the representation of gender. It would be interesting to explore whether the oral tradition of the various authors in this online community is predictive of the gender roles in their stories.

tion. Quantitatively speaking the beta coefficients in columns (1) to (6) are between 0.05 and 0.10. The magnitudes decline slightly but remain significant when we add the individual-level controls.

Ethnic Groups Figure $VIIIa$ depicts the male dominance bias across groups in Berezkin's catalogue. There is significant heterogeneity within countries. India's diversity in languages is also reflected in their oral traditions. For example, among the Khasi people (an indigenous ethnic group in north-eastern India), their folklore shows gender parity in the various stereotypes suggesting a society where gender roles are fluid. On the contrary, the Kashmiri group is among the top 3 groups in India in terms of male bias. Motifs portraying men as dominant exceed images where women appear as such by 33%. In this society, the oral tradition depicts the different roles that men and women occupy in unambiguous terms. Similar within-country differences across oral traditions appear in other continents. In Brazil, for example, the Caduveo and the Palikur among the indigenous groups have oral traditions where the male dominance bias is negative. This means that women dominate men, at least in images and episodes narrated among these groups. On the contrary, in Umotina's folklore, the male dominance bias estimate is 0.45. It is the highest across all groups in Brazil and ranks in the top 10 among oral traditions globally. Within Russia, Chechens top the male dominance bias across indigenous folklore with a score of 0.25, with groups like the Kets (a Yeniseian people in Siberia) being at the bottom with a score of 0.05.

Influential research on gender norms links the historical status of women in agriculture to historical gender roles. Boserup (1970), for example, identifies in the use of the plow a diminished role of women. The agricultural roots of gender norms are clear in the examples described above. According to the EA among the Khasi, who did not use the plow historically, females do appreciably more agriculture, whereas the opposite is true among the Kashmiri. Similarly, Chechens in Russia have historically been a plow-using society with males dominating agriculture, whereas farming was historically absent among the Kets. Finally, although none of the indigenous groups had the plow in Brazil, among the Umotina, only males worked in agriculture. Among the Caduveo and the Palikur, women participated at least as much as males.

We formally explore this pattern in Panel C of Table VII. We use variable $v54$ from the Ethnographic Atlas that records sex differences in agriculture and exclude groups for which agriculture is absent or unimportant, i.e., $v54$ equals 9. The independent variable in columns (1) to (3) is an indicator that takes the value one when males contribute more than females to agriculture. In columns (4) to (6) the independent variable takes the value 1 for groups that had the plow before contact with the Europeans. Across and within continents as well within modern-day countries, historically plow-using societies and those where men did most of the

agricultural work are communities whose traditional stories have a clear gender bias. Men are dominant and physically active, whereas women are submissive and engage in domestic affairs. In the most conservative specifications in columns (3) and (6), the estimated beta coefficients are about 0.10. Figures $VIIIb$ and $VIIIc$ depict the binned residual scatter plots from columns (2) and (4), respectively.

Our work builds on several studies exploring gender norms in specific tales and groups. See, for example, the early contributions by Elder (1968), Upadhyaya (1968), Howe and Hirschfeld (1981), and Weigle (1982) culminating in an exhaustive multi-volume encyclopedia on women's folklore across the globe by Locke, Vaughan, and Greenhill (2009). Our study is the first one in economics and folklore studies to establish econometrically a large-scale association between male bias in folklore and the status of women historically across ethnic groups, and today across countries, and 2nd-generation immigrants. It also offers empirical support to the thesis of the famous American folklorist Lewis (1974) regarding the role traditional narratives play (primarily proverbs and jokes) in supporting and maintaining the negative stereotypes about women's status and teaching such values to the members of society.

6 Concluding Remarks

Narratives are essential building blocks of our society. We introduce to economics a unique catalogue of oral traditions across approximately $1,000$ groups. After validating folklore's content showing that episodes in a group's oral tradition reflect its geographic and social attributes as articulated in the ethnographic record, we undertake a series of applications illustrating how to extract information from folklore. In the first set, we illustrate how to fill in the gaps and expand upon a group's ethnographic record. In the second set, we discuss how machine learning and human-classification methods can help shed light on ancestral norms. Our initial examination indicates a striking consistency between values derived from folklore and contemporary attitudes related to trust, risk-taking, and gender norms. Images and episodes in folklore appear to endure and, possibly, still shape how individuals perceive the world today.

Next Steps We view this study as a springboard for further research. The finding that folklore-based measures of the economy and the polity correspond to what we know from ethnographers suggests that we can obtain more precise estimates of a group's heritage by combining the two sources. Lowering the measurement error in the historical record will allow us to revisit and better understand our societies' legacies and their consequences. One related idea is to use folklore to fill in the EA and SCCS gaps for the *universe* of recorded traits along the lines described in Section 4. Moreover, one can utilize folklore to derive bilateral measures of cultural

proximity, see Spolaore and Wacziarg (2009), or explore how different geographical traits and historical events influence the content of oral traditions. For example, what do groups located in malaria-prone regions, fertile territories, or rugged terrains "talk" about? Similarly, what are the distinctive themes in the folklore of groups that have experienced disruptions from slavery, epidemics, forced migrations, and colonization? This approach would allow testing famous conjectures in anthropology including the "culture of honor" proposed by Goldschmidt and Edgerton in 1971 and "the original affluent society hypothesis" by Sahlins (1972).

There is a long list of contemporary values and attitudes in regional and global surveys that we have not covered, including patience, aspirations, reciprocity, attitudes towards violence, strangers, the elderly, the community, the importance of imagination, obedience, independence, hard work, honesty, etc. We hope that the roadmap provided here can help trace these values in the respective oral traditions. Obtaining folklore-based measures of these attitudes may help us better understand the cultural traits that are stable over time.

Another avenue of future research relates to how motifs and concepts have traveled across oral traditions. Some motifs appear to be universal, whereas others are found in a handful of folklore traditions. Is there a pattern in the content of localized versus universal narratives? Moreover, the multiplicity of characters' attributes in a given motif and oral tradition (at least as classified by humans) may convey important information about the richness and the ambiguity of the characters' personality. This within-oral tradition diversity in attitudes may provide a way to gauge the degree of flexibility in the norms transmitted intergenerationally. It would also be interesting to explore how the individual characteristics of those reading and classifying the motifs may systematically predict how a given motif is perceived. Finally, we posit that the degree of continuity in the narratives between contemporary children's books and the folktales and myths of the respective societies is a direct measure of the rate at which ancestral norms are intergenerationally transmitted.

Given the versatility of folklore as a vehicle for obtaining a unique (and perhaps our only) view of our ancestral cultural heritage, we expect it to be useful to scholars interested in the historical origins of comparative development, social psychology, culture, and beyond.

References

Akerlof, G. A., and D. J. Snower (2016): "Bread and Bullets," *Journal of Economic Behavior and Organization*, 126(Part B), 58–71.

Alesina, A., P. Giuliano, and N. Nunn (2013): "On the Origins of Gender Roles: Women and the Plough," *Quarterly Journal of Economics*, 128(2), 469–530.

Algan, Y., and P. Cahuc (2010): "Inherited Trust and Growth," *American Economic Review*, 100, 2060–2092.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 131(4), 1593–1636.

Bates, R. H. (1983): *Essays on the Political Economy of Rural Africa*. University of California Press, Berkeley, CA.

Becker, A. (2019): "On the Economic Origins of Restrictions on Women Sexuality," mimeo, Harvard University.

Benabou, R., A. Falk, and J. Tirole (2018): "Narratives, Imperatives, and Moral Reasoning," NBER Working Paper No. 24798.

Bentzen, J. S. (2019): "Acts of God: Religiosity and Natural Disasters Across Subnational World Districts," *Economic Journal*, 129(622), 2295–2321.

Berezkin, Y. (2015a): "Folklore and Mythology Catalogue: its Lay-out and Potential for Research," in *The Retrospective Methods Network Newsletter*, ed. by Frog, H. F. Leslie-Jacobsen, and J. S. Hopkins, vol. 10, chap. Between Text and Practice: Mythology, Religion and Research, pp. 58–70. Folklore Studies / Dept. of Philosophy, History, Culture and Art Studies University of Helsinki, Helsinki, Helsinki.

——— (2015b): "Spread of Folklore Motifs as a Proxy for Information Exchange: Contact Zones and Borderlines in Eurasia," *Trames: Journal of Humanities and Social Sciences*, 19(1), 3–13.

——— (2016): "Peopling of the New World in Light of the Data on Distribution of Folklore Motifs," in *Maths Meets Myths: Quantitative Approaches to Ancient Narratives (Understanding Complex Systems)*, ed. by R. Kenna, M. M. Carron, and P. M. Carron, vol. 10, pp. 71–89. Springer Verlag.

Berezkin, Y., and E. Duvakin (2016): "Buried in a Head: African and Asian Parallels to Aesop's Fable," *Folklore*, 127(1), 91–102.

Boas, F. (1916): "The Development of Folk-Tales and Myths," *The Scientific Monthly*, 3(4), 335–343.

Boserup, E. (1970): *Woman's Role in Economic Development*. Allen and Unwin, London.

Brown, M. E. (1989): "Review Women, Folklore and Feminism," *Journal of Folklore Research*, 26(3), 259–264.

Cecil, D. J., D. E. Buechler, and R. J. Blakeslee (2014): "Gridded Lightning Climatology from TRMM-LIS and OTD: Dataset Description," *Atmospheric Research*, 135-136, 404–414.

Chen, K. (2013): "The Effect of Language on Economic Behavior: Evidence from Savings Rates, Health Behaviors, and Retirement Assets," *American Economic Review*, 103(2), 690–731.

Edgerton, R. D. (1971): "The Individual in Cultural Adaptation: a Study of Four East African Peoples," in *Studies in Culture and Ecology in East Africa*, ed. by W. GoldSchmidt. University of California Press, Berkeley and London.

El-Shamy, H. M. (2004): *Types of the Folktale in the Arab World: A Demographically Oriented Tale-Type Index*. Indiana University Press, Bloomington.

Elder, J. D. (1968): "The Male Female Conflict in Calypso," *Caribbean Quarterly*, 14(3), 23–41.

Eldredge, E. (1991): "Women in Production: The Economic Role of Women in Nineteenth-Century Lesotho," *Signs: Journal of Women in Culture and Society*, 16(4), 707–731.

Enke, B. (2020): "Moral Values and Voting," Journal of Political Economy, forthcoming.

Evans, J. A., and P. Aceves (2016): "Machine Translation: Mining Text for Social Theory," *Annual Review of Sociology*, 42, 21–50.

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde (2018): "Global Evidence on Economic Preferences," *Quarterly Journal of Economics*, 133(4), 1645–1692.

Fenske, J. (2014): "Ecology, Trade, and States in Pre-colonial Africa," *Journal of the European Economic Association*, 12(3), 612–640.

Fernandez, R. (2011): "Does Culture Matter?," in *Handbook of Social Economics*, ed. by J. Benhabib, M. O. Jackson, and A. Bisin, vol. 1A, pp. 481–510. North-Holland, Amsterdam.

Galor, O., and O. Ozak (2015): "Land Productivity and Economic Development: Caloric Suitability vs. Agricultural Suitability," mimeo, Brown University.

Galor, O., O. Ozak, and A. Sarid (2016): "Geographical Origins and Economic Consequences of Language Structures," Brown University Working Paper.

Gentzkow, M., B. T. Kelly, and M. Taddy (2019): "Text as Data," *Journal of Economic Literature*, 57(3), 535–574.

Goldberg, C. (1984): "The Historic-Geographic Method: Past and Future," *Journal of Folklore Research*, 21(1), 1–18.

GoldSchmidt, W. (1971): "Introduction: The Theory of Cultural Adaptation, The Individual in Cultural Adaptation: A Study of Four East African Peoples," in *Studies in Culture and Ecology in East Africa*, ed. by W. GoldSchmidt. University of California Press, Berkeley and London.

Grimmer, J., and B. Stewart (2013): "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts," *Political Analysis*, 21(3), 267–297.

Harari, Y. N. (2015): *Sapiens: A Brief History of Humankind*. HarperCollins, New York.

Howe, J., and L. A. Hirschfeld (1981): "The Star Girls' Descent: A Myth about Men, Women, Matrilocality and Singing," *Journal of American Folklore*, 94, 292–322.

Jason, H. (1970): "The Russian Criticism of the Finnish School in Folktale Scholarship," *Norveg*, 14, 285–294.

Karsdorp, F., M. van der Meulen, T. Meder, and A. van den Bosch (2015): "MOMFER: A Search Engine of Thompson Motif-Index of Folk Literature," *Folklore*, 126(1), 37–52.

Lauzen, M. M., D. M. Dozier, and N. Horan (2008): "Constructing Gender Stereotypes Through Social Roles in Prime-Time Television," *Journal of Broadcasting & Electronic Media*, 52(2), 200–214.

Lewis, M. E. (1974): "The Feminists Have Done It: Applied Folklore," *Journal of American Folklore*, 87(343).

Locke, L., T. A. Vaughan, and P. Greenhill (2009): *Encyclopedia of Women's Folklore and Folklife*. Greenwood Press.

Michalopoulos, S., A. Naghavi, and G. Prarolo (2018): "Trade and Geography in the Spread of Islam," *Economic Journal*, 128(616), 3210–3241.

Michalopoulos, S., and E. Papaioannou (2017): *The Long Economic and Political Shadow of History: Vol. 1: A Global View*chap. Introduction: Historical Legacies and Contemporary Development, pp. viii–xxix. VoxEU. PDF e-book.

Mullainathan, S., and J. Spiess (2017): "Machine Learning: An Applied Econometric Approach," *Journal of Economic Perspectives*, 31(2), 87–106.

Murdock, G. P. (1967): *Ethnographic Atlas: A Summary*. University of Pittsburgh Press, Pittsburgh, PA.

Nunn, N. (2020): "The Historical Roots of Economic Development," *Science*, 367(6485), eaaz9986.

Nunn, N., and N. Qian (2020): "Understanding Cultural Persistence and Change," Review of Economic Studies, forthcoming.

Putterman, L., and D. N. Weil (2010): "Post-1500 Population Flows and the Long-Run Determinants of Economic Growth and Inequality," *Quarterly Journal of Economics*, 125(4), 1627–1682.

Qian, N. (2008): "Missing Women and the Price of Tea in China: The Effect of Sex-Specific Earnings on Sex Imbalance," *Quarterly Journal of Economics*, 123, 1251–1285.

Rodrik, D. (2017): "Economics as Narratives," mimeo, Harvard University.

Ross, R. M., S. J. Greenhill, and Q. D. Atkinson (2013): "Population Structure and Cultural Geography of a Folktale in Europe," *Proceedings of the Royal Society B*, 280(1756).

Sahlins, M. (1972): *Stone Age Economics*. Aldine Publishing, Chicago.

Saiz, A., and U. Simonsohn (2013): "Proxying for Unobservable Variables with Internet Document-Frequency," *Journal of the European Economic Association*, 11(1), 137–165.

Schaller, M., A. Norenzayan, S. Heine, T. Yamagishi, and T. Kameda (2010): *Evolution, Culture, and the Human Mind*. Psychology Press.

Shiller, R. J. (2017): "Narrative Economics," *American Economic Review*, 107(4), 967–1004.

Silva, S., and T. Jamshid (2016): "Comparative Phylogenetic Analyses Uncover the Ancient Roots of Indo-European Folktales," *Royal Society Open Science*, 3, 150645.

Speer, R., and J. Lowry-Duda (2017): "Conceptnet at semeval-2017 task 2: Extending word embeddings with multilingual relational knowledge," *arXiv preprint arXiv:1704.03560*.

Spolaore, E., and R. Wacziarg (2009): "The Diffusion of Development," *Quarterly Journal of Economics*, 124(2), 469–529.

Tetlock, P. C. (2007): "Giving Content to Investor Sentiment: The Role of Media in the Stock Market," *Journal of Finance*, 62(3), 1139–1168.

Towbin, M. A., S. A. Haddock, T. S. Zimmerman, L. K. Lund, and L. R. Tanner (2004): "Images of Gender, Race, Age, and Sexual Orientation in Disney Feature-Length Animated Films," *Journal of Feminist Family Therapy*, 15(4), 19–44.

Upadhyaya, H. S. (1968): "Mother-Daughter Relationship Patterns in the Hindu Joint Family A Study Based on the Analysis of the Bhojpuri Folksongs of India," *Folklore*, 79(3), 217–226.

Voigtlander, N., and H.-J. Voth (2012): "Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany," *Quarterly Journal of Economics*, 127(1), 1339–1392.

Weigle, M. (1982): *Spiders and Spinsters Women and Mythology*. University of New Mexico Press.

Whitehouse, H., P. Francois, P. E. Savage, T. E. Currie, K. C. Feeney, E. Cioni, R. Purcell, R. M. Ross, J. Larson, J. Baines, B. ter Haar, A. Covey, and P. Turchin (2019): "Complex Societies Precede Moralizing Gods Throughout World History," *Nature*, 568(3), 226–229.

Xue, M. (2016): "High-Value Work and the Rise of Women: The Cotton Revolution and Gender Equality in China," mimeo.

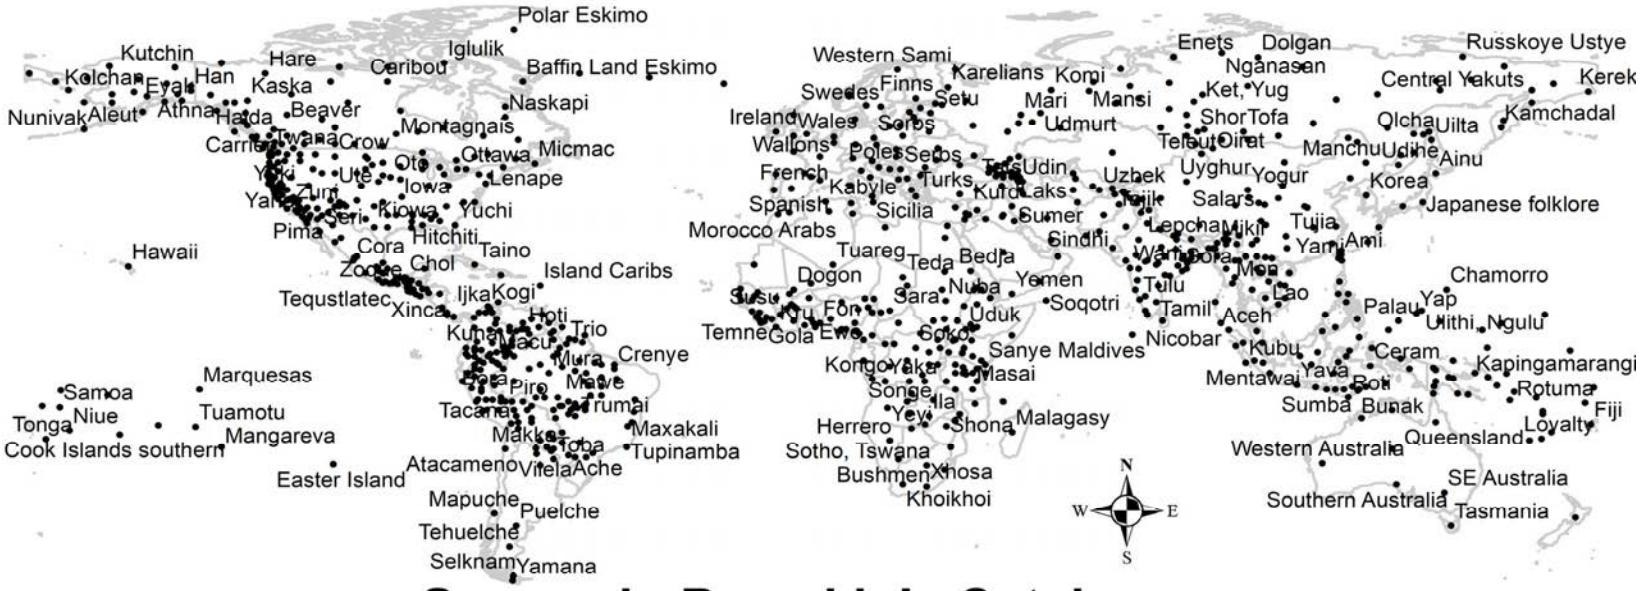

Groups in Berezkin's Catalogue

Figure I

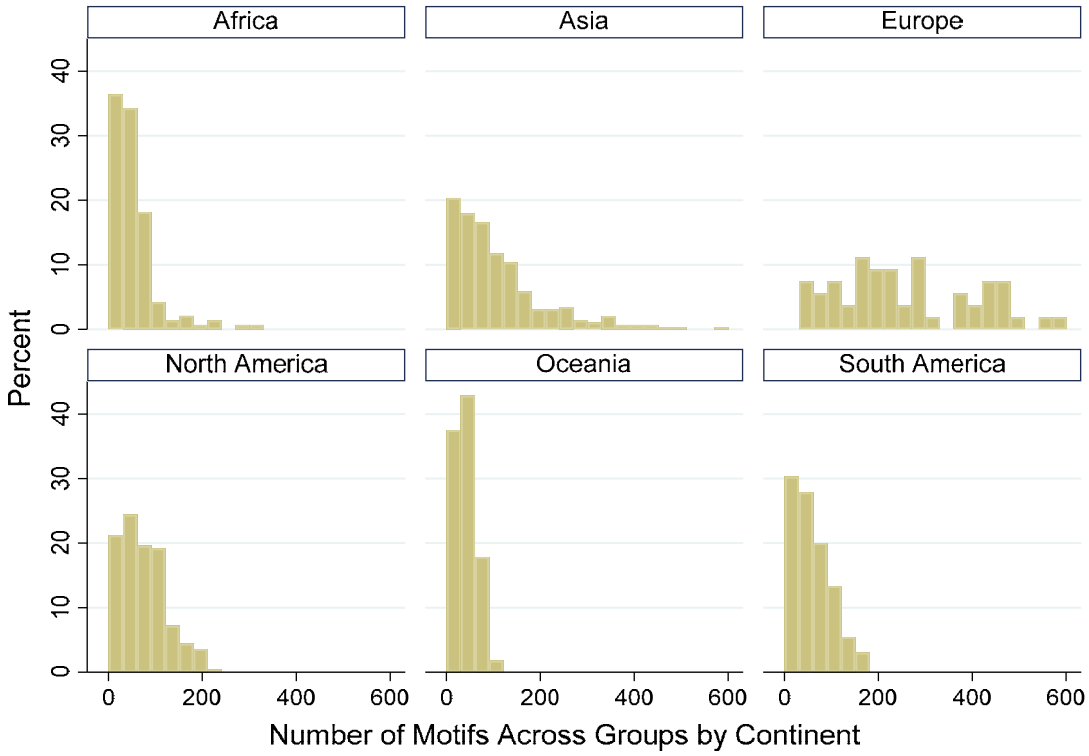

Figure II

Figure IIIa

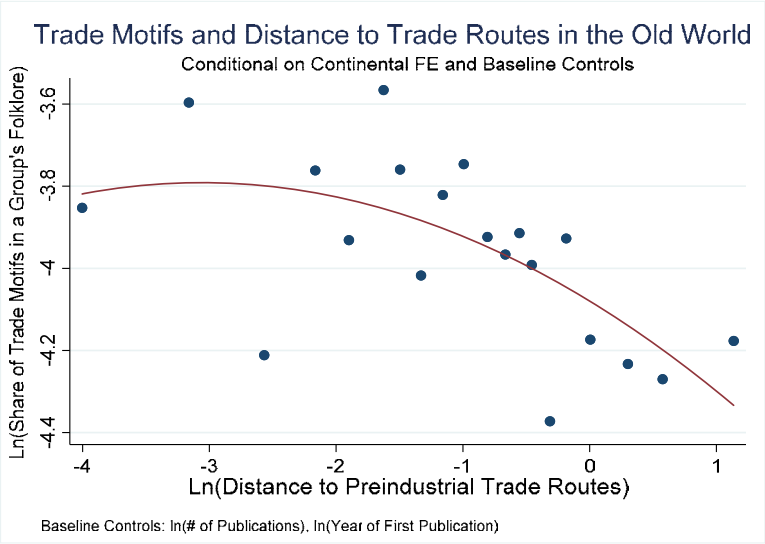

Figure IIIb

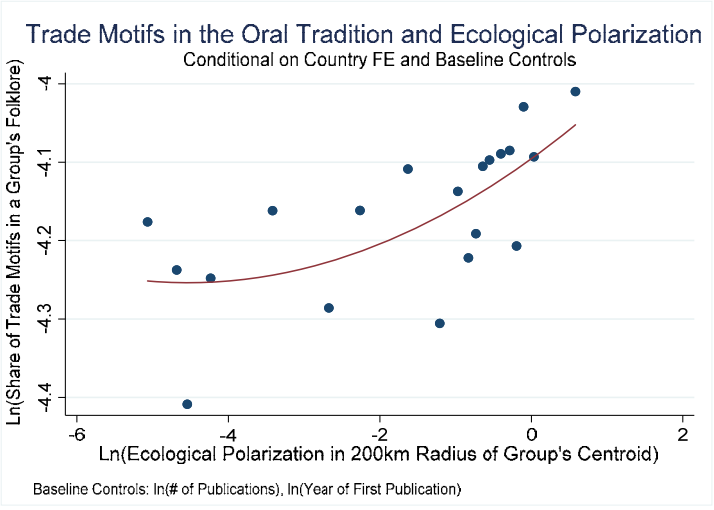

Figure IIIc

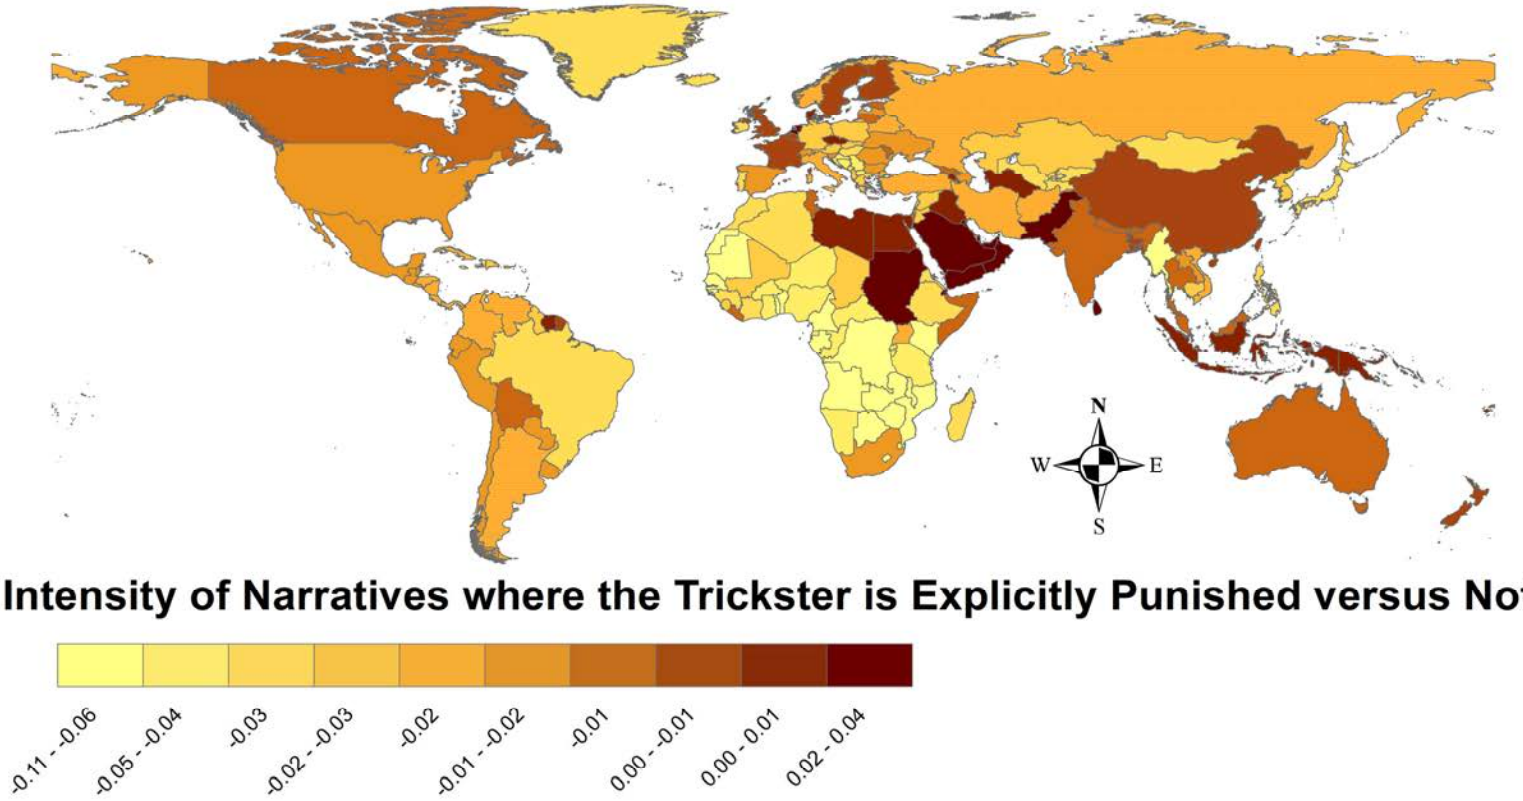

Figure IVa

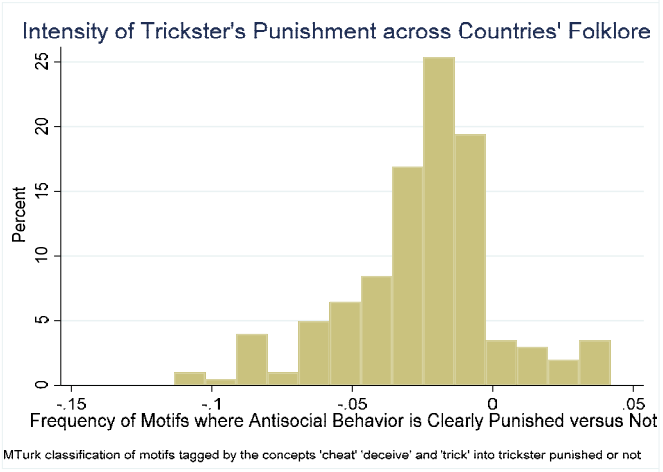

Figure IVb

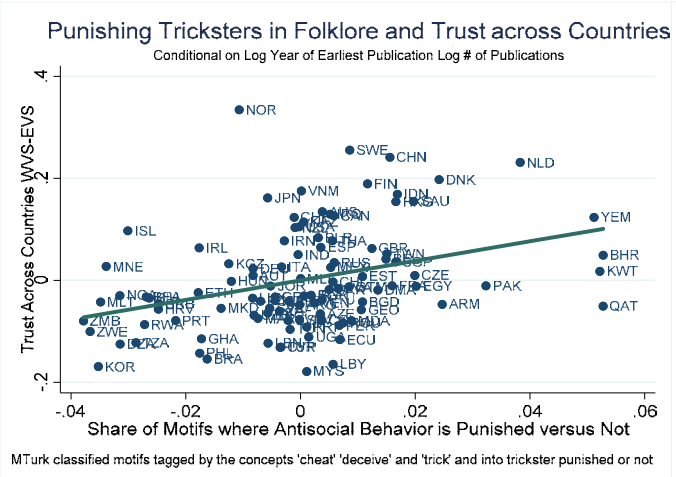

Figure IVc

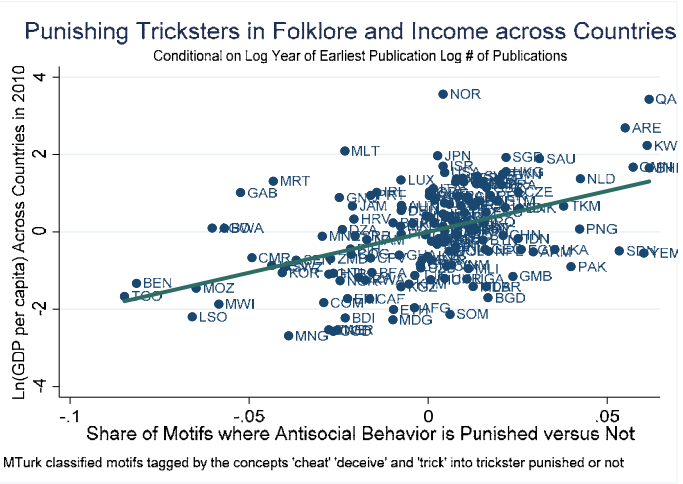

Figure IVd

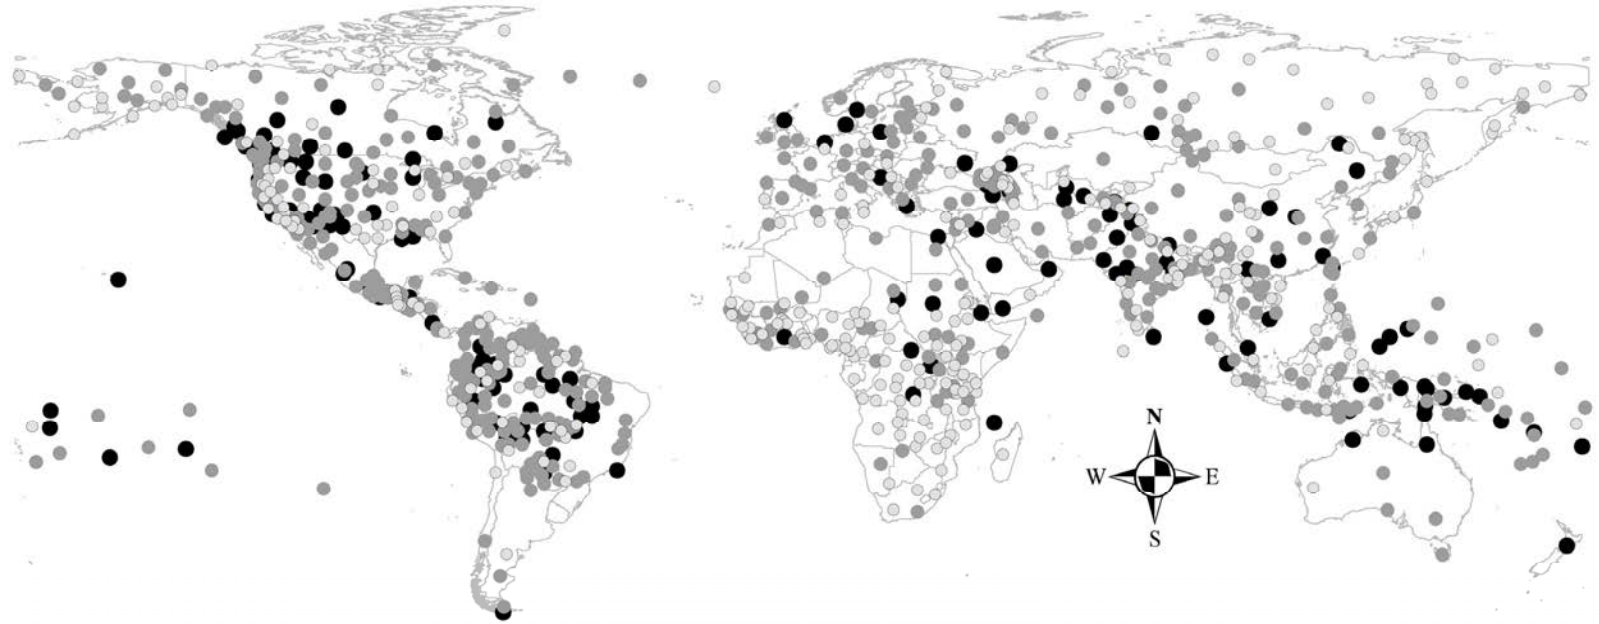

Figure Va

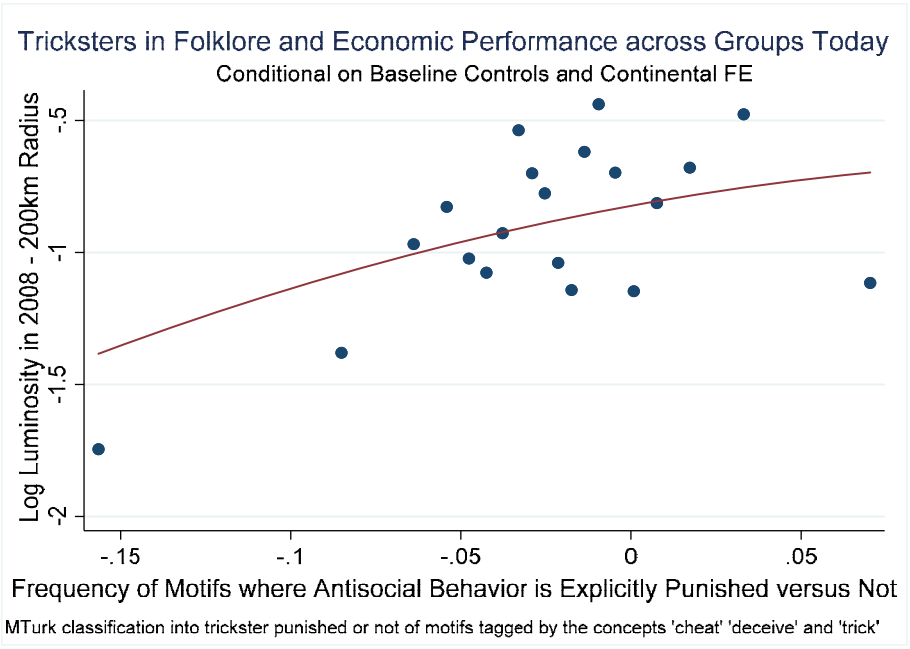

Figure Vb

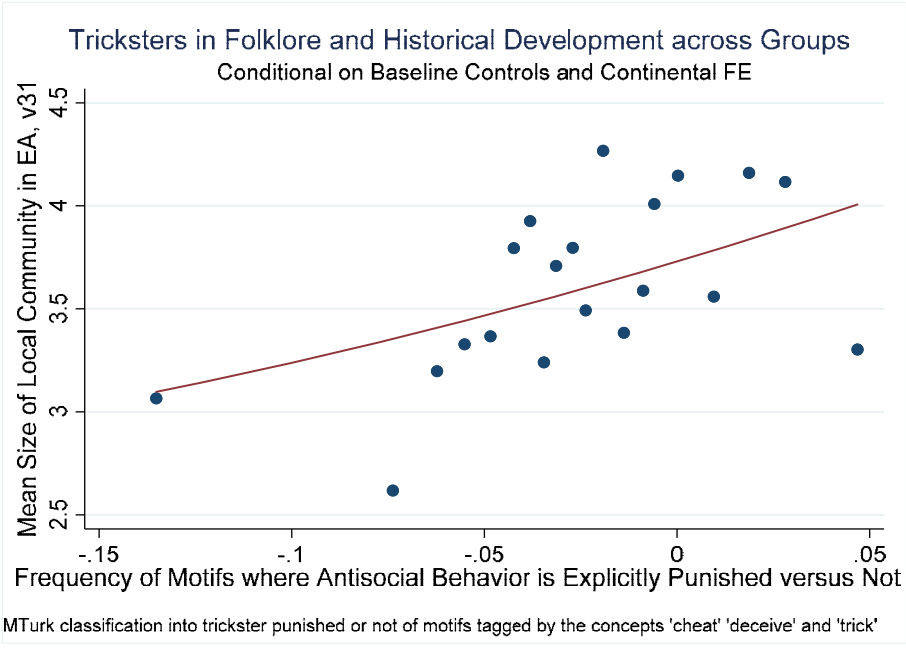

Figure Vc

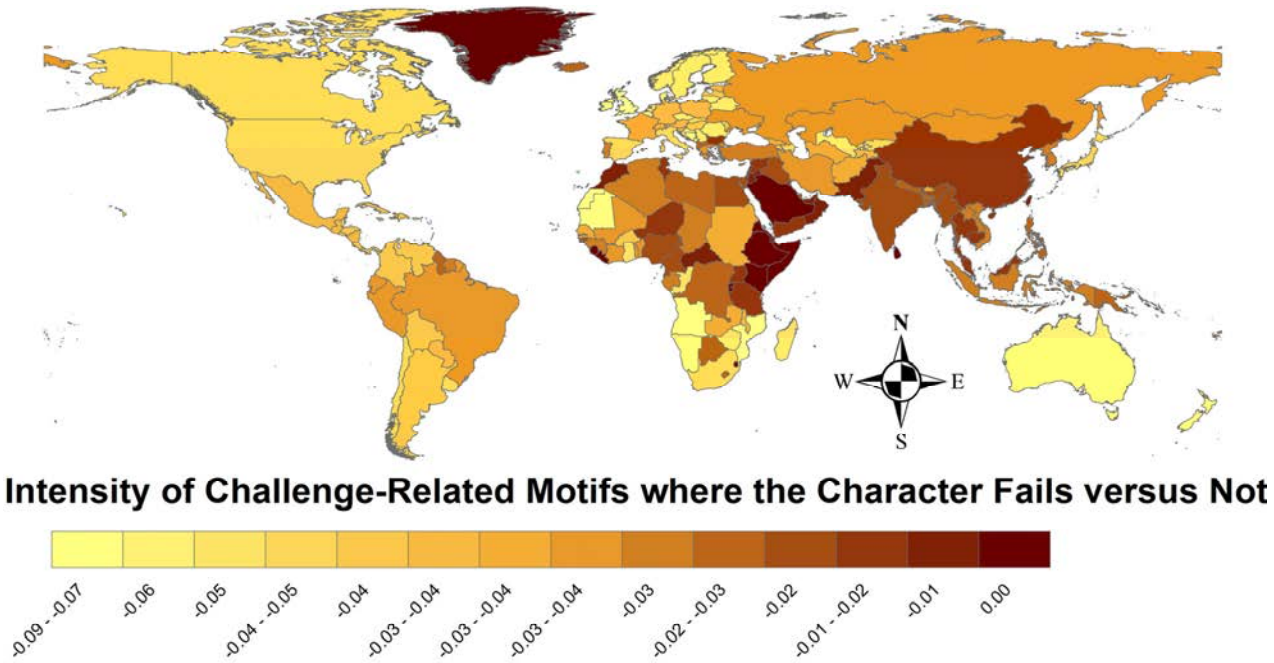

Figure VIa

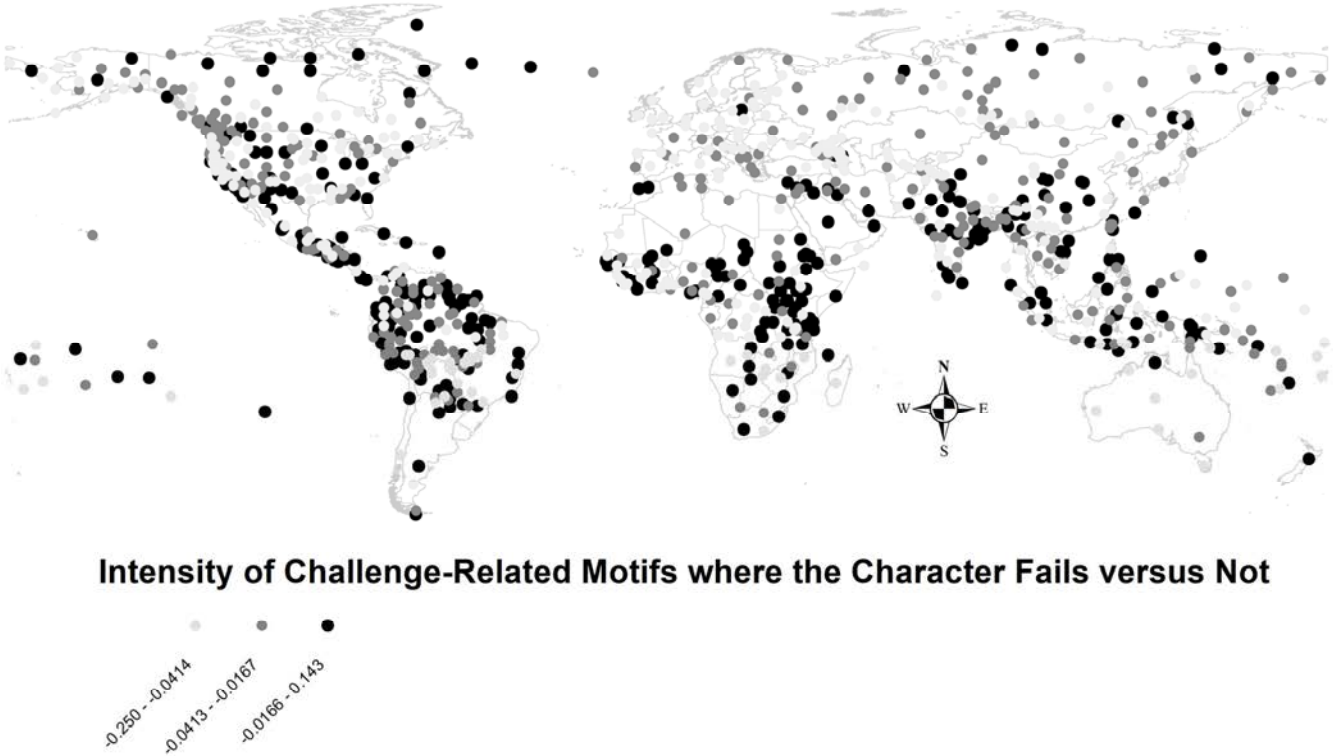

Figure VIb

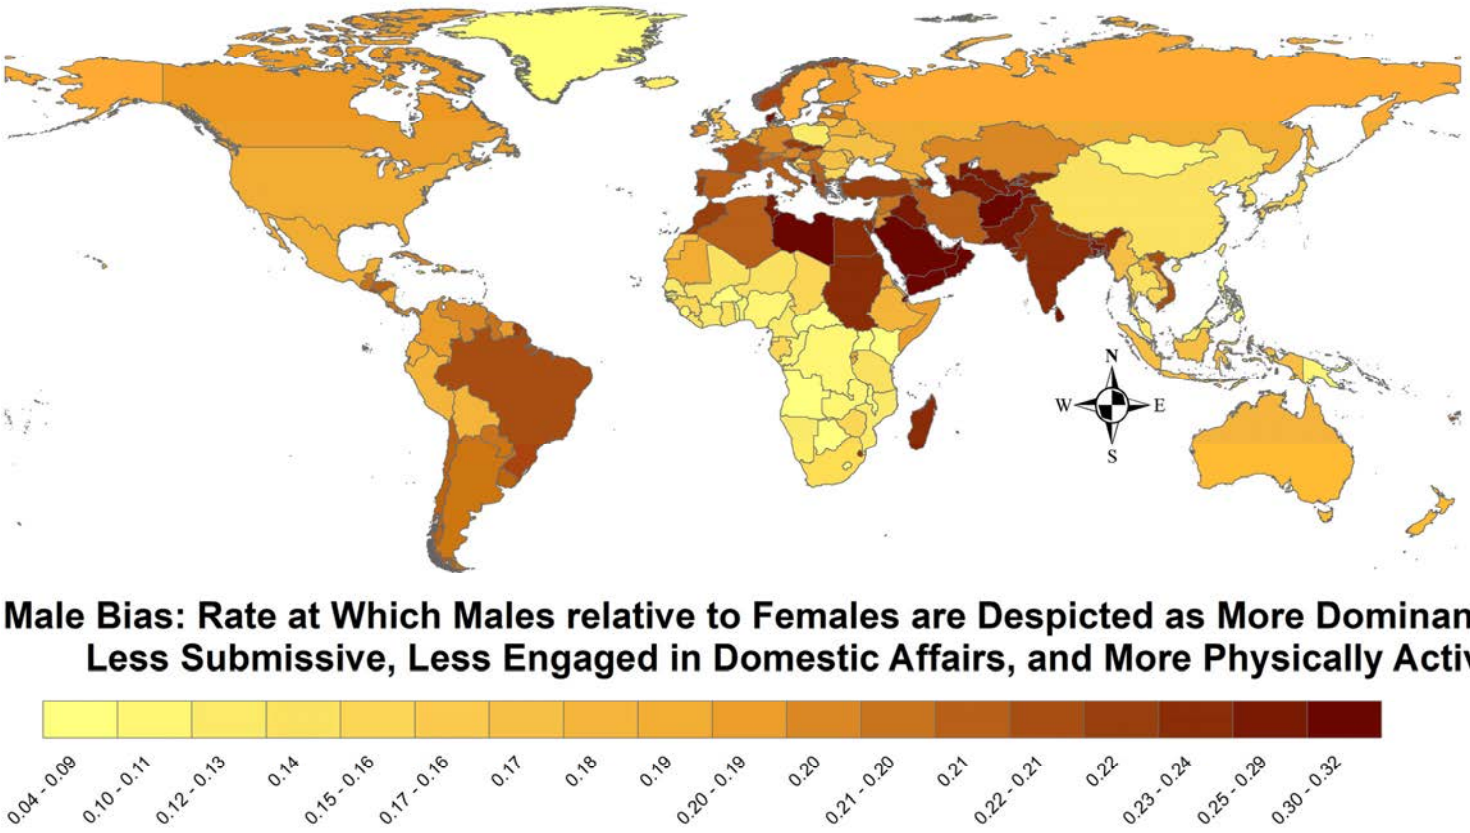

Figure VIIa

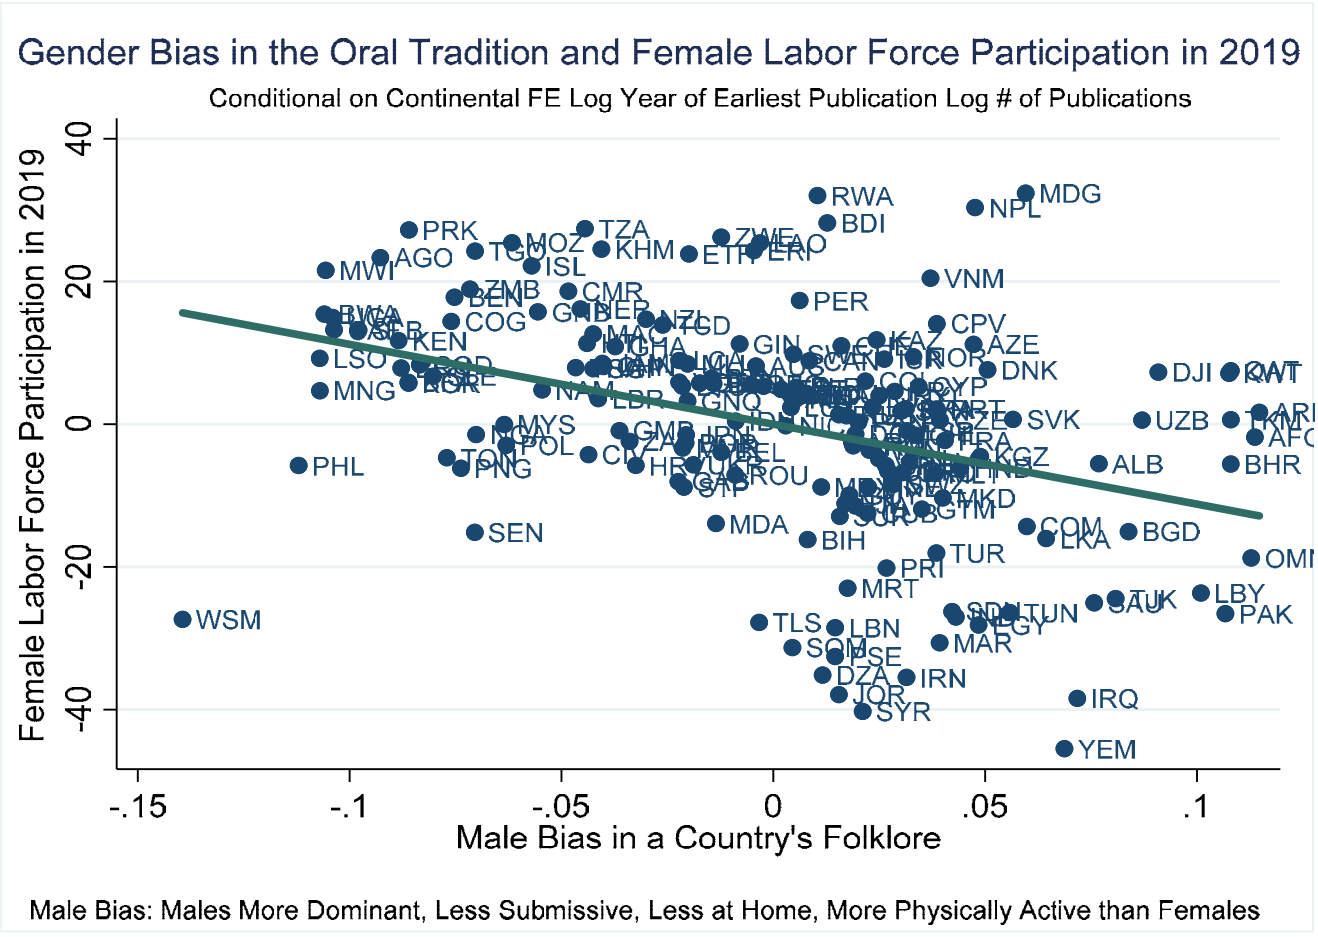

Figure VIIb

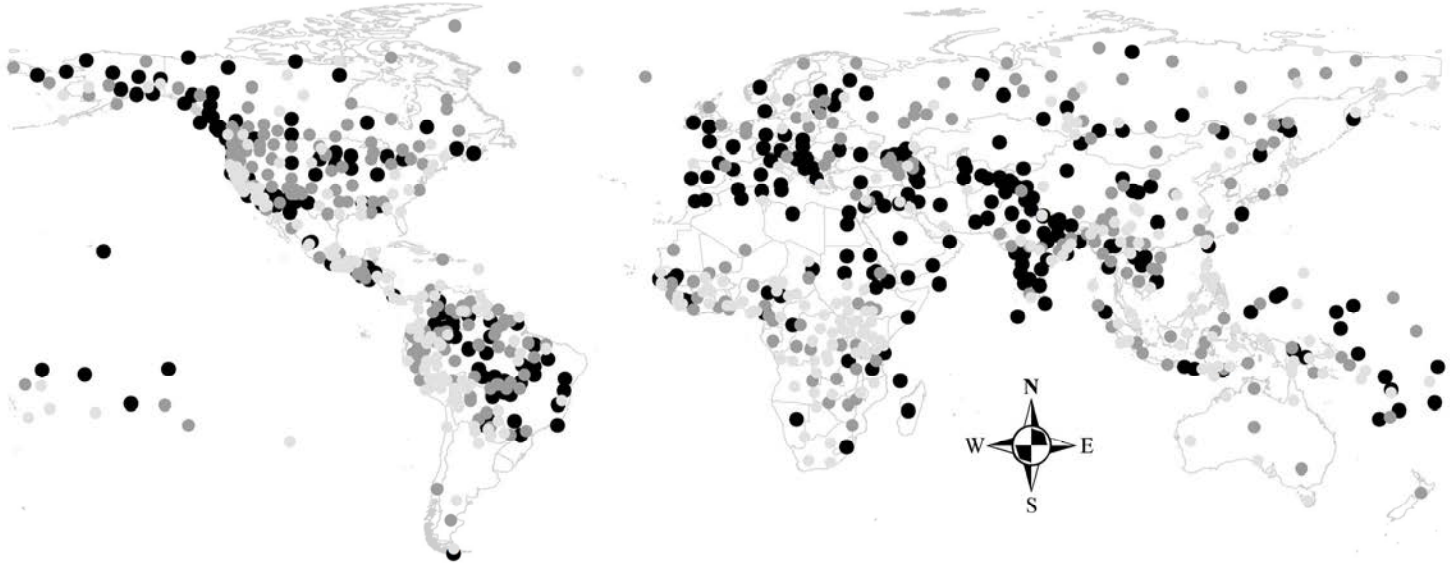

Male Bias: Rate at Which Males relative to Females are Despicted as More Dominant, Less Submissive, Less Engaged in Domestic Affairs, More Physically Active

-0.500 - 0.097

0.098 - 0.187

0.188 - 1.000

Figure VIIIa

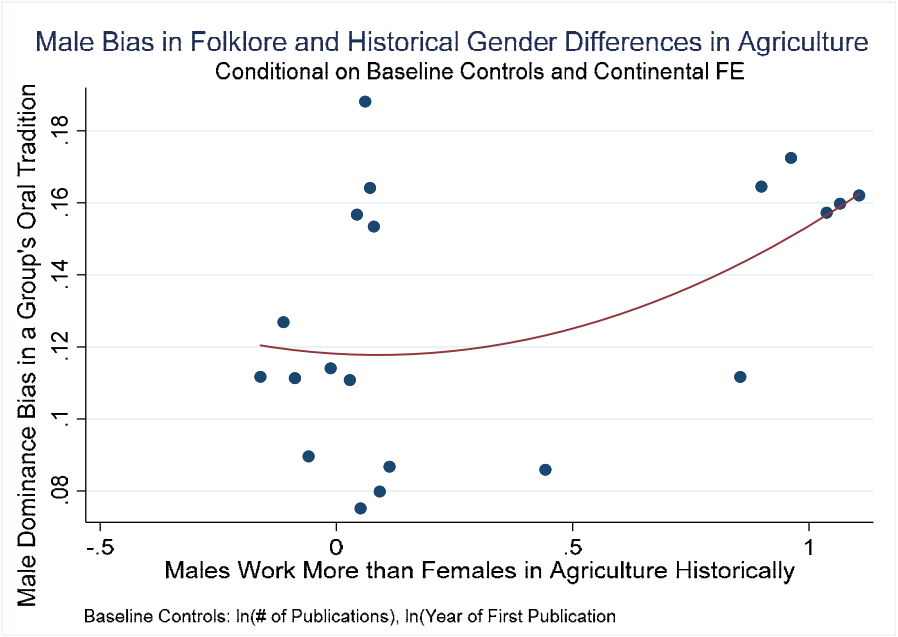

Figure VIIIb

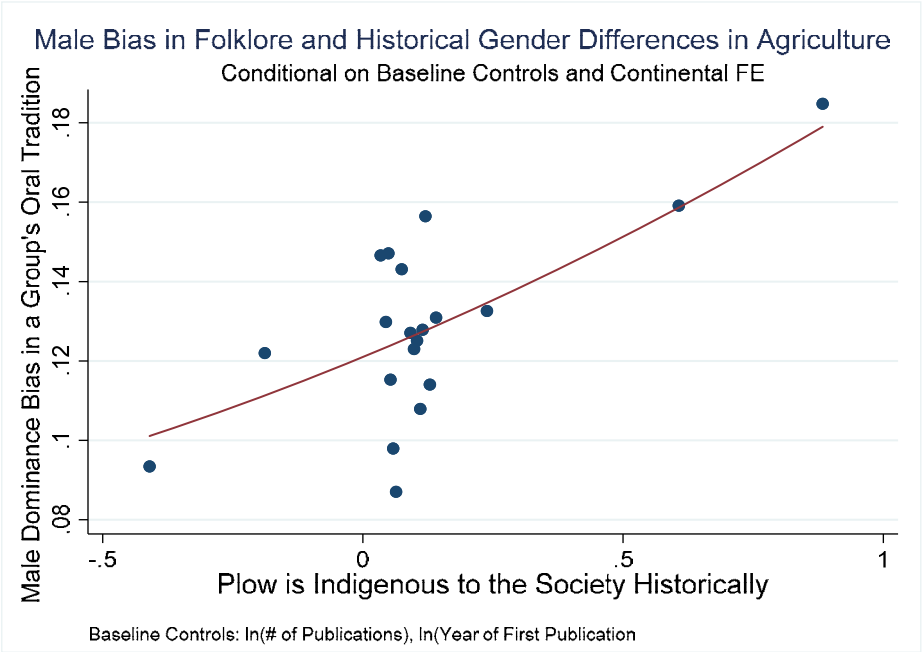

Figure VIIIc

Table I: Folklore and the Physical Environment

	Share of Motifs on									
	Earthquake		Storm		Frozen		Crops		River	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Ln(Distance to Earthquake Zones in 1,000 km's)	-0.1559***	-0.1659***								
	(0.0338)	(0.0508)								
Ln(Mean Lightning Flash Density)			0.0206	0.0412*						
			(0.0195)	(0.0224)						
Mean Yearly Temperature 1900-2000					-0.0194***	-0.0152**				
					(0.0026)	(0.0060)				
Ln(Optimal Agricultural Calories Pre-1500)							0.0594***	0.0550***		
							(0.0140)	(0.0144)		
Ln(Distance to Rivers in Km)									-0.0481*	-0.0536*
									(0.0244)	(0.0313)
Baseline Controls	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Continental FE	yes	no	yes	no	yes	no	yes	no	yes	no
Country FE	no	yes	no	yes	no	yes	no	yes	no	yes
Adjusted R^2	0.038	0.057	0.185	0.203	0.186	0.189	0.183	0.243	0.116	0.129
# of Observations	943	943	942	942	913	913	942	942	943	943

Notes: This table reports OLS estimates. The unit of analysis is a group in Berezkin's catalogue. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). As a dependent variable we use the Ln(0.01 + Share of Concept-Specific Motifs). Standard errors are clustered at the language family level as recorded by Berezkin. ***, **, * denote significance at 1%, 5%, and 10%, respectively. See the main text for variable definition and Appendix Table 2 - Panel A for the summary statistics.

Table II: Folklore, Subsistence, and the Ethnographic Record

	Panel A							
	Share of Motifs on:							
	Crops				**Pastoralism**			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Share of Subsistence from:								
Agriculture	0.0615***	0.0490***					-0.0765***	-0.0686***
	(0.0159)	(0.0124)					(0.0178)	(0.0211)
Animal Husbandry			0.0275	0.0017	0.0722***	0.0685***		
			(0.0200)	(0.0154)	(0.0143)	(0.0176)		
Fishing			-0.0926***	-0.0593***			-0.0677***	-0.0655***
			(0.0215)	(0.0168)			(0.0121)	(0.0146)
Hunting and Gathering			-0.0705***	-0.0546***			-0.0753***	-0.0749***
			(0.0173)	(0.0153)			(0.0154)	(0.0212)
R^2	0.167	0.339	0.175	0.318	0.282	0.449	0.284	0.45

	Panel B							
	Share of Motifs on:							
	Fish				**Hunt**			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Share of Subsistence from:								
Agriculture			-0.0678***	-0.0691***			-0.0857***	-0.0507***
			(0.0166)	(0.0194)			(0.0147)	(0.0168)
Animal Husbandry			-0.0876***	-0.0837***			-0.0398*	-0.0124
			(0.0197)	(0.0293)			(0.0207)	(0.0199)
Fishing	0.0773***	0.0759***					-0.0540**	-0.0411*
	(0.0176)	(0.0211)					(0.0231)	(0.0227)
Hunting and Gathering			-0.0918***	-0.0869***	0.0749***	0.0432***		
			(0.0236)	(0.0283)	(0.0137)	(0.0156)		
Baseline Controls	yes	yes	yes	yes	yes	yes	yes	yes
Continental FE	yes	no	yes	no	yes	no	yes	no
Country FE	no	yes	no	yes	no	yes	no	yes
Adjusted R^2	0.264	0.34	0.267	0.34	0.119	0.235	0.125	0.236
# of Observations	1233	1233	1233	1233	1233	1233	1233	1233

Notes: This table reports OLS estimates. The unit of analysis is a group in the Ethnographic Atlas matched to an oral tradition in Berezkin's catalogue. Baseline Controls include: ln(# of Publications) and ln(Year of First Publication). As a dependent variable we use the Ln(0.01 + Share of Concept-Specific Motifs). Standard errors are clustered at the language family level as classified in the EA, v98. ***, **, * denote significance at 1%, 5%, and 10%, respectively. See the main text for variable definition and Appendix Table 2 - Panel B for summary statistics.

Table III: Comparing Correlations Between the Sample where Ethnographic Values are Observed to those in the Imputed Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A							
Pairwise Correlation Between								
State Centralization &	% Hunting and Gathering	% Fishing	% Animal Husbandry	% Agriculture	Intensive Agriculture	Class Stratification	High Gods	State Centralization
Non-Missing Sample	-0.4538	-0.2635	0.3987	0.3663	0.4245	0.4677	0.4003	1
# of Observations	1129	1129	1129	1129	1129	1032	724	1129
Sample of Missing Values	-0.2509	-0.2132	0.4545	-0.0209	0.4017	0.2605	0.2425	.
# of Observations	128	128	128	128	31	48	22	.
	Panel B							
High Gods &								
Non-Missing Sample	-0.3568	-0.3302	0.5442	0.1735	0.3601	0.2631		0.4003
# of Observations	748	748	748	748	748	697		724
Sample of Missing Values	-0.3899	-0.2777	0.3939	0.354	0.1309	0.1356	.	0.2806
# of Observations	500	500	500	500	403	376	.	22

Notes: Bivariate corrrelations between state centralization in Panel A and high gods in Panel B with 8 ethnographic variables. In the non-missing EA sample, state centralization (high gods) is the one recorded in the EA. In the sample of missing EA values state centralization (high gods) is the folklore-predicted one.

Table IV: Folklore and High Gods

	Share of Motifs in a Group's Oral Tradition Related to:				
	Punishment	Award	Supernatural	1st PC of Ln(Share of Motifs Related to Punishment, Award and Supernatural)	
	(1)	(2)	(3)	(4)	(5)
High Gods	0.1022***	0.0697*	0.0924*	0.1853***	0.1737***
	(0.0318)	(0.0371)	(0.0508)	(0.0490)	(0.0548)
Baseline Controls	yes	yes	yes	yes	yes
Continental FE	yes	yes	yes	yes	no
Country FE	no	no	no	no	yes
Adjusted R^2	0.254	0.309	0.175	0.334	0.489
# of Observations	733	733	733	733	733

Notes: This table reports OLS estimates. In all Columns the unit of analysis is a group in the Ethnographic Atlas, which has been matched to an oral tradition in Berezkin's catalogue. Baseline Controls include: ln(# of Publications) and ln(Year of First Publication). As a dependent variable we use the Ln(0.01 + Share of Concept-Specific Motifs). Standard errors are clustered at the language family level as classified in the EA, v98. ***, **, * denote significance at 1%, 5%, and 10%, respectively. See the main text for the variable definition and Appendix Table 2 - Panel B for the summary statistics.

Table V: Punishment of Tricksters in the Oral Traditions and Trust Today

Panel A: Across Countries

Dependent Variable	Trust (WVS/EVS)		Trust (GPS)		Ln(GDP pc in 2010)	
	(1)	(2)	(3)	(4)	(5)	(6)
Relative Punishment of Antisocial	1.8560***	1.5383***	5.3760***	3.9233*	20.8954***	11.0059***
	(0.4794)	(0.5372)	(1.1388)	(2.1495)	(3.4327)	(3.6698)
Baseline Controls	yes	yes	yes	yes	yes	yes
Continental FE	no	yes	no	yes	no	yes
Adjusted R^2	0.205	0.317	0.188	0.194	0.382	0.55
# of Observations	104	104	76	76	160	160

Panel B: Across 2nd-Generation Immigrants in Europe

Dependent Variable	How Much People Can Be Trusted		Income Decile	
	(1)	(2)	(3)	(4)
Relative Punishment of Antisocial	1.0988***	1.0659***	3.1897	5.7285**
	(0.3689)	(0.4018)	(3.5655)	(2.8876)
Baseline Controls	yes	yes	yes	yes
Individual Controls	no	yes	no	yes
Country - Round FE	yes	yes	yes	yes
Adjusted R^2	0.065	0.072	0.129	0.195
# of Observations	6554	6554	4720	4720

Panel C: Across Groups

Sample	In Berezkin's catalogue		In the Ethnographic Atlas			
Dependent Variable	Ln(Luminosity) in 2008		Historical Mean Size of Local Communities			
			All Groups		Predominantly Farming Societies	
	(1)	(2)	(3)	(4)	(5)	(6)
Relative Punishment of Antisocial	3.2826**	3.1853*	4.8432***	3.9384*	4.3934**	4.4554**
	(1.5809)	(1.6225)	(1.7276)	(2.0552)	(2.1410)	(1.9097)
Baseline Controls	yes	yes	yes	yes	yes	yes
Continental FE	yes	no	yes	no	yes	no
Country FE	no	yes	no	yes	no	yes
Adjusted R^2	0.357	0.487	0.256	0.4	0.289	0.408
# of Observations	942	942	578	578	365	365

Notes: Both Panels report OLS estimates. In Panel A the unit of analysis is a country. In Panel B we use individual level data from the European Social Surveys Rounds 1 to 8 and focus on respondents born in the country of the interview, and whose both parents are born in the *same* country abroad. In Panel C and in columns (1) & (2) an observation is a group in Berezkin's catalogue. In columns (3)-(6) we look at ethnicities in the Ethnographic Atlas. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). Individual Controls include age decile fixed effects, a gender indicator, self-reported health status fixed effects, and religion-specific constants. Standard errors (s.e.), reported in parentheses, are robust in Panel A. In Panel B, s.e. are clustered at the country of origin of the parents of each individual. In Panel C columns (1) & (2), s.e. are clustered at the language family (as classified by Berezkin). In columns (3)-(6), s.e. are clustered at the language family level as classified in the EA, v98. ***, **, * denote significance at 1%, 5%, and 10%, respectively.

Table VI: Challenges and Competitions across Oral Traditions and Risk-Taking & Entrepreneurship Today

	Across Countries								Across 2nd-Generation Immmigrants in Europe	
Dependent Variable	Risk-Taking (GPS)			Ln(# of Patents from Country's Residents Per 100,000 in 2006/2007)			Ln(New business registrations per 1,000 people ages 15-64 in 2006/2007)		Self-Employed	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Share of Motifs on Challenges and Competitions										
All	5.4381***	5.6802**								
	(1.7645)	(2.3076)								
where the character is **successful**			11.6165**							
			(4.5020)							
where the character is **unsuccessful**			-5.1016							
			(4.9413)							
where the character is **neither successful nor unsuccessful**			7.9918*							
			(4.3289)							
where **no challenges are present**			6.6558							
			(13.4506)							
where the character is **relatively unsuccessful**				-9.5625***	-31.0346***	-23.0564**	-30.2746***	-16.7612**	-1.8689***	-1.4029***
				(2.7163)	(11.4762)	(10.0005)	(7.1015)	(7.1745)	(0.3955)	(0.3906)
Baseline Controls	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Individual Controls	no	no	no	no	no	no	no	no	no	yes
Continental FE	no	yes	yes	yes	no	yes	no	yes	no	no
Country-Round FE	no	no	no	no	no	no	no	no	yes	yes
Adjusted R^2	0.13	0.235	0.257	0.274	0.409	0.535	0.263	0.343	0.057	0.077
# of Observations	76	76	76	76	140	140	147	147	6063	6063

Notes: In Columns (1)-(8) the unit of analysis is a country. In Columns (9)-(10) we use individual level data from the European Social Surveys Rounds 1 to 8 and focus on respondents born in the country of the interview, and whose both parents are born in the *same* country abroad. To construct the share of challenge-related motifs where the character is relatively unsuccessful we substract from the share of unsuccessful motifs the share of motifs where the character is either successful or not explicitly harmed. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). Individual Controls include age decile fixed effects, a gender indicator, self-reported health status fixed effects, and religion-specific constants. Standard errors (reported in parentheses) are robust in columns (1)-(8) and are clustered at the country of origin of the parents of each individual in columns (9)-(10). ***, **, * denote significance at 1%, 5%, and 10%, respectively.

Table VII: Gender Norms across Oral Traditions and the Role of Women Today and Historically

Panel A: Across Countries

	Female Labor Force Participation in 2019					
	(1)	(2)	(3)	(4)	(5)	(6)
Male Dominance Bias in the Oral Tradition	-111.9018***			-108.8809***		
	(22.2059)			(25.2061)		
Male Mental Capacity Bias in the Oral Tradition		69.8931			73.4037	
		(46.8905)			(53.2801)	
Male Sexual Bias in the Oral Tradition			73.8563			-5.0315
			(125.1493)			(145.6141)
Baseline Controls	yes	yes	yes	yes	yes	yes
Continental FE	no	no	no	yes	yes	yes
Adjusted R^2	0.136	0.004	-0.011	0.145	0.051	0.034
# of Observations	174	174	174	174	174	174

Panel B: Across 2nd-Generation Immigrants in Europe

	Agree with the Statement:					
	Housewife		Men should have more right to job than women when jobs are scarce		Women should be prepared to cut down on paid work for sake of family	
	(1)	(2)	(3)	(4)	(5)	(6)
Male Dominance Bias in the Oral Tradition	0.5739***	0.4206**	1.8022***	1.7888**	3.3846***	3.0251***
	(0.1799)	(0.1692)	(0.6563)	(0.7274)	(0.9648)	(0.9089)
Baseline Controls	yes	yes	yes	yes	yes	yes
Country - Round FE	yes	yes	yes	yes	yes	yes
Individual Controls	no	yes	no	yes	no	yes
Adjusted R^2	0.025	0.051	0.158	0.182	0.082	0.098
# of Observations	3468	3468	3765	3765	2726	2726

Notes: Both Panels report OLS estimates. In Panel A the unit of analysis is a country. In Panel B we use individual level data from the European Social Surveys Rounds 1 to 8 and focus on respondents born in the country of the interview, and whose both their parents are born in the *same* country abroad. Male Dominance Bias is defined as the share of motifs where males are depicted as more dominant, less submissive, less engaged in domestic affairs, and more physically active than females in a given oral tradition. Male Mental Capacity Bias is the share of motifs where males are depicted as more intelligent and less naive than females. Finally, Male Sexual Bias is the share of motifs where males are depicted as more sexual than females in the oral tradition. Columns (1) & (2) in Panel B focus only on female respondents. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication for the Oral Tradition). Individual Controls include age decile fixed effects, a gender indicator, self-reported health status fixed effects, and religion-specific constants. Standard errors (reported in parentheses) are robust in Panel A and are clustered at the country of origin of the parents of each individual in Panel B. ***, **, * denote significance at 1%, 5%, and 10%, respectively.

Panel C: Ethnic Groups in the Ethnographic Atlas

	Male Dominance Bias in the Oral Tradition					
	(1)	(2)	(3)	(4)	(5)	(6)
Males Contribute more than Females in Agriculture; EA	0.0408***	0.0345***	0.0191**			
	(0.0107)	(0.0111)	(0.0084)			
Plow is Indigenous to the Society; EA				0.0565***	0.0614***	0.0364**
				(0.0132)	(0.0082)	(0.0180)
Baseline Controls	yes	yes	yes	yes	yes	yes
Continental FE	no	yes	no	no	yes	no
Country FE	no	no	yes	no	no	yes
Adjusted R^2	0.056	0.083	0.201	0.049	0.064	0.179
# of Observations	702	702	702	1131	1131	1131

Notes: In Panel C the unit of analysis is a group in the Ethnographic Atlas. Male Dominance Bias is defined as the share of motifs where males are depicted as more dominant, less submissive, less engaged in domestic affairs, and more physically active than females in a given oral tradition; In columns (1) -(3) we focus on groups for which agriculture **is not an unimportant activity** according to variable v54 and the independent variable takes the value 1 if males contribute more to agriculture than females. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). Standard errors (reported in parentheses) are clustered at the linguistic family level (variable v98 in the EA). ***, **, * denote significance at 1%, 5%, and 10%, respectively.

Folklore
Appendix Material

Stelios Michalopoulos
Brown University, CEPR, and NBER

Melanie Meng Xue
NYU Abu Dhabi

January 8, 2021

Abstract

This Appendix consists of the following sections. In the first part, we provide a brief overview of the history of folklore and place Berezkin's catalogue in the context of other works in comparative mythology. Section 2 illustrates the construction of a motif in a given oral tradition. In section 3, we present additional robustness checks on the relationship between folklore and the natural environment. In section 4, we discuss the use of machine learning techniques in the context of folklore and present results implementing shrinkage methods to obtain keywords that predict current beliefs and attitudes across countries, as reflected in the Global Preference Survey (GPS) dataset. In the last section, we provide details on the use of MTurks for motif classification.

Keywords: Folklore, Culture, Development, Values, History

JEL classification Numbers: O10, Z10, Z13.

1 A Short History of Folklore

Folklore studies began in the early 19th century. In 1846, William Thoms invented the word "folklore" to replace existing terms, including "popular antiquities." The terms "folk" and "lore" simply refer to "ordinary people" and "knowledge," respectively. The first generation of folklorists focused on illiterate peasants and groups such as the Romani vagabonds, who would be more or less unaffected by the sweeping social changes of the era, and attempted to document their historical customs and beliefs preserved in the oral traditions. The understanding was that folklore reflected the cultural views of ordinary people in opposition to the elite. With increasing industrialization, urbanization, and literacy in Europe in the 19th century, folklorists were concerned that the oral tradition, the lore of the rural folk, would be lost. It was thought that the stories, beliefs, and customs were surviving fragments of the region's cultural mythology, often predating the spread of Christianity.

From an anthropological perspective, folklore is a critical component of a people's culture, Bascom (1953). It is considered a key mechanism for preserving a group's tradition. According to Bascom, there are four folklore functions: informally teaching cultural attitudes, escaping the accepted limitations of culture, maintaining cultural identity, and validating existing cultural norms.

For over 150 years from the early 19th to the mid-20th century, a vast body of work accumulated from collectors in all parts of the world as they listened to storytellers and published what they heard with improving recording techniques. Hence, the very nature of folklore, that is, its transmission via oral storytelling, might appear to be a source of idiosyncratic variability. According to Barre Toelken, however, this tendency was held in check by the forces of orthodoxy and tradition, which were the "twin laws of folklore performance," Toelken (1996). Audiences expected storytellers to retell familiar stories, and this expectation reined in tendencies to innovate or adapt folklore traditions. To rationalize the stability of the narrative, a famous early 20th-century folklorist, Walter Anderson, posited a double redundancy, that is, a feedback loop between performing and hearing the story multiple times, in order to retain the essential elements of the tale, see Dorst (2016) and Fine (1979).

The natural next step was the development of techniques to categorize this wealth of information for comparative analysis. This was a critical advance in the discipline of folklore, and the indexing of tale types and motifs lies in the heart of its comparative framework. The concept of a tale type was first delineated by the Hungarian folklorist Jànos Honti in 1939. Honti proposed three criteria for considering a tale type as an analytical unit: first, a tale type consists of a specific set of motifs; second, a tale type does not duplicate with other tale types; third, a tale type manifests itself through multiple occurrences (termed variants). According to Thompson

(1946), the motif is defined as "the smallest element in a tale having the power to persist in tradition." Both tale types and motifs are hypothetical archetypes established by comparing a large number of texts that share a common core. The methodology most closely associated with the use of tale types and motifs in comparative mythology is the historic-geographic approach that began in the late 19th century. It focused on establishing an oral tradition, identifying its geographic origin, and its spread, and it has remained popular as a methodological package for classifying textual sources and as a tool for organizing the data according to the degree of similarity. In 1982, Richard Dorson declared the historic-geographic method the dominant force in folklore science.

Folklorists working in the historic-geographic tradition often follow the Aarne-Thompson (AT) classification systems. The latter include the AT motif-index and the AT tale type index, which Uther updated and is now known as the Aarne-Thompson-Uther classification system, Uther (2004). The ATU classifies 2,239 tale types and their variants and the AT index refers to the motif-index of folklore literature created by Stith Thompson in 1955. It is important to keep in mind that the AT/ATU system was originally developed to study European oral traditions, limiting its usefulness for classifying folklore from other parts of the world. Moreover, ethnic attribution is rarely available for folklore found in the non-Western world, even in major regional indices that followed the ATU index. Although Thompson's motif index partially overcomes the lack of non-European coverage, the distribution of motifs remains skewed towards Europe, and obscene-type motifs are intentionally left out, see Dundes (1997).[1] There is also the Encyclopedia of the Folktale (Enzyklopädie des Märchens), an impressive compilation of almost two centuries of international research in the field of folk narrative tradition. However, it focuses on Europe's oral and literary narrative traditions and other countries influenced by Europeans.

Against this background of alternative classification systems, the critical dimension in Berezkin's Folklore and Mythology Catalogue is that he does not limit himself to the European-based tale type classifications and constructs a global database of motifs covering both European and non-European traditional narratives.

1.1 Interpreting Folklore

In the early years of folklore, folklorists' main task was to collect and classify different folklore materials, paying relatively little attention to its interpretation. Indeed, the primary focus was on the "lore" per se. During the 20th century, understanding the relationship between the "folk" and the "lore" became a focal topic for many preeminent scholars of different backgrounds. Summa-

[1] The following example illustrates the lopsided coverage of non-European groups in Thompson's motif index, for example. Take two groups the Irish and the Guarani. In the Thompson (Berezkin) index there are approximately 8,000 (268) motifs in the Irish oral tradition compared to 36 (82) recorded for the Guarani.

rizing the various approaches is not our goal. Below, we mention some of them to illustrate Joseph Campbell's prescient point made in 1968 that "there is no final system for the interpretation of myths, and there never will be any such thing."

There are three dominant approaches: the humanistic, the anthropological, and the psychological. Myths and folktales have been viewed as reflecting: (i) the customs, institutions, and beliefs of a given society, (ii) (distorted) historical facts, (iii) allegories of deeper philosophical issues, (iv) moral truths, (v) explanations of the cosmos, and (vi) metaphors or symbols of the unknown or of religious rituals.

According to the pattern theory of culture, all parts are related and reflect the same values and beliefs, with folklore offering a window into it, see Benedict (1934). Anthropologists have taken an ethnographic approach, a structuralist approach, and, more recently, a symbolic anthropology approach to shed light on the function and meaning of folklore, Green (1997).[2] The use of depth psychology, or psychoanalytic approaches, is the dominant method in interpreting folklore within psychology. Freud and Jung were pioneers in applying this approach. Jung (1968) viewed myths as essentially static symbolism, where recurring themes or motifs can be attributed to some deeply-seated human needs or conflicts, called psychological archetypes. Hence, the "primitive" mentality is what is "hidden in the unconscious and reflected in folkloric symbols."

Postmodernism and Folklore The rise in the popularity of postmodern critical approaches starting in the 1980s and 1990s in cultural studies, linguistics, and humanities, led to a decline in the appeal of comparative mythology in folkloristics. The reason is that relativism, which is in the heart of postmodernism, rejects comparative approaches as intrinsically flawed. This shift in the paradigm explains the scarcity of such works over the last 30 years and why the emergence of digital humanities, which brings together computing technologies and the humanities, has only recently made inroads in folklore studies.

Abello, Broadwell, and Tangherlini (2012) are among the first to discuss a path connecting the traditional close-reading approach in folklore to the distant reading of machine learning. There are two examples of recent efforts to create digital collections of annotated folk narratives. Karsdorp et al. (2015) have launched a search engine called MOMFER, for Thompson's Motif-Index of Folk Literature, whereas Meder et al. (2016) digitize and classify the Folktales in the Dutch Folktale Database. Ours is the first study that conducts text analysis on the monumental corpus of Yuri Berezkin and constructs a catalogue of oral traditions, concept frequencies, and folklore measures of trust, risk-aversion, and gender norms across countries and ethnicities.

[2] Lévi-Strauss, a famous proponent of the structural study of myth, sees in myth's structures elementary contradictions or problems that no society can resolve. Hence, to understand myths, one may focus on the repeated patterns and their structural relations to each other, Levi-Strauss (1955).

2 Going Behind the Scenes of Berezkin's Catalogue

The Construction of a Motif Motif $k27n1$ is titled "Task-giver is a king or a chief," and the description reads: "Person who gives difficult tasks to the hero is a prominent figure in the social hierarchy, i.e., the head of the political unit in the community - or higher level and not a mythical being." This motif is in the top-10 widespread motifs present in 228 societies (Appendix Figure 3*c* depicts its spatial distribution). Going briefly over the sources consulted by Berezkin to map the presence of this motif across oral traditions, it becomes clear that each appearance of $k27n1$ is associated with a different tale/myth. Depending on the context, the "prominent figure" in the tale may be a king, queen, sultan, emir, tsar, rajah, emperor, boyar, khan, prince, aldar, padishah, official, mayor, chief, chieftain, and kurak. This implies that the use of "king" as a seed for capturing the presence of hierarchy across motifs considers many of the alternative configurations of historical political authority worldwide and not only those strictly related to the notion of kingdoms.

Motif $k27n1$ in the Telugu Oral Tradition For the Telugu (15th largest language group worldwide), Berezkin consulted a range of sources to reconstruct its oral tradition. Those include nine different authors from 3 languages typical of the societies in his catalogue. Altogether, 56 motifs are in the Telugu's oral tradition. The particular tale where motif $k27n1$ is detected comes from the book "Folktales from India," written by Mosur Venkataswami and published by the Diocesan Press publishing house in Madras in 1923. Specifically, in the tale: "The adventures of Ratnalpolchetty," a prince gets to perform a series of challenging tasks, including finding the Parijata flower and rare gems upon request of the ruler's daughter in order to convince her that he is worthy. This type of story is very much in the spirit of what Campbell in his famous 1968 work "The Hero with a Thousand Faces" describes as the hero's journey, an archetypical narrative that involves a hero who goes on an adventure wins decisive victories, and returns home transformed.

Other works by Mosur Venkataswami include "Five hundred instructive Tamil proverbs with their English equivalents" and "Uttama neeti". Panels *A* and *B* in Appendix Table 11 describe the underlying sources consulted by Berezkin for the Telugu oral tradition and the resulting motifs, respectively.

3 Notes on Folklore and the Natural Environment

Here, we discuss three aspects of the relationship between folklore and the physical environment. First, the relationship between the local geography and the motifs mentioning this specific geographic trait, besides a direct effect of the group-specific geography, may also reflect the borrowing

of narratives from neighboring groups (whose incidence in turn captures both the borrowing from other groups and the group's own geography). To get at this, in Appendix Table 3 Panel B, we build on the specifications reported in Table I and add as a control the intensity of motifs mentioning the particular geographic feature in the nearest group. Consistent with the idea that narratives travel across space, the intensity of motifs featuring a given geographic attribute in the nearest group's oral tradition (weakly) increases the motif frequency of similar features in the group itself. Nevertheless, the impact of the group's own geography on the share of the geography-related images in the oral tradition remains robust across specifications. Hence, both local factors and interactions with neighboring societies seem to drive the distribution of motifs across oral traditions. Understanding the types of stories that travel further than others is a topic of future research.

Second, clustering standard errors within ethnolinguistic families accounts for the fact that stories are more likely to travel within broadly similar language groups. For example, in Berezkin's dataset, the probability that the nearest group belongs to the same linguistic family is 57%. We now report in Appendix Table 3 Panel C standard errors where we correct for spatial correlation. This requires a cutoff distance beyond which the correlation between the error terms of any two observations is assumed to be zero. The median distance of a group in Berezkin to other groups within the same continent is $3,330$ kilometers. We use this as the threshold. Experimenting with higher cutoff values like $4k$, $5k$, and $6k$, we find consistently smaller standard errors.

Third, which geographic features to explore is an open-ended question. Considering that groups in similar ecologies exhibit limited geographic heterogeneity, we focused on environmental traits displaying both between and within-country variation. We also experimented with features that are spatially concentrated. For example, one can show that the link between proximity to the coast and the intensity of sea-related motifs is strong when comparing oral traditions across continents. This association, nevertheless, is attenuated for groups within the same country. This is perhaps not surprising as differences in the intensity of sea-related motifs in the folklore of groups within landlocked (or coastal) countries may not be linked to their relative proximity to the coast.

4 Machine Learning and Folklore

4.1 A Methodological Digression

The set of values one may extract from folklore is potentially very large. One way to go about it is to let the data "tell" whether there are particular keywords/concepts (if any) that stand out in societies that are, for example, more risk-taking than others. This econometrically is a problem of feature selection. In the context of linear regression models with many predictor

variables, the general solution is to reduce variance at the cost of introducing bias to improve the model's predictive performance. There are various shrinkage techniques, including Lasso, Ridge, and Elastic Net.[3]

The first step is to decide on the underlying set of keywords (the features upon which to implement shrinkage methods). There are many dictionaries one could use, and more are constructed and current ones regularly revised. This situation of an ever-evolving set of dictionaries makes the decision of which one(s) to choose difficult. To avoid this problem, along with entries from the Linguistic Inquiry and Word Count (LIWC) and Moral Foundations Dictionary (MFD), we use ConceptNet to obtain the related words for the top 5k-10k most common words.[4] The downside of incorporating thousands of entries from ConceptNet is that these concepts encompass much more than moral values, psychological states, and personal concerns making the interpretation of the selected features more difficult.

Caveats The use of shrinkage methods and, in particular, that of Lasso is appropriate for a situation like ours where predictors exceed the number of observations and can be useful *as long as* values and lessons are associated with recurrent, common-across-cultures keywords. We recognize that this assumption may apply neither universally across societies nor uniformly across values. For example, different groups may use different keywords and images in their oral tradition to instill the same attitude. Moreover, despite the superior predictive performance of machine learning methods, which exact variables are chosen depends on the specific sample and conditioning set. This problem of uniform model selection consistency, discussed by Mullainathan and Spiess (2017), is further exacerbated in a world of correlated independent variables, which is naturally the case for our concepts. Finally, a word-counting exercise is not fit to capture complex plots like encouraging or discouraging a given behavior/action. For all these reasons, our Lasso regressions are exploratory, meant to provide some insight into the concepts associated with different norms. In Section 5.1, we illustrate how human classifications may overcome many of these shortcomings.

4.1 Machine Learning, Concepts, and Folklore

We construct the share of motifs in each country's oral tradition, that mention a given keyword (along with the related terms). In doing so, we obtain thousands of concept intensities in a coun-

[3] These methods do not touch upon the issue of causality. Nevertheless, they provide an informed way to discipline the choice of potential predictors.

[4] The LIWC and MFD dictionaries have been used extensively by social scientists. The LIWC, for example, offers a list of words related to psychologically meaningful categories, see Pennebaker, Booth, and Francis (2007). The MFD created by Graham, Haidt, and Nosek (2009) aims at capturing the five foundations of morality according to the Moral Foundation Theory developed by Haidt and Joseph (2004) and Graham, Haidt, Koleva, Motyl, Iyer, Wojcik, and Ditto (2013).

try's folklore. To explore whether any keywords predict contemporary values, we take advantage of recent developments in estimating structural parameters in linear, sparse, high-dimensional models with many controls, which is our case, and implement the rigorous Lasso, introduced by Belloni, Chen, Chernozhukov, and Hansen (2012). The latter is a data-driven method of choosing the penalty level that delivers asymptotically valid estimates under non-Gaussian, heteroscedastic disturbances.[5]

We use all variables present in the GPS dataset and supplement it with the commonly-used trust variable from the World and European Values Surveys (WVS/EVS). In Appendix Table 6, we report the variables selected by the Rigorous Lasso. For completeness, we report the concepts picked among increasingly larger sets of keywords. Specifically, we add to the LIWC and MFD categories the frequencies of the top $5,000$ most common words (along with their related words) up to the $10,000$ most frequent ones. We take the log of the share of each concept adding 0.01. We drop the NLTK's list of English common stop-words like "a," "the," etc. We also drop rare terms, i.e., those that in the median country appear in less than 1% of the motifs. We do this to avoid selecting concepts that appear in very few motifs worldwide.[6] Finally, we absorb throughout the baseline controls and show results exploiting between and within-continent variation.

Risk-Taking Let's start with risk-taking. Lasso consistently selects the concept "poke" as a positive predictor. The same applies to the keywords "competition" and "competitive" in three specifications that exploit within-continent variation. There are 113 motifs tagged by words related to "poke." These words are: dig, thrust, pierce, stick, prick, stab, punch, puncture, push, shove, stab. Here is an example (motif $l67$): "Having dug an underground passage to a lying monstrous ungulate, a small animal nibbles wool from the place on the skin where the heart beats; the hero sticks a spear or arrow into this place." There are 32 "competitive"-tagged motifs, and 41 related to "competitions". Here is a competition-related motif ($m185a$) present in 79 societies: "Birds, animals or fish compete as about who is the fastest or can fly higher than others. A weak one imperceptible sticks to the body of the fastest or strongest and wins." These three concepts are all intuitively linked to the attitude of interest, suggesting that societies whose oral tradition features a higher frequency of episodes related to challenges, contests, and physical altercations

[5] There are different approaches for selecting the penalization ('tuning') parameters including those based on information criteria and K-fold cross-validation methods. We implement the Rigorous Lasso for 2 reasons. First, this approach places a high priority on controlling overfitting, thus often producing more parsimonious models than those based on cross-validation methods (see Ahrens, Hansen, and Schaffer (2019), Frandsen (2019), and Chetverikov, Liao, and Chernozhukov (2019)). Second, our attempts to produce results based on cross-validation methods would not converge in finite time. Ahrens, Hansen, and Schaffer (2019) also note that K-fold cross-validation is the slowest method and estimate the Robust Lasso routine to be $1,200$ times faster than the K-fold validation method in their example.

[6] These restrictions shrink the top 5k, 6k, 7k, 8k, 9k, and 10k concepts to $2,820$, $3,239$, $3,672$, $4,071$, $4,464$, and $4,848$, respectively.

exhibit higher levels of risk tolerance today. In Section 5.1.2, we show how we use MTurks to refine these results.

Altruism We repeat the same procedure by looking at the extent of altruism in the sample of 76 countries. In most specifications, Rigorous Lasso selects no predictors. This suggests that single-concept frequency counting exercises may often not be adequate in uncovering robust predictors. Nevertheless, in three specifications, the concept "abusive" is negatively selected. The 41 motifs tagged by words related to "abusive" are: abuse, disrespectful, inappropriate, maltreat, obscene, and sexually. Many of these motifs ($m29a - m29j$) describe different animals engaging in obscene or antisocial behavior. Countries whose folklore has many images of abusive behavior score low in altruism.

Patience Unlike altruism and risk-taking, the Lasso-selected concepts for patience are not intuitively linked to the value of interest. This suggests that complex and abstract concepts like delaying immediate gratification may not be easily captured by counting keywords. It also highlights the limits of shrinkage methods since selecting among covariates to maximize predictive power may generate models that are less straightforward to interpret.

Trust in the GPS and in the WVS For trust, we look at two country-wide proxies, the one from the GPS and the other one from the WVS. When we focus at the WVS measure, Rigorous Lasso picks several concepts. Some can be loosely linked to trusting behavior, including the concepts of "privacy," and "oblige." The most common word related to "privacy" is secret whereas the term tagged more frequently with "oblige" is fulfill. Fulfill-tagged motifs often portray a character that is assigned to complete a task and does so. For example, motif $m182a1$ reads: "A man catches several wild animals but lets them go after they promise to bring him something valuable or to help him. The animals fulfill their promise." Motifs mentioning secrets often describe characters that get punished as a result of not keeping them. For example, motif $m154$ present in 102 societies reads: "A man obtains knowledge of animal languages but if he reveals the secret, he must die. Once he hears animals talking and laughs. His wife thinks that he laughs at her or at her mother. The man is ready to open his secret and either does it and dies or hears how animals (usually a cock) blame him for being so foolish. So he keeps his secret." Several other concepts, nevertheless, are also selected (both negative and positive). These include march, monthly, sting, weekly, and spiritually. These concepts a priori are not easy to link to trusting behavior (either positively or negatively).

When we look at the concepts picked by Lasso using the GPS measure of trust, the keywords selected are rather different and also hard to interpret. This can be explained by the difference

in the sample size and because the correlation between the two proxies (for the sample of 63 countries for which they overlap) is 0.49. This example stresses the sensitivity of machine learning techniques to pinpointing robust predictors for variables that are moderately correlated despite measuring the same outcome (in this case, trust across countries).

Positive & Negative Reciprocity Finally, for the two variables measuring reciprocity in the GPS dataset across countries, Rigorous Lasso selects no keywords for negative reciprocity and terms like "blind," "fleet," "stance," and "logging" for the positive reciprocity.

Summary Overall, the results in Appendix Table 6 paint a mixed picture. For some values, like altruism and risk-taking shrinkage methods select concepts that are intuitively linked to the attitude of interest. Nevertheless, for other contemporary traits, either the chosen keywords are not relevant, or the choice set is empty. Moreover, the instability of the concepts selected, even for alternative measures of the same underlying attitude, suggests that the benefits of shrinkage techniques lie in their exploratory nature. These considerations motivate our preferred method detailed in Section 5, which involves the motifs' classification by multiple individuals.

5 MTurk Details

MTurk is a marketplace for work that requires human intelligence. The Mechanical Turk service offers access to a diverse, on-demand, scalable workforce and gives workers a selection of thousands of tasks to complete whenever it is convenient. See Chan and Holosko (2015) and Buhrmester, Talaifar, and Gosling (2018) for an overview of the extensive use of MTurk in the social sciences. We used the Amazon MTurk interface during January, August, and October of 2020 to classify motifs.[7]

Batch Sizes We started with the classification of the antisocial motifs first and then moved to gender roles. We split the motifs in a given category to smaller batches. Trickster and challenge-related batches contained $26 - 27$ motifs. To speed up the process, we increased the batch sizes for gender classification to $40 - 50$ motifs. After the completion of each batch, we would prevent prior workers from accessing other batches of the same type to avoid the overall classification being influenced by a few MTurks. Hence, the batch size is the upper bound in the number of motifs a single MTurk could categorize. Due to updating times, manual errors, and running multiple surveys at once, a single MTurk may, on occasion, have categorized more than a single batch.

[7] Vafa Behnam and Kush Bavaria were instrumental in the implementation of the MTurk component of the study.

Attention Checks We randomly inserted attention checks throughout the survey batches. This text reads: "In the summer on a snail, in the winter on a bird Moving across the sky, the sun changes the mount animals depending on the season - in the summer it rides on a slow-moving one, in winter - on a fast one. Attention check, please select not antisocial. Either in winter the sun is carried by a young man, and in summer by an old man." In case the attention check was answered incorrectly, all classifications specific to that MTurk were reassigned.

MTurk Qualifications and the Number of Classifications per Motif Initially, we aimed at generating 18 independent classifications per motif, allowing only US-based Master MTurks to perform the tasks. Master MTurks, according to the MTurk website "have consistently demonstrated a high degree of success in performing a wide range of HITs across a large number of Requesters." This assessment is based on statistical models "that analyze all Workers based on several Requester-provided and marketplace data points to make that determination." However, due to the small population of Master MTurk workers, the batches were being completed at a slow pace, so after the 7th batch of antisocial motifs, we allowed non-Master US workers with an accuracy rate of least 95% on previous tasks to participate. At the same time, we adjusted the classifications per motif to 9 instead of 18. For motifs for which both Masters and non-Master workers classified, there does not seem to be a statistically significant difference in the responses. On average, trickster motifs were classified by 15 different MTurks and challenge and gender motifs by 9 MTurks.

Payment We started paying the Masters MTurks 13 cents for a HIT (HIT in the language of MTurk is a task; in our case, this is a classification of a single motif). Once we relaxed the Master's qualification and realized that the median time for completing a HIT was about 25 seconds, we adjusted the payment to 8 cents per HIT (in case of motifs that required classification of both genders, we offered 20 cents per HIT). This payment rate amounts to an hourly rate of 11.5 USD.[8]

[8] Adams et al. (2018) analyzed 3.8 million tasks completed by 2,676 MTurks and estimated hourly earnings of $2 per hour on average. Only 4% of workers earned more than the federal minimum wage of $7.25/hour.

References

Abello, J., P. Broadwell, and T. Tangherlini (2012): "Computational Folkloristics," *Communications of the ACM*, 55(7), 60–70.

Adams, A., K. Hara, K. Milland, S. Savage, C. Callison-Burch, and J. Bigham (2018): "A Data-Driven Analysis of Workers Earnings on Amazon Mechanical Turk," Association for Computing Machinery.

Ahrens, A., C. B. Hansen, and M. E. Schaffer (2019): "lassopack: Model Selection and Prediction with Regularized Regression in Stata," IZA DP No. 12081.

Bascom, W. R. (1953): "Folklore and Anthropology," *Journal of American Folklore*, 66(262), 283–290.

Belloni, A., D. Chen, V. Chernozhukov, and C. Hansen (2012): "Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain," *Econometrica*, 80(6), 2369–2429.

Benedict, R. (1934): *Patterns of Culture*. Houghton Mifflin Harcourt, New York.

Buhrmester, M. D., S. Talaifar, and S. D. Gosling (2018): "An Evaluation of Amazon Mechanical Turk, Its Rapid Rise, and Its Effective Use," *Perspectives on Psychological Science*, 13(2), 149–154.

Campbell, J. (1968): *The Hero with a Thousand Faces*. Princeton University Press, Novato, California.

Chan, C., and M. Holosko (2015): "An Overview of the Use of Mechanical Turk in Behavioral Sciences: Implications for Social Work," *Research on Social Work Practice*.

Chetverikov, D., Z. Liao, and V. Chernozhukov (2019): "On Cross-Validated Lasso," arXiv preprint arXiv:1605.02214.

Dorson, R. M. (1982): *Folklore and Folklife: An Introduction*. University of Chicago Press.

Dorst, J. D. (2016): "Folklore's Cybernetic Imaginary, or, Unpacking the Obvious," *Journal of American Folklore*, 129(512), 127–145.

Dundes, A. (1997): "The Motif-Index and the Tale Type Index: A Critique," *Journal of Folklore Research*, 34(3), 195–202.

Fine, G. A. (1979): "Folklore Diffusion Through Interactive Social Networks: Conduits in a Preadolescent Community," *New York Folklore Society*, 5(1-2).

Frandsen, J. A. B. (2019): "Mashine Labor," NBER Working Paper 26584.

Graham, J., J. Haidt, S. Koleva, M. Motyl, R. Iyer, S. P. Wojcik, and P. H. Ditto (2013): "Moral foundations theory: The pragmatic validity of moral pluralism," in *Advances in experimental social psychology*, vol. 47, pp. 55–130. Elsevier.

Graham, J., J. Haidt, and B. A. Nosek (2009): "Liberals and conservatives rely on different sets of moral foundations.," *Journal of personality and social psychology*, 96(5), 1029.

Green, T. A. (1997): *Folklore: An Encyclopedia of Beliefs, Customs, Tales, Music, and Art*, vol. 1. ABC-CLIO, Santa Barbara, CA.

Haidt, J., and C. Joseph (2004): "Intuitive ethics: How innately prepared intuitions generate culturally variable virtues," *Daedalus*, 133(4), 55–66.

Honti, J. (1939): "Marchenmorphologie und Marchen Typologie," *Folk-Liv*, 3, 307–318.

Jung, G. C. (1968): "The Archetypes and the Collective Unconscious," in *Collected Works of CG Jung*, vol. 9. Princeton University Press, Princeton, NJ.

Karsdorp, F., M. van der Meulen, T. Meder, and A. van den Bosch (2015): "MOMFER: A Search Engine of Thompson Motif-Index of Folk Literature," *Folklore*, 126(1), 37–52.

Levi-Strauss, C. (1955): "The Structural Study of Myth," *The Journal of American Folklore*, 68(270), 428–444.

Meder, T., F. Karsdorp, D. Nguyen, M. Theune, D. Trieschnigg, I. Everhardus, and C. Muiser (2016): "Automatic Enrichment and Classification of Folktales in the Dutch Folktale Database," *Journal of American Folklore*, 129(511), 78–96.

Mullainathan, S., and J. Spiess (2017): "Machine Learning: An Applied Econometric Approach," *Journal of Economic Perspectives*, 31(2), 87–106.

Pennebaker, J. W., R. J. Booth, and M. E. Francis (2007): "Linguistic inquiry and word count: LIWC [Computer software]," *Austin, TX: liwc. net*.

Thompson, S. (1946): *The Folktale*. The Dryden Press, New York, NY.

——— (1955): *Motif-Index of Folk-Literature: a Classification of Narrative Elements in Folktales, Ballads, Myths, Fables, Medieval Romances, Exempla, Fabliaux, Jest-Books and Local Legends*. Indiana University Press.

Toelken, B. (1996): *Dynamics of Folklore*. University Press of Colorado, Boulder, CO.

Uther, H.-J. (2004): *The Types of International Folktales: a Classification and Bibliography, Based on the System of Antti Aarne and Stith Thompson*. Suomalainen Tiedeakatemia, Helsinki, Finland.

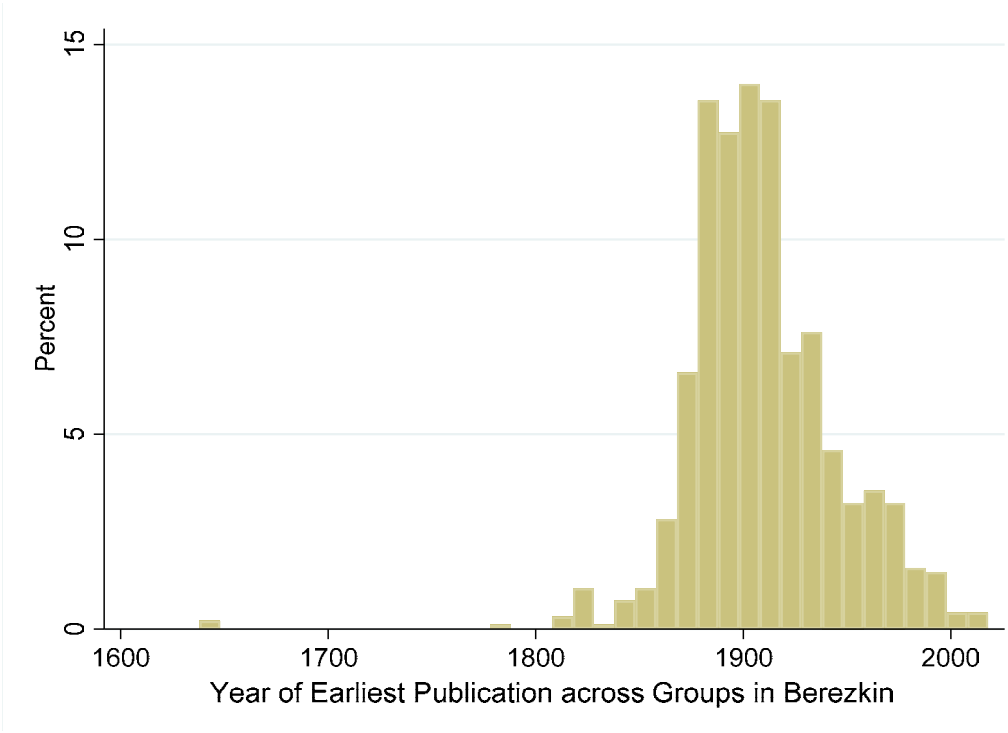

Appendix Figure 1

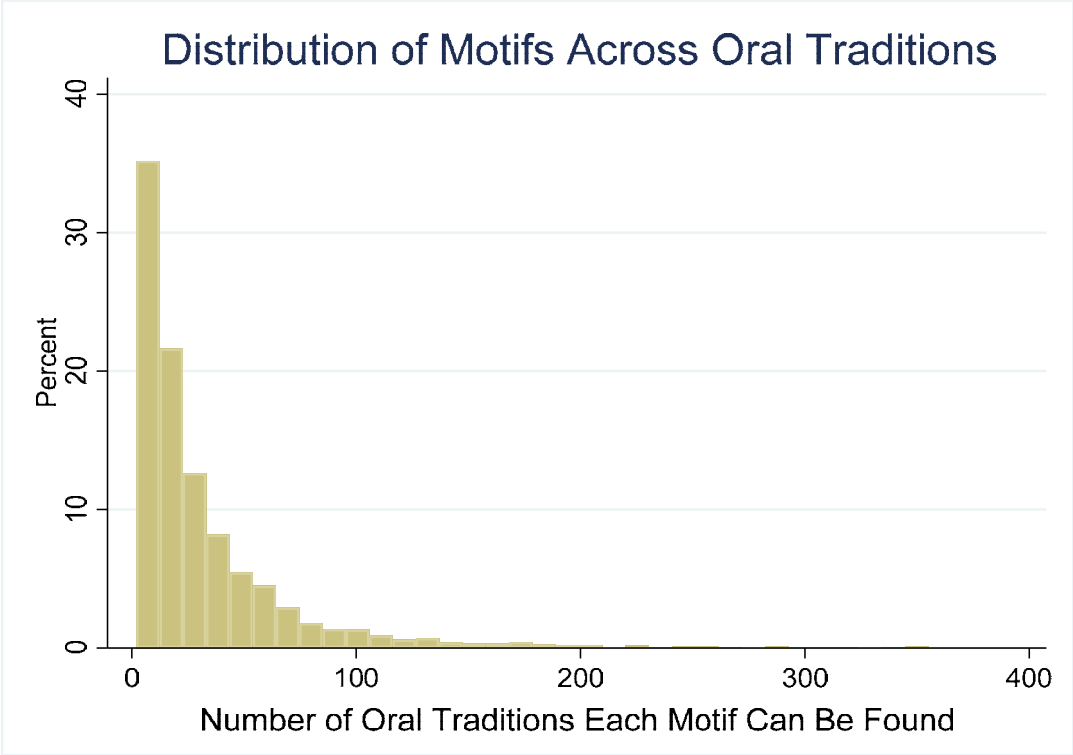

Appendix Figure 2

Appendix Figure 3a, 3b, 3c: Examples of the Most and Least Widespread Motifs in Berezkin's Catalogue

Appendix Figure 3a - Most Common Motif k27n

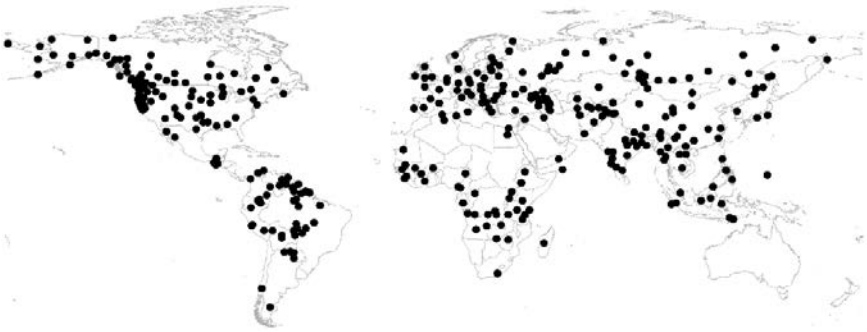

Title: Tasks of the in-laws"; Description "Father or other kinsmen of hero's wife or bride try to kill or test him and/or suggest him difficult tasks"

Appendix Figure 3b - Least Common Motif: b107

Title: "The black oyster"; Description: "During a (world) fire, the oyster burns, so the oyster shells are black". Found in the Oral Tradition of the Flathead people on the East and the Heiltsuk group on the West

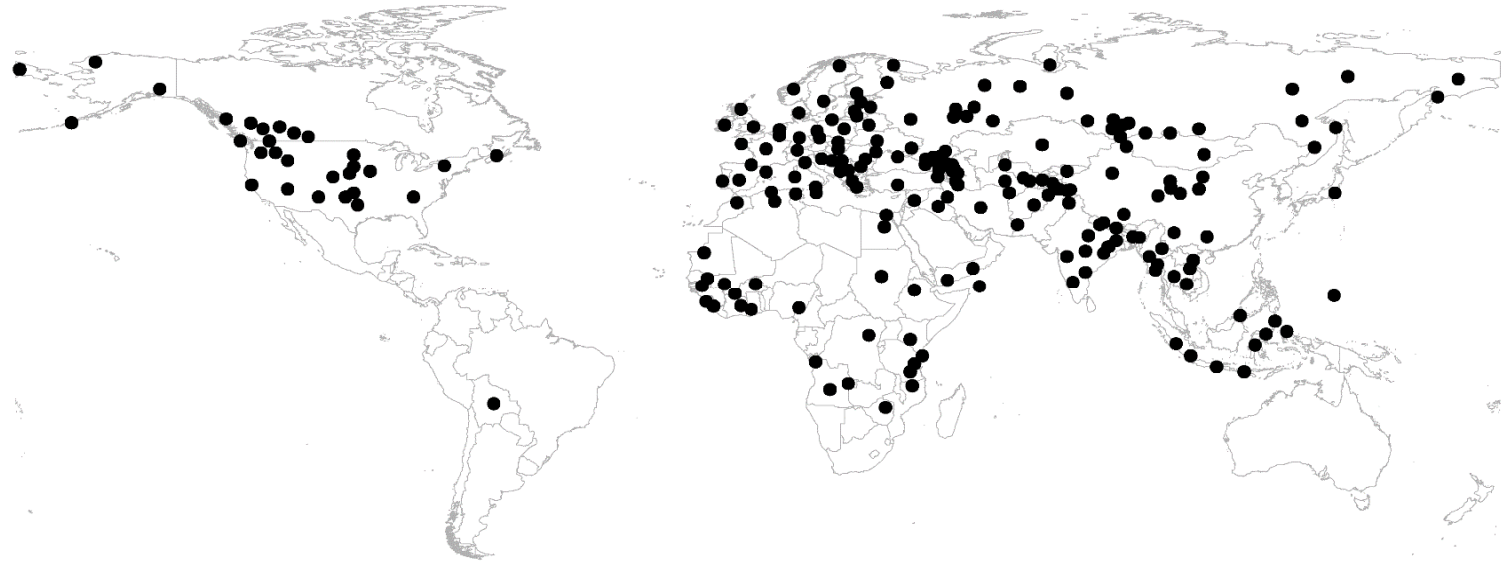

Appendix Figure 3c: The Spatial Distribution of Motif k27n1

Motif Description: “Person who gives difficult tasks to the hero is a prominent figure in social hierarchy, i.e., a head of political unit of community - or higher level and not a mythical being”

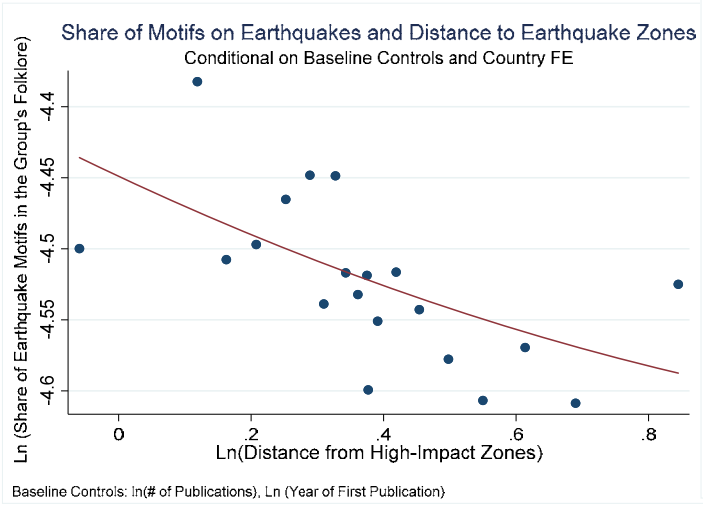

Appendix Figure 4a

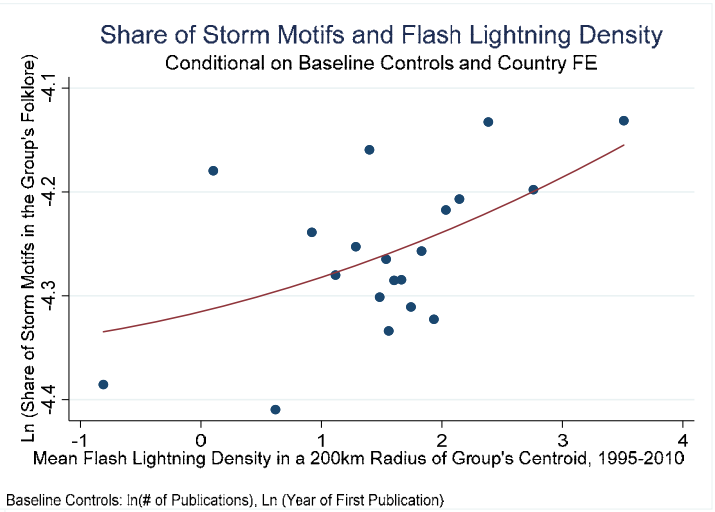

Appendix Figure 4b

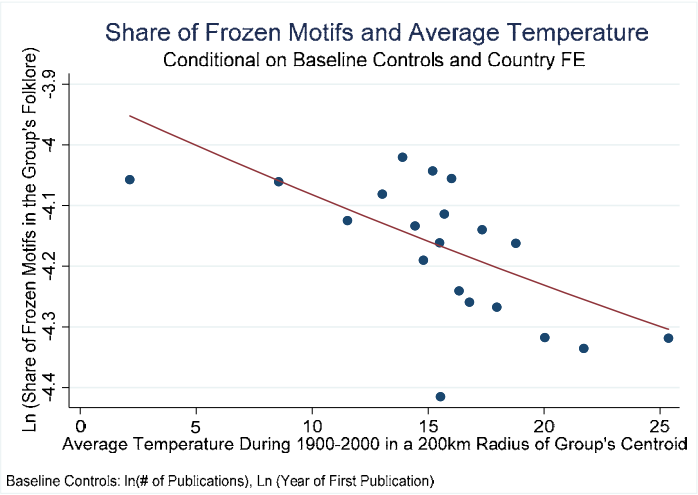

Appendix Figure 4c

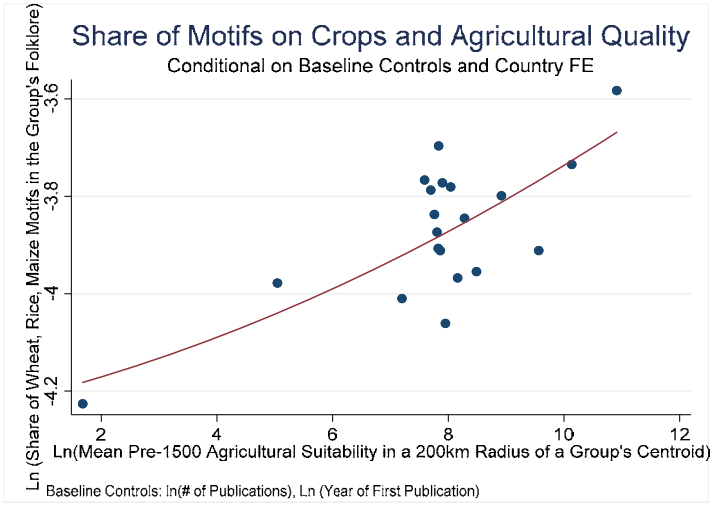

Appendix Figure 4d

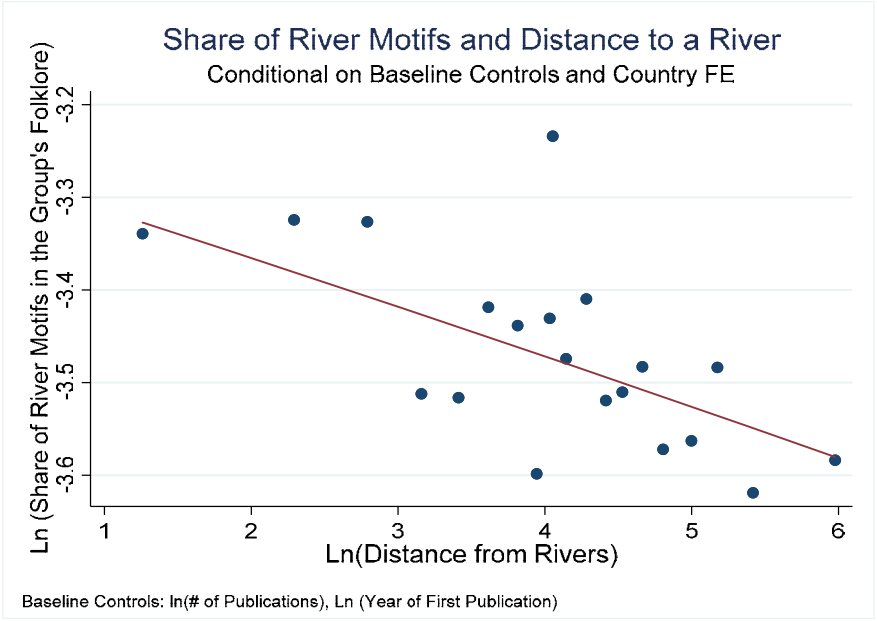

Appendix Figure 4e

Appendix Figure 5a

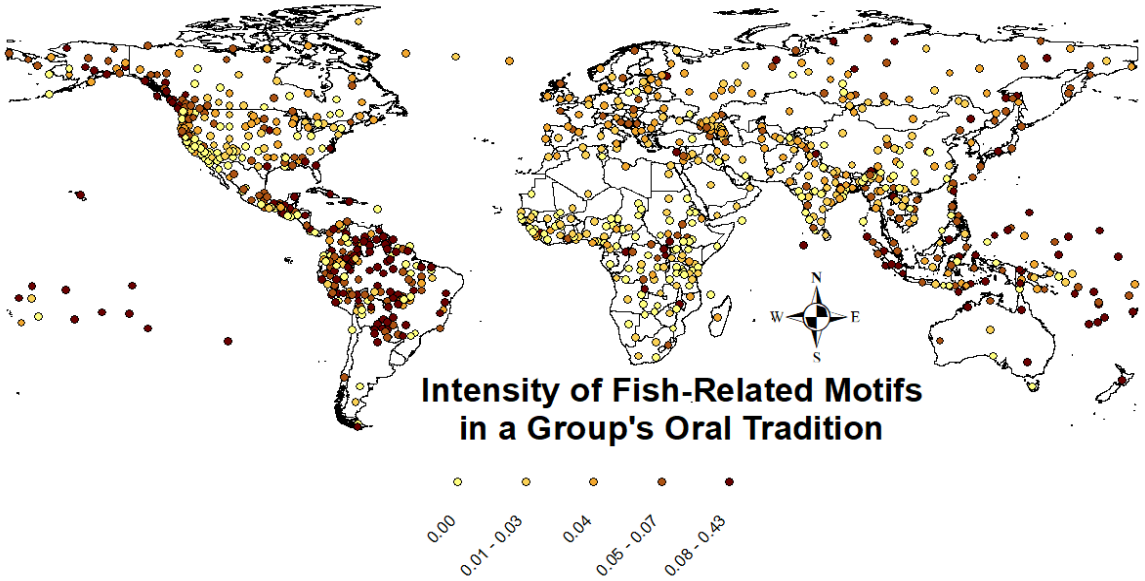

Appendix Figure 5b

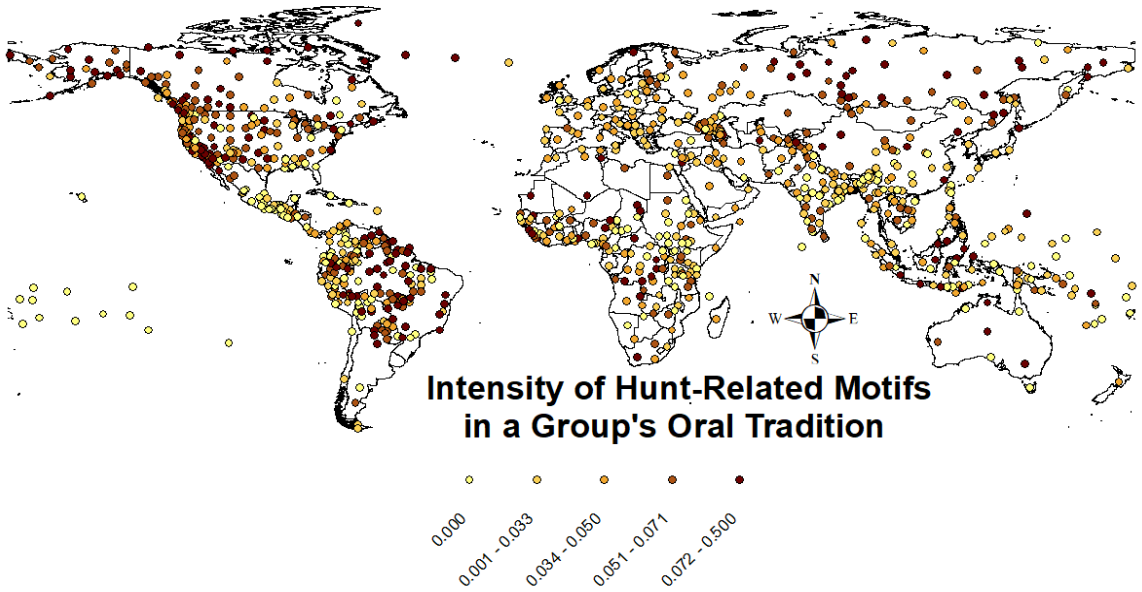

Appendix Figure 5c

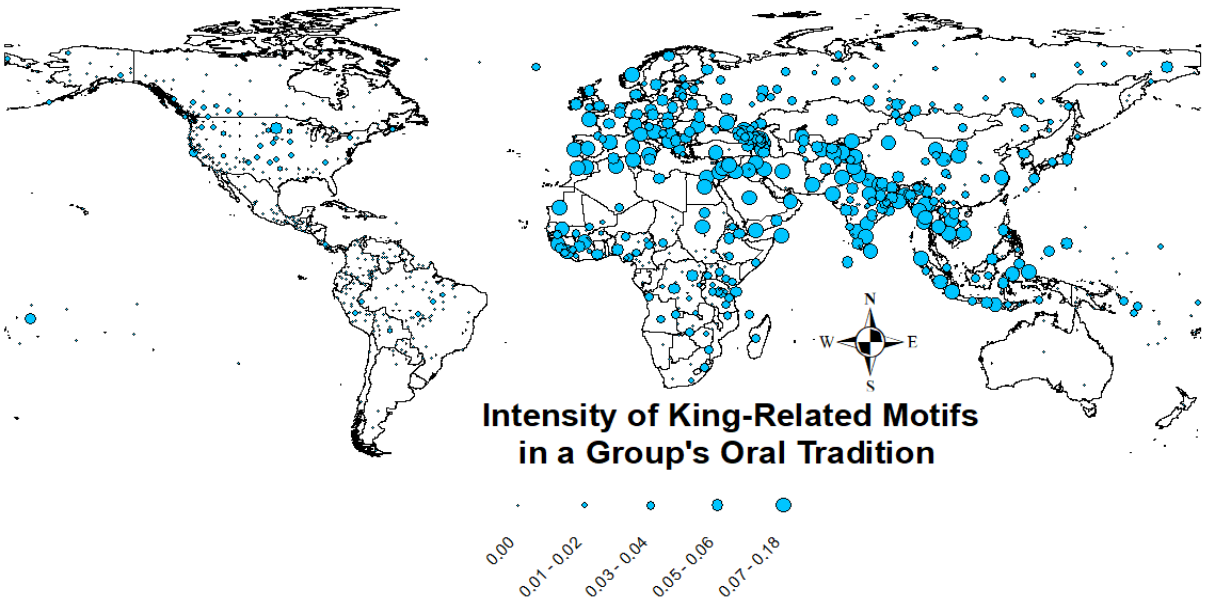

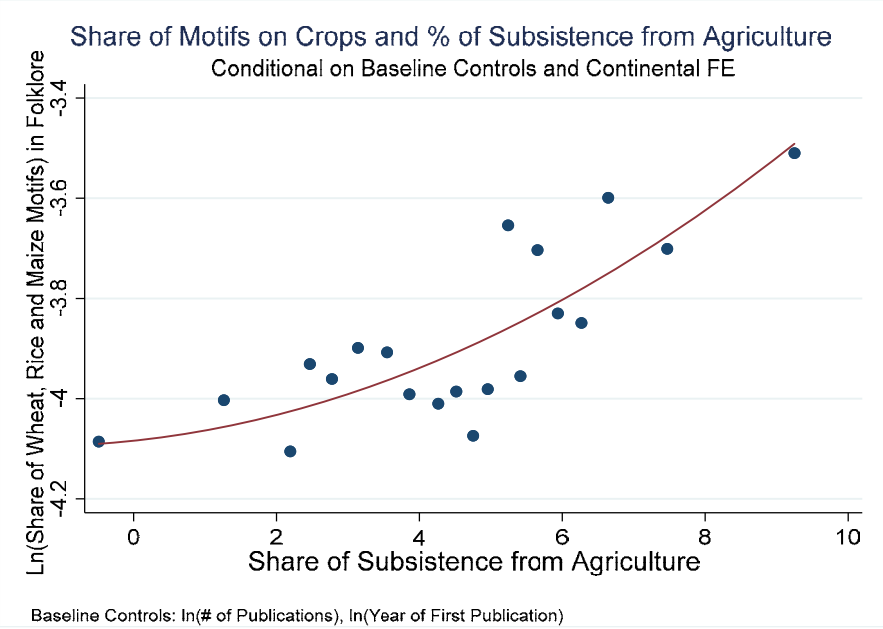

Appendix Figure 6a

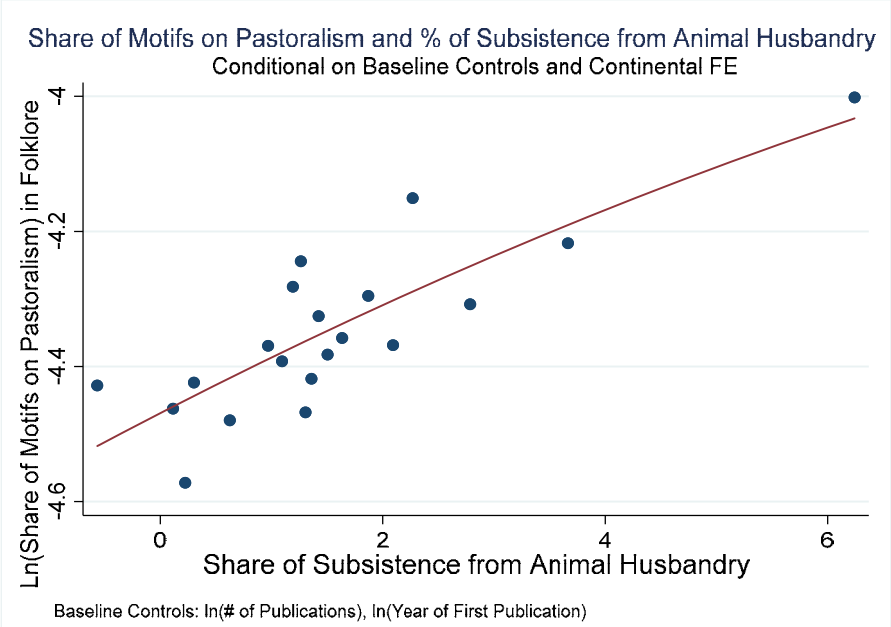

Appendix Figure 6b

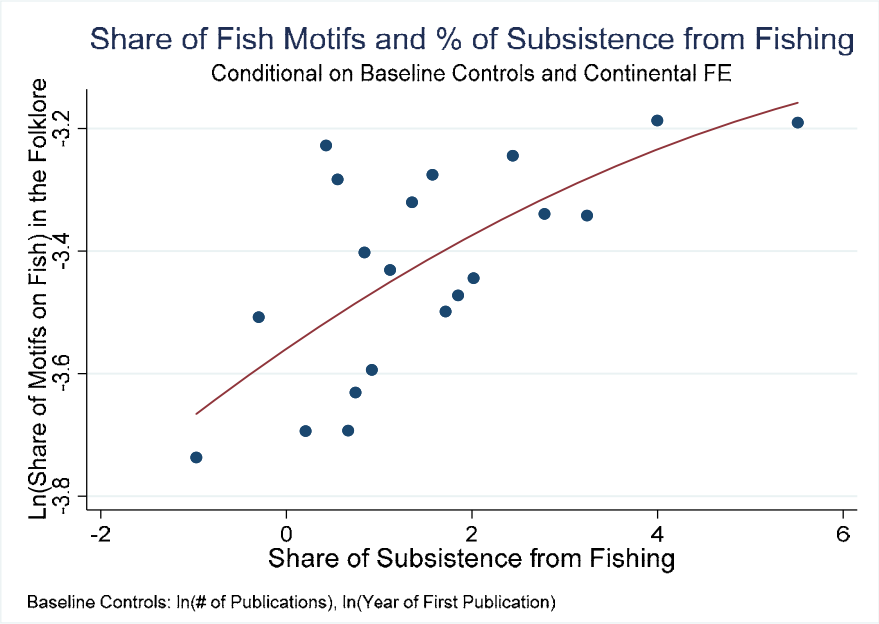

Appendix Figure 6c

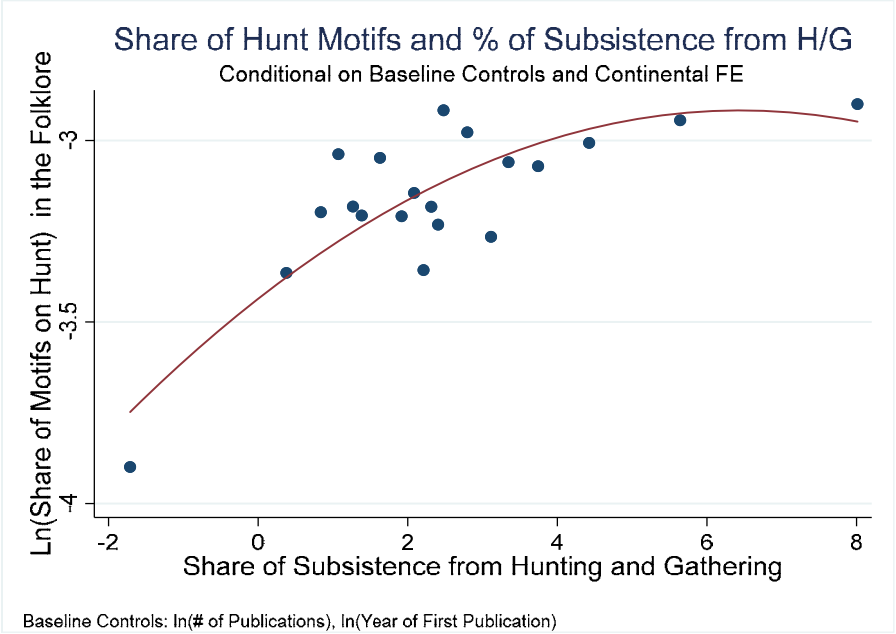

Appendix Figure 6d

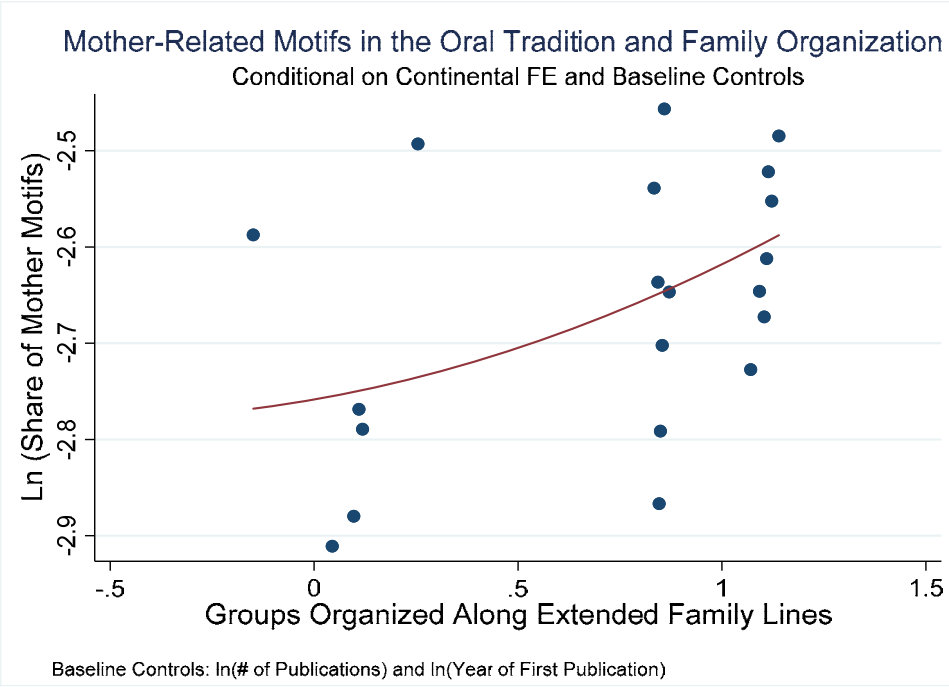

Appendix Figure 7a

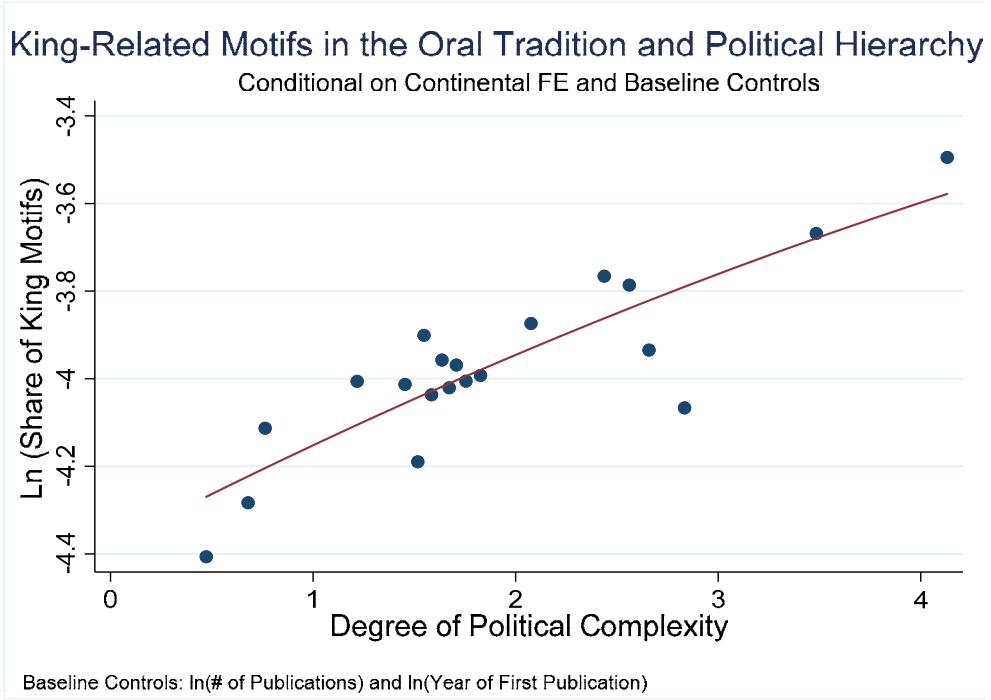

Appendix Figure 7b

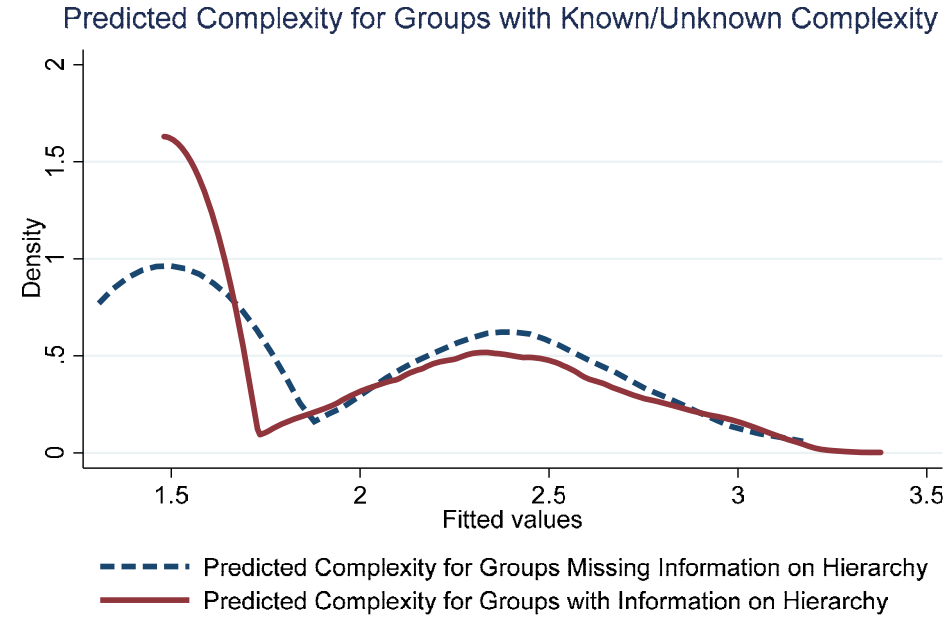

Appendix Figure 7c

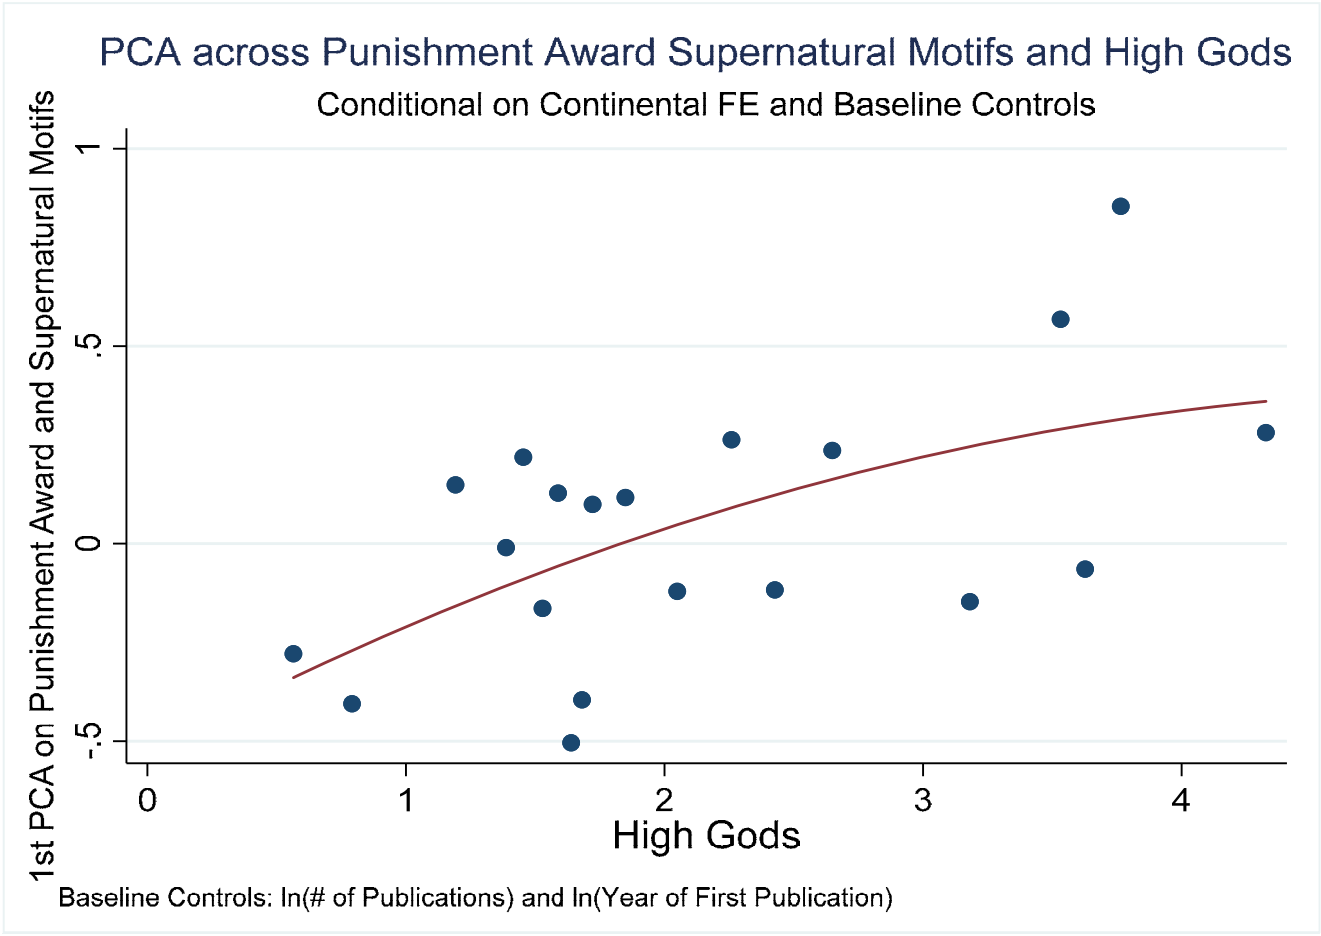

Appendix Figure 8a

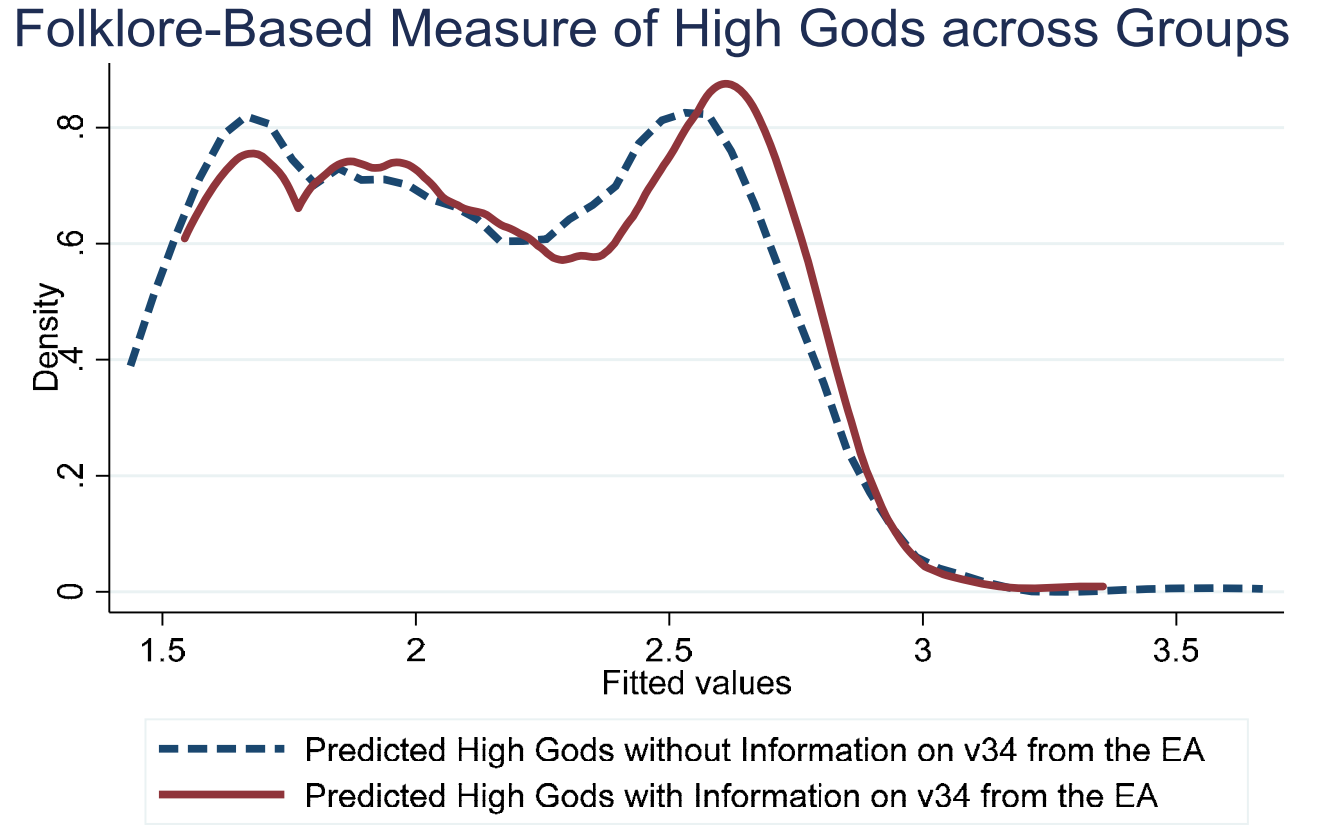

Appendix Figure 8b

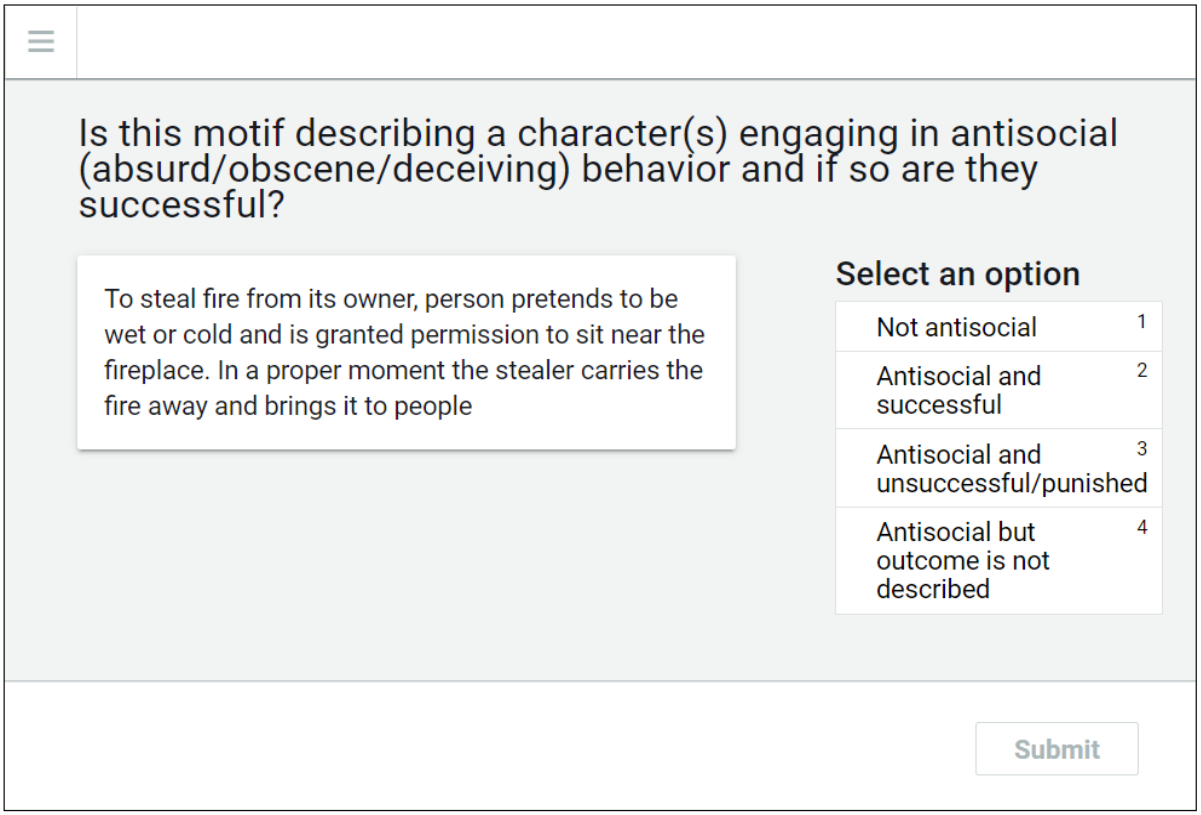

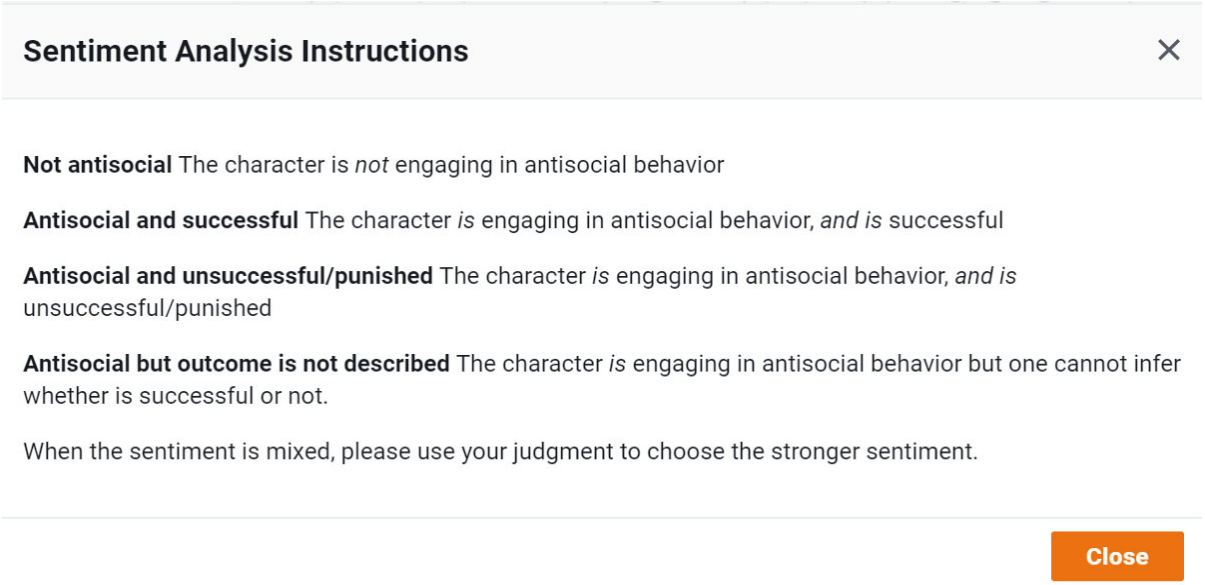

Appendix Figure 9

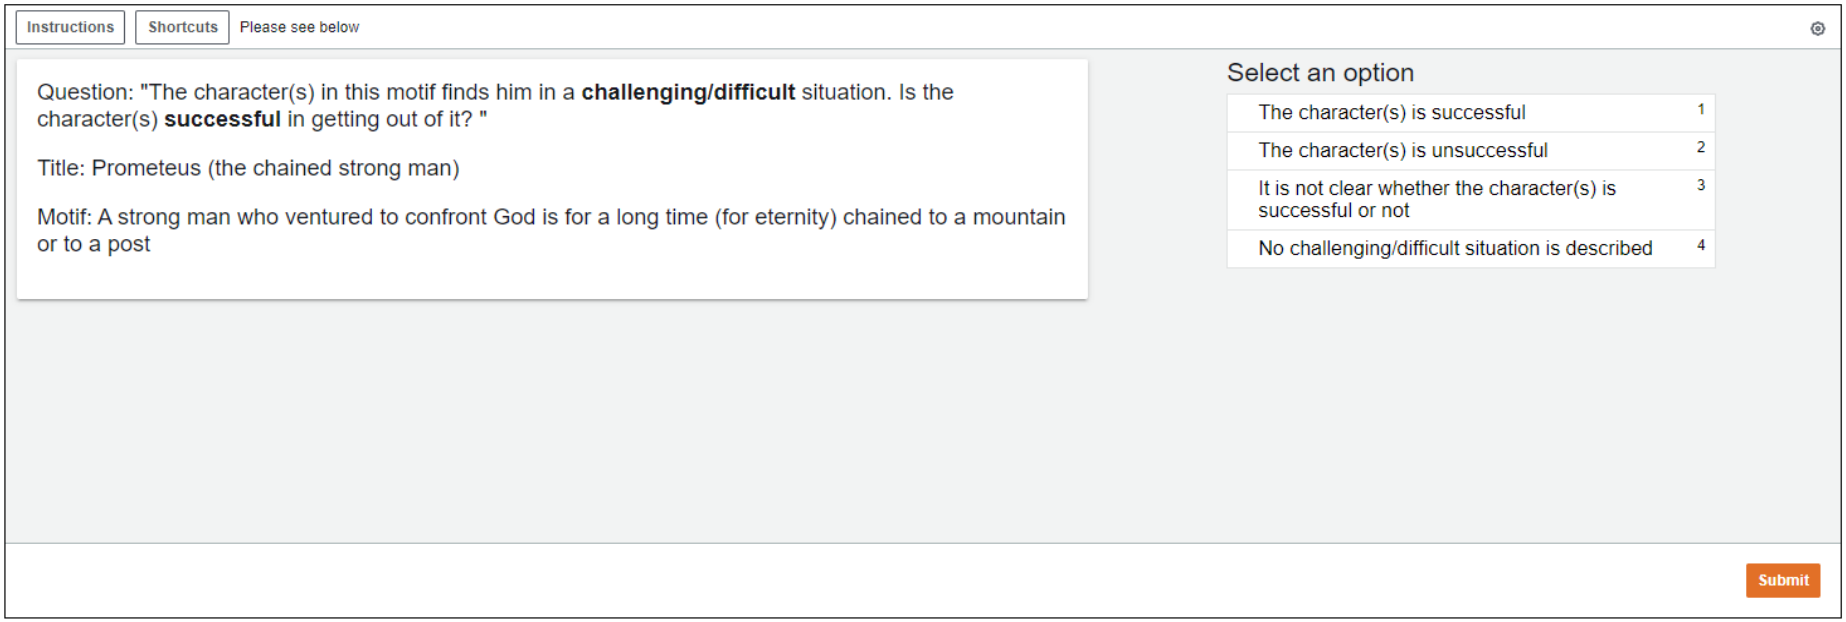

Appendix Figure 10a: Classifying Challenges in a Motif

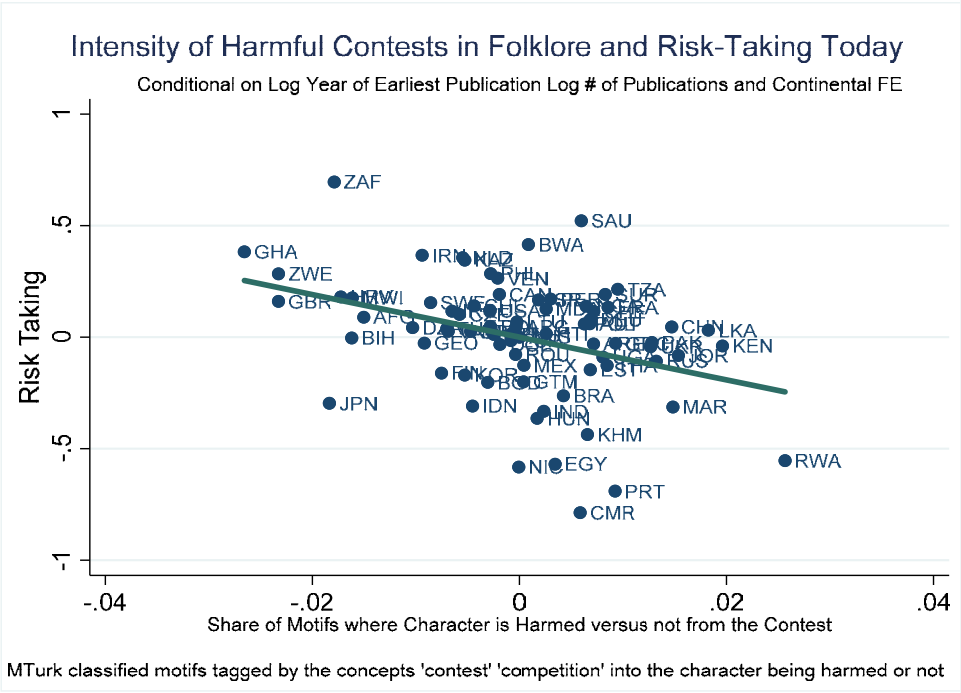

Appendix Figure 10b

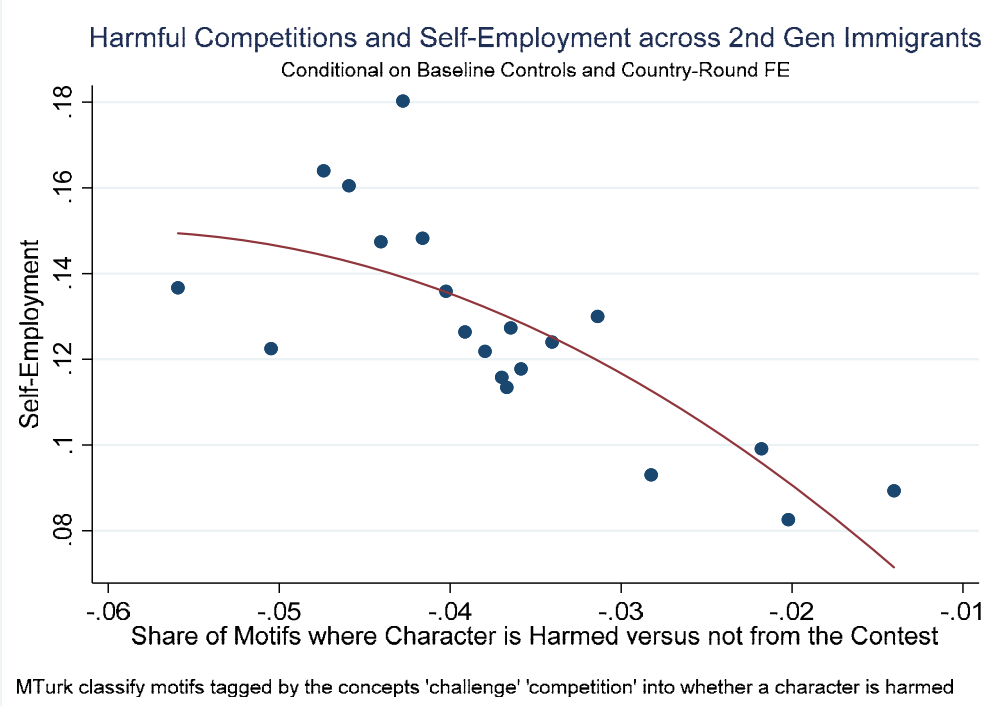

Appendix Figure 10c

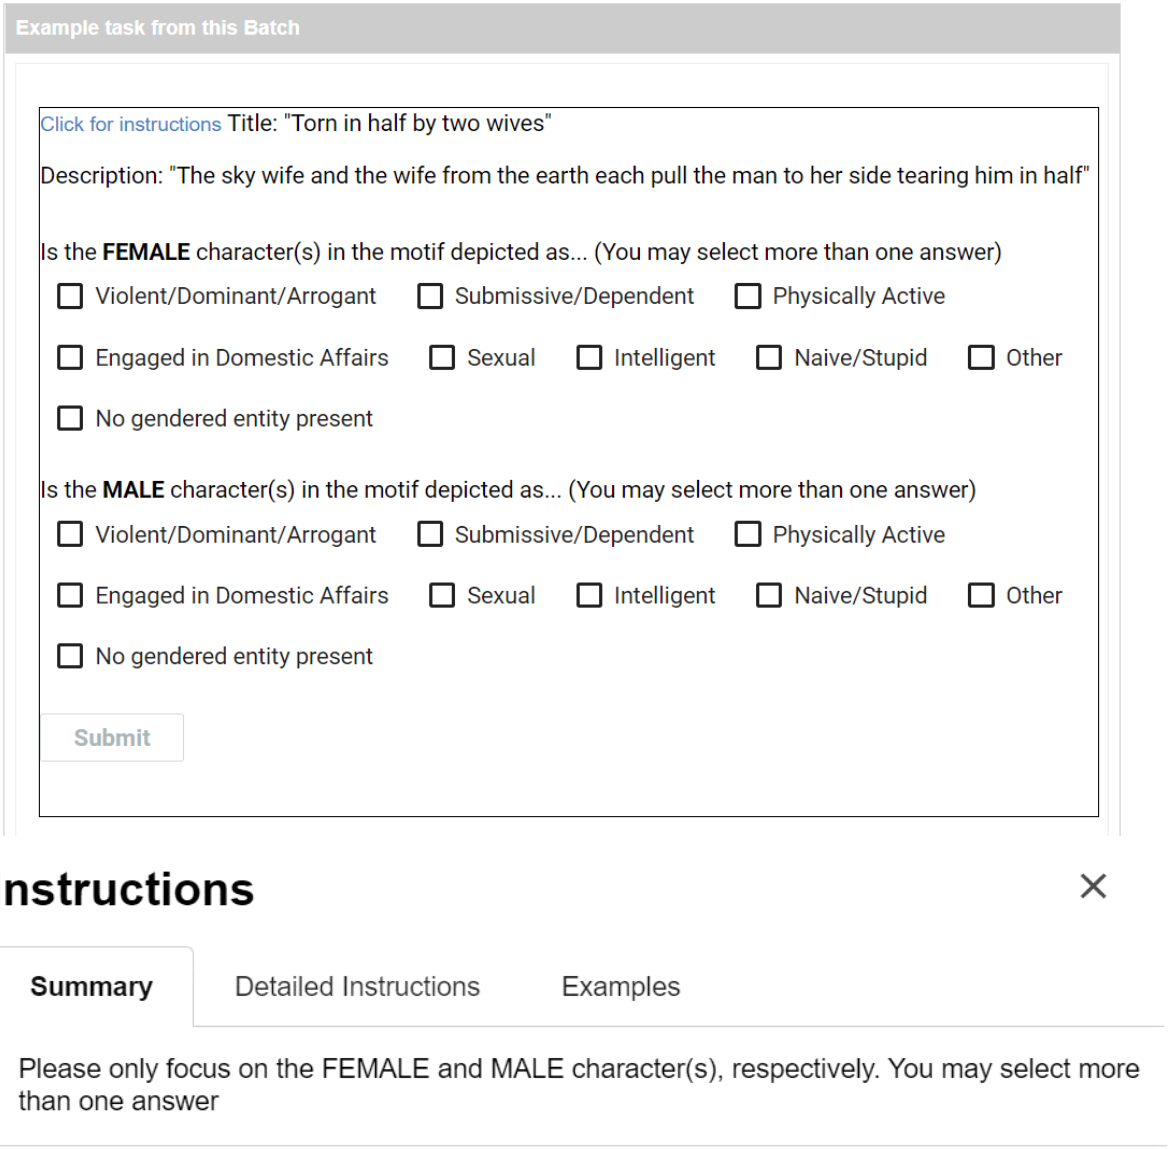

Appendix Figure 11: Classifying Gender Stereotypes in a Motif

Appendix Table 1 - Panel A: Language Breakdown of Publications in Berezkin's Catalogue

Language	# of Publications	Language	# of Publications
English	2527	Belarusian	7
Russian	1844	Scots Gaelic	6
Spanish	558	Czech	4
German	452	Arabic	2
French	389	Galician	2
Portuguese	76	Luxembourgish	2
Estonian	49	Serbian	2
Romanian	42	Ukrainian	2
Hungarian	40	Bulgarian	1
Danish	39	Samoan	1
Dutch	30	Shona	1
Nepali	28	Welsh	1
Catalan	22		
Polish	21		
Latin	18		
Italian	16		
Slovenian	16		
Swedish	16		
Bosnian	13		
Finnish	10		

Appendix Table 1 - Panel B: Top-10 Oral Traditions in Terms of Motifs

Group	# of Motifs
Romanians	454
Germans	458
Greeks (modern)	458
Kazakh	461
Latvians	468
Lithuanians	482
Georgians	488
Bulgaria	565
Ukrainians	588
Russians	598

Appendix Table 2 - Panel A: Summary Statistics for the Oral Traditions in Berezkin's Catalogue

	mean	p50	sd	min	max	N
# of Motifs	87.23	61.50	88.09	1	598	958
# of Publications	14.57	10.00	13.24	1	104	958
# of Different Languages across Publications	3.48	3.00	2.42	1	14	958
# of Authors	12.25	9.00	11.17	1	84	958
Year of First Publication	1909	1904	36	1638	2015	958
Year of Average Publication	1959	1960	20	1882	2015	958
Share of Motifs on Earthquake	0.0016	0	0.0064	0	0.0667	958
Share of Motifs on Storm	0.0074	0	0.0166	0	0.2500	958
Share of Motifs on Frozen	0.0087	0	0.0163	0	0.2000	958
Share of Motifs on Wheat, Rice and Maize	0.0229	0	0.0453	0	0.7500	958
Share of Motifs on River	0.0303	0	0.0341	0	0.5000	958
Distance to Earthquake Zones of Strength 3 and 4 in 1,000km	0.5756	0.3174	0.7133	0	4.6985	958
Mean Lightning Flash Rate (1995-2010)	9.9472	6.3081	10.3328	0.0721	75.3213	957
Mean Annual Temperature (1901-2000)	15.5640	19.5148	10.7498	-17.4485	29.4494	927
Optimal Agricultural Calories Pre-1500	5927	6254	3891	1	18986	957
Mean Distance to Preindustrial Trade Routes in the Old World in 1,000km	0.6801	0.3841	0.7632	0.0063	4.7404	487
Distance in km to a River (as depicted by Natural Earth)	199	56	520	0	5075	958
Ecological Polarization Calculated Across White's Vegetation Zones	0.4611	0.5156	0.3743	0	1.0000	955
Average Luminosity in 2008	1.8006	0.5953	2.7752	0	20.6023	957
Rate At Which Antisocial Behavior is Punished versus Not in the Group's Folklore	-0.0272	-0.0179	0.0664	-1	0.2500	958
Male Dominance Bias in Group's Oral Tradition	0.1429	0.1422	0.1212	-0.5000	1.0000	958

Notes: See the main text for variable definitions. Zonal statistics are calculated in a 200km radius of group's centroid. The latter is recorded in Berezkin's catalogue.

Appendix Table 2 - Panel B: Summary Statistics of Oral Traditions and Ethnographic Traits (in the Sample of the EA)

	mean	p50	sd	min	max	N
# of Motifs	84.57	65.00	73.78	2.00	598.00	1,245
# of Publications	16.30	14.00	11.48	1.00	104.00	1,245
# of Different Languages across Publications	3.80	4.00	2.16	1.00	14.00	1,245
# of Authors	13.77	12.00	9.66	1.00	84.00	1,245
Year of First Publication	1906	1905	29	1638	2015	1245
Year of Average Publication	1957	1957	16	1882	2015	1245
% of Subsistence from Hunting and Gathering, v1+v2 in EA	2.46	1.00	2.73	0.00	10.00	1,264
% of Subsistence from Fishing, v3 in EA	1.53	1.00	1.70	0.00	9.00	1,264
% of Subsistence from Animal Husbandry, v4 in EA	1.56	1.00	1.80	0.00	9.00	1,264
% of Subsistence from Agriculture, v5 in EA	4.45	5.00	2.71	0.00	9.00	1,264
Levels of Jurisdictional Hierachy Beyond the Village level, v33 in EA	1.91	2.00	1.05	1.00	5.00	1,129
High Gods, v34 in EA	2.17	2.00	1.17	1.00	4.00	748
Group Organized Along Extended Family Lines, v8 in EA	0.29	0.00	0.45	0.00	1.00	1,237
The Plow is Aboriginal to Prior Contact; v39 in EA	0.12	0.00	0.33	0.00	1.00	1,156
Males Contribute More in Agriculture; v54 in EA	0.32	0.00	0.47	0.00	1.00	720
Share of Motifs on Wheat, Maize, and Rice	0.0215	0.0000	0.0382	0.00	0.5000	1,245
Share of Motifs on Pastoralism	0.0049	0.0000	0.0099	0.00	0.0755	1,245
Share of Motifs on Fish	0.0346	0.0241	0.0403	0.00	0.4286	1,245
Share of Motifs on Hunt	0.0440	0.0411	0.0366	0.00	0.3333	1,245
Share of Motifs on Mother	0.0747	0.0702	0.0466	0.00	0.5000	1,245
Share of Motifs on King	0.0160	0.0000	0.0233	0.00	0.1667	1,245
Share of Motifs on Punishment	0.0291	0.0227	0.0332	0.00	0.3846	1,245
Share of Motifs on Award	0.0168	0.0065	0.0234	0.00	0.3077	1,245
Share of Motifs on Supernatural	0.0220	0.0182	0.0243	0.00	0.2857	1,245
1st PC Across the Share of Motifs on Supernatural, Punishment and Award	0.0000	-0.0270	1.3017	-1.89	4.4531	1,234

Notes: Each group in the Ethnographic Atlas has been assigned a folklore content by matching it to a group in Berezkin catalogue where possible. See the main text for variable definition.

Appendix Table 3 - Panel A: Folklore and Agriculture

	Share of Motifs on					
	Wheat		Rice		Corn	
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Optimal Agricultural Calories Pre-1500)	0.0316***	0.0266***	0.0349***	0.0274***	0.0406***	0.0378***
	(0.0078)	(0.0080)	(0.0063)	(0.0079)	(0.0088)	(0.0100)
Baseline Controls	yes	yes	yes	yes	yes	yes
Continental FE	yes	no	yes	no	yes	no
Country FE	no	yes	no	yes	no	yes
Adjusted R^2	0.156	0.203	0.15	0.191	0.141	0.215
# of Observations	942	942	942	942	942	942

Notes: This table reports OLS estimates. The unit of analysis is a group in Berezkin's catalogue. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). As a dependent variable we use the Ln(0.01 + Share of Concept-Specific Motifs). Standard errors are clustered at the language family level as recorded by Berezkin. ***, **, * denote significance at 1%, 5%, and 10%, respectively.

Appendix Table 3 - Panel B: The Physical Environment in the Folklore of Neighboring Groups

	Share of Motifs on									
	Earthquake		Storm		Frozen		Crops		River	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Ln(Distance to Earthquake Zones in 1,000 km's)	-0.1370***	-0.1576***								
	(0.0321)	(0.0543)								
Share Motifs on Earthquakes in the Nearest Group	0.1355**	0.0709								
	(0.0559)	(0.0885)								
Ln(Mean Lightning Flash Density)			0.0164	0.0381*						
			(0.0170)	(0.0206)						
Share Motifs on Storms in the Nearest Group			0.1899***	0.0972*						
			(0.0417)	(0.0557)						
Mean Yearly Temperature 1900-2000					-0.0159***	-0.0138**				
					(0.0024)	(0.0057)				
Share of Motifs on Frozen in the Nearest Group					0.1914***	0.1201**				
					(0.0416)	(0.0586)				
Ln(Optimal Agricultural Calories Pre-1500)							0.0525***	0.0523***		
							(0.0119)	(0.0142)		
Share Motifs on Crops in the Nearest Group							0.1786***	0.0677		
							(0.0428)	(0.0419)		
Ln(Distance to Rivers in Km)									-0.0466*	-0.0546*
									(0.0236)	(0.0313)
Share Motifs on River in the Nearest Group									0.0427	-0.0321
									(0.0346)	(0.0361)
Baseline Controls	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Continental FE	yes	no	yes	no	yes	no	yes	no	yes	no
Country FE	no	yes	no	yes	no	yes	no	yes	no	yes
Adjusted R^2	0.052	0.06	0.215	0.209	0.22	0.201	0.213	0.246	0.117	0.129
# of Observations	943	943	942	942	913	913	942	942	943	943

Notes: This table reports OLS estimates. The unit of analysis is a group in Berezkin's catalogue. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). As a dependent variable we use the Ln(0.01 + Share of Concept-Specific Motifs). On the RHS, we use the Ln(0.01 + Share of Concept-Specific Motifs in the folklore of the nearest group). Standard errors are clustered at the language family level as recorded by Berezkin. ***, **, * denote significance at 1%, 5%, and 10%, respectively. See the main text for variable definition and Appendix Table 2 - Panel A for summary statistics.

Appendix Table 3 - Panel C: Folklore and the Physical Environment - Adjusting Standard Errors for Spatial Autocorrelation

	Share of Motifs on									
	Earthquake		Storm		Frozen		Crops		River	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Ln(Distance to Earthquake Zones in 1,000 km's)	-0.1559***	-0.1659***								
	(0.0338)	(0.0508)								
	[0.0349]	[0.0402]								
Ln(Mean Lightning Flash Density)			0.0206	0.0413*						
			(0.0195)	(0.0224)						
			[0.0086]	[0.0246]						
Mean Yearly Temperature 1900-2000					-0.0194***	-0.0152**				
					(0.0026)	(0.0060)				
					[0.0030]	[0.0050]				
Ln(Optimal Agricultural Calories Pre-1500)							0.0594***	0.0550***		
							(0.0140)	(0.0144)		
							[0.0173]	[0.0150]		
Ln(Distance to Rivers in Km)									-0.481*	-0.0536*
									(0.0244)	(0.0313)
									[0.0171]	[0.0249]
Baseline Controls	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Continental FE	yes	no	yes	no	yes	no	yes	no	yes	no
Country FE	no	yes	no	yes	no	yes	no	yes	no	yes
Adjusted R^2	0.038	0.057	0.185	0.203	0.186	0.19	0.183	0.244	0.117	0.129
# of Observations	943	943	942	942	913	913	942	942	943	943

Notes: This table reports OLS estimates. The unit of analysis is a group in Berezkin's catalogue. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). As a dependent variable we use the Ln(0.01 + Share of Concept-Specific Motifs). Standard errors in parentheses are clustered at the language family level as recorded by Berezkin. In brackets we report standard errors adjusted for spatial correlation. The cutoff value beyond which the spatial autocorrelation is assumed to be zero is set to 3,330 kilometers. This number reflects the average distance of a given group to other groups within the same continent. ***, **, * denote significance at 1%, 5%, and 10%, respectively, and correspond to the largest standard errors. See the main text for variable definition and Appendix Table 2 - Panel A for summary statistics.

Appendix Table 4: Folklore, Family Structure, and Political Complexity

	Share of Motifs in a Group's Oral Tradition Related to:			
	Mother		King	
	(1)	(2)	(3)	(4)
Extended Family	0.1458***	0.1598**		
	(0.0426)	(0.0606)		
Levels of Jurisdictional Hierarchy			0.1904***	0.1237***
Beyond the Local Level			(0.0506)	(0.0320)
Baseline Controls	yes	yes	yes	yes
Continental FE	yes	no	yes	no
Country FE	no	yes	no	yes
Adjusted R^2	0.165	0.26	0.42	0.566
# of Observations	1207	1207	1105	1105

Notes: This table reports OLS estimates. The unit of analysis is a group in the Ethnographic Atlas which has been matched to an oral tradition in Berezkin's catalogue. Baseline Controls include: ln(# of Publications) and ln(Year of First Publication). As a dependent variable we use the Ln(0.01 + Share of Concept-Specific Motifs). Standard errors are clustered at the language family level as classified in the EA, v98. ***, **, * denote significance at 1%, 5%, and 10%, respectively. See the main text for variable definition and Appendix Table 2 - Panel B for summary statistics.

Appendix Table 5: Folklore & Trade

	Share of Trade-Related Motifs					
	(1)	(2)	(3)	(4)	(5)	(6)
Ln(Mean Distance to Preindustrial	-0.1005***	-0.1047**			-0.0871***	-0.0944*
Trade Routes)	(0.0242)	(0.0469)			(0.0168)	(0.0462)
Ln(Ecological Polarization within			0.0341**	0.0257***	0.0366*	0.0275*
200km Radius of Group's Centroid			(0.0145)	(0.0092)	(0.0205)	(0.0142)
Baseline Controls	yes	yes	yes	yes	yes	yes
Continental FE	yes	no	yes	no	yes	no
Country FE	no	yes	no	yes	no	yes
Adjusted R^2	0.051	0.105	0.156	0.22	0.059	0.107
# of Observations	477	477	940	940	477	477

Notes: This table reports OLS estimates across groups in Berezkin's Catalogue. Columns (1) and (2) and (5) & (6) focus on groups in the Old World and (3) & (4) on all groups. Baseline Controls include: ln(# of Publications) and ln(Year of First Publication). Standard errors are clustered at the language family level as classified by Berezkin. ***, **, * denote significance at 1%, 5%, and 10%, respectively. As a dependent variable we use the Ln(0.01 + Share of Trade-Related Motifs). Mean Distance to Pre-industrial Trade Routes is the average distance to pre-600 AD and 1700AD trade routes. The pre-600 AD routes also include the historical harbors and ports depicted in the Barrington Atlas and the network of Roman roads. See the main text for variable definition and Appendix Table 2 - Panel B for summary statistics.

Appendix Table 6: Lasso-Selected Concepts Predicting Contemporary Attitudes

Dependent Variable	**Risk Taking (GPS)**		**Altruism (GPS)**		**Patience (GPS)**	
Features	Positive Concepts	Negative Concepts	Positive Concepts	Negative Concepts	Positive Concepts	Negative Concepts
Top 5k Concepts;	"poke"					
Top 5k Concepts + Continental FE Absorbed	"poke," "competition," "competitive"					"trading"
Top 6k Concepts;	"poke"					'"derive"
Top 6k Concepts + Continental FE Absorbed	"poke"					"trading"
Top 7k Concepts	"poke"					'"derive"
Top 7k Concepts + Continental FE Absorbed	"poke"			"abusive"		"trading"
Top 8k Concepts	"poke"					'"derive"
Top 8k Concepts + Continental FE Absorbed	"poke"			"abusive"		"trading"
Top 9k Concepts	"poke"					'"derive"
Top 9k Concepts + Continental FE Absorbed	"poke," "competition," "competitive"	"carcass"				"trading"
Top 10k Concepts	"poke"					'"derive"
Top 10k Concepts + Continental FE Absorbed	"poke," "competition," "competitive"	"carcass"		"abusive"		"trading"

Notes: Along the ConceptNet features we always include the entries from the LIWC & MFD dictionaries. Baseline controls, i.e., ln(# of Publications) and ln(Year of First Publication) are always absorbed.

Appendix Table 6 Cont'd: Lasso-Selected Concepts Predicting Contemporary Attitudes

Dependent Variable	**Trust (GPS)**		**Trust (WVS)**	
Features	Positive Concepts	Negative Concepts	Positive Concepts	Negative Concepts
Top 5k Concepts;	"immediately," "resign," "pour"	"account," "unfold"	"satisfy," "privacy"	
Top 5k Concepts + Continental FE Absorbed		"account"	"arrival," "lose"	"monthly," "march," "beef"
Top 6k Concepts;	"resign"	"account", "walking"	"satisfy," "privacy"	"sting"
Top 6k Concepts + Continental FE Absorbed	"fictional," "outright," "scramble," "downstairs"	"account"		"monthly," "march," "beef," "week"
Top 7k Concepts	"pour," "resign"	"account", "footstep"	"oblige," "privacy"	"sting"
Top 7k Concepts + Continental FE Absorbed	"fictional," "outright," "scramble," "downstairs"	"account"		"monthly"
Top 8k Concepts	"pour," "resign"	"account", "footstep"	"oblige," "privacy"	
Top 8k Concepts + Continental FE Absorbed	"fictional," "outright," "scramble," "downstairs"	"account"		"monthly"
Top 9k Concepts	"pour," "resign"	"account", "footstep"	"privacy"	"march"
Top 9k Concepts + Continental FE Absorbed	"fictional," "outright," "scramble," "downstairs"	"account"		"monthly"
Top 10k Concepts	"pour," "resign," "taunt"	"account", "footstep"	"spiritually"	
Top 10k Concepts + Continental FE Absorbed	"fictional," "outright," "scramble," "downstairs"	"account"		"monthly"

Notes: Along the ConceptNet features we always include the entries from the LIWC & MFD dictionaries. Baseline controls, i.e., ln(# of Publications) and ln(Year of First Publication) are always absorbed.

Appendix Table 7

Panel A: Summary Statistics for Motifs on Challenges Across Countries

	mean	p50	sd	min	max	N
Share of Motifs on Challenges and Competitions						
All	0.0597	0.0603	0.0212	0.0000	0.1247	199
where the character is **successful**	0.0289	0.0293	0.0134	0.0000	0.0602	199
where the character is **unsuccessful**	0.0100	0.0096	0.0070	0.0000	0.0419	199
where the character is **neither** successful **nor** unsuccessful	0.0192	0.0175	0.0107	0.0000	0.0667	199
where **no challenges** present	0.0017	0.0002	0.0029	0.0000	0.0222	199

Panel B: Correlation Structure on Challenges Across Countries

	(1)	(2)	(3)	(4)	(5)
Share of Motifs on Challenges and Competitions					
All	1.0000				
where the character is **successful**	0.7352	1.0000			
where the character is **unsuccessful**	0.5138	0.1198	1.0000		
where the character is **neither** successful **nor** unsuccessful	0.7082	0.1815	0.2205	1.0000	
where **no challenges** present	0.0343	-0.2188	-0.0387	0.0937	1.0000

Notes: This Table reports summary statistics and the correlation structure of the share of challenge-related motifs across countries' oral traditions. We classify the challenge-related motifs as follows: (i) any challenges and competitions (i.e., all motifs tagged by the keywords related to "challenge" and "competiton" according to ConceptNet), (ii) challenges where the character is successful, (iii) challenges where the character is unsuccessful, (iv) motifs where a competition takes places but the outcome for the character is unclear, and (v) motifs that no challenge is present. On average 9 MTurks conducted the classification into categories (ii)-(v). We select the modal answer for each motif allowing for multiple modes in case of a tie.

Appendix Table 8: Gender Roles in the Motifs

Number of Motifs in Berezkin's Catalogue of Oral Traditions where the

Male *Character is:*	N	mean	min	max	sd	***Female*** *Character is:*	N	mean	min	max	sd
Engaged in Domestic Affairs	801	0.0724	0	1	0.259	Engaged in Domestic Affairs	748	0.1297	0	1	0.336
Intelligent	801	0.1511	0	1	0.358	Intelligent	748	0.1096	0	1	0.313
Naïve	801	0.1024	0	1	0.303	Naïve	748	0.0816	0	1	0.274
Other	801	0.2185	0	1	0.413	Other	748	0.2072	0	1	0.406
Physically Active	801	0.1511	0	1	0.358	Physically Active	748	0.0842	0	1	0.278
Sexual	801	0.0674	0	1	0.251	Sexual	748	0.1257	0	1	0.332
Submissive	801	0.1336	0	1	0.340	Submissive	748	0.3035	0	1	0.46
Violent/Dominant /Arrogant	801	0.3283	0	1	0.470	Violent / Dominant / Arrogant	748	0.1457	0	1	0.353

Notes: This Table reports summary statistics on how male and female characters are depicted in Berezkin's motif catalogue. The classification of each motif into stereotypes is done by 9 Mturks on average. We select the modal answer for each motif allowing for multiple modes in case of a tie.

Appendix Table 9a: Gender Roles in the Oral Traditions Across the Globe; Summary Statistics

Share of Motifs across Countries' Oral Traditions where the Character Depicted is:

	N	mean	p50	min	max	sd		N	mean	p50	min	max	sd
Male intelligent	199	0.0824	0.0878	0.0000	0.1579	0.0263	Male Physicaly Active	199	0.0452	0.0467	0.0084	0.0923	0.0134
Female intelligent	199	0.0451	0.0466	0.0000	0.0932	0.0182	Female Physicaly Active	199	0.0182	0.0196	0.0000	0.0623	0.0086
Male Bias in Intelligence Motifs	199	0.0373	0.0374	-0.0251	0.1126	0.0184	**Male Bias in Physical Activity**	199	0.0270	0.0284	-0.0148	0.0739	0.0134
Male Sexual	199	0.0189	0.0172	0.0000	0.0723	0.0102	Male Engaged in Domestic Affairs	199	0.0182	0.0180	0.0000	0.0427	0.0090
Female Sexual	199	0.0260	0.0240	0.0000	0.0900	0.0111	Female Engaged in Domestic Affairs	199	0.0431	0.0444	0.0000	0.0942	0.0168
Male Bias in Sexual Motifs	199	-0.0071	-0.0060	-0.0455	0.0372	0.0097	**Male Bias in Domestic Affairs**	199	-0.0250	-0.0272	-0.0753	0.0206	0.0147
Male Submissive	199	0.0526	0.0502	0.0000	0.1180	0.0186	Male Naïve	199	0.0460	0.0495	0.0000	0.0924	0.0195
Female Submissive	199	0.1099	0.1067	0.0000	0.2134	0.0302	Female Naive	199	0.0355	0.0373	0.0000	0.1035	0.0156
Male Bias in Submissive Motifs	199	-0.0573	-0.0575	-0.1127	0.0222	0.0211	**Male Bias in Naivete**	199	0.0106	0.0121	-0.0773	0.0923	0.0174
Male Dominant	199	0.1132	0.1114	0.0278	0.2121	0.0345	**Male Bias in Dominance, Submissiveness, Engagement in Domestic Affairs and Physical Activity**	199	0.1783	0.1859	0.0444	0.31	0.0539
Female Dominant	199	0.0442	0.0439	0.0000	0.1077	0.0157							
Male Bias in Dominant Motifs	199	0.0690	0.0699	-0.0278	0.1551	0.0263							

Notes: This Table reports summary statistcs on how male and female characters are depicted in the oral traditions of 199 countries.

Appendix Table 9b: Cross-Country Correlation in Gender Stereotypes

Male Bias in Being Portrayed as	Dominant	Submissive	Physically Active	Engaged in Domestic Affairs	Naïve	Intelligent	Sexual	Dominant, Submissive, Engaged in Domestic Affairs and Physically Active
Dominant	1.00							
Submissive	-0.58	1.00						
Physically Active	0.24	-0.17	1.00					
Engaged in Domestic Affairs	-0.28	0.26	-0.23	1.00				
Naïve	-0.03	-0.08	0.09	-0.10	1.00			
Intelligent	0.28	-0.35	0.08	-0.27	-0.22	1.00		
Sexual	-0.03	-0.08	-0.17	-0.13	0.02	0.13	1.00	
Dominant, Submissive, Engaged in Domestic Affairs and Physically Active	0.85	-0.79	0.49	-0.57	0.07	0.37	0.01	1.00

Notes: This Table reports the correlation of the male bias of different stereotypes depicted in the oral traditions of 199 countries. Male bias is the difference in the frequency of male minus female motifs for a given stereotype (normalized by the total number of motifs in a country's oral tradition). The classification of each motif into stereotypes is done by 9 MTurks on average.

Appendix Table 10: Gender Norms in the Oral Traditions and the Role of Women Today

Female Labor Force Participation in 2019 Across Countries							
Male Bias in Being Portrayed as:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dominant	-231.5711***						
	(44.1850)						
Submissive		150.4743**					
		(68.9887)					
Sexual			73.8563				
			(125.1493)				
Intelligent				107.7713*			
				(62.3781)			
Physically Active					-236.3503**		
					(92.3131)		
Engaged Domestic Affairs Motifs						317.7306***	
						(89.3238)	
Naïve							-66.2502
							(86.8053)
Baseline Controls	yes	yes	yes	yes	yes	yes	yes
Adjusted R^2	0.143	0.025	-0.011	0.004	0.024	0.066	-0.008
# of Observations	174	174	174	174	174	174	174

Notes: The unit of analysis is a country. Baseline Controls include: Ln(# of Publications) and Ln(Year of First Publication). Standard errors (reported in parentheses) are robust ***, **, * denote significance at 1%, 5%, and 10%, respectively.

Appendix Table 11 - Panel A: Telugu-related Sources

Author	Year	Source
Elwin, Verrier	1949	Myths of Middle India. Madras: Oxford University Press. p. 532
Payyand, Raghavan	2006	"Religion – native and alien: Interaction, assimilation and annihilation – a study based on worldview". Discourse of Ideology, Religion and Worldview, M.D. Muthukumaraswamy, ed. Chennai: National Folklore Support Center. P. 198-208.
Zograf, Georgy Alexandrovich	1964	Tales of the peoples of India. Translations from Marathi, Punjabi, Tamil, Telugu, Hindi. Preface and notes G. Zograf. M.-L.
Venkataswami, M.N.	1923	Heerámma and Venkataswami or Folktales from India. Madras: Diocesan Press. 230 p.
Hiebert, Paul G.	1971	Konduru. Structure and Integration in a South Indian Village. Minneapolis: University of Minneapolis Press. 192 p.
Pantulu, G.R. Subramiah	1905	Folk-lore of the Telugus, a collection of forty-two highly amusing and instructive tales. Madras: G.A. Natesan & Co. 139 p.
Grierson, George A.	1908	Linguistic Survey of India. Collected and edited by G.A. Grierson. Vol. 9. Part 2. Specimens of Rajastani and Gujarati. Calcutta: Superintendent Government Printing. 477 p.
Jason, Heda	1989	Types of Indic Oral Tales. Supplement. Helsinki: Finnish Academy of Sciences. 100 pp (FF Communications 242).
Kudinova, M.V., A.M. Kudinov	1995	When luck smiles: Indian tales, legends and folk tales. Comp. and trans. from English Mv Kudinova and A.M. Kudinov. M .: Oriental literature, School Press. 320 s. {According to Folk Tales of India (21 Volumes). Sterling Publishers Private Lmt., 1969-1974}.

Appendix Table 11 - Panel B: The Telugu Oral Tradition

Group	ID	Title
Telugu	a12	Eclipses: a monster's attack
Telugu	a12f	Eclipses: creditor
Telugu	b2a	The female earth
Telugu	b3b	Earth grows big
Telugu	c8	The primeval couple of siblings
Telugu	f100	Test of chastity (a queen and a servant girl)
Telugu	h36	The muddled message
Telugu	h36g	Muddled message: how many meals a day
Telugu	h36g1	Bull is a failed messenger
Telugu	h49b	The faithful dog as security for a debt
Telugu	h7	The personified Death
Telugu	h7c	Not finished prayer
Telugu	i127	Ursa major is a bed
Telugu	i40	Rainbow bow
Telugu	j23	A late son kills monsters
Telugu	k101	Night dances of girls
Telugu	k101c	In the palace by day, on the sky by night
Telugu	k116b	The lecherous holy man and the maiden in the box
Telugu	k129	The disenchanted beauty
Telugu	k130a	Girl in the house of several brothers
Telugu	k136c	Killed and revived periodically
Telugu	k144	The predestined death because of an animal
Telugu	k24	Stolen clothes of supernatural woman
Telugu	k25	Woman from sky-world marries mortal man
Telugu	k27n1	Task-giver is a king or a chief
Telugu	k27nn	Envious minister
Telugu	k27x2	To steal an egg from under the bird
Telugu	k27z	Game of chance for life and death
Telugu	k27z3	Cat with a lamp
Telugu	k27z4	The trained animal of the gambler
Telugu	k37	Recognition-test
Telugu	k38b	Snake threatens nestlings
Telugu	k38c	Bird carries hero to his destination
Telugu	k38e	Of copper, of silver, of gold
Telugu	k56b	The worthy man is rewarded, the unworthy punished
Telugu	k77a	Small objects and animals defeat the ogre
Telugu	k77c	Ones who hide in a house frighten dangerous enemy
Telugu	k82	Evil sister-in-law
Telugu	k83	The sons on a quest for a wonderful remedy for their father
Telugu	k92a	The princess responsible for her own fortune
Telugu	k99b	Eloping with the wrong man
Telugu	l19b	Multi-headed being
Telugu	m114j	All women are similar
Telugu	m130a	Bird helps animal to escape from snare
Telugu	m145	The lion in a well

Appendix Table 11 - Panel B: The Telugu Oral Tradition cont'd

Group	ID	Title
Telugu	m152	Why only one wolf?
Telugu	m152a	Animal tied to another for safety
Telugu	m170	Pilgrimage of the animals
Telugu	m197d	The shortened stick
Telugu	m29b	Trickster-fox, jackal or coyote
Telugu	m29w3	The lion is a failure
Telugu	m39e2	The speaking tree
Telugu	m91c2	Put into bag