How Language Shapes Prejudice Against Women: An Examination Across 45 World Languages

David DeFranza¹

Himanshu Mishra²

Arul Mishra³

Forthcoming: *Journal of Personality and Social Psychology (Attitudes and Social Cognition)*

¹ University of Utah, david.defranza@eccles.utah.edu

² University of Utah, himanshu.mishra@utah.edu

³ University of Utah, arul.mishra@utah.edu

Abstract

Language provides an ever-present context for our cognitions and has the ability to shape them. Languages across the world can be gendered (language in which the form of noun, verb, or pronoun is presented as female or male) versus genderless. In an ongoing debate, one stream of research suggests that gendered languages are more likely to display gender prejudice than genderless languages. However, another stream of research suggests that language does not have the ability to shape gender prejudice. In this research, we contribute to the debate by using a Natural Language Processing (NLP) method which captures the meaning of a word from the context in which it occurs. Using text data from Wikipedia and the Common Crawl project (which contains text from billions of publicly facing websites) across 45 world languages, covering the majority of the world's population, we test for gender prejudice in gendered and genderless languages. We find that gender prejudice occurs more in gendered rather than genderless languages. Moreover, we examine whether genderedness of language influences the stereotypic dimensions of warmth and competence utilizing the same NLP method.

Keywords: gender, prejudice, stereotype, natural language processing, word embeddings, text analysis

Language is the ever-present context of our cognitions and provides an insight into our cognition (i.e., thoughts, attitudes, and beliefs; Carnaghi & Maass, 2007; Collins & Clément 2012). It is considered both a vessel for holding cognitions as well as a lens that can direct (or distort) those cognitions (Sutton 2010). Differences in how concepts and objects are defined across languages, therefore, leads to differences in how they are perceived by respective speakers (Casasanto 2008). One dimension where languages differ is gender. World languages can be categorized as gendered (a language in which the form of noun, verb, or pronoun is presented as female or male) versus genderless languages. This raises the following question: if gender is emphasized continuously in language, then is it possible that gendered languages hold differing perceptions of women relative to men compared to genderless languages.

There are conflicting answers to this question. One stream of work shows that people who speak gendered languages do display greater prejudice against women compared to those who speak genderless languages (Wasserman & Weseley, 2009). This stream of research suggests that *genderedness* of language can cause perpetuation of gender perceptions and aligns with findings on linguistic relativity (Casasanto, 2008, 2016; Oh 2003). However, a competing argument is that language cannot influence perceptions against women because language merely reflects the beliefs of a culture and cannot influence or cause gender prejudice (Deutscher, 2010; Fodor, 1975; Pinker, 1994).

We contribute to this debate in three ways. First, we compare 45 world languages covering the majority of the world's nations, to test for gender prejudice across several gendered and genderless languages. Second, instead of testing for prejudice through attitudinal measures or through response time tests, we utilized thousands of gigabytes of data, voluntarily generated

by people across different languages and different contexts, when they may not be specifically discussing gender issues. Wikipedia and a corpus of web crawl data from over five billion web pages, known as the Common Crawl, serve as our data source.

Consistent with past work, we define prejudice as the affective component of intergroup bias where social groups are differentially associated with positively and negatively valenced words (Fiske, Cuddy, Glick, & Xu, 2002; Bhatia, 2017; Todd, Simpson, Thiem, & Neel, 2016). Such valenced associations have been shown to automatically affect subsequent behavior when certain social categories are activated (Allport, 1954; Devine, 1989). Knowing the nature of the diverging association is important (e.g. greater association of positive with women or with men) because only then successful interventions can be designed. Therefore, third, we investigate whether gender prejudice occurs because of a greater prejudice against women or a lesser prejudice for men.

We next present the details of the debate as to whether language can shape cognitions or whether it can only communicate cognitions, and then discuss how we attempt to resolve this debate.

Theoretical Review

The Debate: Does Language Shape or Simply Communicate Thought?

Linguistic relativity. The Sapir-Whorf hypothesis claims that the language one speaks influences that individual's perception of the world. When people speak different languages, they think differently (Casasanto, 2008; Kay & Kempton, 1984). In its strongest formulation, the Sapir-Whorf hypothesis is a type of linguistic determinism, which suggests that the categories made available by a person's language determines that person's thoughts (Kay &

Kempton, 1984; Pinker, 1994). Therefore, it is assumed that language, thought, and culture are intertwined and each language creates a distinctive worldview (Gumperz & Levinson, 1991). A more common and modern version of the Sapir-Whorf hypothesis is known as linguistic relativity (Casasanto, 2008, 2016; Pinker, 1994). It states that “differences among languages cause differences in the thoughts of their speakers” (Pinker, 1994, p. 57) because language provides an ever-present context for each of our cognitions and hence has the ability to shape cognitions (Casasanto, 2008, 2016). In other words, instead of focusing on language and its influence on vague notions of worldview, research on linguistic relativity now focuses on investigating the influence of language on specific cognitive effects like perception of time and motion (Casasanto, 2008).

For example, research has demonstrated that English motion verbs focus on the manner of motion (rolling, climbing, running) while Spanish motion verbs focus on the path of motion (enter versus exit). Such a difference in encoding results in individuals remembering more manner-related versus path-related information depending on which language they use (Casasanto, 2016; Oh, 2003).

One challenge to such research is that it must disentangle cognition from language. In order to ensure that the task used to evaluate cognitions is not related to language, studies have used an *extralinguistic* feature. The extralinguistic feature is orthogonal to language (i.e., it does not have a language tag attached to it) but is still affected depending on which language an individual speaks. Hence, it provides converging, empirical evidence for the ability of language to shape cognition (Casasanto, 2008). For instance, researchers used a combined temporal-spatial assessment task in which participants saw a line being drawn between two points (spatial) over a

period of time (temporal). Participants were then asked to reproduce either the temporal or the spatial aspect of the line. English-speaking participants, for whom time was related to length (e.g., a long versus short time), confounded spatial and temporal tasks. However, Greek-speaking participants, for whom time is associated with volume (e.g., big versus small), made fewer mistakes and were able to disentangle temporal and spatial tasks (Casasanto, 2008, 2016). When the spatial-temporal task was converted into an assessment of volume with time, Greek speakers made more mistakes than English speakers. Therefore, the use of an extralinguistic feature provides support for linguistic relativity. A similar stream of literature examines the role gendered versus genderless languages on gender prejudice.

Gendered versus genderless languages and prejudice against women. Grammatical gender describes the formal rules of a language such as the form of nouns, verbs, or pronouns that reflect gender and distinguish world languages into three broad categories: gendered languages, natural gendered languages, and genderless languages (Prewitt-Freilino, Caswell, & Laakso, 2012; Stahlberg, Braun, Irmen, & Sczesny, 2007). In gendered languages, rules of grammar dictate that most nouns appear with a male or female gender (e.g., *la pelota* or “ball” in Spanish which has a female designation or *le ballon* in French, which has a male designation). Their associated pronouns, adjectives, and verbs reflect the assigned gender. Examples of such languages include Spanish, French, Russian, and Hindi, to name only a few. Generally, in natural gender languages, nouns tend not to have gender tags but gender is indicated through pronouns such as *she*, *he*, or *it*. English is perhaps the best example of a natural gender language.¹ Genderless languages do not have any grammatical rules to distinguish nouns by gender and

their pronouns relating to men and women could be the same (Stahlberg et al., 2007). Examples of genderless languages include Chinese and Finnish.

Gendered languages personify even inanimate objects with a specific male or female gender, and such personification could unconsciously reinforce attitudes about gender (Boroditsky & Schmidt, 2000; Jakobson, 1959). For instance, in languages where the sun is given a male gender (e.g., French and Spanish), it is perceived to be more powerful and courageous but when it is given a female gender (e.g., German), it is considered warm and nurturing (Hofstätter, 1963). Assignment of gender to objects brings the gender marker into the foreground and results in greater salience of gender roles. It follows that such salience would result in perceptions of that gender being reiterated more frequently. Thus, in Spanish and French the moon, which has a female gender assignment, is seen as being weaker and more passive than the sun because of the salience of the perception of women being weaker (Boroditsky & Schmidt, 2000; Hofstätter, 1963).

Gendered (as opposed to genderless) languages influence not only the perception of women but also attitudes about them, which can result in gender prejudice (Boroditsky & Schmidt, 2000). For instance, utilizing EEG, it has been demonstrated that languages can influence cognitions non-consciously, since one uses language continuously and the influence can occur without awareness (Boutonnet, Athanasopoulos, & Thierry, 2012). In countries of the world where gendered language is spoken, women earn lower wages (van der Velde, Tyrowicz, & Siwinska, 2015) and face more barriers to participating in the political arena (Santacreu-Vasut, Shoham, & Gay, 2013).

Language merely reflects thought. A competing argument, which forms the other part of the debate, is that language cannot influence the perpetuation of gender prejudice because it can communicate thought but it cannot shape cognitions (Cela-Conde & Marty, 1998; Pinker, 1994). For instance, researchers from this side of the argument suggest that any cognition can be expressed in any language (Cela-Conde & Marty, 1998; Hunt & Agnoli, 1991; Pinker, 1994). English speakers, for example, may not have a dedicated word for *schadenfreude* but it is easy enough to communicate the experience of relishing in another's misfortune (Deutscher, 2010). Similarly, we can see all languages adjust and adapt to describe new ideas, objects, and innovations (Hunt & Banaji, 1988; Hunt & Agnoli, 1991). Hence, they argue that Whorf's strict relativism, which claims language dictates what its speakers are capable of thinking or understanding cannot be true. This stream of research also argues that it takes far less time to coin and adopt a term that defines a new concept but far longer for the concept to form and acquire an entity of its own (Hunt & Banaji, 1988; Hunt & Agnoli, 1991; January & Kako, 2007; McWhorter, 2016). That is, language provides a label to a construct but does not shape the construct (Deutscher, 2010).

However, it is worth noting that many of these arguments get reduced to strict pronouncements that *only* language can shape cognition or that language can *only* reflect cognition. Some research has argued that a middle ground may be more reflective of reality wherein language is influenced by changing cognition of people but in return language can also affect how people perceive and especially express thoughts (Boroditsky & Schmidt, 2000; Casasanto, 2008; Deutscher, 2010; Kay & Kempton, 1984; Sera, Elieff, Forbes, Burch, Rodríguez, & Dubois, 2002). While language may not constrain cognition, it does dictate the

means through which that cognition can be expressed. For instance, if there is no best-fitting word to express a cognition in a language, then the repeated use of the *only* available word for that cognition could result in the thought becoming modified because of the meaning that word conveys. German speakers, for example, are forced to describe a fork as *die gabel*, which is female (Deutscher, 2010). While the gender marker may not have literal meaning to German speakers, there is evidence that this subtle gender cue leads to associations of a fork with female characteristics (Boutonnet et al., 2012; Hofstätter, 1963; Sera et al., 2002).

In sum, there is an ongoing debate whether language shapes cognitions (linguistic relativity) or merely reflects cognitions. Research has used gendered versus genderless language to test for gender prejudice. This stream of research has demonstrated the greater presence of gender prejudice in gendered compared to genderless language as evidence of language influencing cognition (Boyd & Richerson, 1985; Evans & Levinson, 2009; Konishi, 1993; Liu & Karahanna, 2017; Sera, Berge, & del Castillo-Pintado, 1994; Stubbs, 1996). Such a pattern of results could be because socially assigned gender roles of men (e.g., being stronger) versus women (e.g., being weaker) are emphasized more often in gendered languages. However, research examining the association between gender prejudice and the use of gendered versus genderless language has been criticized on a few fronts.

One criticism is that the evidence rests on studies considering only two or three languages. Since there are more than 100 different languages spoken by 85% of the world population (Simmons & Fennig, 2018), the conclusions drawn from two or three languages might not be generalizable. It is possible that observed effects might disappear when other or more languages are considered. A second criticism is that such research often uses explicit

measures, which can allow people to suppress their perceptions or they may not be aware that they hold a certain perception (Banaji, Hardin, & Rothman, 1993; Greenwald & Banaji, 1995; Greenwald et al., 2002; Greenwald, McGhee, & Schwartz, 1998; Greenwald, Nosek, & Banaji, 2003). Finally, there is very little research that describes the nature of gender prejudice in gendered languages, which makes it difficult to conclude whether the prejudice is predominately against women or in favor of men.

We contribute to this debate by studying 45 world languages in order to examine if gendered languages are more likely to contain gender prejudice. This helps address the criticism of using only one or two languages while testing. We used a natural language processing (NLP) method to test whether genderedness of language is associated with gender prejudice.

If past research showing the relationship of language genderedness and gender prejudice is to hold true, we should find gender prejudice to be more prevalent in gendered rather than genderless languages. Conversely, if it is true that language does not influence gender prejudice, but is simply a reflection of the culture of that language community (Cralley & Ruscher, 2005; Stahlberg et al., 2007), then we should find gender prejudice across both gendered and genderless languages because gender prejudice is known to exist in every part of the world (Matthes, Prieler, & Adam, 2016; Prewitt-Freilino et al., 2012; United Nations Development Programme, 2017; Williams, Satterwhie, & Best, 1999; World Economic Forum, 2017).

Existing studies examining gender prejudice in gendered versus genderless languages have directly queried participants' regarding their gender beliefs. For instance, questions such as whether men and women should be treated similarly or paid the same for the same job or whether women have already been given too much by the government tend to elicit responses

that capture explicit gender prejudice (Wasserman & Weseley, 2009). But some research has suggested that direct questioning cannot reliably capture prejudice or latent traits that inform behavior (Banaji et al., 1993; Deutscher, 2010; Sterling, Jost, & Bonneau 2020). This could happen because individuals tend to respond in a socially desirable manner where they hide gender prejudice (Zerbe & Paulhus, 1987). Secondly, they may engage in self-deception and not believe that they hold any gender prejudice. Finally, sampling from a small set of languages and querying participants directly may lack ecological validity necessary to test psychological constructs (Sterling et al., 2020).

We used an NLP method that takes care of shortcomings of prior investigations. It allowed us to use the voluntarily produced thoughts and feelings of people to examine the association of language genderedness and gender prejudice. The voluntarily produced text comes from large corpora, Wikipedia and Common Crawl. These text corpora reflect the real-world communications and opinions produced by people online. They are a compilation of text generated by individuals on a vast range of topics, and not just when they are specifically discussing gender issues. Briefly, the intuition behind our approach is the following. If a language contains prejudice against women, then words representing women compared to those representing men will appear semantically more associated with negatively valanced words.

The challenge is that it is not always easy to assess semantic association among words. To this end, the NLP method we used can analyze vast quantities of text (in the context of this manuscript it is several thousand gigabytes of data) to capture semantic association among words. Moreover, it is important to note that while text corpora allow us to measure associations between gender and prejudice-exhibiting words, these associations do not suggest a causal link

between genderedness of language and prejudice. However, past research across domains, has demonstrated that cooccurrence studies showing a link between two variables in the real world can be used in deriving many useful insights; we do the same in this research (Maki, 1979; Maki, Maki, & Marsh, 1977).

We next explain this method and how it is able to capture semantic associations from the context of the word and convert it into a numeric format that can be used in further analysis.

Method

In today's world, with people producing vast amounts of digital text, collections of such text, or corpora, have become the holding body of people's cognitions, feelings, and actions across time and regions. Analysis of digitized text has provided insights in the social (e.g., examining the condition of marginalized groups; Ben-David & Matamoros-Fernández, 2016; Garg, Schiebing, Jurafsky, & Zou, 2018; Thelwall, Wilkinson, & Uppal, 2010), business (e.g., analyzing consumer reviews to extract quality inference; Tirunillai & Tellis, 2014), political (e.g., analyzing speeches, legislative addresses, or parliamentary debates to examine influence of laws; Miller, 2013; Schonhardt-Bailey, 2006; Stewart & Zhukov, 2009), and legal (e.g., approval of home loans or using license plates to predict crime; Schrage, 2014) domains. In our research, we use a specific NLP method that creates word embeddings and thus enables study of semantic associations among words.

What are Word Embeddings?

We can easily discern that the colors red and maroon are more similar than the colors red and yellow. However, when trying to discern if "rose" is more similar to "pleasant" or

“unpleasant,” we will hesitate a bit because without a context it is difficult to understand if “rose” is a flower or an individual. This poses a fundamental question: How do words acquire meaning? Saussure (1916/1959) argued that “language is a system of interdependent terms in which the value of each term results solely from the simultaneous presence of others.”

Wittgenstein (1953) asserted that “the meaning of words lies in their use.” Inspired by these assertions linguists developed the distributional hypothesis (Harris, 1954) captured in Firth’s (1957) oft-quoted statement, “you shall know a word by the company it keeps.” That is, words derive their meaning from the context in which they occur. For instance, the meaning of “spring” as a season or a piece of coiled metal can only be ascertained by the context in which it is used. Similarly, pianos are heavy in the context of moving furniture, but they are musical in the context of a virtuoso.

The distributional hypothesis proposes a link between how words are distributed in language and the similarity in their meaning (Jurafsky & Martin, 2019). Consider reading the following sentences: *hikareti* is bad for your lungs; *hikareti* is addictive; *hikareti* contains nicotine. Even if one is not familiar with Maori language, one can infer that *hikareti* means something that is used to inhale smoke like a cigarette or cigar. One can infer the meaning of *hikareti* from words that appeared in its context because these words also appear in the context of cigarette or cigar. Therefore, by examining different words (nicotine, addictive, lungs) that appear in the context of a word (*hikareti*) one can estimate whether it is similar or dissimilar to other words (cigarette vs. apple).

The essence of the distributional hypothesis is that words, which share similar context words are likely to be semantically similar. Note, in this work semantic similarity between words

implies relatedness in their meaning (i.e., likeness of their meaning) because they occur with similar context words, similarity does not imply sharing similar features. When we have large quantities of text, the goal is to let algorithms use the distributional hypothesis to learn semantic associations among words like a human can learn about *hikareti* and *cigarettes/cigar*. Modern NLP algorithms do this by building word embedding models where each word in a large text corpus is represented as a multidimensional vector (commonly referred to as word embeddings or word vectors). That is, word embeddings are the numeric, multidimensional representation of words.

By learning relationships based on the context in which words occur, word embeddings allow complex comparisons, like estimating the semantic similarity between words, to be performed using relatively simple computations. For instance, the cosine of the angle between embeddings of two words informs us how similar they are (Bolukbasi, Chang, Zou, Saligrama, & Kalai, 2016; Caliskan, Bryson, & Narayanan, 2017). If v_{w_1} represents the embedding of word w_1 and v_{w_2} represents the embedding of word w_2 then similarity (S) between w_1 and w_2 is:

$$S(w_1, w_2) = \frac{v_{w_1} \bullet v_{w_2}}{|v_{w_1}| |v_{w_2}|} \quad (\text{Equation 1})$$

In Equation 1, the numerator represents the dot product or inner product of the embeddings v_{w_1} and v_{w_2} , while the denominator is the product of their respective Euclidean lengths.

This similarity calculation offers some interesting insights. When two embeddings are parallel, they have an angle of zero degrees and a cosine of one; they are considered maximally similar in meaning. When two word embeddings are perpendicular, thus having an angle of 90

degrees and a cosine of zero, they are considered orthogonal (unrelated) in meaning. When two word embeddings have an angle of 180 degrees and a cosine of -1, they are considered opposite in meaning. Beyond similarity, word embeddings also preserve linguistic regularity that humans use in answering analogies like “*King:Men::Women:___*”. Thus, embeddings of words can also be used to answer analogy tasks.

Embeddings of words that preserve their semantic relationships align with our goal of examining gender prejudice in gendered versus genderless languages. If a language is documenting prejudice against women then words representing women compared to those representing men will appear semantically more related with negatively valenced words. Therefore, we utilized embeddings created with one of the latest machine learning methods, *fastText*, which accommodates morphological variations across different languages. *fastText* (Bojanowski, Grave, Joulin, & Mikolov, 2016; Grave, Bojanowski, Gupta, Joulin, & Mikolov, 2018; Mikolov, Grave, Bojanowski, Puhrsch, & Joulin, 2017) modifies the Word2Vec algorithm used in prior work (Mikolov, Chen, Corrado, & Dean, 2013; Mikolov, Sutskever, Chen, Corrado, & Dean, 2013).

The intuition behind the method is the following. If we define the context as certain number of words occurring before or after a target word, then the goal is to predict a target word from the words that occur in its context. As an example, take the sentence “farmers grow Arabica coffee beans in Colombia.” If the context window is three then *fastText* uses {Farmers, grow, Arabica, beans, in, Colombia} to predict {coffee}. If we denote target word coffee as w_t and the set of context words as set C_t then probability that w_t is the target word given C_t will be $P(w_t | C_t)$. However, for an irrelevant target word n (e.g., glacier, also called a negative sample),

C_t will not provide a good prediction and the probability of predicting an irrelevant target like n denoted by $P(n | C_t)$ will be lower than predicting $P(w_t | C_t)$.

Extending this example, we see that if two words occur in the same context (e.g., Colombia and coffee) then their similarity will be greater than two words that do not appear in the same context (e.g., Colombia and glacier). Therefore, $P(w_t | C_t)$ will be a function of the similarity between w_t and C_t . Similarly, $P(n | C_t)$ will be a function of similarity between C_t and n . For a relevant target word w_t , the algorithm also samples irrelevant words, like n , and maximizes the similarity between the context and relevant target word while minimizing the similarity between the context and sampled irrelevant words.

Formally, we can describe the above algorithm in the following manner. Given T words $\{w_1, w_2, w_3, \dots, w_t\}$, where C_t is a set of context words and w_t is the target word (e.g., coffee), the similarity (S) will be $S(w_t, C_t) = v_{wt} \bullet v_{ct}$. Here ‘ $\bullet$ ’ is dot product, v_{wt} is the embedding of word w_t , and v_{ct} is average of the embeddings for each word in C_t . The probability that w_t is a target word given context words C_t is a function of similarity S . The function that turns similarity to probability is the sigmoid or logistic function (Jurafsky & Martin, 2019). Therefore,

$$p(w_t | C_t) = \frac{1}{1 + e^{-S(w_t, C_t)}} \quad (\text{Equation 2})$$

Which is equivalent to:

$$\log p(w_t | C_t) = \log \frac{1}{1 + e^{-S(w_t, C_t)}} \quad (\text{Equation 3})$$

Similarly, when word n is drawn from a negative sample (e.g., glacier) then the probability of n occurring in the context window is:

$$p(n | C_t) = \frac{1}{1 + e^{-S(n, C_t)}} \quad (\text{Equation 4})$$

and the probability of n not being a target word in the context of C_t is $1 - p(n | C_t)$ where,

$$1 - p(n | C_t) = 1 - \frac{1}{1 + e^{-S(n, C_t)}} \quad (\text{Equation 5})$$

$$= \frac{e^{-S(n, C_t)}}{1 + e^{-S(n, C_t)}}$$

$$= \frac{1}{1 + e^{S(n, C_t)}}$$

If N_c is the set of negatively sampled words and if we assume these negatively sampled words to be independent, then the probability of all the words in N_c not being a target word in the context of C_t is obtained by multiplying the individual probability of each n :

$$1 - P(n \in N_c | C_t) = \prod \frac{1}{1 + e^{S(n, C_t)}} \quad (\text{Equation 6})$$

$$\log(1 - p(n \in N_c | C_t)) = \sum_{n \in N_c} \log \frac{1}{1 + e^{S(n, C_t)}}$$

Now we have two probabilities, one is the probability that w_t is a target word in the context of C_t and the other is the probability that negatively sampled words n are not target words in the context of C_t . Both of these probabilities are functions of similarity. Thus, the

learning objective becomes maximizing both of these probabilities. For all the words T , this objective can be written as maximizing:

$$\sum_{t=1}^T \left[\log \frac{1}{1 + e^{-S(w_t, C_t)}} + \sum_{n \in N_c} \log \frac{1}{1 + e^{S(n, C_t)}} \right] \quad (\text{Equation 7})$$

This maximization process is governed by a shallow two-layer neural network. Starting with a hidden layer of randomly initialized embeddings, the text corpus is streamed into the network one context window at a time. Embeddings are updated after each instance with the goal of maximizing similarity between context words and relevant target words while minimizing similarity between context words and negative sample words (Bojanowski et al., 2016; Grave et al., 2018).

However, the above approach of creating word embeddings has a limitation: it ignores the internal structure of words. By creating an embedding for each word, the algorithm does not consider that in morphologically rich languages, various forms of a word can occur. For instance, in French or Spanish, most verbs have more than forty different inflected forms, while the Finnish language has fifteen cases for nouns (Bojanowski et al., 2016). In Russian, there are at least 70 word forms for dog, encoding gender, age, number, sentiment and other semantic connotations (Vylomova, Cohn, He, & Haffari, 2016). Many of these forms rarely occur in the text corpus, thus reducing the quality of the resulting word embeddings (i.e., their ability to capture semantic relationship).

fastText addresses this weakness by utilizing subword information to capture common and systematic morphological rules (Bojanowski et al., 2016; Grave et al., 2018; Joulin, Grave,

Bojanowski, & Mikolov, 2016). The word, “where,” for example, will be represented in word2vec by just one set of d -dimensional embeddings. However, in *fastText*, “where” will be decomposed into various subgroups of n contiguous characters known as character n -grams (Bojanowski et al., 2016; Grave et al., 2018; Joulin et al., 2016). When n is set to three, as is common, “where” will be decomposed into subgroups of characters { wh, whe, her, ere, re }. Each of these subgroups will have their own d -dimensional embeddings and the final embedding of the complete word “where” will be averaged to represent the individual word. Formally, a word w_t is represented by:

$$v_{wt} + \frac{1}{N} \sum_{n \in N} x_n \quad (\text{Equation 8})$$

This new representation of w_t essentially changes the similarity calculation, which now becomes

$S(w_t, C_t) = (v_{wt} + \frac{1}{N} \sum_{n \in N} x_n) v_{ct}$. As an added benefit, *fastText* also improves the performance of embeddings when the corpus contains misspelled words or concatenations of words (Bojanowski et al., 2016; Grave et al., 2018; Joulin et al., 2016).

The dimensions of the word embedding actually do not have any meaning on their own. In other words, they don’t represent any specific quality or attribute of a word. These embeddings are numerical values (vectors) assigned to each word that preserve their semantic relationship with other words. Three hundred dimensions of word embeddings, instead of 500 or 200 dimensions, are used since it provides the optimal performance in finding semantic associations (Bojanowski et al., 2016; Grave et al., 2018).

Identifying gender prejudice. To determine whether the *fastText* embeddings are detecting gender prejudice from text corpora, we adapted a method suggested by Caliskan, et al. (2017). Since, *fastText* converts a word into a high dimension numerical vector, we can then find the level of similarity between two words using cosine similarity, as described in Equation 1. The resulting score captures the similarity between two words (i.e., whether a target word is more or less similar to a valenced word).

Consider, for example, how word embeddings can be used to test for gender prejudice. Let us assume that there are two sets of male (e.g., he, him, man) and female (e.g., she, her, woman) words. These gender words are compared to two sets of valenced words: one positive (e.g., freedom, health, love, peace) and the other negative (e.g., abuse, crash, filth, sickness). If a language contains prejudice against women then male words would display more similarity to positive rather than negative words compared to female words. Formally, we can represent the presence of a gender prejudice as a net similarity measure where S represents similarity (as in Equation 1).

$$S(\text{male}, \text{positive}) - S(\text{male}, \text{negative}) > S(\text{female}, \text{positive}) - S(\text{female}, \text{negative}) \quad (\text{Equation 9})$$

The left-hand side of the inequality in Equation 9 measures the net similarity of male words to positively versus negatively valenced words. The term on the right-hand side of the inequality measures the net similarity of female words to positive versus negative words. Estimating these net similarities helps us both detect whether gender prejudice exists and what is the nature of the gender prejudice. To detect gender prejudice, we compare the left-hand side of the inequality with the right-hand side. If the left-hand side is more than the right-hand side then it would indicate that the net association of male words with positive words is more than the net

association of female words with positive words. Another way of representing this is by taking the difference of the right-hand-side from the left-hand side; higher value of the difference would indicate lower prejudice against men. That is, it shows that males are more positively represented in text corpora than females. Such a discrepancy would be surprising in light of the Pollyanna hypothesis (Boucher & Osgood, 1969) also known as linguistic positivity bias (Iliev, Hoover, Dehghani, & Axelrod, 2016). The Pollyanna hypothesis states that positive words are more prevalent, easily learned, and used across languages. People tend to couch communications in positive language, even when the core content of the communication is negative in nature. Studies of large text corpora, using translations across languages, have supported this hypothesis (Dodds et al., 2015). Thus, the finding that despite their prevalence, positive words are semantically more associated with males than females provides evidence of prejudice.

Further, Equation 9 can help in understanding the nature of the prejudice. For instance, the left-hand side of the inequality could be more than the right-hand side if a) male words were roughly equal in similarity to positive and negative words but female words were much more similar to negative words than positive words; b) male words were more similar to positive than negative words but female words were much more similar to negative words than positive words; or c) male words were more similar to positive than negative words but female words were roughly equal in similarity to positive and negative words. Following Equation 9 we can answer what kind of association between gender words and valenced words leads to prejudice.

Details of the measurement of similarity and the statistical tests are provided in Appendix B.

Data Description

We used two text corpora to test if gender prejudice is more likely to exist in gendered languages. The first corpus consisted of data from the free online encyclopedia Wikipedia (Bojanowski et al., 2016) and the second was from the Common Crawl project, which contains snapshots of all the text that can be scraped from the publicly facing Internet, since 2013 (Grave et al., 2018) and contains more than 630 billion words. These two sources are available in several world languages and tend to cover a broad set of topics. Moreover, Wikipedia is considered nominally objective and fact-based (Filipacchi, 2013; Glott, Schmidt, & Ghosh, 2010; Gruwell, 2015; Hill & Shaw, 2013; Pfeil, Zaphiris, & Ang, 2006; Shane-Simpson & Gillespie-Lynch, 2017; Wagner, Garcia, Jadidi, & Strohmaier, 2015; Wagner, Graells-Garrido, Garcia, & Menczer, 2016) and thus provides a conservative setting to test for gender prejudice.

For both data sets, we used pre-trained word embeddings in our analysis. The Wikipedia² and Common Crawl³ embeddings were produced by Bojanowski et al. (2016) and Grave et al. (2018) respectively and are publicly available. Each embedding consisted of a vector for each word, derived from the *fastText* algorithm (Bojanowski et al., 2016; Grave et al., 2018; Joulin et al., 2016) discussed previously. In both data sets, the word embedding for each word consisted of 300 dimensions.

Gender and Valence Words

For our examination of gender prejudice, we used two categories of words: *gender* words and *valence* words. Gender words consisted of 218 noun and pronoun pairs identified and tested by Bolukbasi et al. (2016). For instance: *he, his, him, man, men* versus *she, hers, her, woman, women*; refer to Appendix C for the complete list. Valence words consisted of 25 positively and 25 negatively valenced words commonly used to assess prejudice (Caliskan et al.,

2017; Greenwald et al., 1998, 2003). The relative similarities (see Equation 1), between the gender words and the valenced words were used to measure the strength of association (see Equation 9) and hence, the level of gender prejudice. Details of the lexicon are available in Appendix C.

Language Selection and Validation

We considered several criteria in order to select languages for analysis. First, since one of the embeddings was based on Wikipedia corpora, languages were selected based on their popularity on Wikipedia (Wikipedia, 2018a). Due to the differential properties and behaviors of various Wikipedia communities, there is a trade off in many languages between volume of articles and the depth of the articles. Volume describes the total number of unique articles in the project. Depth is a metric created by Wikipedia to compare the frequency of updates across a project, normalizing for the size and popularity of different projects (Wikipedia, 2018b). We selected languages based on the volume of available data, rather than focusing on those languages that have the most intensely updated and edited articles because we were primarily interested in the natural occurrence of gender prejudice in language. Using these criteria as a guide, the initial search produced 60 languages with 100,000 or more Wikipedia articles.

Since *gender* words and *valence* words used in prior work are in English, we translated them using the Google Translate API using a procedure outlined in prior work (Dodds et al., 2015). While the coverage of Google Translate service is significant, it is not comprehensive. As a result, our list of 60 languages was reduced to 49 to match the coverage of Google Translate. Despite this reduction, the 49 languages spanned various geographic regions and language families. Translation was conducted on every gender and valence word for each

language, as well as on a validation lexicon. The translated words were used to extract vectors from the appropriate language embedding.

Next, a validation method was implemented to test how good the word embeddings were in detecting semantic similarity in the corpora we were using in the analysis. First, a validation word list was constructed consisting of generic positive nouns (e.g., beauty, celebration, smile) and generic negative nouns (e.g., drought, emergency, fatality). The selected words were unambiguous and based on similar validation lexicons used in past research (Caliskan et al., 2017; Nosek, Banaji, & Greenwald, 2002; Nosek et al., 2009), but adapted to avoid associations with specific cultural constructs. The validation word list was translated, both Wikipedia and Common Crawl embeddings were obtained for each translated word, and associations were calculated using cosine similarity.

We hypothesized that a test of our method would require known positive words to have a statistically significant net positive similarity (i.e., known positive words should be more similar to positively valenced words), while known negative words would have a statistically significant net negative similarity (i.e., known negative words should be more similar to negatively valenced words). For example, the association between “hero” and “love” was expected to be greater than the association between “hero” and “filth” (see Appendix C for a full list of the validation words). If the hypothesized relationship (i.e., known positive words are more similar to positively valenced words and known negative words are more similar to negatively valenced words) was not statistically significant, the language was considered to have failed validation and was removed.

Five languages, Cebuano, Japanese, Latin, Vietnamese, and Welsh failed validation and were removed from the Wikipedia data set and four languages, Cebuano, Hebrew, Latin, and Vietnamese, failed validation and were removed from the Common Crawl data set.⁴ Despite the reduction due to translation availability and validation, languages common on every continent, including first or second languages of more than 50% of the world population were included (Simmons & Fennig, 2018). As discussed previously and suggested by past research (Corbett, 1991; Stahlberg et al., 2007), we divided languages based on whether they are gendered or genderless. To assign languages to these groups, we first consulted the *World Atlas of Language Structures* (Haspelmath, 2005), which provided the necessary information for more than half of the languages in our data set. Next, we consulted *Ethnologue: Languages of the World* (Simmons & Fennig, 2018), which accounted for the majority of remaining languages. Finally, any remaining unclassified languages were validated through references in the series *Gender Across Languages* (Hellinger & Bußmann, 2001). After translation, validation, and classification, there were 28 languages in the gendered group and 16 languages in the genderless group for the Wikipedia data set and 28 languages in the gendered group and 17 languages in the genderless group for the Common Crawl data set. Across both data sets, we considered 45 unique languages (see Appendix A for more details on the validation and classification procedures).⁵ Moreover, we also controlled for various geographic, demographic, and cultural covariates in our analysis (see Appendix E).

Results

Gender Prejudice

We tested for gender prejudice in the two corpora, Wikipedia and Common Crawl, where we compared the similarity of target male words (e.g., *he, him*) and female words (e.g., *she, her*) with positive (e.g., *love, pleasure, lucky*) as well as negatively valenced words (e.g., *abuse, sickness, ugly*).

Wikipedia. The analysis revealed a statistically significant gender prejudice in the Wikipedia data set for 21 of the 44 validated languages, $\chi^2(1) = 3.8724, p = 0.0491$. Specifically the analysis revealed that 17 gendered languages (60%) exhibited gender prejudice while only four (25%) genderless languages did so.

We also examined whether net gender prejudice for males or females interacted with gendered versus genderless languages. Recall from Equation 9, gender prejudice for male is $S(\text{male, positive}) - S(\text{male, negative})$ and for female is $S(\text{female, positive}) - S(\text{female, negative})$. As mentioned earlier, a higher value of the difference of the right-hand-side from the left-hand side would indicate lower prejudice against men. A one-way MANOVA with 2 (type of prejudice: male vs. female) dependent variables x 2 (genderedness of languages: gendered vs. genderless) independent variables demonstrated a significant interaction $F(2, 41) = 6.1939, p = 0.0045$. On decomposing this interaction, we find that in genderless languages, there is no difference in net prejudice for male words ($M = 0.007$) or for female words ($M = 0.0048; t = 0.6057, p > 0.05$). However, among gendered languages, the net prejudice was lower against male words ($M = 0.0111$) than against female words ($M = 0.0048; t = 0.6057, p < 0.001$). Therefore, gender prejudice occurs in gendered language because of the greater association of male words (compared to female words) with positively valenced words (see Figure 1).

Figure 1. Nature of gender prejudice in the Wikipedia data set.

Please refer to Table S2 in Appendix D where similarity estimates ($S(\text{male}, \text{positive})$, $S(\text{male}, \text{negative})$, $S(\text{female}, \text{positive})$, and $S(\text{female}, \text{negative})$) across all languages are reported.

Common Crawl

We conducted the same analysis using Common Crawl word embeddings. As with the Wikipedia data set, we find that gendered languages exhibited significant gender prejudice compared to genderless languages (see Figure 2). 19 gendered languages exhibited gender prejudice (67%), however, none of the genderless languages exhibited gender prejudice, $\chi^2(1) = 17.281, p \approx 0$.

Similar to the Wikipedia analysis, the interaction of net prejudice against male or female words with gendered versus genderless languages was significant $F(2, 42) = 9.212, p < 0.001$. On decomposing this interaction, we find that in genderless languages, there is no difference in net prejudice against male words ($M = 0.0165$) or against female words ($M = 0.0145; t = 0.4744, p > 0.05$). However, among gendered languages the net prejudice was lower against male words ($M = 0.0229$) than against female words ($M = 0.0072; t = 5.504, p < 0.001$). This pattern is consistent with the Wikipedia analysis and again indicates that male words (compared to female words) had a higher net similarity to positive words in gendered but not in genderless languages (see Figure 2). Therefore, our results demonstrate that not only is gender prejudice more prevalent in gendered languages, it is caused by a higher association of male words with positive words (see Figure 3 and Figure 4).

Figure 2. Nature of the gender prejudice in the Common Crawl data set.

Figure 3. Gender prejudice in languages in Wikipedia corpora

Figure 4. Gender prejudice in languages in Common Crawl corpora

Additional tests

It is worth noting that languages can belong to different language families (e.g. Indo-European, Sino-Tibetan). Therefore, one could argue that the difference we are observing between gendered and genderless languages is not because of genderedness but because of the

language family. To address this account, we conducted an analysis using a Hierarchical Linear model (HLM) that controlled for variance in each language family, as well as across. The results indicated that genderedness still had a significant influence on net gender prejudice in both the Wikipedia and Common Crawl data (detailed analysis are reported in Appendix E).

One could also argue that there may be specific features of the region where languages are spoken that could affect gender prejudice over and above genderedness of languages. In order to address such a concern we ran our analysis controlling for geographic, cultural, and demographic covariates that have been shown to influence gender equality. The results indicate that for both Wikipedia and Common Crawl corpora genderedness significantly predicted gender prejudice for each language (detailed analysis are reported in Appendix E).

Our examination so far focused on prejudice, which is the affective component of the intergroup bias. Our findings provide some insights into the debate as to whether gendered languages display prejudice against women. However, there is another component of the intergroup bias, the cognitive component, known as stereotyping (Fiske et al., 2002; Bhatia, 2017; Todd et al., 2016). Stereotypes have been defined as specific beliefs about a group, such as descriptions of what members of a particular group look like, how they behave, or their abilities. They include cognitive representations of how members of a group are similar to one another and different from members of other groups. Hence, a gender stereotype would be an association of male or female with different semantic traits (e.g., home or work, art or science, warmth or competence).

We next examine whether genderedness of languages can affect the associations of men versus women with the traits of warmth and competence.

Gender Stereotypes

Past research has categorized gender stereotypes along two general dimensions: warmth and competence (Abele et al., 2016; Durante et al., 2013; Durante et. al. 2017; Fiske, Cuddy, & Glick, 2007; Rosenberg, Nelson, & Vivekananthan, 1968; Rudman, Greenwald, & McGhee, 2001). Although both are positive in nature, research has considered whether women might be stereotypically more associated with warmth compared to men while men are believed to be more associated with competence. Moreover, it has been demonstrated that the warmth-competence dimensions occur across cultures and regions (Durante et. al., 2017). Given the broad presence of the warmth-competence dimensions, we examined whether women were associated more with warmth and men more with competence. And importantly whether such stereotypic associations were stronger in gendered languages compared to genderless languages.

Similar to the gender prejudice analysis, this analysis utilized a lexicon of words to represent gender constructs (male and female), validated in previous research (Bolukbasi et al., 2016). Since, no single lexicon has been developed to clearly represent the warmth-competence constructs, we drew from established descriptions from Rosenberg et al. (1968), Rudman et al. (2001), Fiske et al. (2007), and Abele et al. (2016), to build a representative construct dictionary for warmth (e.g. sincere, warm, caring) and competence (e.g. capable, competent, skillful). Details of this lexicon can be found in Appendix C.

We tested for gender stereotypes in the two corpora, Wikipedia and Common Crawl, by comparing the similarity of target male and female words (e.g., he, she, her, him) to warmth words (e.g. *friendly*, *warm*, *good natured*) as well as competence words (e.g. *skillful*, *confident*, *competent*). It is important to note that when we analyzed gender prejudice, we had positively

and negatively valenced words. Hence, we were able to calculate a net similarity as described in Equation 9. Since competence and warmth are both positively valenced, creating a net difference score like $S(\text{male}, \text{warmth}) - S(\text{male}, \text{competence})$ or $S(\text{female}, \text{warmth}) - S(\text{female}, \text{competence})$ would not be meaningful. Therefore, we analyzed the competence and warmth dimensions separately.

Common Crawl.

For the construct of competence, a MANOVA revealed a significant interaction between male or female words with gendered versus genderless languages, $F(2, 42) = 13.077, p < 0.001$). Decomposing the interaction across gendered and genderless languages, we find that in genderless languages there is no difference in similarity or association of competence words with male words ($M = 0.219$) or female words ($M = 0.211; t = -0.1997, p > 0.05$). However, among gendered languages, there is a greater association of competence words with male words ($M = 0.182$) than female words ($M = 0.157; t = -4.6437, p < 0.001$) showing that men are associated more with competence than women in gendered languages (see Figure 5).

Figure 5. Association of male and female words with the construct of competence in gendered versus genderless languages in Common Crawl.

Similarly for the construct of warmth in the Common Crawl data, a MANOVA revealed a significant interaction between male or female words with gendered versus genderless languages, $F(2, 42) = 13.599, p < 0.001$. On decomposing the interaction, we find that in genderless languages, there is no difference in association with warmth words for male words ($M = 0.233$) or for female words ($M = 0.234; t = 0.0445, p > 0.05$). However, among gendered languages, male words ($M = 0.194$) are more associated with warmth words than female words ($M = 0.182; t = -2.1421, p = 0.0367$; see Figure 6). This means that male words are associated more with warmth than female words in gendered languages.

Figure 6. Association of warmth with male versus female words in Common Crawl.

Across both warmth and competence dimensions we see an interesting pattern. Males are more associated with both competence and warmth in gendered languages, which could be because both are positively valenced.

Wikipedia.

For the construct of competence, results were similar to those from Common Crawl. A MANOVA revealed a significant interaction $F(2, 41) = 3.8329, p = 0.0298$. On decomposition, we find that there is no difference in association with competence words for male words ($M = 0.238$) or for female words ($M = 0.229; t = -0.1527, p > 0.05$) in genderless languages. However, among gendered languages, male words were more associated with

competence words ($M = 0.201$) than female words ($M = 0.184$; $t = -2.2439$, $p = 0.0291$; see Figure 7).

Figure 7. Association of competence with male and female words in Wikipedia data in gendered versus genderless languages.

For the construct of warmth, a MANOVA revealed a significant interaction $F(2, 41) = 5.1579$, $p = 0.01$. Specifically on decomposing the interaction, we find that there is no difference in association with warmth for male words ($M = 0.254$) or for female words ($M = 0.258$; $t = 0.0545$, $p > 0.05$) in genderless languages. Among gendered languages, there was no significant difference in association with warmth with male words ($M = 0.217$) or for female words ($M = 0.213$; $t = -0.4966$, $p > 0.05$; see Figure 8).

Figure 8. Association of male and female words with the construct of warmth in gendered versus genderless languages in Wikipedia.

Similar to the analysis we conducted with gender prejudice, we examine the influence of gendered versus genderless languages by controlling for language families in a HLM model. As in the previous analysis of prejudice, we see a significant influence of genderedness when the influence of individual language families are considered. We present the details of the analysis in Appendix F.

Discussion

Our findings contribute to the debate of whether the genderedness of language is associated with gender prejudice. One stream of research suggests that the genderedness of

language is associated with gender prejudice while another stream of research suggests that language merely reflects the beliefs of a culture and cannot be associated with gender prejudice. If the latter argument were to hold true, we should have seen the prevalence of gender prejudice equally across languages. Instead, in our examination of large text corpora of 45 world languages, we find that gender prejudice appears more in gendered rather than genderless languages.

Our findings indicate that when gender is made more salient through the genderedness of language, it results in certain associations of that gender being reiterated more frequently. Unfortunately, we find that associations of men tend to be more positive (i.e., there is a closer association of male words with positive words) than women. Moreover, past research has examined gender prejudice in English (Caliskan et al., 2017) and across time (Garg et al., 2018), suggesting that languages can effectively be used to test for gender prejudice. However, to date, this approach has not been used comparatively across a large group of common world languages. Using languages that span continents and cultures, we find evidence that gender prejudice exists more in gendered languages.

Second, instead of testing for gender prejudice through self-reports or direct questioning, we tested for gender prejudice using millions of pages of text data. We used Wikipedia and Common Crawl corpora that contains text (e.g., facts, thoughts, opinions) voluntarily produced by millions of people. By measuring the association of words voluntarily produced across 45 different world languages when people were not specifically being queried about gender prejudice, gives us a glimpse into how prevalent gender prejudice is. We demonstrate that gender prejudice exists cross-culturally. We also show, in line with past

research, albeit using more languages and using text data, that prejudice present in language has the potential to influence future thoughts. Wikipedia and Common Crawl are not just repositories of people's thought but are used by them to form future opinions and hence, can result in perpetuation of gender prejudice in gendered languages.

Third, knowing the nature of the prejudice can help in developing successful interventions. Therefore, it is important to know whether gender prejudice exists in certain languages but it is equally important to know whether the prejudice occurs because of a greater prejudice against women or a lesser prejudice against men. Our results across both data sets indicate that the gender prejudice is driven more by the greater association of positive words with men, which would result in their being evaluated favorably, in many significant decisions. For instance, two equally qualified individuals, one female and one male, may be treated differently just because of their gender. As mentioned earlier, such a finding is troubling in light of the "Pollyanna" effect (Boucher & Osgood, 1969; Dodd et al., 2015), in which individuals have a higher tendency to use positive words, because it shows that the association of positive words with women is weaker. Figure 3 and Figure 4 demonstrate that the positive-negative association of men and women varies significantly across languages. Understanding the nature of this association will be critical for future efforts and policies focused on eliminating gender prejudice around the world.

Limitations and Future Research

In our research we tested for the presence of gender stereotypes using the dimensions of warmth and competence. Our findings indicate that men are associated more with both warmth and competence than women in gendered languages, despite the fact that past

research has shown that women are stereotypically associated with warmth. Moreover, we do not find stereotypic associations in genderless languages. Our conjecture is that such a pattern of results could occur because of the following. It is possible that what we are observing in the text corpora is distinct from commonly studied psychological stereotyping and is instead a specific construct that is encoded linguistically by which males are consistently associated with positive attributes. That is, although women are typically reported to be more closely associated with warmth when such responses are directly elicited from participants, language usage shows a different relationship. Language usage might be denigrating women such that positively valenced words are more strongly associated with men than women. In this scenario, our analysis is revealing a linguistic pattern of association which is stronger in gendered languages. Future research could examine the association of men with positive constructs, generally, across different languages and cultures.

We examined only warmth and competence but other stereotypic dimensions such as weakness versus power or home versus career that have been stereotypically associated with one gender can also be examined. Future research can also consider other dimensions such as perceptions of emotional versus rational decision-making and test whether gendered versus genderless languages show a differing pattern of association.

It is also worth noting that we find evidence consistent with predictions made in prior work on linguistic relativity. However, since we are using associations to infer relationship, our evidence does not directly test for the tenet of linguistic relativity: that language can affect cognitions. Future research can consider other extra-linguistic perception to cast additional light on the language-shapes-cognition versus language-reflects-cognition debate. The findings of this

research focus on written text. Given the prevalence of spoken along with written words, future research should consider whether genderedness of language affects prejudice or stereotype in spoken language.

On a broader level, social psychology research can use the method employed in this manuscript in various ways. First, the large text corpora of Wikipedia and Common Crawl can be used to replicate many existing findings in social psychology. As an example, one could examine whether promotion-focused language does indeed occur more with aspirational words while prevention-focused language occurs more with responsibility words. This would help establish robustness of existing findings in an ecologically valid data set in which people were not directly responding to a question but were instead providing their opinion on a diverse set of topics. Second, other text corpora exist that range across time such as news articles or twitter feeds. These can be used to gain novel insights into how events occurring across time are influencing different psychological constructs. For example, such a method could be used to chart the rise of populist discourse and associated changes in the perception of minorities or women. Finally, word embedding methods represent just one type of text analysis tool capable of modeling semantic associations. Hence, word embeddings are best suited to testing for the association between psychological constructs. However, simpler text analysis methods such as topic modelling (e.g., Latent Dirichlet Allocation; LDA) that rely on word frequency can also be very useful to examine what topics are most likely in a document. For instance, LDA analysis of right and left leaning media outlets might reveal a difference in topics popular on such venues. Subsequent word embedding analysis could offer insight into the association these topics might have with specific individuals, events, or groups. In this way, word embedding analysis becomes

an important part of an ensemble of methods that together provide deeper insight into, greater generalizability of, and a more robust understanding of psychological phenomena.

References

- Abele, A. E., Hauke, N., Peters, K., Louvet, E., Szymkow, A., & Duan, Y. (2016). Facets of the fundamental content dimensions: Agency with competence and assertiveness—communion with warmth and morality. *Frontiers in Psychology*, 7, 1-17.
- Allport, G. W. (1954). *The nature of prejudice*. Cambridge, MA: Addison-Wesley Publishing Company.
- Banaji, M. R., Hardin, C., & Rothman, A. J. (1993). Implicit stereotyping in person judgment. *Journal of Personality and Social Psychology*, 65(2), 272–281.
- Ben-David, A., & Matamoros-Fernández, A. (2016). Hate speech and covert discrimination on social media: Monitoring the Facebook pages of extreme-right political parties in Spain. *International Journal of Communication*, 10, 1167–1193.
- Bhatia, S. (2017). The semantic representation of prejudice and stereotypes. *Cognition*, 164, 46-60.
- Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2016). Enriching word vectors with subword information. *ArXiv:1607.04606 [Cs]*. Retrieved from <http://arxiv.org/abs/1607.04606>
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* (Vol. 29, pp. 4349–4357). Retrieved from <http://papers.nips.cc/paper/6228-man-is-to-computer-programmer-as-woman-is-to-homemaker-debiasing-word-embeddings.pdf>
- Boroditsky, L., & Schmidt, L. A. (2000). Sex, syntax, and semantics. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 22, 1–6.

- Boucher, J., & Osgood, C. E. (1969). The Pollyanna hypothesis. *Journal of Verbal Learning and Verbal Behavior*, 8(1), 1–8.
- Boutonnet, B., Athanasopoulos, P., & Thierry, G. (2012). Unconscious effects of grammatical gender during object categorisation. *Brain Research*, 1479, 72–79.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago, IL: University of Chicago Press.
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186.
- Carnaghi, A., & Maass, A. (2007). In-group and out-group perspectives in the use of derogatory group labels: Gay versus fag. *Journal of Language and Social Psychology*, 26(2), 142–156.
- Casasanto, D. (2008). Who’s Afraid of the Big Bad Whorf? Crosslinguistic differences in temporal language and thought. *Language Learning*, 58(s1), 63–79.
- Casasanto, D. (2016). Linguistic relativity. In N. Riemer (Ed.), *The Routledge Handbook of Semantics* (pp. 158–174). New York, NY: Routledge.
- Cela-Conde, C. J., & Marty, G. (1998). Noam Chomsky’s minimalist program and the philosophy of mind. An interview. *Syntax*, 1(1), 19–36.
- Collins, K. A., & Clément, R. (2012). Language and prejudice: Direct and moderated effects. *Journal of Language and Social Psychology*, 31(4), 376–396.
- Corbett, G. G. (1991). *Gender*. New York: Cambridge University Press.
- Cralley, E. L., & Ruscher, J. B. (2005). Lady, girl, female, or woman: Sexism and cognitive busyness predict use of gender-biased nouns. *Journal of Language and Social Psychology*, 24(3), 300–314.
- Deutscher, G. (2010). *Through the language glass: Why the world looks different in other languages*. New York, NY: Metropolitan Books.

- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56(1), 5–18.
- Dodds, P. S., Clark, E. M., Desu, S., Frank, M. R., Reagan, A. J., Williams, J. R.,...Danforth, C. M. (2015). Human language reveals a universal positivity bias. *Proceedings of the National Academy of Sciences*, 112(8), 2389-2394.
- Durante, F., Fiske, S. T., Gelfand, M. J., Crippa, F., Suttora, C., Stillwell, A.,...Teymoori, A. (2017). Ambivalent stereotypes link to peace, conflict, and inequality across 38 nations. *Proceedings of the National Academy of Sciences*, 114(4), 669-674.
- Durante, F., Fiske, S. T., Kervyn, N., Cuddy, A. J. C., Akande, A., Adetoun, B. E. (2013). Nations' income inequality predicts ambivalence in stereotype content: How societies mind the gap. *British Journal of Psychology*, 52(4), 726-746.
- Evans, N., & Levinson, S. C. (2009). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, 32(05), 429.
- Filipacchi, A. (2013, April 24). Wikipedia's sexism toward female novelists. *The New York Times*. Retrieved from <https://www.nytimes.com/2013/04/28/opinion/sunday/wikipedias-sexism-toward-female-novelists.html>
- Firth, J. R. (1957). A synopsis of linguistic theory, 1930-1955. *Studies in Linguistic Analysis*, 1957, 1–32.
- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77–83.
- Fiske, S. T., Cuddy, A. J. C., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, 82(6), 878-902.
- Fodor, J. A. (1975). *The Language of Thought*. New York, NY: Crowell.

- Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16), E3635–E3644.
- Glott, R., Schmidt, P., & Ghosh, R. (2010). *Wikipedia survey: Overview of results* (pp. 1–11). Retrieved from http://www.ris.org/uploadi/editor/1305050082Wikipedia_Overview_15March2010-FINAL.pdf
- Grave, E., Bojanowski, P., Gupta, P., Joulin, A., & Mikolov, T. (2018). Learning word vectors for 157 languages. *ArXiv:1802.06893 [Cs]*. Retrieved from <http://arxiv.org/abs/1802.06893>
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4–27.
- Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Mellott, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review*, 109(1), 3–25.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74(6), 1464–1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85(2), 197–216.
- Gruwell, L. (2015). Wikipedia's politics of exclusion: Gender, epistemology, and feminist rhetorical (in)action. *Computers and Composition*, 37, 117–131.
- Gumperz, J. J., & Levinson, S. C. (1991). Rethinking linguistic relativity. *Current Anthropology*, 32(5), 613–623.
- Harris, Z. S. (1954). Distributional structure. *Word*, 10(2–3), 146–162.

- Haspelmath, M. (2005). *World Atlas of Language Structures*. Oxford, UK: Oxford University Press.
- Hellinger, M. (2001). English: Gender in a global language. In M. Hellinger & H. Bußmann (Eds.), *Gender across languages: The linguistic representation of women and men* (pp. 105-113). Amsterdam: J Benjamins Publishing Company.
- Hellinger, M., & Bußmann, H. (Eds.). (2001). *Gender across languages: The linguistic representation of women and men*. Philadelphia, PA: J. Benjamins Publishing Company.
- Hill, B. M., & Shaw, A. (2013). The Wikipedia gender gap revisited: Characterizing survey response bias with propensity score estimation. *PLOS ONE*, 8(6), e65782.
- Hofstätter, P. R. (1963). Über sprachliche Bestimmungsleistungen: Das Problem des grammatikalischen Geschlechts von Sonne und Mond. [On linguistic performances of connotation: The problem of the grammatical gender of sun and moon.]. *Zeitschrift Für Experimentelle Und Angewandte Psychologie*, 10(1), 91–108.
- Hunt, E., & Banaji, M. R. (1988). The Whorfian hypothesis revisited: A cognitive science view of linguistic and cultural effects on thought. In J. W. Berry, S. H. Irvine, & E. B. Hunt (Eds.), *Indigenous Cognition: Functioning in Cultural Context* (pp. 57–84). Dordrecht, Netherlands: Martinus Nijhoff Publishers.
- Hunt, Earl, & Agnoli, F. (1991). The Whorfian hypothesis: A cognitive psychology perspective. *Psychological Review*, 98(3), 377–389.
- Iliev, R., Hoover, J., Dehghani, M., & Axelrod, R. (2016). Linguistic positivity in historical texts reflects dynamic environmental and psychological factors. *Proceedings of the National Academy of Sciences*, 113(49), E7871–E7879.
- Jakobson, R. (1959). On the linguistic aspects of translation. In R. A. Brower (Ed.), *On Translation* (pp. 232–239). Cambridge, MA: Harvard University Press.

- January, D., & Kako, E. (2007). Re-evaluating evidence for linguistic relativity: Reply to Boroditsky (2001). *Cognition*, 104(2), 417–426.
- Joulin, A., Grave, E., Bojanowski, P., & Mikolov, T. (2016). Bag of tricks for efficient text classification. *ArXiv:1607.01759 [Cs]*. Retrieved from <http://arxiv.org/abs/1607.01759>
- Jurafsky, D., & Martin, J. H. (2019). *Speech and language processing* (3rd ed.). Retrieved from <https://web.stanford.edu/~jurafsky/slp3/>
- Kay, P., & Kempton, W. (1984). What is the Sapir-Whorf hypothesis? *American Anthropologist, New Series*, 86(1), 65–79.
- Konishi, T. (1993). The semantics of grammatical gender: A cross-cultural study. *Journal of Psycholinguistic Research*, 22(5), 519–534.
- Liu, Q. B., & Karahanna, E. (2017). The dark side of reviews: The swaying effects of online product reviews on attribute preference construction. *MIS Quarterly*, 41(2), 427-A6.
- Maki, W. S. (1979). Discrimination learning without short-term memory: Dissociation of memory processes in pigeons. *Science*, 204(4388), 83-85.
- Maki, R. H., Maki, W. S., & Marsh, L. G. (1977). Generalization of above/below, right/left effects to compass directions and natural stimuli. *Bulletin of the Psychometric Society*, 10(4), 307-310.
- Matthes, J., Prieler, M., & Adam, K. (2016). Gender-role portrayals in television advertising across the globe. *Sex Roles*, 75(7–8), 314–327.
- McWhorter, J. H. (2016). *The language hoax: Why the world looks the same in any language*. New York, NY: Oxford University Press.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *ArXiv:1301.3781 [Cs]*. Retrieved from <http://arxiv.org/abs/1301.3781>

- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26(NIPS 2013), 9.
- Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., Joulin, A. (2017). Advances in pre-training distributed word representations. *ArXiv:1712.09405 [Cs]*. Retrieved from <https://arxiv.org/abs/1712.09405>
- Miller, M. C. (2013). *Wronged by empire: Post-imperial ideology and foreign policy in India and China*. Stanford, CA: Stanford University Press.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice*, 6(1), 101–115.
- Nosek, B. A., Smyth, F. L., Sriram, N., Lindner, N. M., Devos, T., Ayala, A., ... Kesebir, S. (2009). National differences in gender: Science stereotypes predict national sex differences in science and math achievement. *Proceedings of the National Academy of Sciences*, 106(26), 10593-10597.
- Oh, K. J. (2003). *Language, cognition, and development: Motion events in English and Korean*. University of California, Berkeley.
- Pew Research Center. (2014). *Religious diversity index scores by country*. Retrieved from <http://www.pewforum.org/2014/04/04/religious-diversity-index-scores-by-country/>
- Pfeil, U., Zaphiris, P., & Ang, C. S. (2006). Cultural differences in collaborative authoring of Wikipedia. *Journal of Computer-Mediated Communication*, 12(1), 88–113.
- Pinker, S. (1994). *The Language Instinct*. New York: W. Morrow and Co.
- Prewitt-Freilino, J. L., Caswell, T. A., & Laakso, E. K. (2012). The gendering of language: A comparison of gender equality in countries with gendered, natural gender, and genderless languages. *Sex Roles*, 66(3–4), 268–281.

- Rosenberg, S., Nelson, C., & Vivekananthan, P. S. (1968). A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology*, 9(4), 283–294.
- Rudman, L. A., Greenwald, A. G., & McGhee, D. E. (2001). Implicit self-concept and evaluative implicit gender stereotypes: Self and ingroup share desirable traits. *Personality and Social Psychology Bulletin*, 27(9), 1164–1178.
- Santacreu-Vasut, E., Shoham, A., & Gay, V. (2013). Do female/male distinctions in language matter? Evidence from gender political quotas. *Applied Economics Letters*, 20(5), 495–498.
- Saussure, F. de (1959). *Course in general linguistics*. (W. Baskin, Trans.). New York, NY: McGraw-Hill. (Original work published 1916).
- Schonhardt-Bailey, C. (2006). *From the corn laws to free trade: Interests, ideas, and institutions in historical perspective*. Cambridge, MA: MIT Press.
- Schrage, M. (2014, January 29). Big data's dangerous new era of discrimination. Retrieved from <https://hbr.org/2014/01/big-datas-dangerous-new-era-of-discrimination>
- Sera, M. D., Berge, C. A. H., & del Castillo-Pintado, J. (1994). Grammatical and conceptual forces in the attribution of gender by English and Spanish speakers. *Cognitive Development*, 9(3), 261–292.
- Sera, M. D., Elieff, C., Forbes, J., Burch, M. C., Rodríguez, W., & Dubois, D. P. (2002). When language affects cognition and when it does not: An analysis of grammatical gender and classification. *Journal of Experimental Psychology: General*, 131(3), 377–397.
- Shane-Simpson, C., & Gillespie-Lynch, K. (2017). Examining potential mechanisms underlying the Wikipedia gender gap through a collaborative editing task. *Computers in Human Behavior*, 66, 312–328.

- Simmons, G. F., & Fennig, C. D. (2018). *Ethnologue: Languages of the World* (21st ed.). Retrieved from <https://www.ethnologue.com/>
- Stahlberg, D., Braun, F., Irmen, L., & Sczesny, S. (2007). Representation of the sexes in language. In K. Fiedler (Ed.), *Social Communication* (pp. 163–187). New York, NY: Psychology Press.
- Sterling, J., Jost, J. T., & Bonneau, R. (2020). Political psycholinguistics: A comprehensive analysis of the language habits of liberal and conservative social media users. *Journal of Personality and Social Psychology*.
- Stewart, B. M., & Zhukov, Y. M. (2009). Use of force and civil–military relations in Russia: An automated content analysis. *Small Wars & Insurgencies*, 20(2), 319–343.
- Stubbs, M. (1996). *Text and corpus analysis: Computer-assisted studies of language and culture*. Oxford, OX, UK ; Cambridge, MA: Blackwell Publishers.
- Sutton, R. M. (2010). The creative power of language in social cognition and intergroup relations. In G. Howard, S. Reid, & J. Hardwood (Eds.), *Dynamics of intergroup communication* (pp. 105–115). New York, NY: Peter Lang.
- Tajick, S. (2018). *STC climate index*. Retrieved from <http://globalresidenceindex.com/hnwi-index/climate-index/>
- Thelwall, M., Wilkinson, D., & Uppal, S. (2010). Data mining emotion in social network communication: Gender differences in MySpace. *Journal of the American Society for Information Science and Technology*, 61(1), 190–199.
- Tirunillai, S., & Tellis, G. J. (2014). Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent dirichlet allocation. *Journal of Marketing Research*, 51(4), 463–479.

- Todd, A. R., Simpson, A. J., Thiem, K. C., & Neel, R. (2016). The generalization of implicit racial bias to young black boys: Automatic stereotyping or automatic prejudice? *Social Cognition*, 34(4), 306-323.
- Transparency International. (2016). *Corruption perceptions index*. Retrieved from https://www.transparency.org/news/feature/corruption_perceptions_index_2016
- United Nations Development Programme. (2017). *Human Development Report 2016: Human Development for Everyone*.
- van der Velde, L., Tyrowicz, J., & Siwinska, J. (2015). Language and (the estimates of) the gender wage gap. *Economics Letters*, 136, 165–170.
- Vylomova, E., Cohn, T., He, X., & Haffari, G. (2016). Word representation models for morphologically rich languages in neural machine translation. *ArXiv:1606.04217 [Cs]*. Retrieved from <http://arxiv.org/abs/1606.04217>
- Wagner, C., Garcia, D., Jadidi, M., & Strohmaier, M. (2015). It's a man's Wikipedia? Assessing gender inequality in an online encyclopedia. *ArXiv:1501.06307 [Cs]*. Retrieved from <http://arxiv.org/abs/1501.06307>
- Wagner, C., Graells-Garrido, E., Garcia, D., & Menczer, F. (2016). Women through the glass ceiling: gender asymmetries in Wikipedia. *EPJ Data Science*, 5(1), 5.
- Wasserman, B. D., & Weseley, A. J. (2009). ¿Qué? Quoi? Do languages with grammatical gender promote sexist attitudes? *Sex Roles*, 61(9–10), 634–643.
- Wikimedia Foundation. (2018). Wikimedia traffic analysis report: Page views per Wikipedia language. Retrieved from https://stats.wikimedia.org/archive/squid_reports/2018-05/SquidReportPageViewsPerLanguageBreakdown.htm
- Wikipedia. (2018a). List of Wikipedias. In *Wikipedia*. Retrieved from https://en.wikipedia.org/w/index.php?title=List_of_Wikipedias&oldid=860318950

- Wikipedia. (2018b). Wikipedia article depth. Retrieved from
https://meta.wikimedia.org/wiki/Wikipedia_article_depth
- Williams, J. E., Satterwhite, R. C., & Best, D. L. (1999). Pancultural gender stereotypes revisited: The five factor model. *Sex Roles*, 40(7–8), 513–525.
- Wittgenstein, L. (1973) *Philosophical investigations* (3rd ed.; G.E.M. Anscombe, Trans.). Oxford, UK: Pearson.
- World Economic Forum. (2017). *The Global Gender Gap Report: 2017*. Geneva: World Economic Forum.
- Zerbe, W. J., & Paulhus, D. L. (1987). Socially desirable responding in organizational behavior: A reconception. *Academy of Management Review*, 12(2), 250–264.

Footnotes

¹ For the purposes of this research, we classify languages based on the assignment of gender markers to nouns, verbs, and adjectives. As such, English is classified as a genderless language, which is consistent with past literature that utilizes a binary taxonomy (Hellinger, 2001).

² <https://github.com/facebookresearch/fastText/blob/master/docs/pretrained-vectors.md> (Wikipedia)

³ <https://github.com/facebookresearch/fastText/blob/master/docs/crawl-vectors.md> (Common Crawl)

⁴ While it is difficult to assess exactly why these languages failed validation, it could be due to failed or imprecise translations. Future work could address this challenge by developing specific lexicons for these languages which are drawn from, translated, and validated by native speakers of the language.

⁵ A power analysis confirmed that the resulting group sizes would be sufficient to identify medium-sized effects, if such effects existed.

SUPPLEMENTAL MATERIAL

Appendix A: Validation

A validation method was implemented to verify the ability of the word embedding algorithm to correctly identify semantic associations in each language. A validation lexicon was constructed consisting of generic positive nouns (e.g., beauty, fun, hero) and generic negative nouns (e.g., casualty, destruction, emergency). An attempt was made to select words that were unambiguous. For this reason, validation lexicons commonly for testing prejudice (Greenwald, McGhee, & Schwartz, 1998; Greenwald, Nosek, & Banaji, 2003) and past word embedding research (e.g., Caliskan, Bryson, & Narayanan, 2017), for example lists of insects and weapons, were avoided. The words used in the validation lexicon are listed in Appendix C.

Once the validation lexicon was translated, word embeddings were obtained for each word. Next, a search for exact duplicates was performed. In some languages, for example, the word for man and men is the same. In other languages, semantically similar words like boy and girl may be combined, resulting in a single word for both concepts, for example, child. Both of these scenarios have a confounding effect on the results of the analysis and as such, duplicates were removed.

Association was then calculated using cosine similarity. We hypothesized that a strong validation of the method would result in a stronger association between known positive nouns and positively valenced words on the one hand and known negative nouns and negatively valenced words on the other. Moreover, these observations should be statistically significant. Formally, a strong test of validation would satisfy both propositions:

$$\cos(\text{noun}_+, \text{valence}_+) - \cos(\text{noun}_+, \text{valence}_-) > 0, p < 0.05$$

$$\cos(\text{noun}_-, \text{valence}_-) - \cos(\text{noun}_-, \text{valence}_+) > 0, p < 0.05$$

Five languages, Cebuano, Japanese, Latin, Vietnamese, and Welsh failed validation and were removed from the Wikipedia data set. Four languages, Cebuano, Hebrew, Latin, and Vietnamese failed validation and were removed from the Common Crawl data set. The remaining language data sets outlined in Table S1 were analyzed utilizing the method that follows.

Table S1

Languages included in analysis after translation, classification, and validation.

Language	Genderedness (Source)	Prejudice Analysis
Arabic	Gendered (Haspelmath, 2005)	Wikipedia ^{**} , Common Crawl [*]
Armenian	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Azerbaijani	Genderless (Simmons & Fennig, 2018)	Wikipedia, Common Crawl
Basque	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Belarusian	Gendered (Simmons & Fennig, 2018)	Wikipedia ^{***} , Common Crawl ^{***}
Bulgarian	Gendered (Simmons & Fennig, 2018)	Wikipedia ^{***} , Common Crawl ^{***}
Catalan	Gendered (Simmons & Fennig, 2018)	Wikipedia, Common Crawl
Chinese (Mandarin)	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Czech	Gendered (Hellinger & Bußmann, 2001)	Wikipedia ^{***} , Common Crawl ^{***}
Danish	Gendered (Hellinger & Bußmann, 2001)	Wikipedia [*] , Common Crawl
Dutch	Gendered (Hellinger & Bußmann, 2001)	Wikipedia ^{**} , Common Crawl ^{**}
English	Genderless (Hellinger & Bußmann, 2001)	Wikipedia [*] , Common Crawl
Esperanto	Genderless (Hellinger & Bußmann, 2001)	Wikipedia ^{***} , Common Crawl
Estonian	Genderless	Wikipedia [*] . Common Crawl

Finnish	(Hellinger & Bußmann, 2001) Genderless	Wikipedia, Common Crawl
French	(Haspelmath, 2005) Gendered	Wikipedia ^{***} , Common Crawl ^{***}
Galician	Gendered (Simmons & Fennig, 2018)	Wikipedia, Common Crawl
Georgian	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
German	Gendered (Haspelmath, 2005)	Wikipedia, Common Crawl
Greek	Gendered (Haspelmath, 2005)	Wikipedia ^{**} , Common Crawl ^{***}
Hebrew	Gendered (Corbett, 2013)	Wikipedia [*]
Hindi	Gendered (Haspelmath, 2005)	Wikipedia ^{**} , Common Crawl ^{***}
Hungarian	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Indonesian	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Italian	Gendered (Hellinger & Bußmann, 2001)	Wikipedia, Common Crawl ^{***}
Japanese	Genderless (Hellinger & Bußmann, 2001)	Common Crawl
Korean	Genderless (Simmons & Fennig, 2018)	Wikipedia [*] , Common Crawl
Lithuanian	Gendered (Simmons & Fennig, 2018)	Wikipedia, Common Crawl
Malay	Genderless (Simmons & Fennig, 2018)	Wikipedia, Common Crawl
Norwegian	Gendered (Hellinger & Bußmann, 2001)	Wikipedia, Common Crawl
Persian	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Polish	Gendered (Hellinger & Bußmann, 2001)	Wikipedia ^{**} , Common Crawl ^{***}
Portuguese	Gendered (Hellinger & Bußmann, 2001)	Wikipedia, Common Crawl ^{***}
Romanian	Gendered (Hellinger & Bußmann, 2001)	Wikipedia ^{***} , Common Crawl ^{***}
Russian	Gendered (Haspelmath, 2005)	Wikipedia ^{***} , Common Crawl ^{***}

Serbian	Gendered (Hellinger & Bußmann, 2001)	Wikipedia ^{***} , Common Crawl ^{***}
Slovak	Gendered (Simmons & Fennig, 2018)	Wikipedia ^{***} , Common Crawl ^{***}
Slovenian	Gendered (Hellinger & Bußmann, 2001)	Wikipedia ^{***} , Common Crawl ^{***}
Spanish	Gendered (Haspelmath, 2005)	Wikipedia [*] , Common Crawl [*]
Swedish	Gendered (Hellinger & Bußmann, 2001)	Wikipedia, Common Crawl
Tamil	Gendered (Haspelmath, 2005)	Wikipedia, Common Crawl [*]
Thai	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Turkish	Genderless (Haspelmath, 2005)	Wikipedia, Common Crawl
Ukrainian	Gendered (Haspelmath, 2005)	Wikipedia, Common Crawl ^{***}
Urdu	Gendered (Simmons & Fennig, 2018)	Wikipedia, Common Crawl
Welsh	Gendered (Hellinger & Bußmann, 2001)	Common Crawl

^{*} $p < 0.05$, ^{**} $p < 0.01$, ^{***} $p < 0.001$.

Appendix B: Testing for Gender Prejudice

To identify gender prejudice in text corpora, we adapted a method used extensively in social psychology to study the presence and influence of prejudice and stereotypes (Greenwald et al., 1998, 2003). It suggests that if a target word is more strongly associated with a positively than a negatively valenced word in a person's memory, she will respond faster to the pairing of the target word with a positive than a negative word. After employing the *fastText* algorithm on the text corpora (Bojanowski, Grave, Joulin, & Mikolov, 2016; Grave, Bojanowski, Gupta, Joulin, & Mikolov, 2018) a word is represented as a point in vector space with 300 dimensions. Given these dimensions, we can then find the level of similarity between two points (words), as suggested by Caliskan et al. (2017). The resulting score captures the strength of

association between words. Therefore, we use semantic similarity between words as a measure of association.

Formally, let X represent a set of female words and Y a set of male words. Let A be the set containing positively valenced words and B the set containing negatively valenced words. Cosine similarity is used to measure the similarity between two words where each word has an associated vector. Thus, for target word $x \in X$ and a valence word $a \in A$, $\text{Cos}(x, a) = \frac{x \cdot a}{\|x\| \|a\|}$ and $x \cdot a$ is the dot product of vectors x and a . Hence, we can calculate the similarity between x and each of the members of A such that the $\text{mean}_{a \in A} \text{Cos}(x, a)$, denotes the average similarity between x and all the positive words. Similarly, the estimated average similarity of x to all negative words in set B would be given by $\text{mean}_{b \in B} \text{Cos}(x, b)$.

The difference between $\text{mean}_{a \in A} \text{Cos}(x, a)$ and $\text{mean}_{b \in B} \text{Cos}(x, b)$ provides a single measure of relative similarity between gender word x and all positive and negative words. Intuitively, such a difference will provide the net similarity to positive words. If $\text{mean}_{a \in A} \text{Cos}(x, a) - \text{mean}_{b \in B} \text{Cos}(x, b) > 0$, then it demonstrates that x is more closely associated with positive words than negative words. Conversely, if $\text{mean}_{a \in A} \text{Cos}(x, a) - \text{mean}_{b \in B} \text{Cos}(x, b) < 0$, it demonstrates x is more closely associated with negative words. If $S(x, A, B) = \text{mean}_{a \in A} \text{Cos}(x, a) - \text{mean}_{b \in B} \text{Cos}(x, b)$, then $\sum_{x \in X} S(x, A, B)$ captures the sum of the net similarities for all members of x . Similarly, $\sum_{y \in Y} S(y, A, B)$ captures the sum of the net similarities for all members of Y . The main measure of prejudice would then be:

$$S(X, Y, A, B) = \sum_{x \in X} S(x, A, B) - \sum_{y \in Y} S(y, A, B) \quad (\text{Equation 1})$$

A positive value indicates the words in X are more closely associated with positive words than the words in set Y . However, a negative value indicates words in Y are more closely associated with positive words than those in X .

Permutation test. However, $S(X, Y, A, B)$ is a single measure of relative similarity and could thus be the result of random error. To estimate the probability that the obtained similarity was not observed due to random allocations in X and Y , words are shuffled between groups. Put another way, words from the original sets X and Y are randomly shuffled between groups to create X^* and Y^* for the purposes of calculating $S(X^*, Y^*, A, B)$. If the null hypothesis is true in such a non-parametric permutation test, the result of $S(X^*, Y^*, A, B)$ should be no different than $S(X, Y, A, B)$.

Formally, if (X_i, Y_i) represents the potential random shuffling of words in X and Y , then the probability of the observed score being the result of a random process will be:

$$\text{Probability} = \frac{\text{Number of times } (|S(X_i, Y_i, A, B)| > |S(X, Y, A, B)|)}{\text{Number of all possible shuffles of } X \text{ and } Y} \quad (\text{Equation 2})$$

Appendix C: Gender and Attribute Words

The Gender-specific lexicon was adapted from Bolukbasi, Change, Zou, Saligrama, and Kalai (2016) and consists of 216 items separated by gender.

1. *Male words*: he, his, him, man, men, spokesman, himself, son, father, chairman, husband, guy, boy, boys, brother, male, brothers, dad, sons, boyfriend, king, businessman, grandfather, uncle, males, congressman, grandson, bull, businessmen, nephew, fathers, lads, lion, gentleman, fraternity, bachelor, bulls, husbands, prince, colt, salesman, dude, beard, councilman, gentleman, stepfather, monks, lad, sperm, testosterone, nephews, daddy, fiancé, kings, dads, sir, stud, lions, czar, countrymen, penis, bloke, spokesmen,

suitors, monastery, brethren, prostate, schoolboy, brotherhood, stepson, uncles, monks, viagra, paternity, suitor, macho, statesman, fathered, hubby, blokes, dudes, strongman, grandsons, studs, godfather, boyhood, baritone, grandpa, boyfriends, countryman, stallion, semen, fella, chap, salesmen, widower, beau, beards, handyman, horsemen, fatherhood, princes, colts, fraternities, pa, fellas, councilmen, barbershop, fraternal, prostate cancer, younger brother, ex boyfriend, twin brother, eldest son, estranged husband, elder brother

2. *Female words*: her, she, women, woman, wife, mother, daughter, girls, girl, spokeswoman, female, sister, herself, actress, mom, girlfriend, daughters, lady, sisters, mothers, grandmother, deer, ladies, queen, wives, widow, bride, females, aunt, lesbian, chairwoman, moms, maiden, granddaughter, niece, hers, princess, lesbians, actresses, maiden, mare, fiancée, waitress, maternal, heroine, nieces, girlfriends, mistress, womb, grandma, maternity, estrogen, widows, diva, nuns, nun, brides, housewife, menopause, motherhood, stepmother, hostess, fillies, congresswoman, witch, sorority, businesswoman, gal, schoolgirl, goddess, stepdaughter, uterus, mama, hens, hen, mommy, grandmothers, feminism, heiress, queens, witches, aunts, granddaughters, convent, vagina, maids, gals, housewives, obstetrics, councilwoman, matriarch, ma, dowry, ballerina, ex girlfriend, estranged wife, ovarian cancer, teenage girl, twin sister

The positive and negative (pleasant and unpleasant) lexicon was adapted from Caliskan et al. (2017).

Positive words: caress, freedom, health, love, peace, cheer, friend, heaven, loyal, pleasure, diamond, gentle, honest, lucky, rainbow, diploma, gift, honor, miracle, sunrise, family, happy, laughter, paradise, vacation

Negative words: abuse, crash, filth, murder, sickness, accident, death, grief, poison, stink, assault, disaster, hatred, pollute, tragedy, divorce, jail, poverty, ugly, cancer, kill, rotten, vomit, agony, prison

The warmth and competence lexicon was adapted from Rosenberg, Nelson, and Vivekananthan (1968), Rudman, Greenwald, and McGhee (2001), Fiske, Cuddy, and Glick (2007), and Abele et al. (2016).

- *Warmth words:* friendly, good natured, sincere, warm, tolerant, caring, just, fair, considerate, trustworthy, reliable, appreciative, honest, sentimental, happy, humorous, sociable, supportive, nurture, nice, gentle, accepting, pleasant, loving, open
- *Competence words:* capable, competent, confident, skillful, independent, competitive, intelligent, clever, efficient, decisive, persistent, determined, industrious, imaginative, potent, commanding, assertive, successful, bold, dynamic, logical, triumphant, dominant, respectable, accomplished

The positive and negative validation lexicon was constructed by the authors.

- *Positive validation words:* accomplishment, beauty, celebration, comfort, delight, encouragement, fun, generous, genius, hero, holiday, humor, innovation, kiss, music, nature, relief, safety, satisfaction, smile, success, sweetness, tranquility, victory, welcome
- *Negative validation words:* arson, avalanche, blackout, blizzard, bushfire, casualty, cyclone, destruction, drought, earthquake, emergency, erosion, fatality, flood, hailstorm,

hurricane, mishap, sandstorm, storm, tornado, tremor, tsunami, volcano, whirlpool,
windstorm

Appendix D: Range of prejudice across languages

Across the 45 languages, the means of each of the four comparisons in the Wikipedia data set are reported in Table S2. The table shows average similarity of female (or male) words with positively (or negatively) valenced words for each of the languages. The same relationships are detailed in Table S3 for the Common Crawl data set.

Organizing the data in this way allows us to test the three-way interaction 2 (genderedness of languages: gendered vs. genderless) x 2 (gender words: male vs. female) x 2 (valence words: positive vs. negative). For the Wikipedia data set, this interaction is not significant, $F(7, 168) = 1.237, p > 0.05$. However, for the Common Crawl data set, the three-way interaction is significant, $F(7, 172) = 6.086, p < 0.001$. Decomposing this interaction across gender words suggests that female words are more associated with positive words ($M = 0.1816$) when a language is genderless than when a language is gendered ($M = 0.1629; t = -2.9695, p < 0.01$) even as negative words are relatively more associated with female words in genderless languages ($M = 0.1671$) than gendered languages ($M = 0.1557; t = 2.0203, p > 0.05$). At the same time, male words are associated with positive words in genderless ($M = 0.1772$) and gendered ($M = 0.1710; t = -0.8178, p > 0.05$) languages to a relatively similar degree, but are much more strongly associated with negative words in genderless ($M = 0.1607$) as opposed to gendered ($M = 0.1480; t = -1.9681, p > 0.05$) languages.

Table S2

Range of prejudice in the Wikipedia data set.

Language	Female-Negative	Female-Positive	Male-Negative	Male-Positive
Arabic	0.1848	0.1880	0.1772	0.1901
Armenian	0.2145	0.2177	0.2039	0.2069
Azerbaijani	0.1781	0.1929	0.1704	0.1815
Basque	0.2250	0.2291	0.2077	0.2113
Belarusian	0.2228	0.2115	0.1931	0.2005
Bulgarian	0.2065	0.1961	0.1921	0.1942
Catalan	0.2133	0.2237	0.2046	0.2156
Chinese	0.8852	0.8837	0.8781	0.8780
Czech	0.1969	0.1975	0.1790	0.1961
Danish	0.2255	0.2176	0.2020	0.2039
Dutch	0.2407	0.2353	0.2249	0.2296
English	0.2137	0.2220	0.1936	0.2100
Esperanto	0.2374	0.2185	0.2141	0.2113
Estonian	0.2078	0.2004	0.1857	0.1866
Finnish	0.2107	0.2391	0.1869	0.2096
French	0.2188	0.2102	0.1965	0.2105
Galician	0.1999	0.2089	0.1912	0.1954
Georgian	0.2226	0.2187	0.1969	0.1969
German	0.1882	0.2092	0.2304	0.2592
Greek	0.2102	0.2085	0.1956	0.2047
Hebrew	0.1803	0.1869	0.1603	0.1839
Hindi	0.2505	0.2457	0.2371	0.2448
Hungarian	0.2007	0.2168	0.1940	0.2134
Indonesian	0.2321	0.2277	0.2180	0.2216
Italian	0.2115	0.2230	0.2049	0.2227
Korean	0.2777	0.2864	0.2684	0.2657
Lithuanian	0.2126	0.2189	0.2007	0.2053
Malay	0.2325	0.2343	0.2219	0.2265
Norwegian	0.2151	0.2175	0.2032	0.2110
Persian	0.2389	0.2374	0.2281	0.2303
Polish	0.2071	0.2147	0.1842	0.2014
Portuguese	0.2068	0.2212	0.2078	0.2218
Romanian	0.2310	0.2197	0.2162	0.2218
Russian	0.2010	0.2076	0.1747	0.2022
Serbian	0.2547	0.2520	0.2317	0.2470
Slovak	0.2158	0.2103	0.1948	0.2060
Slovenian	0.2081	0.2071	0.1775	0.1896
Spanish	0.2091	0.2123	0.2022	0.2141
Swedish	0.2766	0.2783	0.2688	0.2724
Tamil	0.2000	0.2053	0.1877	0.1942
Thai	0.2323	0.2523	0.2223	0.2408

Turkish	0.1952	0.2041	0.1856	0.1966
Ukrainian	0.2131	0.2206	0.1878	0.2001
Urdu	0.2802	0.2764	0.2725	0.2713
	Mean	Mean	Mean	Mean
	(SD)	(SD)	(SD)	(SD)
Gendered	0.2171	0.2187	0.2035	0.2146
	(0.0243)	(0.0219)	(0.0261)	(0.0242)
Genderless	0.2628	0.2676	0.2485	0.2555
	(0.1675)	(0.1658)	(0.1695)	(0.1673)

Table S3

Range of prejudice in the Common Crawl data set.

Language	Female-Negative	Female-Positive	Male-Negative	Male-Positive
Arabic	0.1306	0.1519	0.1285	0.1578
Armenian	0.1689	0.1955	0.1646	0.1875
Azerbaijani	0.1451	0.1578	0.1298	0.1464
Basque	0.1343	0.1496	0.1235	0.1377
Belarusian	0.1631	0.1548	0.1420	0.1557
Bulgarian	0.1465	0.1450	0.1300	0.1605
Catalan	0.1762	0.1925	0.1840	0.1962
Chinese	0.1820	0.1867	0.1787	0.1887
Czech	0.1555	0.1487	0.1356	0.1582
Danish	0.1693	0.1772	0.1533	0.1641
Dutch	0.1671	0.1698	0.1680	0.1814
English	0.1454	0.1518	0.1383	0.1461
Esperanto	0.1863	0.1967	0.1723	0.1901
Estonian	0.1510	0.1607	0.1416	0.1461
Finnish	0.1791	0.2199	0.1758	0.2194
French	0.1676	0.1738	0.1570	0.2024
Galician	0.1571	0.1705	0.1583	0.1640
Georgian	0.1544	0.1757	0.1555	0.1769
German	0.1748	0.2070	0.1748	0.2255
Greek	0.1595	0.1655	0.1506	0.1801
Hindi	0.1379	0.1284	0.1284	0.1431
Hungarian	0.1553	0.1906	0.1451	0.1854
Indonesian	0.2064	0.2123	0.1980	0.2097
Italian	0.1540	0.1711	0.1537	0.1886
Japanese	0.1939	0.1918	0.1822	0.1805
Korean	0.1717	0.1661	0.1647	0.1570
Lithuanian	0.1625	0.1764	0.1683	0.1887
Malay	0.1746	0.1919	0.1712	0.1939

Norwegian	0.1742	0.1967	0.1718	0.1949
Persian	0.1657	0.1769	0.1521	0.1646
Polish	0.1499	0.1601	0.1470	0.1723
Portuguese	0.1481	0.1578	0.1506	0.1756
Romanian	0.1420	0.1469	0.1355	0.1628
Russian	0.1461	0.1477	0.1160	0.1600
Serbian	0.1613	0.1618	0.1488	0.1752
Slovak	0.1515	0.1487	0.1290	0.1469
Slovenian	0.1522	0.1538	0.1265	0.1451
Spanish	0.1450	0.1555	0.1482	0.1679
Swedish	0.1863	0.2029	0.1924	0.2153
Tamil	0.1504	0.1633	0.1448	0.1679
Thai	0.1760	0.1851	0.1868	0.2094
Turkish	0.1508	0.1774	0.1511	0.1723
Ukrainian	0.1829	0.1745	0.1676	0.1866
Urdu	0.1383	0.1475	0.1325	0.1479
Welsh	0.1088	0.1106	0.1010	0.1028
	Mean	Mean	Mean	Mean
	(SD)	(SD)	(SD)	(SD)
Gendered	0.1558	0.1629	0.148	0.1709
	(0.01665)	(0.0212)	(0.0208)	(0.0246)
Genderless	0.1671	0.1816	0.1607	0.1772
	(0.0194)	(0.0199)	(0.0209)	(0.0246)

Appendix E: Additional Analysis for Gender Prejudice

It is worth noting that languages can belong to different language families (e.g., Indo-European, Sino-Tibetan). Therefore, one could argue that the observed difference between gendered versus genderless languages is not because of genderedness but because of the language family. To address this account, we conducted a Hierarchical Linear model (HLM) with language family as the second-level variable. We classified the 45 languages using data from Ethnologue (Simmons & Fennig, 2018; see Table S4). An analysis using a random-intercept HLM model that controls for variance in each language family as well as across, showed that genderedness still has a significant influence on net gender prejudice in both the Wikipedia and Common Crawl data. For the Wikipedia data language family had a positive influence on net

prejudice ($\beta=0.0096$, $p < 0.001$) and classification as a genderless language reduced the net prejudice ($\beta=-0.0074$, $p < 0.01$). The analysis in the Common Crawl data also indicated that language family had a positive influence on net prejudice ($\beta=0.0158$, $p < 0.001$) and classification as a genderless language reduced the net prejudice ($\beta=-0.0137$, $p < 0.001$).

Table S4

Languages classified by their macro-family.

Language	Language Family
Arabic	Afro-Asiatic
Armenian	Indo-European
Azerbaijani	Turkic
Basque	Isolate
Belarusian	Indo-European
Bulgarian	Indo-European
Catalan	Indo-European
Chinese (Mandarin)	Sino-Tibetan
Czech	Indo-European
Danish	Indo-European
Dutch	Indo-European
English	Indo-European
Esperanto	International Auxiliary Languages
Estonian	Uralic
Finnish	Uralic
French	Indo-European
Galician	Indo-European
Georgian	Caucasian
German	Indo-European
Greek	Indo-European
Hebrew	Afro-Asiatic
Hindi	Indo-European
Hungarian	Uralic
Indonesian	Austronesian
Italian	Indo-European
Japanese	Japonic
Korean	Koreanic

Lithuanian	Indo-European
Malay	Austronesian
Norwegian	Indo-European
Persian	Indo-European
Polish	Indo-European
Portuguese	Indo-European
Romanian	Indo-European
Russian	Indo-European
Serbian	Indo-European
Slovak	Indo-European
Slovenian	Indo-European
Spanish	Indo-European
Swedish	Indo-European
Tamil	Dravidian
Thai	Tai-Kadai
Turkish	Turkic
Ukrainian	Indo-European
Urdu	Indo-European
Welsh	Indo-European

One could also argue that there may be a systematic variation such that prosperous countries are more likely to have genderless language while less prosperous countries have gendered languages or vice versa, and that prosperity of the country in which a language is spoken affects the pattern of our findings. However, the content contributed online (e.g., on Wikipedia) in a certain language can come from people in different countries (e.g., people in any country can contribute to online English, Spanish, or French content if they speak that language). As a result, we do not conclusively observe such a systematic pattern in our data.

Finally, one could argue that there may be specific features of the region where languages are spoken that could affect gender prejudice over and above the genderedness of languages. In order to address such a concern we ran our analysis controlling for geographic, cultural, and demographic covariates that have been shown to influence gender equality (Prewitt-

Freilino, Caswell, & Laakso, 2012; United Nations Development Programme, 2017; World Economic Forum, 2017).

We used the Global Gender Gap index developed by World Economic Forum (2017), the Corruption Perceptions Index developed by Transparency International (2016), the Religious Diversity Index developed by Pew Research Center (2014), the STC Climate Index compiled from public weather data (Tajick, 2018), and the Human Development Index developed by the United Nations Development Programme (2017). Since most global indices are calculated per country, not per language, we utilized data from Wikipedia outlining the relative contribution of residents from each country to each language project (Wikimedia Foundation, 2018). Country-level index values were then weighted and averaged for each language, based on that country's relative contribution to the Wikipedia language project. This same distribution was used as a proxy to analyze both Wikipedia and Common Crawl data. Covariate values were centered to make interpretation of coefficients easier.

For the Wikipedia data, a linear regression indicated that the predictor of genderedness significantly predicted gender prejudice for each language. The results indicated that the predictor explained 21.3% of the total variance ($R^2 = 0.2131$, $F(1,42) = 11.37$, $p < 0.01$). Gendered languages had more gender prejudice than genderless languages ($\beta = 0.0074$, $p < 0.01$). A more complex model including the covariates of gender parity, corruption, religious diversity, development, and climate did not improve predictive ability ($R^2 = 0.2132$, $F(6,34) = 1.535$, $p > 0.05$). Therefore, even after controlling for covariates gendered languages displayed more gender prejudice than genderless languages.

A similar pattern of results emerged for the Common Crawl data. A linear regression indicated a significant effect of genderedness on gender prejudice; it explained 30.2% of the total variance ($R^2 = 0.3017$, $F(1,43) = 18.57$, $p < 0.001$). Gendered language had more gender prejudice than genderless languages ($\beta = 0.0137$, $p < 0.001$). In the Common Crawl data, one covariate, corruption, along with genderedness improved the predictive ability of the model ($R^2 = 0.5128$, $F(6,35) = 6.14$, $p < 0.001$). Specifically, genderedness ($\beta = 0.0142$, $p < 0.001$) and corruption ($\beta = -0.0006$, $p < 0.01$)¹ predicted gender prejudice.

In sum, the regression model results (see Table S5) suggest that the presence of genderedness alone is responsible for between 20% and 30% of the variance in observed gender prejudice across 45 languages in two different data sets, even after accounting for geographic, demographic, and cultural confounds.

Table S5

Summary of the regression analysis including geographic, demographic and cultural control variables.

	Wikipedia	Wikipedia	Common Crawl	Common Crawl
	Gender Model	Complete Model	Gender Model	Complete Model
Variable	β (SE)	β (SE)	β (SE)	β (SE)
Intercept	0.0022 (0.0018)	-0.0068 (0.0248)	0.0021 (0.0025)	-0.0130 (0.0313)
Genderedness	0.0074** (0.0022)	0.0075* (0.0028)	0.0137*** (0.0032)	0.0142*** (0.0034)
Global Gender Gap		-0.0090 (0.0420)		0.0473 (0.0494)
Corruption Perceptions		-0.0001 (0.0309)		-0.0006** (0.0002)
Religious Diversity		0.0005 (0.0010)		0.0019 (0.0012)

STC Climate Index		0.0008 (0.0121)		-0.0121 (0.0153)
Human Development Index		0.0104 (0.0310)		0.0250 (0.0388)
R^2	0.2131	0.2132	0.3017	0.5128
F	11.37**	1.5350	18.57***	6.14***

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix F: Additional Analysis for Gender Stereotype

A similar analysis was conducted for the stereotypic associations of gender with warmth and competence. Since warmth and competence are not binary variables like positive and negative, the two constructs are considered separately. A random-intercept HLM that controls for the variance in each language family as well as across each language showed that the observed patterns held. For the Wikipedia data, language family did not have a significant influence on associations between gender and warmth words ($p = 0.1693$), but the presence of genderedness ($\beta = 0.0076$, $p < 0.05$) had a positive influence on the difference in association between gender words and warmth words. The analysis in the Common Crawl data again indicated that language family did not have an influence on associations between gender and warmth words ($p = 0.451$); however, genderedness did have a significant positive influence on the difference in association between gender words and warmth words ($\beta = 0.0134$, $p < 0.05$). For the Wikipedia data, both language family ($\beta = 0.0094$, $p < 0.001$) and genderedness ($\beta = 0.0079$, $p < 0.05$) had a significant influence on the difference between the association of gender words and competence. Similarly, for the Common Crawl data, both language family ($\beta = 0.0079$, $p < 0.01$) and genderedness ($\beta = 0.0168$, $p < 0.001$) had a significant influence on the difference between the association of gender words and competence. This shows that even after taking onto account language family, the influence of genderedness of language persists.

Again, to account for the influence of covariates, we used the Global Gender Gap index developed by World Economic Forum (2017), the Corruption Perceptions Index developed by Transparency International (2016), the Religious Diversity Index developed by Pew Research Center (2014), the STC Climate Index compiled from public weather data (Tajick, 2018), and the Human Development Index developed by the United Nations Development Programme (2017), this time with the results for the warmth-competence constructs weighted by the proportional contribution of each country to each language to Wikipedia, to account for possible cultural and geographic covariates (Prewitt-Freilino et al., 2012).

For the Wikipedia data, a linear regression indicated that the predictor of genderedness significantly predicted gender stereotypes associated with warmth. The results indicated that the predictor explained 16.8% of the total variance ($R^2 = 0.168$, $F(1, 42) = 8.748$, $p < 0.01$). Gendered languages showed a greater difference in association between gender words and warmth words ($\beta = 0.0071$, $p < 0.01$). A more complex model including the covariates of gender parity, corruption, religious diversity, development, and climate improved predictive ability ($R^2 = 0.4681$, $F(6, 34) = 4.986$, $p < 0.001$), with a similar relationship between gendered languages and the difference in association between gender and warmth words ($\beta = 0.0072$, $p < 0.01$) and none of the other covariates showing a significant influence (see Table S6).

The same analysis was conducted for the Common Crawl data. A linear regression indicated a significant effect of genderedness on gender stereotypes; it explained 34% of the total variance ($R^2 = 0.3404$, $F(1, 43) = 22.19$, $p < 0.001$). Gendered languages showed a greater difference in the association between gender words and warmth words ($\beta = 0.0134$, $p < 0.001$). With the more complex model suggested by Prewitt-Freilino et al. (2012), improved predictive

ability ($R^2 = 0.5658$, $F(6, 35) = 7.6$, $p < 0.001$), with a similar relationship between gendered languages and the difference in association between gender and warmth words ($\beta = 0.0143$, $p < 0.001$) and none of the other covariates showed a significant influence (see Table S6).

Table S6

Summary of the regression analysis for the stereotype of warmth, including geographic, demographic and cultural control variables.

	Wikipedia	Wikipedia	Common Crawl	Common Crawl
	Warmth	Warmth	Warmth	Warmth
	Gender Model	Complete Model	Gender Model	Complete Model
Variable	β (SE)	β (SE)	β (SE)	β (SE)
Intercept	-0.0032 (0.002)	-0.0212 (0.0205)	-0.0017 (0.0023)	0.0201 (0.0252)
Genderedness	0.0071** (0.0024)	0.0072** (0.0025)	0.0134*** (0.0029)	0.0143*** (0.003)
Global Gender Gap		-0.0415 (0.0359)		-0.0166 (0.0417)
Corruption Perceptions		-0.0001 (0.0001)		-0.0003 (0.0002)
Religious Diversity		-0.0001 (0.0001)		0.0001 (0.001)
STC Climate Index		-0.0001 (0.0107)		-0.0179 (0.0131)
Human Development Index		-0.0275 (0.0252)		-0.0152 (0.0309)
R^2	0.168	0.4681	0.3404	0.5128
F	8.478**	4.986***	22.19***	6.14***

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The analysis was repeated for competence words in each data set. For the Wikipedia data set, a linear regression indicated genderedness was a significant predictor of gender stereotypes associated with competence. The results indicated that the predictor explained 14%

of the total variance ($R^2 = 0.1435$, $F(1, 42) = 7.034$, $p < 0.05$). Gendered languages showed a greater difference in association between gender words and competence words ($\beta = 0.0079$, $p < 0.05$). A more complex model including the covariates of gender parity, corruption, religious diversity, development, and climate improved predictive ability ($R^2 = 0.3076$, $F(6, 34) = 2.517$, $p < 0.05$), with a similar relationship between gendered languages and the difference in association between gender and competence words ($\beta = 0.0085$, $p < 0.05$) and none of the covariates showing significant influence (see Table S6). For the Common Crawl data, a linear regression indicated a significant effect of genderedness on gender stereotypes of competence; it explained 14% of the total variance ($R^2 = 0.1435$, $F(1, 42) = 7.034$, $p < 0.05$). Gendered languages showed a greater difference in the association between gender words and competence words ($\beta = 0.0079$, $p < 0.05$). The more complex model improved predictive ability ($R^2 = 0.3076$, $F(6, 34) = 2.517$, $p < 0.05$), with a similar relationship between gendered languages and the difference in association between gender and competence words ($\beta = 0.0085$, $p < 0.05$) and none of the other covariates showing a significant influence (see Table S7).

Table S7

Summary of the regression analysis for the stereotype of warmth, including geographic, demographic and cultural control variables.

	Wikipedia	Wikipedia	Common Crawl	Common Crawl
	Competence	Competence	Competence	Competence
	Gender Model	Complete Model	Gender Model	Complete Model
Variable	β (SE)	β (SE)	β (SE)	β (SE)
Intercept	0.0094 (0.0024)	-0.0292 (0.0285)	0.0094 (0.0024)	-0.0292 (0.0285)

Genderedness	0.0079 ** (0.0029)	0.0085* (0.0035)	0.0079* (0.0029)	0.0085* (0.0035)
Global Gender Gap		-0.0233 (0.0499)		-0.0232 (0.0499)
Corruption Perceptions		-0.0003 (0.0002)		-0.0003 (0.0002)
Religious Diversity		-0.0019 (0.0012)		-0.0019 (0.0012)
STC Climate Index		-0.0226 (0.0149)		-0.0226 (0.0149)
Human Development Index		0.0595 (0.0349)		0.0595 (0.03499)
R^2	0.1435	0.3076	0.1435	0.3076
F	7.034*	2.517*	7.034*	2.517*

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

References

- Abele, A. E., Hauke, N., Peters, K., Louvet, E., Szymkow, A., & Duan, Y. (2016). Facets of the fundamental content dimensions: Agency with competence and assertiveness—communion with warmth and morality. *Frontiers in Psychology*, 7.
- Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2016). Enriching word vectors with subword information. *ArXiv:1607.04606 [Cs]*. Retrieved from <http://arxiv.org/abs/1607.04606>
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 29, pp. 4349–4357). Retrieved from <http://papers.nips.cc/paper/6228-man-is-to-computer-programmer-as-woman-is-to-homemaker-debiasing-word-embeddings.pdf>
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186.
- Corbett, G. G. (1991). *Gender*. New York, NY: Cambridge University Press.
- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77–83.
- Grave, E., Bojanowski, P., Gupta, P., Joulin, A., & Mikolov, T. (2018). Learning word vectors for 157 languages. *ArXiv:1802.06893 [Cs]*. Retrieved from <http://arxiv.org/abs/1802.06893>
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, 74(6), 1464–1480.

- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85(2), 197–216.
- Haspelmath, M. (2005). *World Atlas of Language Structures*. Oxford, UK: Oxford University Press.
- Hellinger, M., & Bußmann, H. (Eds.). (2001). *Gender across languages: The linguistic representation of women and men*. Philadelphia, PA: John Benjamins Publishing Company.
- Pew Research Center. (2014). *Religious diversity index scores by country*. Retrieved from <http://www.pewforum.org/2014/04/04/religious-diversity-index-scores-by-country/>
- Prewitt-Freilino, J. L., Caswell, T. A., & Laakso, E. K. (2012). The gendering of language: A comparison of gender equality in countries with gendered, natural gender, and genderless languages. *Sex Roles*, 66(3–4), 268–281.
- Rosenberg, S., Nelson, C., & Vivekananthan, P. S. (1968). A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology*, 9(4), 283–294.
- Rudman, L. A., Greenwald, A. G., & McGhee, D. E. (2001). Implicit self-concept and evaluative implicit gender stereotypes: Self and ingroup share desirable traits. *Personality and Social Psychology Bulletin*, 27(9), 1164–1178.
- Simmons, G. F., & Fennig, C. D. (2018). *Ethnologue: Languages of the World*. Retrieved from <https://www.ethnologue.com/>
- Tajick, S. (2018). *STC climate index*. Retrieved from <http://globalresidenceindex.com/hnwi-index/climate-index/>

- Transparency International. (2016). *Corruption perceptions index*. Retrieved from https://www.transparency.org/news/feature/corruption_perceptions_index_2016
- United Nations Development Programme. (2017). *Human development report 2016: Human development for everyone*. New York, NY: United Nations Development Programme.
- Wikimedia Foundation. (2018). Wikimedia traffic analysis report: Page views per Wikipedia language. Retrieved from https://stats.wikimedia.org/archive/squid_reports/2018-05/SquidReportPageViewsPerLanguageBreakdown.htm
- World Economic Forum. (2017). *The global gender gap report: 2017*. Geneva: World Economic Forum.

Footnotes

¹ Note that the corruption index is scaled such that a number closer to zero indicates more public corruption and a number closer to one indicates less public corruption (Transparency International, 2016).