

Stereotypical Gender Associations in Language Have Decreased Over Time

Jason J. Jones, Mohammad Ruhul Amin, Jessica Kim, Steven Skiena

Stony Brook University

Abstract: Using a corpus of millions of digitized books, we document the presence and trajectory over time of stereotypical gender associations in the written English language from 1800 to 2000. We employ the novel methodology of word embeddings to quantify male gender bias: the tendency to associate a domain with the male gender. We measure male gender bias in four stereotypically gendered domains: career, family, science, and arts. We found that stereotypical gender associations in language have decreased over time but still remain, with career and science terms demonstrating positive male gender bias and family and arts terms demonstrating negative male gender bias. We also seek evidence of changing associations corresponding to the second shift and find partial support. Traditional gender ideology is latent within the text of published English-language books, yet the magnitude of traditionally gendered associations appears to be decreasing over time.

Keywords: gender; stereotypes; gender ideology; implicit bias; word embeddings; natural language processing

TRADITIONAL gender ideology, defined as the “level of support for a division of paid work and family responsibilities that is based on the notion of separate spheres” (Davis and Greenstein 2009:88), normalizes men’s and women’s differences¹ and justifies their separation into the public and private spheres, respectively, as well as the inequalities that arise from this arrangement. Surveys show that although many still endorse traditional gender ideology, support for it has generally declined over time (Burns et al. 2015; Cherlin and Walters 1981; Cotter, Hermsen, and Vanneman 2011; Swim et al. 1995) as gender equality and women’s participation in the public sphere has increased (Bianchi et al. 2000; Burns et al. 2015; Weitzman 1978). Yet traditional gender ideology also extends beyond the individual or their awareness. People not only harbor implicit attitudes consistent with traditional gender ideology often outside of their overtly declared sentiments toward women (Nosek, Banaji, and Greenwald 2002a; Nosek et al. 2007; Steele and Ambady 2006) but they also embed traditional gender ideology within the cultural artifacts they use and produce, such as language and books (Madera, Hebl, and Martin 2009; Ng 2007; Pratto, Korchmaros, and Hegarty 2007; Schmader, Whitehead, and Wysocki 2007; Silveira 1980; Trix and Psenka 2003).

Citation: Jones, Jason J., Mohammad Ruhul Amin, Jessica Kim, and Steven Skiena. 2019. “Stereotypical Gender Associations in Language Have Decreased Over Time.” *Sociological Science* 7: 1-35.

Received: August 13, 2019

Accepted: October 31, 2019

Published: January 7, 2019

Editor(s): Jesper Sørensen, Sarah Soule

DOI: 10.15195/v7.a1

Copyright: © 2019 The Author(s). This open-access article has been published under a Creative Commons Attribution License, which allows unrestricted use, distribution and reproduction, in any form, as long as the original author and source have been credited. 

In turn, it has therefore become necessary to generate accompanying strategies for detecting these subtler manifestations of traditional gender ideology. Recent developments in computer science that document the capacity for computers and artificial intelligence to inherit and reflect humanlike biases, such as traditional gender ideology (Barocas and Selbst 2016; Caliskan, Bryson, and Narayanan 2017; Ensign et al. 2018; Friedler, Scheidegger, and Venkatasubramanian 2016; O’Neil 2016), offer an innovative solution. Although some have begun to make use of this

method for studying social life (Garg et al. 2018; Kozlowski, Taddy, and Evans 2019; Ornaghi, Ash, and Chen 2018), it remains a largely understudied yet promising new area of research. We add to this budding literature and build upon previous qualitative content analyses of gender bias in language and books to quantitatively examine the presence and trajectory of gender ideology in the written English language across 200 years.

Specifically, we use word embedding representations of the English language constructed from published books from every decade from 1800 to 2000 to quantify the association between gender and stereotypically gendered domains (career, family, science, and arts) latent within language patterns. In doing so, we expand upon existing work that demonstrates stereotypical gender bias in current writing by employing a historical approach that not only detects the presence of traditional gender ideology in published books but also captures how it has changed over time, both across and within domains. Our findings suggest that gender bias in language has decreased over time but is still present in written work to varying degrees. We further examine the relative “stickiness” of stereotypical associations (England 2010) and find that although male–family and female–career terms are becoming increasingly associated, male–career and female–family term associations remain quite stable over time. This suggests that the overall decrease in traditional gender ideology in language is driven by an imbalanced rather than symmetrical shift whereby the concepts of male and female may diversify in their associated roles, but their primary domain prescriptions remain strongly intact.

We compare our findings to existing societal and public opinion trends to conclude that the traditional gender ideology present in the corpus of books written in English roughly mirrors social sentiments, both in terms of broader trends (Burns et al. 2015; Cherlin and Walters 1981; Cotter et al. 2011; Swim et al. 1995) and as a partial reflection of the second shift (Hochschild and Manchung 2003). We posit that this approach is a novel and effective method for detecting latent traditional gender ideology within written work and, more broadly, for making inferences about human culture and social life over time. In the following sections, we review research on the changing nature of traditional gender ideology over time, discuss the presence of implicit biases in humans and the things they manufacture, and develop hypotheses regarding the temporal trajectories of gender and domain associations.

Traditional Gender Ideology Over Time

Defining Traditional Gender Ideology

At its core, gender as a social institution establishes differences between men and women and justifies social inequalities between them on the basis of those differences (Connell 1987; Lorber 1994; Martin 2004; Risman 2004). Embedded within several levels of social organization—individual, interactional, and institutional—gender not only produces social structures that funnel men and women into unequal life trajectories but it also socializes them to draw upon accepted notions of their gender differences when engaging in interactions and constructing their personal

identities (Bem 1993; Fagot, Rodgers, and Leinbach 2000; Risman 2004). Moreover, as a social institution, gender is upheld by a legitimating ideology that “proclaims the rightness and necessity of [the institution’s] arrangements, practices, and social relations” (Berger and Luckmann 1966; Martin 2004:1257). This legitimating ideology—gender ideology—justifies men’s and women’s specified social roles and expectations according to corresponding beliefs about each group. Although gender ideology can manifest in many forms (e.g., traditional, egalitarian, liberal, feminist) and by extension render various organizational outcomes based on gender (Kroska 2007), arguably the most fundamental and enduring gender ideology in modern society is traditional gender ideology.

Traditional gender ideology—also referred to as separate spheres ideology—draws upon beliefs about dispositional gender differences to justify the separation of men and women into the public and private spheres, respectively, with men’s rightful social position being that of the “breadwinner” and women’s that of the “homemaker” (Davis and Greenstein 2009). Specifically, traditional gender ideology draws upon cultural beliefs, or gender stereotypes,² that describe men as agentic (Eagly and Steffen 1984), rational (Broverman et al. 1972), “tough, aggressive, and assertive” (Huddy and Terkildsen 1993:121) and women as expressive (Broverman et al. 1972), “warm, gentle, kind, and passive” (Huddy and Terkildsen 1993:121) to validate the idea that men are better suited for work and should earn paid employment and women should uphold the domestic and familial obligations of the household (Gerstel and Gross 1987; Kerber 1988; Mitchell 1971; Thorne and Yalom 1982; Welter 1966). Conversely, women are expected to be unemployed and thus financially dependent on men, and men are excused from involvement in the home and family. Women’s employment is considered problematic in that it detracts from their caregiving, and men’s contributions to the household are deemed unnecessary or inappropriate (Acker 1988; DeVault 1990; Feldberg and Glenn 1979; Pleck 1977). Although the arrangement prescribed by traditional gender ideology represents an ideal type that was achieved by relatively few families historically (Kerber 1988; Kessler-Harris 1982), it has nevertheless enjoyed widespread acceptance and perceived validity.

Historical Trends in Traditional Gender Ideology

Although national-level surveys capturing individuals’ traditional gender ideology were not utilized until the second half of the twentieth century, historical evidence points to its operation within society as far back as the early 1800s. Perhaps most famously, Welter’s “Cult of True Womanhood” was prominent throughout the Victorian era of the 1800s and characterized the typical American man as “a busy builder of bridges and railroads, at work long hours,” and the American woman a “hostage in the home” marked by her submissiveness and domesticity (1966:151). Early women’s rights activists into the 1920s expressed similar dismay at their restricted access to the public sphere and overall dependence on men, citing women’s prohibition from voting, education, property ownership, and wage labor (Stanton 1848). And, despite significant gains for women by way of the 19th Amendment in 1920 (U.S. Constitution) and wartime demand for women’s labor during World

War II (Feree and Hess 2000; Rupp 1979), women were still restricted in the type of employment available to them, popular conceptions about what constituted “men’s work” versus “women’s work” (Milkman 2016), and the continued notion that women’s primary obligation was to the family and household (Kessler-Harris 1990).

Yet despite the potency of traditional gender ideology, by the 1970s, further increases in women’s labor force participation (Bianchi et al. 2000), increasing divorce rates, decreasing fertility (Weitzman 1978), and the resurgent women’s rights movement set the stage for increased attention to gender inequality and changing attitudes regarding the roles of men and women (Burns et al. 2015). To capture the extent to which the public supported these changing work and family roles, several public opinion surveys³ began incorporating items that quantified and measured traditional gender ideology into their questionnaires (Cherlin and Walters 1981; Mason, Czajka, and Sara Arber 1976). Typical questions included, “It is much better for everyone concerned if the man is the achiever outside the home and the woman takes care of the home and family,” “A working mother can establish just as warm and secure a relationship with her children as a mother who does not work,” and “A wife should not expect her husband to help around the house after he comes home from a hard day’s work” (Davis and Greenstein 2009).

Analyses of these early surveys demonstrate sizable increases in support for gender equality by the mid-1970s (Feree 1974; Mason et al. 1976; Tallichet and Willits 1986; Thornton and Freedman 1979). For instance, from 1964 to 1974, women’s attitudes toward their role in both the home and labor market became more egalitarian, with a particularly large decrease in their belief that children will suffer when their mother works outside the home (Mason et al. 1976). Studies focusing on the years 1972 to 1978 found similar results, with both men and women demonstrating a sizeable shift away from belief in traditional gender ideology (Cherlin and Walters 1981). Later studies continued to map this general trend as it continued into the 1980s (Mason and Lu 1988; Thornton, Alwin, and Camburn 1983) and the early 1990s (Cotter et al. 2011).

By the mid-1990s, however, the pace of gains in egalitarian attitudes set in the 1970s and 1980s significantly slowed, and according to some accounts, even slightly reversed (Cotter et al. 2011; Thornton and Young-DeMarco 2001). Despite a small surge in egalitarian attitudes in the early 2000s, the post-mid-1990s era is generally characterized by “small and shifting changes” away from traditional gender ideology (Cotter et al. 2011:260), with a sizable proportion of Americans still endorsing it (Burns et al. 2015).

Some point to the changing nature of traditional gender ideology in light of recent societal changes. Compared to earlier eras, Americans are exchanging blatant “old-fashioned” stereotypes easily captured by conventional survey instruments for more elusive and subtle “modern” (Swim et al. 1995), “ambivalent” (Glick and Fiske 1996), “hostile,” or “benevolent” (Glick and Fiske 1996, 2001) forms of sexism marked by denial of gender inequality as a problem, antagonism toward women pursuing additional rights, resentment of women who use their sexuality to gain power, and feelings of protective paternalism towards women (Glick and Fiske 1996, 2001; Swim et al. 1995). Yet traditional gender ideology underlies these

more modern beliefs, with high-ranking “modern” sexists being more likely to hold traditional beliefs about women (Swim et al. 1995).

Others posit this tempering to be the result of a cultural backlash against feminism to produce an “egalitarian but traditional” outlook on women’s proper place (Charles and Grusky 2004; Cotter et al. 2011:284). This frame suggests that in the context of increasing norms of childcare and the concomitant “ideology of intensive mothering⁴” presented by pop culture (Douglas and Michaels 2004; Hays 1996:9x), women attempting to balance both their careers and families simply “can’t have it all” and should choose to focus their efforts solely on their families (Blair-Loy 2003:66; Cotter et al. 2011; Stone 2007). Indeed, studies show that people continue to draw upon traditional gender stereotypes, emphasizing women’s maternal instincts and predisposition to family work to structure parental involvement in and division of childcare (Deutsch 1999; Gaunt 2006).

Gender stereotypes also continue to permeate expectations about men’s and women’s appropriate careers and hobbies, with men being seen as rational and thus better at math and science and women being perceived as verbal and emotional and therefore better suited for arts and literature (Eccles, Jacobs, and Harold 1990; Jacobs et al. 2002). These expectations are borne out through parents’ perceptions of their children (Eccles et al. 1990), children’s school subject preferences and performance (Coles and Hall 2002; Ely and Ryan 2008; Jones, Howe, and Rua 2000; Mullis et al. 2007; Sainsbury and Schagen 2004; Sasson et al. 2010), adults’ career choices and outcomes (Beede et al. 2011; Ginther and Khan 2006), and through self-fulfilled stereotype threat⁵ (Estes and Felker 2012; Kiefer and Sekaquaptewa 2007; Spencer, Steele, and Quinn 1999).

It seems then that although traditional gender ideology has adapted and waned over time, it continues to be a relevant factor in how we shape our identities, beliefs, and expectations about the appropriate social roles of men and women. However, this does not necessarily imply that expectations for men and women have changed symmetrically. In fact, there is evidence to suggest that women have a difficult time extracting themselves from arrangements prescribed by traditional gender ideology (Diekman and Goodfriend 2006; Hochschild and Manchung 2003).

Differential Rates of Change: The Second Shift

Although endorsement of traditional gender ideology has decreased over time (Burns et al. 2015; Swim et al. 1995), some stereotypes are “stickier” than others (England 2010). Despite women’s significant gains in the paid workforce (Fullerton 1999; Toossi and Morisi 2017) as well as increasing expectations that women be agentic, powerful, and capable of breadwinning (Brewster and Padavic 2000; Diekman and Eagly 2000; Diekman and Goodfriend 2006; England 2010), the perception of women as integral to the home and family persists (Deutsch 1999; Diekman and Goodfriend 2006; Gaunt 2006). In one study, although participants rated women’s adoption of traditionally masculine characteristics positively, they rated women’s possession of traditionally feminine characteristics most positively, suggesting higher approval of traditional as compared to nontraditional women (Diekman and Goodfriend 2006). It appears that women are making strides toward equality but

cannot seem to escape their prescriptive ties to domesticity. As Rudman and Phelan (2010) aptly suggest, the “insistence on female nurturance may signal a desire to keep women stoking the home fires even as they blaze a path to glory in public spheres” (P. 192).

One major consequence of this trend is the double standard it sets for women: they must simultaneously be career driven yet family focused; agentic yet communal; confident yet feminine; in short, they must strive to be “supermoms” (Diekman and Goodfriend 2006; Hochschild and Machung 2003:22). At the same time, expectations of men’s roles have remained relatively unchanged (Diekman and Eagly 2000; Diekman, Goodfriend, and Goodwin 2004). That is, by and large, women’s inclusion in the workforce fails to be matched by men’s increased participation in the home (Coltrane 2000; Hochschild and Manchung 2003).

This phenomenon is most succinctly captured by Arlie Hochschild’s seminal work, *The Second Shift*. As women significantly increase their participation in paid employment, men’s participation in the home has only increased marginally, leaving women to shoulder the dual burdens of both work and home in what Hochschild coins as their second shift. As the name implies, following the completion of their first eight-hour shift in paid employment, women arrive home to begin their second shift of unpaid household labor (Hochschild and Manchung 2003). England (2010) makes similar claims of asymmetry in women’s progress, suggesting that the path toward gender equality has been, by and large, unidirectional: women experience strong incentives to enter male jobs, but men do not experience similar incentives to pursue female jobs.

A wealth of evidence supports this claim. Although men do contribute more to the household than they have in the past (Sayer 2005) and the number of hours dedicated to household labor is decreasing overall (Bianchi et al. 2000), women and mothers still perform the majority of it (Sayer 2005; Shelton and John 1996; Yavorsky, Kamp Dush, and Schoppe-Sullivan 2015). Moreover, when domestic needs become more pressing, mothers (compared to fathers) are far more likely to leave their paid jobs in order to focus on the home, whereas fathers tend to work additional paid hours and/or contribute no additional labor to the household (Gjerdengen and Center 2005; Klerman and Leibowitz 1999; Stone 2007). And yet in practice, mothers often maintain steady work hours while still bearing the brunt of childcare (Yavorsky et al. 2015).

The effects of the second shift are not just normative. They also produce tangible and quantifiable consequences for both men and women. Women find themselves “mommy tracked” in their careers, experiencing a wage penalty upon having children (Budig and England 2001; Epstein et al. 1995:296; Rose and Hartmann 2008), and operate along a “leaky pipeline” of employment, missing opportunities for promotion and failing to break through into senior management and top-level earning positions (Padavic and Reskin 2002; Stone 2007:13). At the same time, men experience a “fatherhood premium,” where their salaries either do not change after having a child or they increase (Glauber 2008; Hodges and Budig 2010; Lundberg and Rose 2000, 2002:264). Experimental research reinforces these tropes, finding that hiring preference was highest among candidates identified as fathers and lowest among those identified as mothers (Correll, Bernard, and Paik 2007).

Taking the trajectory of traditional gender ideology as a whole, it seems that although traditional gender ideology has generally decreased over time—particularly in the 1970s and 1980s—it still remains a central model through which gender operates in modern society (Cotter et al. 2011). Yet how it manifests and the tools necessary for measuring it have undoubtedly changed. This is especially true given the significant gains that women have made over the last century and the corresponding belief among some that gender inequality has been completely eradicated (Swim et al. 1995:201). Unlike earlier eras marked by conventional and easy-to-measure beliefs, more subtle expressions of traditional gender ideology, such as the “egalitarian but traditional,” “modern,” or “ambivalent” forms of sexism as well as the documented second shift, have become more prevalent. Alternatively, individuals may be unaware of their own biases or unwilling to express their support of traditional gender ideology for fear of social retribution (Gawronski and Bodenhausen 2007; Greenwald and Banaji 1995), which further complicates measurement. It is perhaps these trends that have made detecting implicit and unconscious manifestations of traditional gender ideology so useful in recent years.

Implicit Attitudes and Unconscious Gender Ideology

Implicit attitudes, defined as “introspectively unidentified (or inaccurately identified) traces of past experience that mediate favorable or unfavorable feeling, thought, or action toward social objects,” suggest that individuals may possess underlying biases that operate and exist much outside of their awareness or control (Greenwald and Banaji 1995:8). One of the most ubiquitous methods for measuring implicit attitudes is the Implicit Association Test (IAT), which captures the ease with which respondents are able to make automatic associations across categories of a given social group. For instance, respondents who can more quickly associate “White + pleasant” versus “Black + pleasant” may possess unconscious anti-black attitudes (Greenwald, McGhee, and Schwartz 1998). This strategy is especially useful for drawing out intentionally concealed attitudes, therefore undercutting social desirability bias typically endemic to sensitive topics (Dovidio et al. 1997; Gawronski and Bodenhausen 2007; Greenwald et al. 1998; Maccoby and Maccoby 1954). Research employing IATs has demonstrated implicit bias across many attributes, including race and ethnicity (Baron and Banaji 2006; Dovidio et al. 1997; McConnell and Leibold 2001; McGhee, and Schwartz 1998; Ziegert and Hanges 2005), sexuality (Banse, Seise, and Zerbes 2001; Jellison, McConnell, and Gabriel 2004), and weight (Nosek et al. 2007; Schwartz et al. 2006), among others (Nosek et al. 2007).

Traditional gender ideology is no exception to the issue of implicit bias. Gender-specific IATs regularly find that individuals implicitly endorse traditional gender ideology even despite self-proclaimed support for equality (Nosek et al. 2002a, 2007; Steele and Ambady 2006). Akin to the breadwinner/homemaker dichotomy embedded within traditional gender ideology, the Gender–Career IAT finds that participants associate “male” and “career” terms as well as “female” and “family” terms more quickly than their reverse counterparts. Similar stereotypes play out for academic disciplines in the Gender–Science IAT, with participants associating

“math” and “science” terms with male terms more quickly than with female terms, and “humanities” terms (e.g., arts and literature) with “female” terms more quickly than with “male” terms (Nosek et al. 2002a, 2007). Indicative of the latent nature of implicit bias, analyses of the Gender–Career tests find that traditional gender ideology is more easily detected at the implicit as opposed to explicit level, as indicated by the marked increase in bias using the implicit measure (Nosek et al. 2002a).

Experimental studies examining these biases confirm this phenomenon. They find, for instance, that traditional gendered priming increases implicit gender bias and decreases female interest in male-associated careers and areas of study (Nosek, Banaji, and Greenwald 2002b; Rudman and Phelan 2010; Steele and Ambady 2006). Others find in experimental studies that gender alone is qualification enough for arithmetic-based tasks, with “employers” preferring men over women, even when men display poor past performance (Reuben, Sapientza, and Zingales 2014). Moreover, when students are asked to “draw a scientist,” children are far more likely to draw a male over female scientist, even across categories of the child’s own race, gender, and age group (Barman 1997; Chambers 1983:256; Finson 2002; Fort and Varney 1989; Jones, Howe, and Rua 2000; Mead and Metraux 1957). Analogous findings have been made for children’s drawings of mathematicians (Picker and Berry 2000; Rock and Shaw 2000). Implicit gender biases are so deeply embedded that research subjects even demonstrate gender stereotypical responses to computers with nothing more than gendered voice output (Nass, Moon, and Green 1997).

Although this provides compelling evidence to suggest that individuals subliminally draw upon traditional gender ideology to inform their decisions and guide their behaviors, this is by no means the only way implicit gender bias can manifest. Individuals are also capable of imbuing everyday objects and entities with traditional gender ideology.

Nonhuman Entities Inheriting Human Biases

Language

Although implicit biases originate and exist within individuals, they can and do extend beyond the individual. One such space in which traditional gender ideology is particularly rife is language. Given that language is a cultural artifact of human creation, it is not surprising that dominant human ideas and values, such as traditional gender ideology, are expressed through it. Previous research on this topic suggests this is indeed the case.

Broadly speaking, scholars find abundant evidence for sexism in the written English language. This is true both for English as a standalone construct and in its practical use. Perhaps most fundamentally, English contains grammatical rules and standard conventions that reinforce gender biases (Ng 2007; Pratto et al. 2007; Silveira 1980). From a descriptive perspective, the English language contains more words that refer to men than to women (Maass and Arcuri 1996). Of those that refer to women, many have experienced derogation or reduction in status in the

process of feminization (e.g., governess has lost its meaning as a political leader, a counterpart to governor, and has come to mean a woman responsible for childcare) (Schulz 2000). Conventional rules also promote “discriminatory speech acts” that favor the explicit designation of atypical or socially inferior identities and the absence of their normatively default and privileged counterparts (e.g., specifying someone’s gender if they are female but not if they are male) (Pratto et al. 2007). This coincides with the assumption of maleness that generic terms confer (Silveira 1980), the use of masculine generics to refer to both men and women (e.g., mankind) (Ng 2007), and the placement of gender indicators before stereotypically male or female nouns (e.g., male nurse or madam chairman) that reinforces the traditional gender coding of words (He 2010; Sklar 1983; Zuber and Reed 1993).

Traditional gender stereotypes also emerge in analyses of language use, and they permeate across various social institutions and types of written work. Dictionary examples, for instance, rely on traditional gender stereotypes to illustrate word meanings (e.g., using the phrase “rabid feminist” as an exemplar of the word “rabid,” or using male pronouns to illustrate examples of doctors or researchers) (Oman-Regan 2016). In the educational and developmental setting, research on children’s books also finds traditional gender stereotypes embedded in both illustration and text, from dispositional traits, to role centrality and agency, to occupations—though stereotypes do appear to be decreasing (Benjamin and Irwin-DeVitis 1998; Collins, Ingoldsby, and Dellmann 1984; Gooden and Gooden 2001; Richardson 1986; Weitzman et al. 1972). Similar gender stereotypes have been found in elementary school textbooks (Evans and Davies 2000; Lee and Collins 2009), English as foreign language textbooks (Amini and Birjandi 2012; Otlowski 2003), and teaching materials (Amare 2007). Furthermore, studies find that writers invoke traditional gender stereotypes when crafting recommendation letters for undergraduate students (LaCroix 1985), graduate students (Watson 1987), and academic faculty positions (Madera et al. 2009; Schmader et al. 2007; Trix and Psenka 2003). Specifically, writers are more likely to describe women as communal and men as agentic even when controlling for ability and credentials (Madera et al. 2009), and they describe male candidates with “standout” adjectives (e.g., outstanding, exceptional, superb) (Schmader et al. 2007; Trix and Psenka 2003).

Perhaps most consequentially, exposure to gender-biased language also reinforces and perpetuates traditional gender stereotypes as well as gender inequality. With regard to hiring decisions, applicants described using female-coded traits, compared to male-coded traits, are rated as less hirable (Madera et al. 2009). Additionally, women asked to read job overviews utilizing masculine pronouns—compared to gender-inclusive or gender-neutral pronouns—reported lower interest and investment in the job and higher feelings of exclusion (Stout and Dasgupta 2011). Similarly, job advertisements that employ masculine language increased participants’ perceptions that the career in question was occupied more by men, thus decreasing female participants’ interest in those jobs (Friesen, Gaucher, and Kay 2011). Experiments that employ subliminal stereotypic gender priming render comparable results, with participants correctly classifying gender pronouns more quickly when presented with gender-congruent primes compared to gender-noncongruent primes (Banaji and Hardin 1996). Even exposure to generic masculine

words or gender-neutral language produces cognitive male biases, though deliberate interventions are capable of moderating this (Bauer and Baltes 2002; Lassonde and O'Brien 2013; McConnell and Fazio 1996; Ng 1990; Wilson and Ng 1988).

Taken together, this evidence suggests that the use of gender-stereotypical language and conventions perpetuates individuals' automatic gender stereotypes as well as gender stereotypical social and institutional arrangements, sometimes even outside of the knowledge or intention of those involved (Banaji and Hardin 1996; Gaucher et al. 2011). Although substantial, the influence of traditional gender stereotypes in language extends far beyond written work alone; gender stereotypic language also has far-reaching consequences for computer algorithms trained on biased materials.

Computers and Artificial Intelligence

Current research demonstrates that computers and artificial intelligence are also capable of inheriting human biases, such as traditional gender ideology. In fact, some assert that computer algorithms are inherently biased because their training materials are created by imperfect humans who hold prejudices and stereotypes (O'Neil 2016). That is, because language is a human creation that reflects their values, beliefs, and customs, it so happens that when language is used to train computer algorithms, computers will also inherit those human cultural tendencies embedded within it.

Research demonstrates the power of this algorithmic bias to impact real-life outcomes. Biased algorithms have been shown to generate runaway feedback loops, which use prior information to maximize future predictions. If prior information is biased and the loop is iterative, it produces an increasingly biased baseline that generates predictions through compounding sampling bias, thus reinforcing stereotypical beliefs. For instance, if you train a computer to predict crime hotspots using past crime data, it will reinforce its expectation for crime in places where it knows crimes have already taken place (Ensign et al. 2018). Machines have also been shown to replicate existing inequalities or discriminatory associations by race, gender, and employment opportunities (Barocas and Selbst 2016; Caliskan et al. 2017).

Yet, much like implicit bias, this inherited bias emerges outside of the intention or awareness of the programmer. This is true even despite deliberate strategies used to avoid biased outcomes (Caliskan et al. 2017), such as diversity measures (Sweeney 2013). That is, algorithmic bias is not a product of a biased programmer or programming but of biased inputs (Barocas and Selbst 2016). It appears that "machine prejudice derives so fundamentally from human culture... *semantics*, the meaning of words, necessarily reflects regularities latent in our culture, some of which we now know to be prejudiced" (Caliskan et al. 2017:2). These findings suggest that algorithmic fairness is difficult, if not impossible, to achieve (Friedler et al. 2016).

To the extent that this is true, social scientists can co-opt the capacity for computer code to reflect human biases to investigate and make claims about existing and historical biases in social life. For instance, machines trained on data from the

World Wide Web have also replicated known human stereotypes gleaned through traditional IATs (Caliskan et al. 2017). Just as implicit bias in humans renders quicker associations between men and career/science terms and women with family/arts, so too do computer models come to learn stereotypical gender associations, among others (Caliskan et al. 2017). Other researchers have used this method to examine the clustering of race, class, and gender terminology to conceptually map American culture (Kozlowski et al. 2019), the presence of implicit racial and gender bias in written judicial opinions (Ornaghi et al. 2018), and the correlation of gender and ethnic stereotypes with demographic changes, occupational shifts, and social movements in American society (Garg et al. 2018). Similar racial and gender biases have been demonstrated in computer-analyzed visual images (Buolamwini and Gebru 2018; Zhao et al. 2017).

We build upon these studies to perform the first sociological analysis of stereotypically gendered domains of social life (career, science, family, and arts) as they manifest in writing. We move beyond the use of case studies and present-day text to instead characterize change over time. Using a corpus of English-language books published from 1800 to 2000 as training material, we estimate the magnitude of traditional gender ideology over 200 years of time. By doing so, we capture a comprehensive picture of traditional gender ideology as it manifests through language and its trajectory over time. Moreover, we use documented trends in public opinion about traditional gender ideology and its manifestation in social structures to inform our analysis. Specifically, given (1) the decreasing but remaining presence of traditional gender ideology today, particularly through implicit attitudes associations, and (2) evidence to suggest that computer models inherit human biases to replicate implicit attitudes, we anticipate:

H1: Stereotypical gender associations reflecting gender ideology in English printed text have steadily decreased over time but still exist.

At the same time, we expect that traditional gender ideology, however lessened, continues to operate in stereotypically gendered ways. Specifically, we anticipate:

H2a: Career and science terms are more closely associated with male terms, whereas family and arts terms are more closely associated with female terms.

H2b: These stereotypical gender associations are decreasing over time.

We further posit the translation of this second shift ideology in English printed text, therefore building upon prior gender research to suggest that the second shift ideology is so entrenched in society that it permeates the written word. That is, just as we expect our corpus to reflect the general trends of gender bias in society, so too do we expect that it will reflect the imbalanced change in expectations of men and women. Simply, we anticipate that a quantitative second shift will emerge. Specifically, we expect:

H3a: Female–family associations and male–career associations will remain relatively stable over time.

H3b: Female–career associations and male–family associations will increase, but female–career associations will increase more than male–family associations.

Data and Methods

We will illustrate changes in gender-domain associations through analysis of word embeddings. A word embedding places words in a high-dimensional latent space that encodes many dimensions of meaning (Mikolov et al. 2013a, 2013b). Words that are similar in meaning are placed close together in this space. The distance between words can thus be interpreted as their conceptual dissimilarity. This distance can be measured numerically and continuously, thus allowing direct comparisons of the relative association between pairs of words.

Word embeddings are compact representations of language. A word embedding is constructed algorithmically; a large quantity of text is the input, and the output is the vocabulary from that text arranged in a space. The text data constrain the output such that words are bound more tightly together the more often they appear in the same document and the closer they are to each other within the text. Secondary associations also constrain the output: two words that appear in similar contexts but not together (think of synonyms) will also be placed close together in the representational space.

This method relies crucially upon the distributional hypothesis (Firth 1957; Harris 1954). The distributional hypothesis asserts that the context surrounding a word is a good representation of the meaning of that word. In short, “a word is characterized by the company it keeps” (Firth 1957:11). Most techniques of natural language processing (the computer science subdiscipline concerned with textual data) assume the distributional hypothesis. In the specific case of word embeddings, the distributional hypothesis is instantiated by clustering words together when they frequently and closely co-occur in text and when the words appear in similar contexts. Specifically, words are represented as vectors in space, one word to a vector. It has been demonstrated that the relative position of words in the constructed space reflect syntactic and semantic relationships recognizable to human readers (Mikolov et al. 2013a). For example, consider how one should complete *France* : *Paris* as *Italy* : _____. In a word embedding, subtracting two word vectors results in a third vector that can be thought of as a relationship: *France* – *Paris* = a vector representing *capital city of*. Adding that vector to the vector of another country is a movement in representation space in the direction of the computed relationship. Word embeddings correctly complete analogies such as these (*France* : *Paris* as *Italy* : *Rome*) for both semantic relationships (*king* : *man* as *queen* : *woman*) and syntactic relationships (*walking* : *walked* as *swimming* : *swam*).

The primary use of the word embeddings technique is in applications of artificial intelligence. Its advantage is that it converts text (and specifically words) from a symbolic form opaque to mathematical manipulation into numerical values perfectly suited to algebraic algorithms. For our purposes, however, we merely rely on the fact that distances in the representational space reflect the level of association between words in the source texts.

Our aim is to characterize the association between genders and domains reflected in published English texts over time. To do so we used *HistWords: Word Embeddings for Historical Text* (Hamilton, Leskovec, and Jurafsky 2016). These public domain data sets represent word embeddings created from corpora of published books, subset by decade of publication. The intermediate source of the text is the Google Books Ngram Corpus, an aggregated compendium of all words and phrases as they appear in millions of scanned and digitized books (Lin et al. 2012). The Google Books Ngram Corpus was derived from the Google Books project, which comprehensively digitized the book collections of more than 40 university libraries (Michel et al. 2011). These collections represent a substantial portion of all English text ever produced; the previous citation specifically states that 5,195,769 books are included, which corresponds to more than 361 billion words of text and an estimated coverage of 4 percent of all books ever published. The books are those selected for inclusion in university library collections, and one might speculate as to what selection bias this might introduce. However, a pure random selection of all English text from the past two centuries is not otherwise available. The Google Books corpus is available, not obviously deficient in any particular way, and importantly is the source material for much of computational modeling of language. For our purposes, the HistWords embeddings based on this corpus provide the best available cross-sectional estimates of the semantic association between words over time. See the online supplement for more detail on the construction of the HistWords embeddings.

To measure associations between genders and domains using this method, we must operationalize gender, domains, and association. We represent the concepts of male and female gender with lists of words (all word lists appear in Table 1). These lists are drawn from previous studies of implicit association conducted in human samples (Nosek et al. 2002a; 2002b). Similarly, we represent our domains of interest—arts, career, family, and science—with word lists previously employed in these studies. We define association as the complement of the mean distance in a word embedding between all combinations of words between two lists. In other words, we estimate the distance between a gender and a domain by choosing every possible pairing of a gender word and a domain word, calculating the distance between the corresponding word vectors and averaging these distances. Because words are represented as vectors in the embedding, the natural measure of distance is the cosine distance. A cosine distance varies from -1 to 1. Two vectors overlaid on top of each other will have distance zero, and the distance will move away from zero as the two vectors become more misaligned. Similarity is the complement of the distance—it is one for two indistinguishable concepts and zero for concepts so dissimilar they are orthogonal in the latent space. Informally, one may consider the cosine similarity as a score indicating the strength of the “belief” present in the word embedding model that two concepts belong close together.

We take one further step to quantify stereotypical gender associations. We define male gender bias as the association between a domain and the male gender minus the association between that domain and the female gender. Positive values thus indicate a domain is associated more closely with the male gender, and negative values indicate a domain is associated more closely with the female

Table 1: List of Implicit Association Test words

Female	Male	Family	Career	Science	Arts
female	male	home	executive	science	poetry
woman	man	parents	management	technology	art
girl	boy	children	professional	physics	Shakespeare
sister	brother	family	corporation	chemistry	dance
she	he	cousins	salary	Einstein	literature
her	him	marriage	office	NASA	novel
hers	his	wedding	business	experiment	symphony
daughter	son	relatives	career	astronomy	drama

Notes: NASA, National Aeronautics and Space Administration.

gender. Formal definitions of association are presented in the online supplement, and we have made the code we produced for this research available online at <https://github.com/ruhulbsu/StereotypicalGenderAssociationsInLanguage>.

Unlike supervised machine learning methods that use “ground truth” or human-supplied labels to guide the machine’s learning and attempt to mimic what has been provided as “correct” output, the word embedding approach is an unsupervised method. Unsupervised methods recover statistical regularities in the input data (published text in this application) without guidance. Therefore, should the method place a domain closer to one gender than another, this should be interpreted as an outcome conditioned by the input data and not the intentions of the algorithm’s creators or of the personified algorithm itself. As Barocas and Selbst (2016) aptly caution, “an algorithm is only as good as the data it works with” (P. 671).

In sum, we quantify the extent to which implicit gender stereotypical associations known to be expressed by humans are reproduced in the latent structure of written English. In so doing, we effectively use word embeddings to administer an algorithmic Gender IAT test to books written in English between the years 1800 and 2000. We expand upon previous work that uses machine learning to examine the gender bias in written English (Caliskan et al. 2017; Garg et al. 2018; Kozlowski et al. 2019; Ornaghi et al. 2018) to trace the various trajectories of several different gender domains and compare them to advances made toward gender equality in social life.

Results

The primary result is depicted in Figure 1. This chart places a composite measure of stereotypical gender bias on the *y* axis and time on the *x* axis. Specifically, the mean of the absolute values of the male gender bias score for each of the four domains is presented for each decade. The trend is clearly toward decreasing stereotypical gender bias over time, supporting H1. By this measure, stereotypical gender bias observable in language has decreased nearly threefold from 1800 to 2000 ($0.054/0.019 = 2.84$). A linear regression provides an estimate for the per-decade decrease as 0.0014 (95 percent confidence interval: 0.0009 to 0.0018).

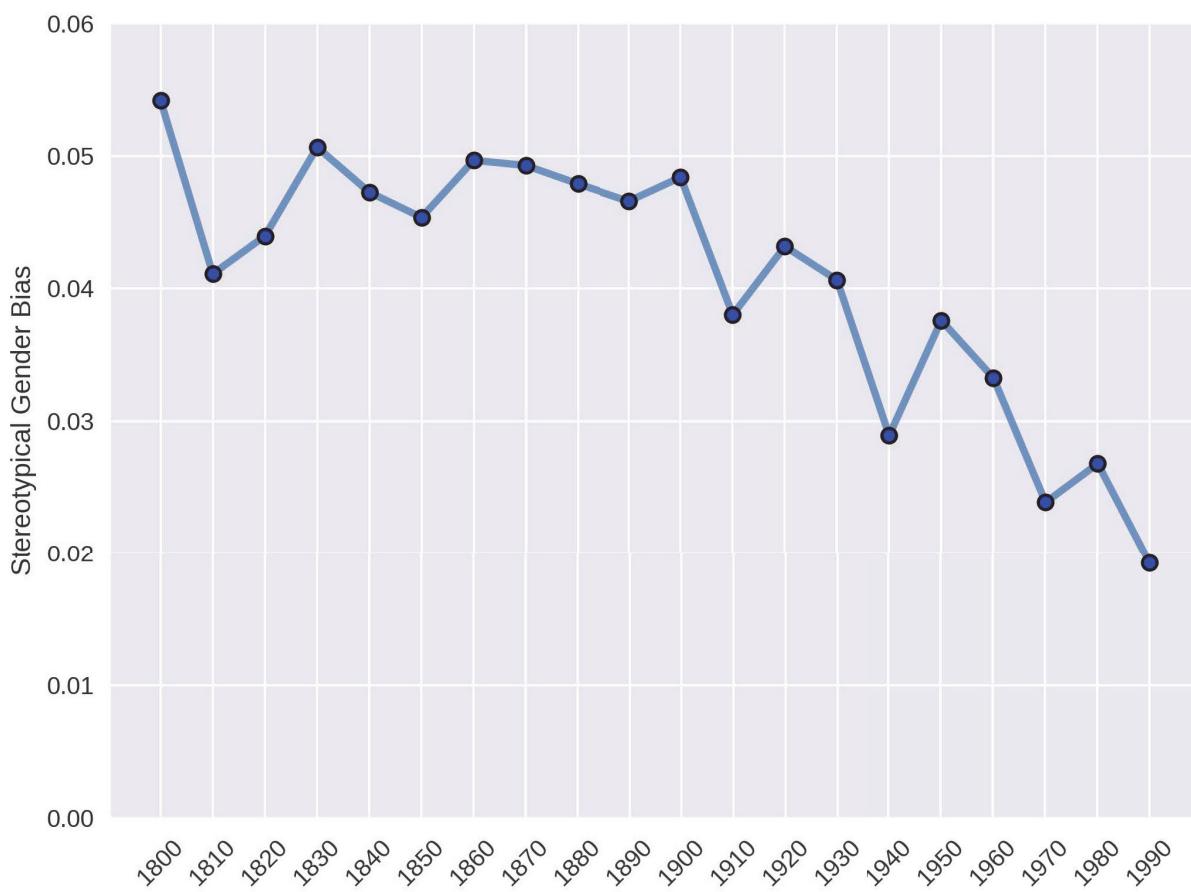


Figure 1: Composite stereotypical gender bias measured over time. Each data point represents the average of the absolute values of male gender bias for the four domains of Family, Career, Science, and Arts. The x axis is labeled with the first year of each decade (e.g., the label “1910” is used to mark the measurement based on the word embedding trained on text from 1910 up to but not including 1920).

Disaggregated trends for each domain are depicted in Figure 2. Each data point represents the domain’s male gender bias (male association minus female association) measured within a decade’s worth of published texts. Overall, the domains of Career and Science are male-associated, whereas Family and Arts are female-associated. In every case, the extent of the stereotypical association is decreasing over time. This is made clear in Figure 2 by the regression lines overlaid on the data points and their convergence toward zero.

Table 2 contains estimates of a linear regression of male gender bias over decades for each domain. Even with an overly conservative Bonferroni correction (Dunn 1961; Hochberg 1988), we can reject the null hypothesis of a zero-slope linear fit in every case except Science. (In the Family, Career and Arts regressions, the p values for a test of reliable difference from zero for the decade coefficients are all less than $0.05/4 = 0.0125$. The test is significant at the 0.05 level in isolation for the Science

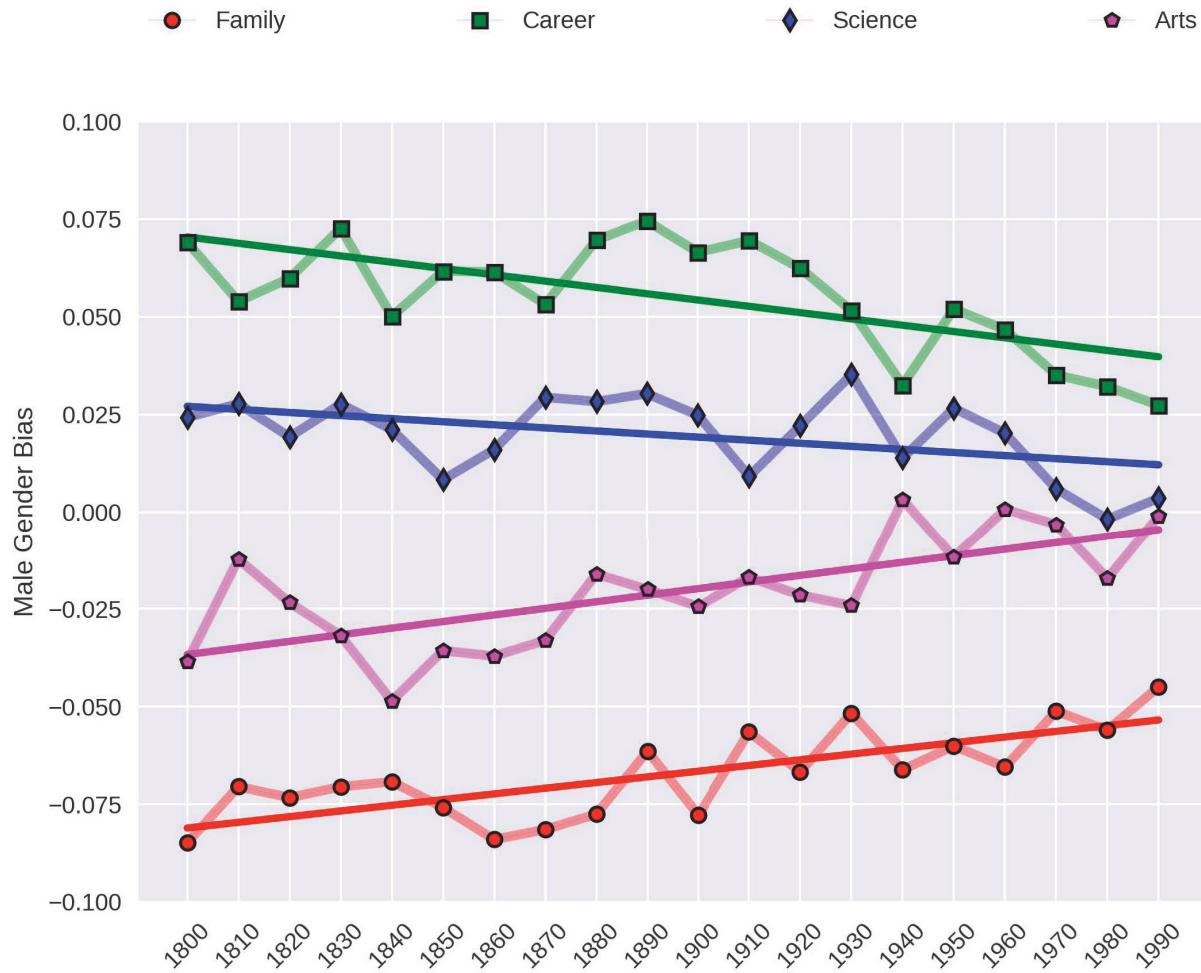


Figure 2: Male gender bias as measured by word embeddings in each decade of text for the four domains of Family, Career, Science, and Arts. Male gender bias is calculated by subtracting female cosine similarity from the male cosine similarity per each domain. The x axis is labeled with the first year of each decade (e.g., the label “1910” is used to mark the measurement based on the word embedding trained on text from 1910 up to but not including 1920).

regression, but does not meet the overly conservative Bonferroni correction cutoff.) These results support H2a and H2b and indicate that although gender stereotypical associations are decreasing over time, they still operate in predictably stereotypical ways. Career and science terms remain associated with men, and family and arts with women.

The results may be even further disaggregated to examine male association and female association within each domain separately. Figures 3 through 6 present the trends in these measures over time. Focusing on the data within Figures 3 and 4 will allow for testing of H3a and H3b. Specifically, we may examine the results of separate linear regressions for each series. These results are presented in Table 3. The coefficients have a straightforward interpretation: the estimate represents the

Table 2: Linear regression model of male gender bias over decade by domain

Domain	Per-Decade Estimated Change
Family	0.001456 [†] (0.000299)
Career	-0.001614 [†] (0.000428)
Science	-0.000788 [*] (0.000359)
Arts	0.001690 [†] (0.000400)
Number of Observations	20

Notes: The first number represents the regression coefficient, and the second number (in parentheses) represents the standard error. All domains remain significant after a Bonferroni correction except Science.

* $p < 0.05$, [†] $p < 0.01$.

average change in association per decade. The female–family and male–career coefficient estimates are consistent with H3a. The estimates for the magnitude of change in these associations are small compared to effects found elsewhere in these results. The coefficients do not meet the criterion for significant difference from zero. Specifically, 95 percent confidence intervals for the per-decade change in association are (-0.000715, 0.001939) for female–family and (-0.001070, 0.000965) for male–career. Simply, there has been no significant change in women’s association with family or men’s association with career.

H3b finds only partial support in the results. Confidence intervals for the per-decade change in association are (0.000810, 0.002313) for female–career and (0.000877, 0.003259) for male–family. Both are reliably positive, which means that female and career terms became more closely associated over time while concurrently, male and family terms became more associated. However, the supposition that the rate of increase would be greater for the female–career association was not borne out. A Wald test—in which equality of the coefficients is the null hypothesis—failed to reject that null with a p value of 0.22.

Discussion

Patterns of male and female associations with domains emerged as expected. Overall, gender bias decreased, as predicted by H1, as we sweep through the text published in books from 1800 to 2000. H2a and H2b are also supported by trends in the data. Career and Science terms are placed in conceptual space closer to male terms, and Family and Arts terms are placed closer to female terms. The magnitude of bias decreases over time, with all regression lines converging toward but not reaching zero. We take these results as evidence for changes to traditional gender ideology. Every temporal trend in the data points toward converging representations of the concepts of male and female. Over time, the domains of career and science are less

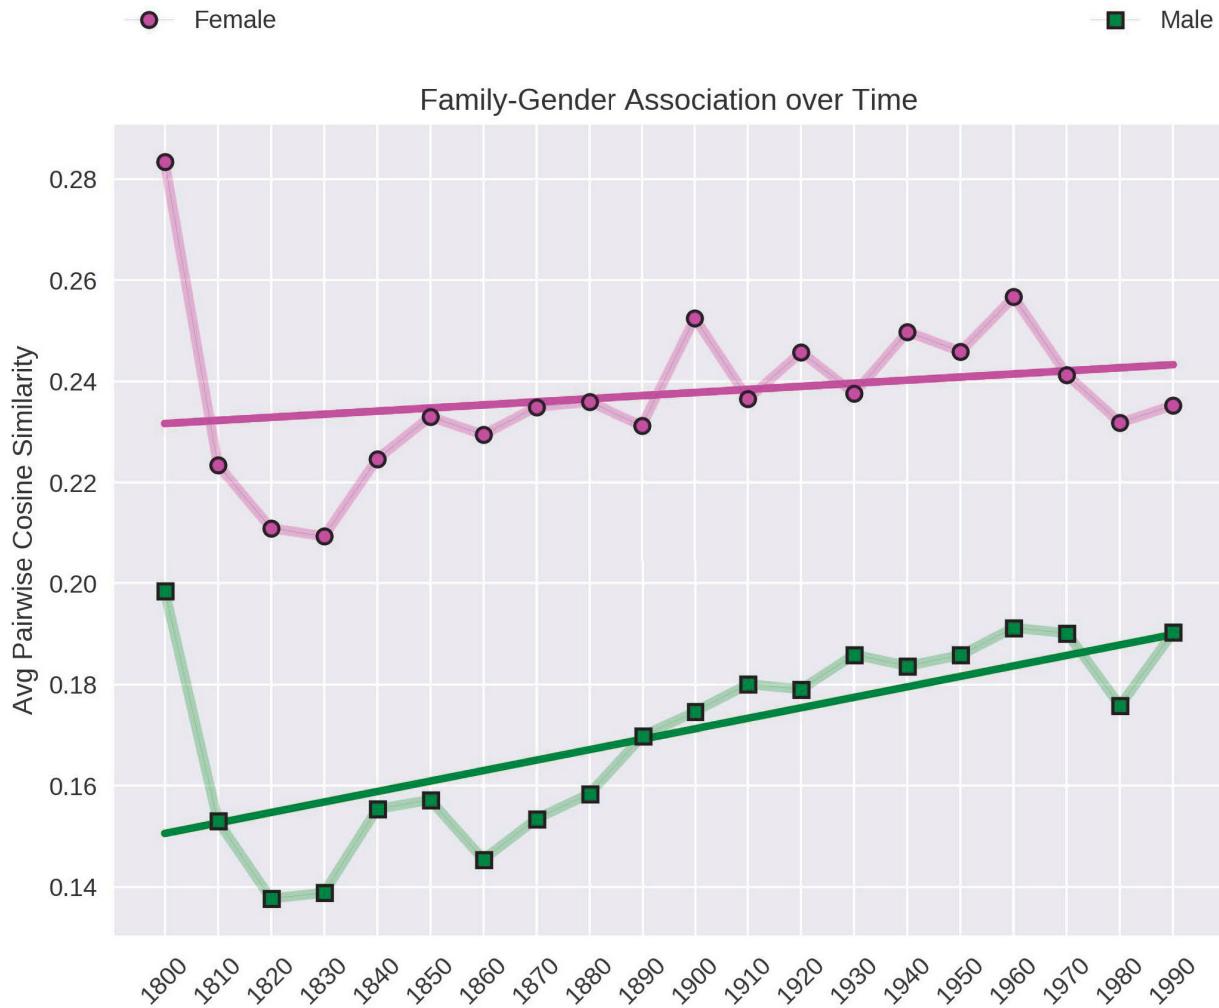


Figure 3: Female association and male association with the Family domain measured by decade. The x axis is labeled with the first year of each decade (e.g., the label “1910” is used to mark the measurement based on the word embedding trained on text from 1910 up to but not including 1920). Avg, average.

exclusively coupled to maleness in authors’ minds (as revealed by their language use), whereas family and arts are less exclusively linked to femaleness.

Yet these results challenge the notion that differential gender associations have disappeared. Despite general evidence for decreasing male bias over time, stereotypical associations remain. If we assume these millions of books are a relatively accurate representation of social mores, then we can assume that people still put stock in traditional gender ideology, perhaps outside of individual (or even societal) consciousness. Even when adopting a more conservative level of inference, we can safely conclude that the sum representation of our society through writing continues to be stereotypically biased. It is crucial to note here that the process of putting gender ideology into books is not necessarily a deliberate effort. In fact, it is

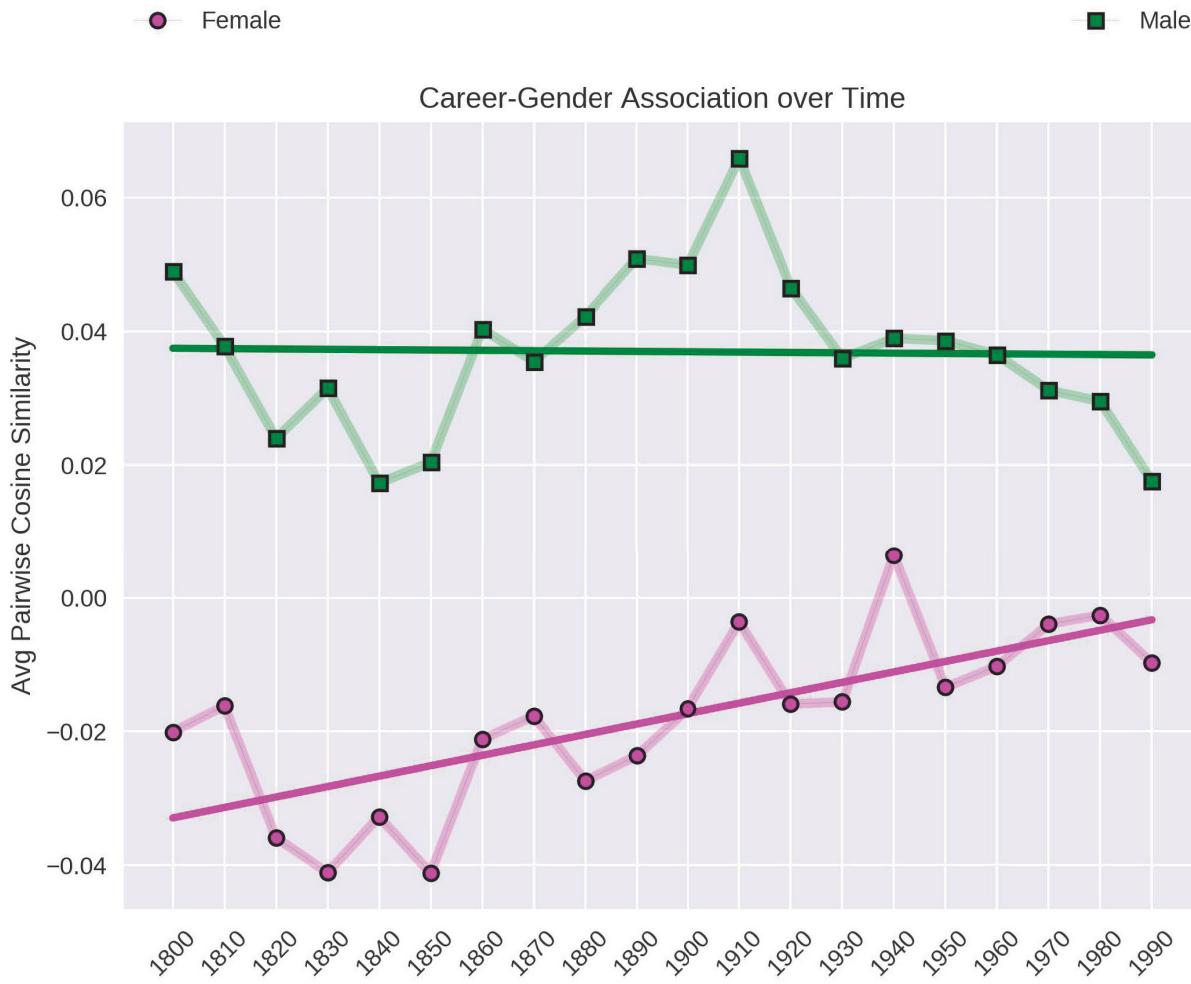


Figure 4: Female association and male association with the Career domain measured by decade. The x axis is labeled with the first year of each decade (e.g., the label “1910” is used to mark the measurement based on the word embedding trained on text from 1910 up to but not including 1920). Avg, average.

much more likely these biases are subtle and unintentionally incorporated, which perhaps makes their presence all the more worrisome.

This actuality has several implications for society. Although some may feel that we as a society are beyond issues of gender inequality (Fingerhut 2016), these word embeddings contribute evidence to suggest that we are not, even if we fail to realize it. This makes it especially crucial to expose persisting inequalities where they remain. These results raise the natural follow-up question: where else does traditional gender ideology operate without explicit intention or knowledge? Of course, this question is impossible to answer without further investigation. But these results should caution against assuming the impartiality of not explicitly gendered entities and spaces and should motivate us to question the extent to which we receive latent exposure to stereotypically gendered content.

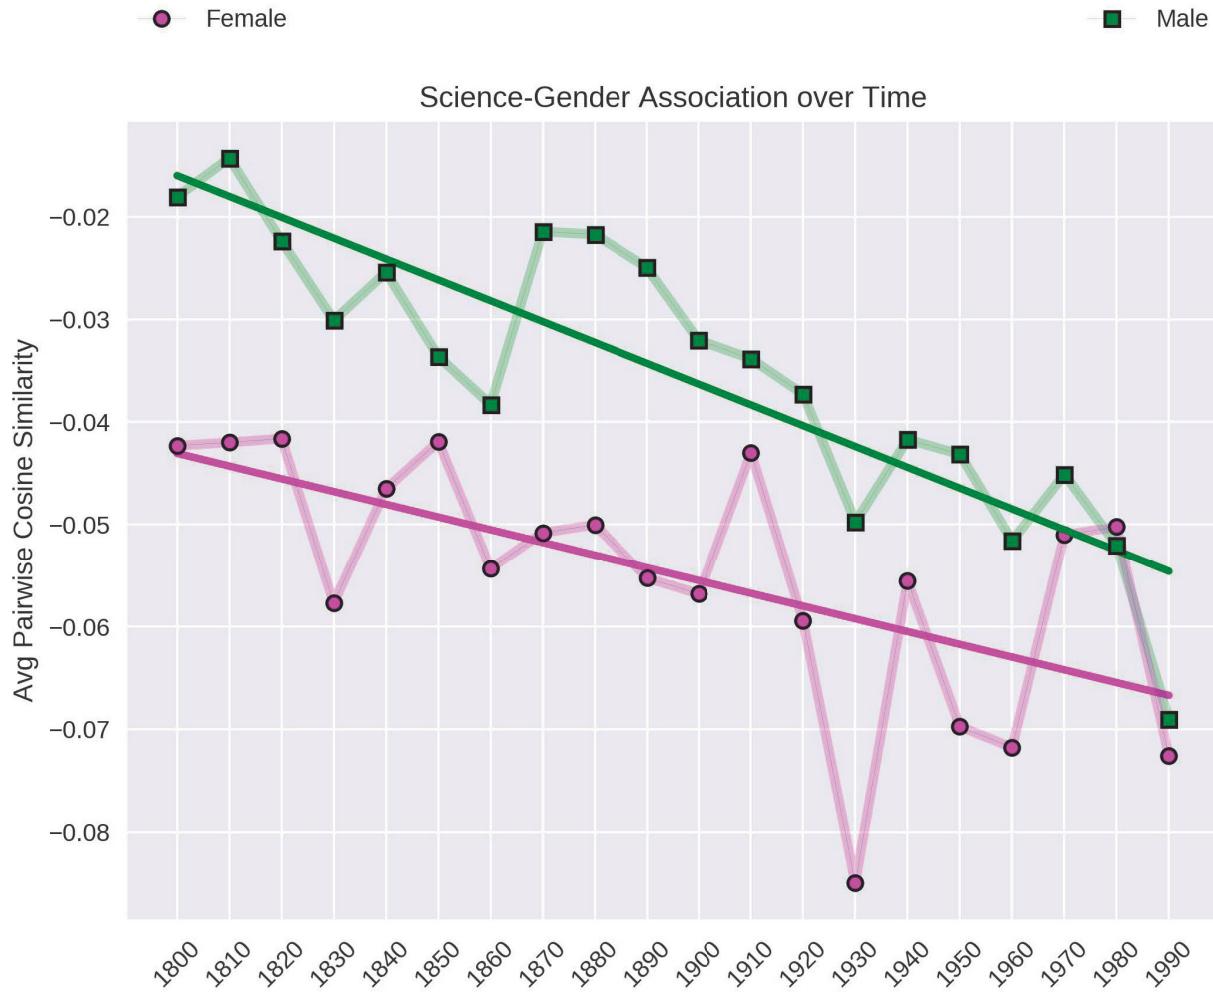


Figure 5: Female association and male association with the Science domain measured by decade. The x axis is labeled with the first year of each decade (e.g., the label “1910” is used to mark the measurement based on the word embedding trained on text from 1910 up to but not including 1920). Avg, average.

This is especially consequential given the high accessibility of books and their capacity to influence our beliefs. Books do not exist in a void; they are produced to be read. A large literature demonstrates that the things people read not only contain gender bias (Amare 2007; Amini and Birjandi 2012; Benjamin and Irwin-DeVitis 1998; Gooden and Gooden 2001; Schmader et al. 2007; Trix and Psenka 2003) but are also capable of reinforcing traditional gender stereotypes (Banaji and Hardin 1996; Friesen et al. 2011; Lassonde and O’Brien 2013; Madera et al. 2009; Stout and Dasgupta 2011). Together, these findings indicate the power of language to reflect and promote traditional gender ideology. If left unaddressed, extreme levels of male bias within published books may serve to perpetuate traditional gender ideology, potentially creating a cyclical pattern of ideology expression and reinforcement.

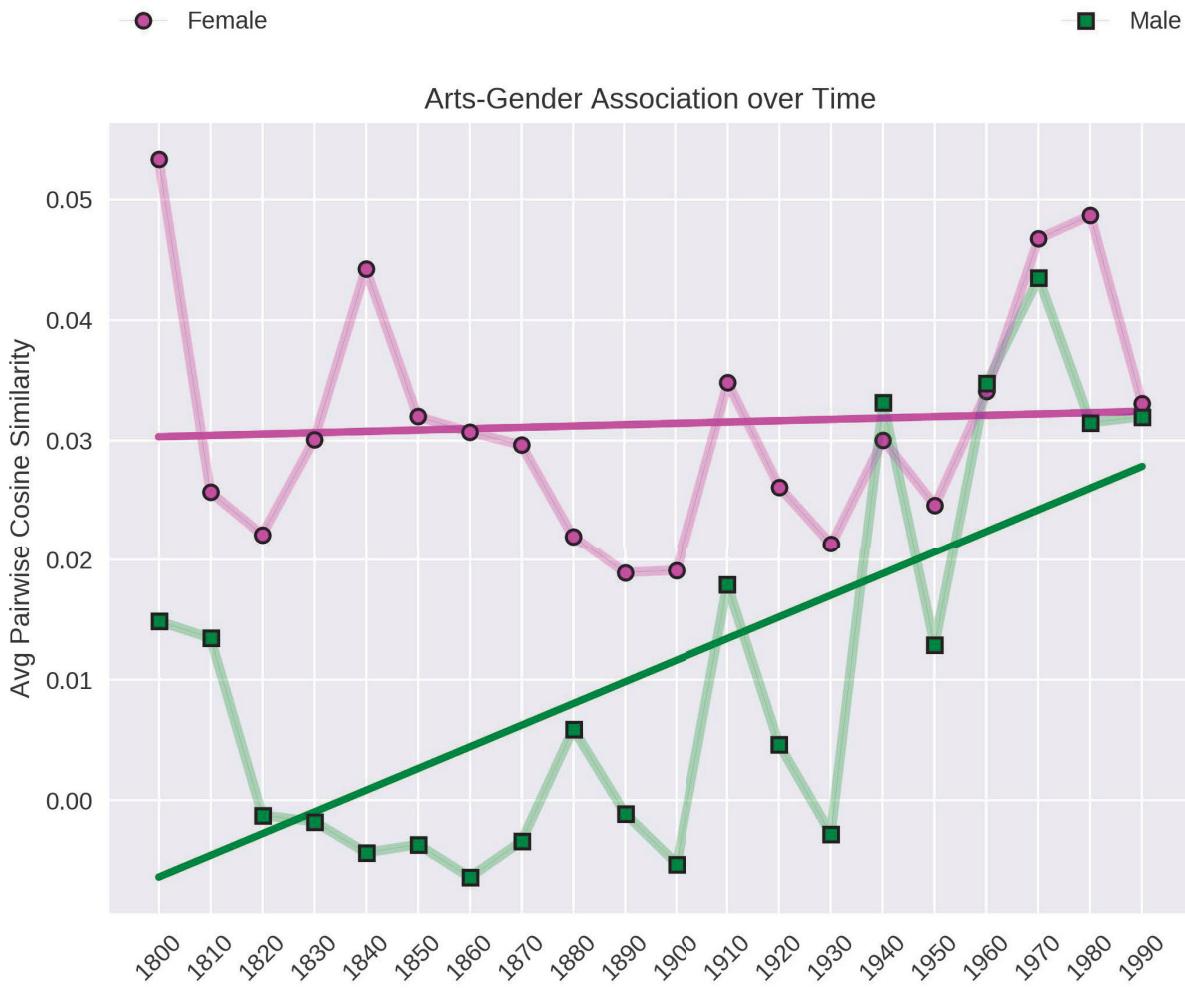


Figure 6: Female association and male association with the Arts domain measured by decade. The x axis is labeled with the first year of each decade (e.g., the label “1910” is used to mark the measurement based on the word embedding trained on text from 1910 up to but not including 1920). Avg, average.

Moreover, the results presented in Table 3 and Figures 3 and 4 demonstrate that female–family associations and male–career associations do remain relatively stable over time. Combined with the partial support of H3b—female–career and male–family associations increase—an interesting interpretation suggests itself. The overall change in gender bias in the work–home dichotomy is driven by the formerly dissociated gender. It has become more common to perceive women as belonging in the sphere of work and men belonging in the sphere of home. Rather than existing stereotypical associations fading, perceptions of both men and women trend toward the opposite gender’s fixed position. This suggests that although the content of books allows men and women to associate with traditionally feminine and masculine coded things, respectively, it is changing little in the way of moving away from status quo expectations.

Table 3: Linear regression model of male gender bias over decade by gender-domain

Domain	Per-Decade Estimated Change
Female–Family	0.000612 (0.000631)
Female–Career	0.001561 [†] (0.000358)
Male–Family	0.002068 [†] (0.000567)
Male–Career	–0.000052 (0.000484)
Number of Observations	20

Notes: The first number represents the regression coefficient, and the second number (in parentheses) represents the standard error. [†] $p < 0.01$.

Additionally, we expected to observe differential rates of change for female–career and male–family associations, but the data do not support that prediction. That is not to say that expectations of women to be more involved at work are matched by expectations of men to be more involved at home. It may be the case that this expectation is intentionally unstated or understated in published text. Or it may be that this three-way interaction of time, gender, and domain stretches this analysis technique and data set past its breaking point.

Finally, we draw conclusions regarding the problem of algorithmic bias. Previous work has raised the alarm that automated systems trained on historical data will inherit the biases of the historical systems that created those data (Caliskan et al. 2017; O’Neil 2016). The current findings do not obviate that concern, but they do lead to a more hopeful conclusion. If historical data increasingly reflects less biased associations—as we argue it does—then automated systems will become increasingly unbiased. Data shape machine learning outputs. That is what makes them powerful and what necessitates paying careful attention to the input data. Yet we must also remain attentive to the fact that gender bias in language is only becoming less biased in some ways: men and women remain affixed to their traditional domains despite increasing associations between each gender and their respective nontraditional domains.

Indeed, it is useful to imagine the word-embedding systems trained here as subjects in an IAT experiment. The more recent the training text, the weaker a gender-stereotypical implicit attitude effect the system would display. Cross-sectional studies of gender IAT effects provide an interesting set of comparisons to these results. Humans acquire language, knowledge, and attitudes over years, and if books play a meaningful part in shaping this acquisition, we would expect to see similar trends over age. In Nosek et al. (2007), the authors report the results of many different implicit attitude tests based on 2.5 million self-administered online sessions. Consistent with the current results, stereotypical gender associations are strongest for older adults and weakest among younger adults. The more recent the human or machine’s “socialization,” the weaker the expression of traditional gender ideology.

Limitations

Several limitations apply to the results presented here. Foremost is ambiguity in interpreting the observed changes in association strength over time. The process of constructing word embeddings is a process of finding patterns in the observed texts. When a pattern shows that women are less associated with science, this can be due to collectively held stereotypes (as we emphasize) but also empirically true associations. It is likely true that women became increasingly represented in the population of scientists over the span of our data set, so some amount of increasing association of women and science in the embeddings would be present even if the text was produced by a completely bias-free neutral observer. If a book about famous scientists was published in 1850, and all the scientists discussed were male, this could be due to collectively held stereotypes at that moment, the individual author's bias or particular gender ideology, or social structures that made it harder to be a female than a male recognized scientist at that time. Interpreting the observed trends in conceptual association is complicated by this mix of potential causal processes. However, the empirical results remain. Language use is shifting toward less pronounced gender bias.

Another limitation comes from the corpus. Although the data set is large, it represents published books, and obviously most humans live their lives never having authored a book. Thus, the population whose gender associations this method may measure is atypical and not representative of a general population in terms of demographic background. At the same time, books and reading are practically impossible to avoid. Even if those producing books are not representative of the population, the population at large is not immune from the effects of consuming their material. Furthermore, the books were selected for their presence in a university library. This means the source material is both mixed (fiction and nonfiction are both included, for example) and curated. One might speculate that scientific writing would be overrepresented in comparison to the books one might find by randomly sampling bookstores or the bookshelves of individuals. Such alternative samples (or other pristine random samples of text) are not readily available at all, much less in the quantities needed to properly train word embedding models.

The corpus is also limited in that data are only available until the year 2000. Although 1800 to 2000 offers rich insight into the history of gender ideology in English, many social changes have occurred since then. This may help to explain why we achieved only partial support for H3b: the associations of the second shift may still be in the process of being written and simply aren't detectable in data prior to 2000.

Finally, this analysis does not speak to the causal mechanism for the reduction in stereotypical gender associations. Over 200 years of published books, many factors are changing. Perhaps the gender ideology of all English speakers is changing, and the published record reflects this. Or perhaps gender ideology remains relatively stable but the set of English speakers who come to publish books changes, and the new entrants are less extreme in their gender bias.

Further research could address the limitations of this work. Tagging books by author demographics (particularly gender) would be a challenging task but would

allow for contrasting within-demographic trends that are more likely to represent attitude change to between-demographic differences that may drive the overall trend through replacement. Similarly, in this work we include all books in search of broad trends, but we encourage future authors to investigate the role of genre on these results. In addition, these same techniques applied to other data sets would deepen the context in which to view the current results. It is an open question whether a similar trend would appear in historical corpora in other languages or from sources other than books. Finally, work on the convergent validity of implicit attitudes measured through separate means would be welcome. Specifically, it would be useful to measure the correlation between individuals' or groups' implicit attitudes measured through the IAT instrument and by word embedding analysis of their writing.

Concluding Remarks

Stereotypical gender associations latent in the text of published books have decreased over time. Extrapolating from the current trend implies a future in which gender biases that are currently taken for granted will shrink to zero or even reverse. Future studies of traditional gender ideology should further assess the validity and velocity of this trend and its applicability in different settings. Additionally, changes in conceptual similarity, revealed through word embeddings conditioned by historical text, should be studied far beyond the domain of gender. The unspoken relationships that infuse the collective thinking of two centuries of authors and readers are available for inspection using these data and techniques. We look forward to more research that undertakes mining of large digitized data sets to examine historical trends.

Notes

- 1 Although we acknowledge that gender identities expand far beyond the male-female binary, we restrict our discussion and analysis of gender ideology here to be most consistent with those dominant in Western societies that characterize gender using two gender categories, male and female, and understand these categories as "distinct and opposite" (Chatillon, Charles, and Bradley 2018:217).
- 2 Defined as "cultural constructs, shared at the societal level, that describe what men and women are "known" to be like" (Fiske 1998:377–8.).
- 3 Current surveys that include gender ideology questions include the General Social Survey; High School and Beyond; Intergenerational Panel Study of Parents and Children; International Social Survey Program; Marital Instability over the Life Course; National Longitudinal Survey of Youth, 1979 Cohort and Child/Young Adult Sample; National Study of the Changing Workforce; National Study of Families and Households; and the World Values Survey. This is not intended as a complete list.
- 4 Defined by Hayes (1996) as "a gendered model that advises mothers to expend a tremendous amount of time, energy, and money in raising their children" (P.x).
- 5 Defined as "being at risk of confirming, as self-characteristic, a negative stereotype about one's group" (Steele and Aronson 1995:797). In other words, people fearful that their

performance or behavior will confirm a stereotype about a group to which they belong, they perform poorly in light of this anxiety. In this case, women fearful of confirming stereotypes about women being bad at math tend to perform badly on standardized math tests (Smith and White 2002).

References

Acker, Joan. 1988. "Class, Gender, and the Relations of Distribution." *Signs* 13:473–97. <https://doi.org/10.1086/494429>.

Amare, Nicole. 2007. "Where is She? Gender Occurrences in Online Grammar Guides." *Research in the Teaching of English* 42:163–87.

Amini, Mohadeseh, and Parviz Birjandi. 2012. "Gender Bias in the Iranian High School EFL Textbooks." *English Language Teaching* 5:134–47. <https://doi.org/10.5539/elt.v5n2p134>.

Banaji, Mahzarin R., and Cutis D. Hardin. 1996. "Automatic Stereotyping." *Psychological Science* 7:136–41. <https://doi.org/10.1111/j.1467-9280.1996.tb00346.x>.

Banse, Rainer, Jan Seise, and Nikola Zerbes. 2001. "Implicit Attitudes towards Homosexuality: Reliability, Validity, and Controllability of the IAT." *Experimental Psychology* 48:145–60. <https://doi.org/10.1026/0949-3946.48.2.145>.

Barman, Charles R. 1997. "Students' Views of Scientists and Science: Results from a National Study." *Science and Children* 35:18–23.

Barocas, Solon, and Andrew D. Selbst. 2016. "Big Data's Disparate Impact." *California Law Review* 104:671–732. <https://doi.org/10.2139/ssrn.2477899>.

Baron, Andrew Scott, and Mahzarin R. Banaji. 2006. "The Development of Implicit Attitudes: Evidence of Race Evaluations from Ages 6 and 10 and Adulthood." *Psychological Science* 17:53–58. <https://doi.org/10.1111/j.1467-9280.2005.01664.x>.

Bauer, Cara C., and Boris B. Baltes. 2002. "Reducing the Effects of Gender Stereotypes on Performance Evaluations." *Sex Roles* 47:465–76. <https://doi.org/10.1023/A:1021652527696>.

Beede, David, Tiffany Julian, David Langdon, George McKittrick, Beethika Khan, and Mark Doms. 2011. "Women in STEM: A Gender Gap to Innovation." *Economics and Statistics Administration* Issue Brief No. 04-11. <https://doi.org/10.2139/ssrn.1964782>.

Bem, Sandra Lipsitz. 1993. *The Lenses of Gender: Transforming the Debate on Sexual Inequality*. New Haven, CT: Yale University Press.

Benjamin, Beth, and Linda Irwin-DeVitis. 1998. "Censoring Girls' Choices: Continued Gender Bias in English Language Arts Classrooms." *The English Journal* 87:64–71. <https://doi.org/10.2307/821553>.

Berger, Peter L., and Thomas Luckmann. 1966. *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. New York, NY: Vintage Books.

Bianchi, Suzanne M., Melissa A. Milkie, Liana C. Sayer, and John P. Robinson. 2000. "Is Anyone Doing the Housework? Trends in the Gender Division of Household Labor." *Social Forces* 79:191–228. <https://doi.org/10.1093/sf/79.1.191>.

Blair-Loy, Mary. 2003. *Competing Devotions: Career and Family among Women Executives*. Cambridge, MA: Harvard University Press.

Brewster, Karin L., and Irene Padavic. 2000. "Change in Gender-Ideology, 1977–1996: The Contributions of Intracohort Change and Population Turnover." *Journal of Marriage and Family* 62:477–87. <https://doi.org/10.1111/j.1741-3737.2000.00477.x>.

Broverman, Inge K., Susan Raymond Vogel, Donald M. Broverman, Frank E. Clarkson, and Paul S. Rosenkrantz. 1972. "Sex-Role Stereotypes: A Current Appraisal." *Journal of Social Issues* 28:59–78. <https://doi.org/10.1111/j.1540-4560.1972.tb00018.x>.

Budig, Michelle J., and Paula England. 2001. "The Wage Penalty for Motherhood." *American Sociological Review* 66:204–25. <https://doi.org/10.2307/2657415>.

Buolamwini, Joy, and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." *Proceedings of Machine Learning Research* 81:77–91.

Burns, Nancy, Ashley E. Jardina, Donald Kinder, and Molly E. Reynolds. 2015. "The Politics of Gender." Pp. 124–45 in *New Directions in Public Opinion*, 2nd ed, edited by Adam J. Berinsky. New York, NY: Routledge.

Caliskan, Aylin, Joanna J. Bryson, and Arvind Narayanan. 2017. "Semantics Derived Automatically from Language Corpora Contain Human-Like Biases." *Science* 356:183–6. <https://doi.org/10.1126/science.aal4230>.

Chambers, David Wade. 1983. "Stereotypic Images of the Scientist: The Draw-a-Scientist Test." *Science Education* 67:255–65. <https://doi.org/10.1002/sce.3730670213>.

Charles, Maria, and David B. Grusky. 2004. *Occupational Ghettos: The Worldwide Segregation of Women and Men*. Stanford, CA: Stanford University Press.

Chatillon, Anna, Maria Charles, and Karen Bradley. 2018. "Gender Ideologies." Pp. 217–26 in *Handbook of the Sociology of Gender: Second Edition*, edited by B. J. Risman, C. M. Froyum, and W. J. Scarborough. New York, NY: Springer.

Cherlin, Andrew, and Pamela Barnhouse Walters. 1981. "Trends in United States Men's and Women's Sex-Role Attitudes: 1972 to 1978." *American Sociological Review* 46:453–60. <https://doi.org/10.2307/2095264>.

Coles, Martin, and Christine Hall. 2002. "Gendered Readings: Learning from Children's Reading Choices." *Journal of Research in Reading* 25:96–108. <https://doi.org/10.1111/1467-9817.00161>.

Collins, Laura J., Bron B. Ingoldsby, and Mary M. Dellmann. 1984. "Sex-Role Stereotyping in Children's Literature: A Change From the Past." *Childhood Education* 60:278–85.

Coltrane, Scott. 2000. "Research on Household Labor: Modeling and Measuring the Social Embeddedness of Routine Family Work." *Journal of Marriage and Family* 62:1208–33. <https://doi.org/10.1111/j.1741-3737.2000.01208.x>.

Connell, Raweyn W. 1987. *Gender and Power: Society, the Person and Sexual Politics*. Stanford, CA: Stanford University Press.

Correll, Shelley J., Stephen Bernard, and In Paik. 2007. "Getting a Job: Is There a Motherhood Penalty?" *American Journal of Sociology* 112:1297–339. <https://doi.org/10.1086/511799>.

Cotter, David, Joan M. Hermsen, and Reeve Vanneman. 2011. "The End of the Gender Revolution? Gender Role Attitudes from 1977 to 2008." *American Journal of Sociology* 117:259–89. <https://doi.org/10.1086/658853>.

Davis, Shannon N., and Theodore N. Greenstein. 2009. "Gender Ideology: Components, Predictors, and Consequences." *Annual Review of Sociology* 35:87–105. <https://doi.org/10.1146/annurev-soc-070308-115920>.

Deutsch, Francine M. 1999. *Halving it All: How Equally Shared Parenting Works*. Cambridge, MA: Harvard University Press.

DeVault, Marjorie L. 1990. "Conflict Over Housework: A Problem that (Still) Has No Name." Pp. 189–202 in *Research in Social Movements, Conflict and Change*, edited by L. Kriesberg. Greenwich, CT: JAI Press.

Diekman, Amanda B., and Alice H. Eagly. 2000. "Stereotypes as Dynamic Constructs: Women and Men of the Past, Present, and Future." *Personality and Social Psychology Bulletin* 26:1171–88. <https://doi.org/10.1177/0146167200262001>.

Diekman, Amanda B., and Wind Goodfriend. 2006. "Rolling with the Changes: A Role Contiguity Perspective on Gender Norms." *Psychology of Women Quarterly* 30:369–83. <https://doi.org/10.1111/j.1471-6402.2006.00312.x>.

Diekman, Amanda B., Wind Goodfriend, and Stephanie Goodwin. 2004. "Dynamic Stereotypes of Power: Perceived Change and Stability in Gender Hierarchies." *Sex Roles* 50:201–15. <https://doi.org/10.1023/b:sers.0000015552.22775.44>.

Douglas, Susan J., and Meredith W. Michaels. 2004. *The Mommy Myth: The Idealization of Motherhood and How It Has Undermined Women*. New York, NY: Free Press.

Dovidio, John F., Kerry Kawakami, Craig Johnson, Brenda Johnson, and Adaiah Howard. 1997. "On the Nature of Prejudice: Automatic and Controlled Processes." *Journal of Experimental Social Psychology* 33:510–40. <https://doi.org/10.1006/jesp.1997.1331>.

Dunn, Olive Jean. 1961. "Multiple Comparisons among Means." *Journal of the American Statistical Association* 56:52–64. <https://doi.org/10.2307/2282330>.

Eagly, Alice H., and Valerie J. Steffen. 1984. "Gender Stereotypes Stem from the Distribution of Women and Men into Social Roles." *Journal of Personality and Social Psychology* 46:735–54. <https://doi.org/10.1037/0022-3514.46.4.735>.

Eccles, Jacquelynne S., Janis E. Jacobs, and Rena D. Harold. 1990. "Gender Role Stereotypes, Expectancy Effects, and Parents' Socialization of Gender Differences." *Journal of Social Issues* 46:183–201. <https://doi.org/10.1111/j.1540-4560.1990.tb01929.x>.

Ely, Richard, and Elizabeth Ryan. 2008. "Remembering Talk: Individual and Gender Differences in Reported Speech." *Memory* 16:395–409. <https://doi.org/10.1080/09658210801949869>.

England, Paula. 2010. "The Gender Revolution: Uneven and Stalled." *Gender and Society* 24:149–66. <https://doi.org/10.1177/0891243210361475>.

Ensign, Danielle, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. "Runaway Feedback Loops in Predictive Policing." *Proceedings of Machine Learning Research* 81:1–12.

Epstein, Cynthia Fuchs, Robert Sauté, Bonnie Oglensky, and Martha Gever. 1995. "Glass Ceilings and Open Doors: Women's Advancement in the Legal Profession." *Fordham Law Review* 64:291–449.

Estes, Zachary, and Sydney Felker. 2012. "Confidence Mediates the Sex Difference in Mental Rotation Performance." *Archives of Sexual Behavior* 41:557–70. <https://doi.org/10.1007/s10508-011-9875-5>.

Evans, Lorraine, and Kimberly Davies. 2000. "No Sissy Boys Here: A Content Analysis of the Representation of Masculinity in Elementary School Reading Textbooks." *Sex Roles* 42:255–70. <https://doi.org/10.1023/A:1007043323906>.

Fagot, Beverly I., Carie S. Rodgers, and Mary D. Leinbach. 2000. "Theories of Gender Socialization." Pp. 65–89 in *The Developmental Social Psychology of Gender*, edited by T. Eckes and H. M. Trautner. Mahwah, NJ: Lawrence Erlbaum Associates Publishers.

Feldberg, Roslyn L., and Evelyn Nakano Glenn. 1979. "Male and Female: Job Versus Gender Models in the Sociology of Work." *Social Problems* 26:524–38. <https://doi.org/10.1525/sp.1979.26.5.03a00050>.

Feree, Myra Marx. 1974. "A Woman for President? Changing Responses: 1958-1972." *Public Opinion Quarterly* 38:390-9. <https://doi.org/10.1086/268175>.

Feree, Myra Marx, and Beth B. Hess. 2000. *Controversy and Coalition: The New Feminist Movement across Three Decades of Change*. New York, NY: Routledge.

Fingerhut, Hannah. 2016. "In Both Parties, Men and Women Differ Over Whether Women Still Face Obstacles to Progress." *Pew Research Center*, August 16. Retrieved August 8, 2018, <http://pewrsr.ch/2b9DdPJ>.

Finson, Kevin D. 2002. "Drawing a Scientist: What We Do and Do Not Know After Fifty Years of Drawings." *School Science and Mathematics* 102:335-45. <https://doi.org/10.1111/j.1949-8594.2002.tb18217.x>.

Firth, John Rupert. 1957. *A Synopsis of Linguistic Theory, 1930-1955*. Oxford, UK: Basil Blackwell.

Fiske, Susan T. 1998. "Stereotyping, Prejudice, and Discrimination." Pp. 357-411 in *Handbook of Social Psychology*, edited by D. T. Gilbert, S. T. Fiske, and G. Lindzey. New York, NY: Oxford.

Fort, Deborah C., and Heather L. Varney. 1989. "How Students See Scientists: Mostly Male, Mostly White, and Mostly Benevolent." *Science and Children* 26:8-13.

Friedler, Sorelle A., Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. "On the (Im)possibility of Fairness." Retrieved August 8, 2018 (<https://arxiv.org/abs/1609.07236>).

Friesen, Justin, Danielle Gaucher, and Aaron C. Kay. 2011. "Evidence that Gendered Wording in Job Advertisements Exists and Sustains Gender Inequality." *Journal of Personality and Social Psychology* 101:109-28. <https://doi.org/10.1037/a0022530>.

Fullerton, Howard N, Jr. 1999. "Labor Force Participation: 75 Years of Change, 1950-98 and 1998-2025." *Bureau of Labor Statistics Monthly Labor Review*, December. Retrieved September 9, 2018, <https://www.bls.gov/opub/mlr/1999/12/art1full.pdf>.

Garg, Nikhil, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. "Word Embeddings Quantify 100 Years of Gender and Ethnic Stereotypes." *Proceedings of the National Academy of Sciences* 115:E3635-E3644. <https://doi.org/10.1073/pnas.1720347115>.

Gaunt, Ruth. 2006. "Biological Essentialism, Gender Ideologies, and Role Attitudes: What Determines Parents' Involvement in Child Care." *Sex Roles* 55:523-33. <https://doi.org/10.1007/s11199-006-9105-0>.

Gawronski, Bertram, and Galen V. Bodenhausen. 2007. "What Do We Know about Implicit Attitude Measures and What Do We Have to Learn?" Pp. 265-86 in *Implicit Measures of Attitudes*, edited by B. Wittenbrink and N. Schwarz. New York, NY: The Guilford Press.

Gerstel, Naomi, and Harriet Gross. 1987. *Families and Work*. Philadelphia, PA: Temple University Press.

Ginther, Donna K., and Shulamit Khan. 2006. "Does Science Promote Women? Evidence from Academia 1973-2001." Working Paper, National Bureau of Economic Research. <https://doi.org/10.3386/w12691>.

Gjerdigen, Dwenda K., and Bruce A. Center. 2005. "First-time Parents' Postpartum Changes in Employment, Childcare, and Housework Responsibilities." *Social Science Research* 34:103-16. <https://doi.org/10.1016/j.ssresearch.2003.11.005>.

Glauber, Rebecca. 2008. "Race and Gender in Families and at Work: The Fatherhood Wage Premium." *Gender and Society* 22:8-30. <https://doi.org/10.1177/0891243207311593>.

Glick, Peter, and Susan T. Fiske. 1996. "The Ambivalent Sexism Inventory: Differentiating Hostile and Benevolent Sexism." *Journal of Personality and Social Psychology* 70:491–512. <https://doi.org/10.1037/0022-3514.70.3.491>.

Glick, Peter, and Susan T. Fiske. 2001. "An Ambivalent Alliance: Hostile and Benevolent Sexism as Complementary Justifications for Gender Inequality." *American Psychologist* 56:109–18. <https://doi.org/10.1037/0003-066x.56.2.109>.

Gooden, Angela M., and Mark A. Gooden. 2001. "Gender Representation in Notable Children's Picture Books: 1995–1999." *Sex Roles* 45:89–101. <https://doi.org/10.1023/A:1013064418674>.

Greenwald, Anthony G., and Mahzarin R. Banaji. 1995. "Implicit Social Cognition: Attitudes, Self-Esteem, and Stereotypes." *Psychological Review* 102:4–27. <https://doi.org/10.1037/0033-295x.102.1.4>.

Greenwald, Anthony G., Debbie E. McGhee, and Jordan L. K. Schwartz. 1998. "Measuring Individual Differences in Implicit Cognition: The Implicit Association Test." *Journal of Personality and Social Psychology* 74:1464–80. <https://doi.org/10.1037/0022-3514.74.6.1464>.

Hamilton, William L., Jure Leskovec, and Dan Jurafsky. 2016. "Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change." Pp. 1489–501 in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Stroudsburg, PA: Association for Computational Linguistics. <https://doi.org/10.18653/v1/p16-1141>.

Harris, Zellig S. 1954. "Distributional Structure." *Word* 10:146–62. <https://doi.org/10.1080/00437956.1954.11659520>.

Hays, Sharon. 1996. *The Cultural Contradictions of Motherhood*. New Haven, CT: Yale University Press.

He, Guimei. 2010. "An Analysis of Sexism in English." *Journal of Language Teaching and Research* 1:332–5. <https://doi.org/10.4304/jltr.1.3.332-335>.

Hochberg, Yosef. 1988. "A Sharper Bonferroni Procedure for Multiple Tests of Significance." *Biometrika* 75:800–802. <https://doi.org/10.2307/2336325>.

Hochschild, Arlie, and Anne Manchung. 2003. *The Second Shift: Working Families and the Revolution at Home*. New York, NY: Penguin Books.

Hodges, Melissa J., and Michelle J. Budig. 2010. "Who Gets the Daddy Bonus? Organizational Hegemonic Masculinity and the Impact of Fatherhood on Earnings." *Gender and Society* 24:717–45. <https://doi.org/10.1177/0891243210386729>.

Huddy, Leonie, and Nayda Terkildsen. 1993. "Gender Stereotypes and the Perception of Male and Female Candidates." *American Journal of Political Science* 37:119–47. <https://doi.org/10.2307/2111526>.

Jacobs, Janis E., Stephanie Lanza, D. Wayne Osgood, Jacquelynne S. Eccles, and Allan Wigfield. 2002. "Changes in Children's Self-Competence and Values: Gender and Domain Differences across Grades One through Twelve." *Child Development* 73:509–27. <https://doi.org/10.1111/1467-8624.00421>.

Jellison, William A., Allen R. McConnell, and Shira Gabriel. 2004. "Implicit and Explicit Measures of Sexual Orientation Attitudes: In Group Preferences and Related Behaviors and Beliefs Among Gay and Straight Men." *Personality and Social Psychology Bulletin* 30:629–42. <https://doi.org/10.1177/0146167203262076>.

Jones, M. Gail, Ann Howe, and Melissa J. Rua. 2000. "Gender Differences in Students' Experiences, Interests, and Attitudes toward Science and Scientists." *Science Education*

ucation 84:180–92. [https://doi.org/10.1002/\(sici\)1098-237x\(200003\)84:2<180::aid-sce3>3.0.co;2-x](https://doi.org/10.1002/(sici)1098-237x(200003)84:2<180::aid-sce3>3.0.co;2-x).

Kerber, Linda K. 1988. "Separate Spheres, Female Worlds, Women's Place: The Rhetoric of Women's History." *Journal of American History* 75:9–39. <https://doi.org/10.1515/9783110968859.173>.

Kessler-Harris, Alice. 1982. *Out to Work: A History of Wage-Earning Women in the United States*. New York, NY: Oxford University Press.

Kessler-Harris, Alice. 1990. *A Woman's Wage: Historical Meanings and Social Consequences*. Lexington, KY: University Press of Kentucky.

Kiefer, Amy K., and Denise Sekaquaptewa. 2007. "Implicit Stereotypes, Gender Identification, and Math-Related Outcomes: A Prospective Study of Female College Students." *Psychological Science* 18:13–18. <https://doi.org/10.1111/j.1467-9280.2007.01841.x>.

Klerman, Jacob Alex, and Arleen Leibowitz. 1999. "Job Continuity among New Mothers." *Demography* 36:145–55. <https://doi.org/10.2307/2648104>.

Kozlowski, Austin C., Matt Taddy, and James A. Evans. 2019. "The Geometry of Culture: Analyzing Meaning through Word Embeddings." *American Sociological Review* 84:905–49. <https://doi.org/10.1177/0003122419877135>.

Kroska, Amy. 2007. "Gender Ideology and Gender Role Ideology." Malden, MA: Blackwell.

LaCroix, Patricia P. 1985. "Sex in Recs: Gender Bias in Recommendation Writing." *Journal of College Admissions* 109:24–26.

Lassonde, Karla A., and Edward J. O'Brien. 2013. "Occupational Stereotypes: Activation of Male Bias in a Gender-Neutral World." *Journal of Applied Social Psychology* 43:387–96. <https://doi.org/10.1111/j.1559-1816.2013.01008.x>.

Lee, Jackie F. K., and Peter Collins. 2009. "Australian English-Language Textbooks: The Gender Issues." *Gender and Education* 21:353–70. <https://doi.org/10.1080/09540250802392257>.

Lin, Yuri, Jean-Baptiste Michel, Eerez Lieberman Aiden, Jon Orwant, Will Brockman, and Slav Petrov. 2012. "Syntactic Annotations for the Google Books Ngram Corpus." Paper presented at the Proceedings of the ACL 2012 System Demonstrations, July 10, Jeju, Korea. Retrieved June 1, 2018 (<https://www.aclweb.org/anthology/P12-3029/>).

Lorber, Judith. 1994. *Paradoxes of Gender*. New Haven, CT: Yale University Press.

Lundberg, Shelly, and Elaina Rose. 2000. "Parenthood and the Earnings of Married Men and Women." *Labour Economics* 7:689–710. [https://doi.org/10.1016/s0927-5371\(00\)00020-8](https://doi.org/10.1016/s0927-5371(00)00020-8).

Lundberg, Shelly, and Elaina Rose. 2002. "The Effects of Sons and Daughters on Men's Labor Supply and Wages." *The Review of Economics and Statistics* 84:251–68. <https://doi.org/10.1162/003465302317411514>.

Maass, Anne, and Luciano Arcuri. 1996. "Language and Stereotyping." Pp. 192–226 in *Stereotypes and Stereotyping*, edited by C. N. Macrae, C. Stangor, and M. Hewstone. New York, NY: Guilford Press.

Maccoby, Eleanor E., and Nathan Maccoby. 1954. "The Interview: A Tool of Social Science." Pp. 449–87 in *Handbook of Social Psychology*. Vol. 1, *Theory and Method*, edited by G. Lindzey. Reading, MA: Addison-Wesley.

Madera, Juan M., Michelle R. Hebl, and Randi C. Martin. 2009. "Gender and Letters of Recommendation for Academia: Agentic and Communal Differences." *Journal of Applied Psychology* 94:1591–9. <https://doi.org/10.1037/a0016539>.

Martin, Patricia Yancey. 2004. "Gender as Social Institution." *Social Forces* 82:1249–73. <https://doi.org/10.1353/sof.2004.0081>.

Mason, Karen Oppenheim, John L. Czajka, and Sara Arber. 1976. "Change in U.S. Women's Sex-Role Attitudes, 1964–1974." *American Sociological Review* 41:573–96. <https://doi.org/10.2307/2094837>.

Mason, Karen Oppenheim, and Yu-Hsia Lu. 1988. "Attitudes toward Women's Familial Roles: Changes in the United States, 1977–1985." *Gender and Society* 2:39–57. <https://doi.org/10.1177/089124388002001004>.

McConnell, Allen R., and Russell H. Fazio. 1996. "Women as Men and People: Effects of Gender-Marked Language." *Personality and Social Psychology Bulletin* 22:1004–13. <https://doi.org/10.1177/01461672962210003>.

McConnell, Allen R., and Jill M. Leibold. 2001. "Relations among the Implicit Association Test, Discriminatory Behavior, and Explicit Measures of Racial Attitudes." *Journal of Experimental Social Psychology* 37:435–42. <https://doi.org/10.1006/jesp.2000.1470>.

Mead, Margaret, and Rhoda Métraux. 1957. "Image of the Scientist among High-School Students: A Pilot Study." *Science* 126:384–90. <https://doi.org/10.1126/science.126.3270.384>.

Michel, John-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, The Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden. 2011. "Quantitative Analysis of Culture Using Millions of Digitized Books." *Science* 331:176–82. <https://doi.org/10.1126/science.1199644>.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. "Efficient Estimation of Word Representations in Vector Space." Retrieved June 1, 2018 ([arXiv:1301.3781](https://arxiv.org/abs/1301.3781)).

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. "Distributed Representations of Words and Phrases and Their Compositionality." Presented at the Proceedings of the 26th International Conference on Neural Information Processing Systems, December 5–10, Lake Tahoe, NV.

Milkman, Ruth. 2016. *On Gender, Labor, and Inequality*. Champaign, IL: University of Illinois Press.

Mitchell, Juliet. 1971. *Woman's Estate*. New York, NY: Pantheon Books.

Mullis, Ina V. S., Michael O. Martin, Ann M. Kennedy, and Pierre Foy. 2007. *PIRLS 2006 International Report: IEA's Progress in International Reading Literacy Study in Primary Schools in 40 Countries*. Chestnut Hill, MA: Boston College. Retrieved January 31, 2019 (https://timss.bc.edu/pirls2006/intl_rpt.html).

Nass, Clifford, Youngme Moon, and Nancy Green. 1997. "Are Machines Gender Neutral? Gender-Stereotypic Responses to Computers with Voices." *Journal of Applied Social Psychology* 27:864–76. <https://doi.org/10.1111/j.1559-1816.1997.tb00275.x>.

Ng, Sik Hung. 1990. "Androcentric Coding of Man and His in Memory Language by Users." *Journal of Experimental Social Psychology* 26:455–64. [https://doi.org/10.1016/0022-1031\(90\)90069-x](https://doi.org/10.1016/0022-1031(90)90069-x).

Ng, Sik Hung. 2007. "Language-Based Discrimination: Blatant and Subtle Forms." *Journal of Language and Social Psychology* 26:106–22. <https://doi.org/10.1177/0261927x07300074>.

Nosek, Brian A., Mahzarin R. Banaji, and Anthony G. Greenwald. 2002a. "Harvesting Implicit Group Attitudes and Beliefs from a Demonstration Web Site." *Group Dynamics* 6:101–15. <https://doi.org/10.1037/1089-2699.6.1.101>.

Nosek, Brian A., Mahzarin R. Banaji, and Anthony G. Greenwald. 2002b. "Math = Male, Me = Female, Therefore Math \neq Me." *Journal of Personality and Social Psychology* 83:44–59. <https://doi.org/10.1037/0022-3514.83.1.44>.

Nosek, Brian A., Frederick L. Smyth, Jeffrey J. Hansen, Thierry Devos, Nicole M. Lindner, Kate A. Ranganath, Colin Tucker Smith, Kristina R. Olson, Dolly Chugh, Anthony G. Greenwald, and Mahzarin R. Banaji. 2007. "Pervasiveness and Correlates of Implicit Attitudes and Stereotypes." *European Review of Social Psychology* 18:36–88. <https://doi.org/10.1080/10463280701489053>.

O'Neil, Cathy. 2016. *Weapons of Math Destruction*. New York, NY: Crown Publishing.

Oman-Regan, Michael. 2016. "Sexism in the Oxford Dictionary of English." *Medium*. Retrieved November 22, 2019, <https://medium.com/space-anthropology/sexism-in-the-oxford-dictionary-of-english-6d335c6a77b5#.f7113isan>.

Ornaghi, Arianna, Elliot Ash, and Daniel L. Chen. 2019. "Stereotypes in High Stakes Decisions: Evidence from U.S. Courts." Unpublished Manuscript. ([http://users.nber.org/\\$sim\\$dlchen/papers/Implicit_Bias_in_the_Judiciary.pdf](http://users.nber.org/simdlchen/papers/Implicit_Bias_in_the_Judiciary.pdf)).

Otlowski, Marcus. 2003. "Ethnic Diversity and Gender Bias in EFL Textbooks." *Asian EFL Journal* 5:1–15.

Padavic, Irene, and Barbara Reskin. 2002. *Women and Men at Work*. 2nd ed. Thousand Oaks, CA: Pine Forge Press.

Picker, Susan H., and John S. Berry. 2000. "Investigating Pupils' Images of Mathematicians." *Educational Studies in Mathematics* 43:65–94. <https://doi.org/10.1023/A:1017523230758>.

Pleck, Joseph H. 1977. "The Work-Family Role System." *Social Problems* 24:417–27. <https://doi.org/10.2307/800135>.

Pratto, Felicia, Josephine D. Korchmaros, and Peter Hegarty. 2007. "When Race and Gender Go Without Saying." *Social Cognition* 25:221–47. <https://doi.org/10.1521/soco.2007.25.2.221>.

Reuben, Ernesto, Paola Sapienza, and Luigi Zingales. 2014. "How Stereotypes Impair Women's Careers in Science." *Proceedings of the National Academy of Sciences of the United States of America* 111:4403–8. <https://doi.org/10.1073/pnas.1314788111>.

Richardson, Robin. 1986. "The Hidden Messages of Schoolbooks." *Journal of Moral Education* 15:26–42. <https://doi.org/10.1080/0305724860150104>.

Risman, Barbara J. 2004. "Gender as a Social Structure: Theory Wrestling with Activism." *Gender and Society* 18:429–50. <https://doi.org/10.1177/0891243204265349>.

Rock, David, and Jean M. Shaw. 2000. "Exploring Children's Thinking about Mathematicians and Their Work." *Teaching Children Mathematics* 6:550–5.

Rose, Stephen J., and Heidi Hartmann. 2008. *Still a Man's Labor Market: The Long-Term Earnings Gap*. Washington, DC: Institute for Women's Policy Research. Retrieved September 25, 2018 (<https://iwpr.org/publications/still-a-mans-labor-market-the-long-term-earnings-gap/>).

Rudman, Laurie A., and Julie E. Phelan. 2010. "The Effect of Priming Gender Roles on Women's Implicit Gender Beliefs and Career Aspirations." *Social Psychology* 41:192–202. <https://doi.org/10.1027/1864-9335/a000027>.

Rupp, Leila. 1979. "Women's Place is in the War: Propaganda and Public Opinion in the U.S. and Germany, 1939–1945." Pp. 342–59 in *Women of America: A History*, edited by C. R. Berkin and M. B. Norton. Boston, MA: Houghton Mifflin.

Sainsbury, Marian, and Ian Schagen. 2004. "Attitudes to Reading at Ages Nine and Eleven." *Journal of Research in Reading* 27:373–86. <https://doi.org/10.1111/j.1467-9817.2004.00240.x>.

Sasson, Noah J., Amy E. Pinkham, Jan Richard, Paul Huggett, Raquel E. Gur, and Ruben C. Gur. 2010. "Controlling for Response Biases Clarifies Sex and Age Differences in Facial Affect Recognition." *Journal of Nonverbal Behavior* 34:207–21. <https://doi.org/10.1007/s10919-010-0092-z>.

Sayer, Liana C. 2005. "Gender, Time and Inequality: Trends in Women's and Men's Paid Work, Unpaid Work and Free Time." *Social Forces* 84:285–303. <https://doi.org/10.1353/sof.2005.0126>.

Schmader, Toni, Jessica Whitehead, and Vicki H. Wysocki. 2007. "A Linguistic Comparison of Letters of Recommendation for Male and Female Chemistry and Biochemistry Job Applicants." *Sex Roles* 57:509–14. <https://doi.org/10.1007/s11199-007-9291-4>.

Schulz, Muriel R. 2000. "The Semantic Derogation of Women." Pp. 82–91 in *The Routledge Language and Cultural Theory Reader*, edited by L. Burke, T. Crowley, and A. Girvin. New York, NY: Routledge.

Schwartz, Marlene B., Lenny R. Vartanian, Brian A. Nosek, and Kelly D. Brownell. 2006. "The Influence of One's Own Body Weight on Implicit and Explicit Anti-Fat Bias." *Obesity* 14:440–7. <https://doi.org/10.1038/oby.2006.58>.

Shelton, Beth Anne, and Daphne John. 1996. "The Division of Household Labor." *Annual Review of Sociology* 22:299–322. <https://doi.org/10.1146/annurev.soc.22.1.299>.

Silveira, Jeanette. 1980. "Generic Masculine Words and Thinking." Pp. 165–78 in *The Voices and Words of Women and Men*, edited by C. Kramarae. Oxford, UK: Pergamon Press.

Sklar, Elizabeth S. 1983. "Sexist Grammar Revisited." *College English* 45:348–58. <https://doi.org/10.2307/376543>.

Smith, Jessi L., and Paul H. White. 2002. "An Examination of Implicitly Activated, Explicitly Activated, and Nullified Stereotypes on Mathematical Performance: It's Not Just a Woman's Issue." *Sex Roles* 47:179–91. <https://doi.org/10.1023/A:1021051223441>.

Spencer, Steven J., Claude M. Steele, and Diane M. Quinn. 1999. "Stereotype Threat and Women's Math Performance." *Journal of Experimental Psychology* 35:4–28. <https://doi.org/10.1006/jesp.1998.1373>.

Stanton, Elizabeth Cady. 1848. "The Declaration of Sentiments." Seneca Falls, NY: Seneca Falls Convention.

Steele, Claude M., and Joshua Aronson. 1995. "Stereotype Threat and the Intellectual Test Performance of African Americans." *Attitudes and Social Cognition* 69:797–811. <https://doi.org/10.1037/0022-3514.69.5.797>.

Steele, Jennifer R., and Nalini Ambady. 2006. "'Math is Hard!' The Effect of Gender Priming on Women's Attitudes." *Journal of Experimental Psychology* 42:428–36. <https://doi.org/10.1016/j.jesp.2005.06.003>.

Stone, Pamela. 2007. *Opting Out? Why Women Really Quit Careers and Head Home*. Berkeley, CA: University of California Press.

Stout, Jane G., and Nilanjana Dasgupta. 2011. "When He Doesn't Mean You: Gender-Exclusive Language as Ostracism." *Personality and Social Psychology Bulletin* 37:757–69. <https://doi.org/10.1177/0146167211406434>.

Sweeney, Latanya. 2013. "Discrimination in Online Ad Delivery." *ACM Queue* 11:1–19.

Swim, Janet K., Kathryn J. Aikin, Wayne S. Hall, and Barbara A. Hunter. 1995. "Sexism and Racism: Old-Fashioned and Modern Prejudices." *Journal of Personality and Social Psychology* 68:199–214. <https://doi.org/10.1037/0022-3514.68.2.199>.

Tallichet, Suzanne E., and Fern K. Willits. 1986. "Gender-Role Attitude Change of Young Women: Influential Factors from a Panel Study." *Social Psychology Quarterly* 49:219–27. <https://doi.org/10.2307/2786804>.

Thorne, Barrie, and Marilyn Yalom, eds. 1982. *Rethinking the Family: Some Feminist Questions*. New York, NY: Longman.

Thornton, Arland, and Deborah Freedman. 1979. "Changes in the Sex Role Attitudes of Women, 1962–1977: Evidence from a Panel Study." *American Sociological Review* 44:831–42. <https://doi.org/10.2307/2094530>.

Thornton, Arland, Duane F. Alwin, and Donald Camburn. 1983. "Causes and Consequences of Sex-Role Attitudes and Attitude Change." *American Sociological Review* 48:211–27. <https://doi.org/10.2307/2095106>.

Thornton, Arland, and Linda Young-DeMarco. 2001. "Four Decades of Trends in Attitudes toward Family Issues in the United States: The 1960s through the 1990s." *Journal of Marriage and Family* 63:1009–37. <https://doi.org/10.1111/j.1741-3737.2001.01009.x>.

Toossi, Mitra, and Teresa L. Morisi. 2017. *Women in The Workforce Before, During, And After The Great Recession*. Washington, DC: U.S. Bureau of Labor Statistics.

Trix, Frances, and Carolyn Psenka. 2003. "Exploring the Color of Glass: Letters of Recommendation for Female and Male Medical Faculty." *Discourse and Society* 14:191–220. <https://doi.org/10.1177/0957926503014002277>.

U.S. Constitution, Amendment 19.

Watson, Carol. 1987. "Sex-Linked Differences in Letters of Recommendation." *Women and Language* 10:26–28.

Weitzman, Murray S. 1978. "Finally the Family." Pp. 61–82 in *America in the Seventies: Some Social Indicators*, edited by C. Taeuber. Philadelphia, PA: The Annals of the American Academy of Political and Social Science.

Weitzman, Lenore J., Deborah Eifler, Elizabeth Hokada, and Catherine Ross. 1972. "Sex-Role Socialization in Picture Books for Preschool Children." *American Journal of Sociology* 77:1125–50. <https://doi.org/10.1086/225261>.

Welter, Barbara. 1966. "The Cult of True Womanhood: 1820–1860." *American Quarterly* 18:151–74. <https://doi.org/10.2307/2711179>.

Wilson, Elizabeth, and Sik Hung Ng. 1988. "Sex Bias in Visual Images Evoked by Generics: A New Zealand Study." *Sex Roles* 18:159–68. <https://doi.org/10.1007/bf00287786>.

Yavorsky, Jill E., Claire M. Kamp Dush, and Sarah J. Schoppe-Sullivan. 2015. "The Production of Inequality: The Gender Division of Labor across the Transition to Parenthood." *Journal of Marriage and Family* 77:662–79. <https://doi.org/10.1111/jomf.12189>.

Zhao, Jieyu, Tianlu Wang, Mark Yatskar, Vicente Ordóñez, and Kai-Wei Chang. 2017. "Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints." Retrieved August 7, 2018 (<https://arxiv.org/abs/1707.09457v1>).

Ziegert, Jonathan C., and Paul J. Hanges. 2005. "Employment Discrimination: The Role of Implicit Attitudes, Motivation, and a Climate for Racial Bias." *Journal of Applied Psychology* 90:553–62. <https://doi.org/10.1037/0021-9010.90.3.553>.

Zuber, Sharon, and Ann M. Reed. 1993. "The Politics of Grammar Handbooks: Generic He and Singular They." *College English* 55:515–30. <https://doi.org/10.2307/378587>.

Acknowledgments: This material is based upon work supported by the National Science Foundation under grants IIS-1546113 and IIS-1927227. The authors would like to thank Stony Brook Research Computing and Cyberinfrastructure as well as the Institute for Advanced Computational Science at Stony Brook University for access to the high-performance SeaWulf computing system, which was made possible by a \$1.4 million National Science Foundation grant (#1531492).

Jason J. Jones: Department of Sociology and Institute for Advanced Computational Science, Stony Brook University. E-mail: Jason.J.Jones@stonybrook.edu.

Mohammad Ruhul Amin: Department of Computer Science, Stony Brook University. E-mail: moamin@cs.stonybrook.edu.

Jessica Kim: Department of Sociology, Stony Brook University. E-mail: jessica.a.kim@stonybrook.edu.

Steven Skiena: Department of Computer Science, Stony Brook University. E-mail: skiena@cs.stonybrook.edu.