

# A Latent Semantic Analysis of Gender Stereotype-Consistency and Narrowness in American English

Alison P. Lenton · Constantine Sedikides ·  
Martin Bruder

Published online: 20 September 2008  
© Springer Science + Business Media, LLC 2008

**Abstract** Using latent semantic analysis, we examined gender stereotypes in American English by submitting over 100 masculine, neutral, and feminine role-words and trait-words to pair-wise semantic similarity comparisons with masculine (*man, he, him*) and feminine (*woman, she, her*) referents separately. We expected to find: (a) Stereotyping—roles and traits would be more semantically similar to the ostensible ‘matching’ than ‘mismatching’ gender category referent; (b) Categorical narrowness—both categories would be less semantically similar to counterstereotypical than to neutral or stereotypical characteristics; but this would be especially so for the male category, indicating its relatively greater narrowness. Results supported these hypotheses, but only among role-words. American English reflects and reinforces gender stereotypes regarding gender roles at a level beyond that recognized previously.

**Keywords** Gender stereotypes · Stereotype content · Stereotype breadth · Latent semantic analysis

---

The research reported herein was supported by Economic and Social Research Council grant #RES-000-22-0253.

---

A. P. Lenton (✉)  
Department of Psychology, University of Edinburgh,  
7 George Square,  
Edinburgh EH8 9JZ, UK  
e-mail: a.lenton@ed.ac.uk

C. Sedikides  
University of Southampton,  
Southampton, UK

M. Bruder  
Cardiff University,  
Cardiff, UK

## Introduction

In 2006, Jaques Chirac, then-president of France, walked out of an European Union summit meeting upon hearing one of his fellow countrymen speaking English, in an apparent protest over the proliferation of the use of English (and the concomitant shrinkage of the use of French) across Europe and the world (Watt and Gow 2006). English—and arguably, English of the American variety—indeed has become the international language for trade, commerce, and science, not to mention film and television (Graddol 2000), as well as the internet where nearly 70% of the content is in English (Global Reach 2004). In the present study, we set out to determine whether gender stereotypes are embedded in the very meaning of American English. We propose that they are, and that this will be evident in a latent semantic analysis (Landauer et al. 1998a) of the degree to which stereotypically masculine, neutral, and feminine role-words and trait-words are similar in meaning to the most common category referents for ‘man’ and ‘woman.’ In light of the increasing numbers of people learning (American) English across the world, it is important to understand to what extent these people are exposed to gender stereotypes.

Of course, we are not the first to suggest that language reveals stereotypes more generally. Children’s literature is rife with stereotype-reinforcing depictions (e.g., boys described as *strong*, girls as *sweet*; Ernst 1995). Even a content analysis of male and female business leaders’ obituaries (Rodler et al. 2001) has revealed gender stereotypes (e.g., men described as *expert*, women described as *loyal*). More generally, speakers use relatively abstract terminology (e.g., adjectives such as “he was *aggressive*”) when conveying unfavorable information about members of other groups, but use relatively concrete terminology (e.g., descriptive action verbs such as “he *ran*”) when conveying

favorable information about members of those same groups (Maass et al. 1989). These linguistic differences subtly contribute to the persistence of stereotypes. Stereotypes also impact language comprehension: If the word *he* follows a sentence describing a secretary's actions, reading times are slower than when the pronoun's and the antecedent's stereotypic gender match (Kennison and Trofe 2003).

It is clear that language communicates, and that our understanding of it is affected by, stereotypes. What is not clear, however, is the extent to which the very meaning of a social category label overlaps with that of stereotype-consistent versus stereotype-inconsistent words. Again, we investigate this proposition in American English, given its increasing use across the world. Latent semantic analysis lends itself well to examining semantic overlap, as it is both a “theory and method for extracting and representing the contextual-usage meaning of words by statistical computations applied to a large corpus of text” (Landauer et al. 1998a, p. 259). In latent semantic analysis, the semantic similarity of any two words depends on whether they appear in similar contexts (i.e., surrounded by a similar array of other words; Landauer and Dumais 1997). Consider this sentence:

The nurse retrieved a handkerchief and dabbed the child's nose.

To determine whether the word *woman* is more similar in meaning to the word *nurse* than the word *man* is, one could ask how likely it is that *man* appears amongst the words *retrieved*, *handkerchief* and the other words surrounding *nurse* across a variety of texts. And how likely is *man* to appear amongst the words that typically surround those words and so on? In comparison, how likely is the word *woman* to be found amongst this same array of words? If the latter probability is higher than the former, then *woman* is more similar in meaning to *nurse* than is *man* to *nurse*. To facilitate understanding of this comparison, the reader might try to imagine the average meaning of all of the sentences in which the word *woman* appears in a given book, and then imagine the average meaning of all of the sentences in which the word *man* appears in that same book (Landauer et al. 1998a). Now the reader should compare each of these averages to the average meaning of all of the sentences in which the word *nurse* appears in that book. How similar or different are they?

The exact mathematics of how latent semantic analysis computes the degree of shared meaning are beyond this article's scope, but suffice it to say that it relies upon singular value decomposition, which has commonalities with factor analysis and multidimensional scaling (Landauer and Dumais 1997; see also: Anderson and Sedikides 1991; Schiffman et al. 1981). Latent semantic analysis represents each word as a vector in multidimensional

semantic space (the limits of which are defined by the corpus selected for analysis). The analysis of shared meaning yields a similarity value for a given word-pair, which literally is the cosine of the angle between the two vectors representing the words under comparison (the following website can be used to extract LSA scores: <http://lsa.colorado.edu/>).

Importantly, latent semantic analysis does not merely measure the simple, first-order co-occurrence between words; that is, it does not only measure whether the words appear together in the same local context, which is usually a paragraph. In fact, in one study, 99% of the word-pairs whose similarity was assessed never co-occurred in the same paragraph (Dennis et al. 2003). Thus, *man* and *engineer* need not ever appear in the same local context for latent semantic analysis to assess them as being similar in meaning or as possessing a high cosine. Indeed, direct co-occurrence does not necessarily yield a large cosine, and not all large cosines stem from directly co-occurring words (Deerwester et al. 1990). With respect to the latter, for example, Lemaire and Denhière (2006) report that, in a 24 million word French corpus, the words *internet* and *web* never co-occurred, though they are, of course, strongly related to one another. Latent semantic analysis is based upon the assumption that linguistic meaning is derived from an irreducibly high dimensional space (Landauer et al. 2004), and this is why it takes into account higher-order or indirect associations in addition to first-order co-occurrence. For example, one study found that latent semantic analysis can take into account up to fifth-order co-occurrences (Kontostathis and Pottenger 2002).

Whereas co-occurrences of such order are difficult for most of us to imagine, it might help the reader to fix in mind the meaning of co-occurrences of a lower, more easily interpreted order. Again, first-order co-occurrence means that the word-pairs can be found within the same local context. Second-order co-occurrence, on the other hand, means that the word-pairs do *not* ever appear together in the same context; instead, at least some of the words that surround one of the pair, also surround the other of the pair. Still higher-order co-occurrences continue analogously. Consequently, although latent semantic analysis does not specifically exclude direct or first-order co-occurrence, it is more than this: It is a measure of the larger pattern of co-occurrence across a vast number of local contexts (Landauer et al. 1998a). At the same time, however, we must note that it is difficult to assess the precise extent of the role played by higher-order (versus first-order) co-occurrence in producing any given set of LSA scores (Lemaire and Denhière 2006).

Nevertheless, given that latent semantic analysis models word-meaning through both direct and indirect associations, it possibly offers a naturalistic representation of how

people learn language: that is, inductively, based on experience (Landauer and Dumais 1997; Landauer 2002). Language learners encounter words in close spoken and/or written temporal proximity (e.g., the same sentence or paragraph), and from this proximity, they infer semantic similarity. In addition to relying on these first-order co-occurrences, however, learners also make use of higher-order associations to infer semantic meaning. For example, the synonyms *provide* and *supply* are unlikely to occur together in the same local context (because each word replaces, rather than extends, the meaning of the other), but people deduce their shared meaning from the fact that they are regularly found amongst the same set of other words (e.g., *services*, *goods*, *assistance*). It is this critical process of assigning each word a place in the mesh of prior knowledge that latent semantic analysis attempts to model.

To investigate the feasibility of such a model, one can examine the latent semantic analysis scores between words after uploading various corpora to its database. For example, Landauer and Dumais (1997) found that—after ‘training’ on an American encyclopedia—latent semantic analysis had ‘learned’ English synonyms to a degree akin to that of non-native English speakers taking the *Test of English as a Foreign Language* and, further, that approximately 75% of its understanding of word meaning stemmed from indirect induction. This feature makes latent semantic analysis distinct from and more powerful than a content analysis focusing on direct or first-order co-occurrence only. This is because latent semantic analysis models language and knowledge acquisition, in addition to (post-acquisition) meaning representation (Landauer et al. 1998b).

An earlier study showed that latent semantic analysis is more likely to match an exemplar (*apple*) to its superordinate category (*fruit*) than to any of 13 other superordinate categories (Laham 1997). We aimed to achieve something similar in the domain of gender stereotypes. We examined the degree to which the snapshot of the average American English-speaker’s semantic network provided by latent semantic analysis (see “Method”) reflects modern gender stereotypes. We hypothesized that stereotypically masculine, feminine, and gender-neutral words would have distinct patterns of semantic overlap with the most commonly used gender category referents: *man*, *he*, and *him* versus *woman*, *she*, and *her*. In particular, with Hypothesis 1, we expected to find evidence of stereotyping: Stereotypically masculine words would share meaning with *man/he/him* more so than with *woman/she/her*, whereas stereotypically feminine words would share meaning with *woman/she/her* more so than with *man/he/him*. Such a finding would indicate that many American English-speakers’ understanding of words is founded upon a knowledge-base laden with gender stereotypes, as it would demonstrate

that stereotypes may permeate language at a more indirect and, therefore, insidious level (via higher order co-occurrence in addition to first-order co-occurrence) than previously considered. A corollary is that American English-speakers’ very understanding of *he*, for example, is inextricably bound up with their understanding of gender stereotypical words (e.g., *engineer*) and their joint higher-order associates. This corollary is not without controversy, however, and it is therefore a point to which we will return in the Discussion.

Hypothesis 2 stemmed from our conceptualization of *stereotype breadth* (Lenton et al., unpublished data), which is an indication of the degree to which a stereotype representation can be ‘stretched’ to include stereotype-inconsistent aspects. These aspects are counterstereotypes, or characteristics that have been traditionally associated with the contrasting category. To the extent that these stereotype representations do not stretch very far, they can be called ‘narrow.’ We predicted that *man/he/him* and *woman/she/her* would be narrow. That is, each representation would share less semantic overlap with counterstereotypical than with other characteristics, including those that are stereotype-neutral or stereotype-consistent.

We examined both stereotype-inconsistency (breadth) and stereotype-consistency, because there is evidence showing that the two are not inversely related (i.e., represented on a single bipolar continuum). For example, Blair et al. (2001) found that counterstereotype mental imagery sometimes impacted stereotype-consistent (only) responding and sometimes stereotype-inconsistent (only) responding. Diekmann and Eagly (2000) found that participants perceived women to have become more stereotypically masculine since the 1950s, but not simultaneously less stereotypically feminine. More generally, connectionist models posit (Smith and Conroy 2007; Smith and DeCoster 1998) and simulations based thereon show (Queller 2002; Queller and Smith 2002) that counterstereotypic information can be represented in long-term memory alongside stereotype-consistent information. Indeed, lay conceptions of gender are represented by (at least) two unipolar scales: degree of masculinity and degree of femininity (Helgeson 1994). The negative correlations between masculinity and femininity that have been found (e.g., Biernat 1991) are likely due to current input—such as the methods used (e.g., ‘imagine meeting a woman who is pretty, delicate, and soft’)—activating internally consistent subtypes (the ‘female flower’) rather than the general category (‘women’).

We anticipated a potential qualification of Hypothesis 2, however. With Hypothesis 2a (‘differential narrowness’), we predicted that *man/he/him* would be more narrow than *woman/she/her*. Close examination of Diekmann and Eagly’s (2000) ‘present time’ condition suggests that women are perceived to possess masculine characteristics to a greater

degree than are men perceived to possess feminine characteristics. Similarly, Prentice and Carranza's (2002) research indicates that there are relaxed prescriptions regarding the desirability of counterstereotypical traits for women, but not men. The developmental literature also supports the notion that there is a stronger prohibition against males exhibiting traditionally feminine characteristics than the converse (Burn 2000; Maccoby 1998). If we obtain simultaneous support for Hypotheses 1 and 2a, it will confirm the utility of the stereotype breadth construct, as it will demonstrate that stereotype inconsistency is an independent aspect of stereotype representations (in that its results do not perfectly mirror those in tests of Hypothesis 1) and, as such, ought to be disentangled from stereotype consistency. Moreover, support for Hypothesis 2 and 2a would suggest that the word-meaning embedded in the American English language reflects and reinforces cultural representations of gender in ways not previously considered.

To recap, in this study we test several hypotheses: (1) Stereotyping—roles and traits will be more semantically similar to the ostensible 'matching' than 'mismatching' gender category referent; (2) Categorical narrowness—both categories will be less semantically similar to counterstereotypical than to neutral or stereotypical characteristics; but (2a) this will be especially so for the male category, indicating relatively greater narrowness thereof. Our study also explores whether the type of attribute—role versus trait—matters. According to social role theory (Eagly et al. 2000), gender-based expectations regarding traits stem from observations of the differential distribution of men and women across various types of roles and occupations: Traits (e.g., *caring*) are deduced from roles (e.g., *nurse*). For example, a now-classic study showed that people perceive men and women as being similar to one another when it is clear that they share the same role (e.g., full-time employee, homemaker; Eagly and Steffen 1984): Homemakers—whether male or female—are thought to be especially kind, whereas those in full-time employment—again, whether male or female—are thought to be competitive. But when men's and women's roles are not made explicit, people assume that their roles are distinct from one another, and only then are differential trait ascriptions applied to the sexes. Because roles are primary in social role theory, it would seem to follow that gender stereotyping and (differential) breadth could be stronger among the role-words than the trait-words (Hypothesis 3).

## Method

### Materials

We obtained the words used in this study from three primary sources. First, we included the 20 masculine, 20

feminine, and 20 neutral trait-words comprising the Bem Sex Role Inventory (BSRI; Bem 1974), with two exceptions: We excluded the items *feminine* and *masculine* because of their explicitly gendered nature.

One of the first author's previous studies (Lenton and Webber 2006) provided a second source. In that study, 40 role-words assessed participants' *gender diagnosticity*, Lippa and Connelly's (1990) reformulation of gender role orientation. Nearly 200 UK participants, predominantly students from the University of Edinburgh, rated the extent to which they would like to engage in each role (e.g., 'I would like to be a *pilot*' or 'I would like to be a *librarian*'; 1 = *strongly disagree*, 7 = *strongly agree*). For the purposes of the present research, a role was categorized as masculine if the male participants ( $n=89$ ) showed significantly greater interest in it, a role was categorized as feminine if the female participants ( $n=93$ ) showed significantly greater interest in it, and a role was categorized as gender-neutral if the male and female participants showed equivalent interest. However, some of these initial categorizations were notably inconsistent with traditional occupational gender stereotypes and thus they were excluded (e.g., 'biologist' and 'chemist' as feminine roles; 'minister' and 'librarian' as gender-neutral roles). The final set of words from this source comprised 14 masculine, eight neutral, and nine feminine roles.

The third source of both role-words and trait-words was our own prior research (Lenton et al., unpublished data). One hundred forty-three online participants (92 women, 51 men) rated 154 words with respect to their masculinity–femininity (1 = *very masculine*; 7 = *very feminine*) and valence (1 = *very negative*; 7 = *very positive*). Each word's frequency (The British National Corpus 2001) and length in number of letters was also recorded. The purpose of that study was to obtain sets of masculine, feminine, and gender-neutral words that were matched in several respects (valence, word length, word frequency) other than their gendered nature. The final set consisted of 20 words in each category, with 12 roles and eight traits in each of the masculine and feminine categories, and seven roles and 13 traits in the neutral category. There was overlap across the three sources of 160 words (e.g., two contained *farmer*), which meant that the initial pool for the current study contained 134 distinct roles and traits.

Before conducting the analyses, we transformed several phrases to their single-word synonyms. For example, we replaced *willing to take a stand*, with *bold*. In some cases (e.g., *automobile sales person*), there was no good single-word alternative; we dropped these expressions. Also, we excluded one word, because two traits (*leadership ability* and *acts as a leader*) effectively resulted in the same single-word transform: *authoritative*. Finally, given that the LSA corpus cannot identify compound words (e.g., *wage-*



earner), we replaced all four of these with synonyms (e.g., *worker*). Thus, the final set comprised 118 words, 42 of which were masculine (20 roles), 37 were feminine (17 roles), and 39 were gender-neutral (ten roles). “Appendix” lists these words.

We selected the category referents—*man/he/him* and *woman/she/her*—on the basis of their high frequency in American English (American National Corpus 2007). Thus, our results can be generalized across the most commonly used pairs of referents for the primary gender categories.

### Procedure

To extract the semantic similarity scores, we submitted each trait- or role-word to LSA six times, once per category referent (*man, he, him, woman, she, and her*). Subsequently, we averaged across all referents of one category to produce for each word one average LSA score for the male and female referents separately. This was warranted by the high internal consistency of the LSA scores for each of the three referents within a gender category,  $\alpha=.95$  and  $\alpha=.93$  for male and female referents, respectively. We selected *General Reading through First Year of College* as our corpus for analysis. This corpus consists of nearly 11 million words, and is based on a representative sample of texts (e.g., textbooks, novels, newspapers) read by students from grade three through the first year of university in the United States (Landauer et al. 1998b). The corpus was originally put together by Touchstone Applied Science and Associates (TASA) for the purposes of developing The Educator’s Word Frequency Guide, the largest study of word frequency completed up to that point (TASA Inc. 2006). Accordingly, it has been suggested that the corpus is representative of American English-speakers’ *world* knowledge, in addition to their *word* knowledge (Wolfe and Goldman 2003). Because of these unique features, latent semantic analysis may provide researchers with a relatively realistic snapshot of the average American English-speakers’ semantic network in a way that an analysis of a more restricted corpus (such as storybooks or magazines) cannot.

In conducting the analyses, we allowed the number of factors to reach the maximum (300), due to lack of a priori theory that would inform us otherwise. Given that the number of factors extracted influences the similarity scores, one should not assign much weight to the particular scores, but, instead, consider their relative values (Berry et al. 1995).

### Results

We initially tested all of the hypotheses within the confines of one analysis of variance (ANOVA). We ran a 2 (referent

gender: male referents average vs. female referents average)  $\times$  2 (word gender: masculine vs. feminine)  $\times$  2 (word type: role vs. trait) mixed-model ANOVA on the LSA scores, with repeated measures on the first factor and word gender and word type as between-subjects factors. Evidence for stereotyping or (differential) narrowness would be shown by an interaction between referent gender and word gender (with word gender being based on a priori classification of roles and traits into masculine, feminine, and neutral as described in the methods section). This interaction was significant,  $F(1, 75)=37.33, p=.001, \eta_p^2=.33$ , but was qualified by a three-way interaction involving word type,  $F(1, 75)=14.84, p=.001, \eta_p^2=.17$  (see Table 1 for the means and standard errors).


We broke down the three-way interaction by re-running the ANOVA twice (excluding the word type factor), once for each type of word. Using Bonferroni-corrected  $p$ -values to account for this test’s redundancy, and confirming Hypothesis 3 (that the expected effects would be stronger for the role-words than the trait-words), the Referent Gender  $\times$  Word Gender interaction was statistically significant for roles,  $F(1, 35)=55.40, p=.001, \eta_p^2=.61$ , but nonsignificant for traits,  $F(1, 40)=2.39, p=.260, \eta_p^2=.06$ . We thus examined stereotyping (Hypothesis 1) and (differential) narrowness (Hypothesis 2 and 2a) only amongst the role-words (Fig. 1).

Again, we expected to find evidence for stereotyping (Hypothesis 1): Roles will be more semantically similar to the ostensible ‘matching’ than ‘mismatching’ gender category referent. In other words, stereotyping is the extent to which the gendered words, on average, show greater semantic similarity to the stereotypic than to the counterstereotypic referent gender. To test for this, we conducted simple effects analyses of: (a) Masculine stereotyping—*masculine* words with the *male* referents versus with the *female* referents; (b) feminine stereotyping—*feminine* words with the *female* referents

**Table 1** LSA score means (standard errors) by word gender, referent gender, and word type (unadjusted for masculine generic effect).

	Roles		Traits	
	Male referent	Female referent	Male referent	Female referent
Word gender				
Feminine	.07 (.03)	.12 (.02)	.21 (.02)	.20 (.02)
Neutral	.15 (.03)	.12 (.03)	.20 (.02)	.16 (.02)
Masculine	.17 (.02)	.08 (.02)	.15 (.02)	.10 (.02)

Again, one should not assign much weight to the particular scores, but should, instead, consider their relative values (Berry et al. 1995). Nevertheless, it might be useful for the reader to know that cosine values for randomly paired words typically range between  $-.01$  and  $+.03$  (Dennis et al. 2003)



**Fig. 1** Word gender  $\times$  referent gender interaction for roles (unadjusted for masculine generic effect).

versus with the *male* referents; and (c) stereotyping of neutral words—*neutral* words with the *male* referents versus with the *female* referents. We examined stereotyping of neutral words in order to test for *masculine norming* (i.e., the male being seen as normative or generic; Gastil 1990) and, if necessary, correct for it.

Among the roles, there was evidence for both masculine stereotyping,  $M_{\text{difference}}=.092$ ,  $SD=.059$ ,  $t(19)=6.90$ ,  $p=.001$ , and feminine stereotyping,  $M_{\text{difference}}=.047$ ,  $SD=.053$ ,  $t(16)=3.68$ ,  $p=.002$ . As can be seen in Table 1 and Fig. 1, the masculine role-words were more similar in meaning to the male referents than to the female referents, whereas the feminine role-words were more similar in meaning to the female referents than to the male referents. Stereotyping of neutral words ( $M_{\text{difference}}=.035$ ,  $SD=.046$ ) was significantly different from 0 however,  $t(9)=2.41$ ,  $p=.039$ , suggesting that the neutral role-words were more similar to the masculine than to the feminine role-words. If even relatively neutral roles are somewhat masculine in meaning, then perhaps the observed male stereotyping effect is inflated and the female stereotyping effect deflated. To control for this possibility, we re-ran the stereotyping comparisons: This time, we ‘discounted’ (subtracted from) the masculine stereotyping contrast by half of the mean-difference between the neutral words’ LSA score for the masculine referents versus the feminine referents, and ‘reimbursed’ (added to) the feminine stereotyping contrast by this same value (.0175). Upon doing so, feminine stereotyping of course became stronger ( $M_{\text{difference}}=.065$ ,  $SD=.052$ ),  $t(16)=5.04$ ,  $p=.001$ . More importantly, the masculine stereotyping effect remained significant ( $M_{\text{difference}}=.074$ ,  $SD=.059$ ),  $t(19)=5.58$ ,  $p=.001$ .

To test Hypotheses 2 (categories will be less semantically similar to counterstereotypical than to neutral or stereotypical characteristics) and 2a (more so for the male category), we conducted two one-way ANOVAs. The first examined male narrowness by testing the contrast between the semantic similarity to *man/he/him* of the feminine

words (+2) versus the average of the neutral (–1) and masculine (–1) words. The second examined female narrowness by testing the contrast between the semantic similarity to *woman/she/her* of the masculine (+2) versus the average of the neutral (–1) and feminine (–1) words (see Table and Fig. 1). The greater the difference, the narrower the gender categories (neutral words were included in order to conduct the most conservative tests). Male narrowness was significant ( $M_{\text{difference}}=-.09$ ,  $SE=.010$ ),  $t(45)=-3.01$ ,  $p=.004$ ,  $pr=-.413$ , whereas female narrowness was not ( $M_{\text{difference}}=-.042$ ,  $SE=.008$ ),  $t(45)=-1.82$ ,  $p=.076$ ,  $pr=-.264$ . Thus, *man/he/him* is significantly less semantically similar to feminine than to other roles (masculine + neutral), whereas *woman/she/her* is only marginally less similar to masculine than to other roles (feminine + neutral). The male stereotype is narrower than the female stereotype. Our finding that the neutral roles were relatively more masculine than feminine in meaning renders this interpretation problematic, however, because the female representation may appear broader than the male representation simply because masculine-related words (masculine and neutral words) are on both sides of the former contrast but on the same side of the latter contrast. To control for this possibility, we re-ran the analyses, where we (a) adjusted upward the relationship of the neutral role-words to the female referents (i.e., *woman/she/her*), and (b) adjusted downward the relationship of the neutral role-words to the male referents (i.e., *man/he/him*). In both cases, we again adjusted by half of the mean difference between the LSA scores of the neutral role-words to the male versus female referents. This time, both tests of narrowness were significant,  $t(45)=-2.71$ ,  $p=.009$ ,  $pr=-.379$  for male ( $M_{\text{difference}}=-.081$ ,  $SE=.010$ ), and  $t(45)=-2.12$ ,  $p=.033$ ,  $pr=-.314$  for female ( $M_{\text{difference}}=-.051$ ,  $SE=.008$ ). Still, the partial correlations suggest that male narrowness is stronger than female narrowness.

## Discussion

The principle of linguistic relativity (Carroll 1956), alongside research showing that the language use conveys stereotypic information (Ernst 1995; Kennison and Trofe 2003; Maass et al. 1989), led us to expect that gender stereotypes would be evident in American English-language semantics. We hypothesized that stereotypically masculine and feminine words would share more semantic meaning with their matching (*man/he/him*, *woman/she/her*, respectively) than mismatching (*woman/she/her*, *man/he/him*, respectively) category referents. We also hypothesized that stereotype narrowness—whereby a category referent is distinctly dissimilar to words stereotypical of the contrasting category—would be evinced in American English semantics, with this pattern expected to

be more pronounced for the male than the female category because of differential cultural injunctions against men versus women engaging in counterstereotypical behavior (Bosson et al. 2005; Burn 2000; Maccoby 1998).

The results supported these hypotheses. Our research shows that gender stereotypes are inherent in the very meaning of the most common social category referents for *man* and *woman*. Stated differently, because lexical acquisition is inductive (Landauer and Dumais 1997), American English-speakers' understanding of the words *man*, *he*, or *him* and *woman*, *she*, or *her* is fundamentally tied to their understanding of stereotype-relevant words. Thus, to understand *he*, for example, one also must understand gender stereotypical role-words (e.g., *engineer*) and their joint higher-order associates. These results demonstrate that stereotypes permeate language at a very deep level, as LSA is carried out via the inclusion of indirect semantic associations.

In line with our differential narrowness hypothesis, our research also shows that *woman/she/her* is a broader concept than *man/he/him*. After taking into account the somewhat masculine nature of the 'neutral' role-words, it was apparent that the male and female categories are both narrowly construed: Compared to their association with stereotypic and neutral role-words combined, they show distinctly less semantic overlap with counterstereotypic role-words. The results further indicated that the male category is likely to be even more narrow than the female category. These findings possibly reflect the relative success of measures taken to broaden the concept of *woman*. The focus might now need to shift toward extending the concept *man*. These findings are also notable, because they demonstrate that the existence of counterstereotypic characteristics in the mental representation of *woman* does not depend on the possession of an implicit role theory about the relationship between women's changing roles and, thus, traits (Diekmann and Eagly 2000). LSA, of course, possesses no such theory, and it still suggests that the female representation is construed somewhat more broadly than the male one. Finally, the findings support our argument that stereotype-consistency and stereotype-inconsistency both ought to be taken into account when investigating gender stereotypes.

We used a powerful tool for assessing the semantic similarity of words: latent semantic analysis (Landauer et al. 1998b). This technique identifies the semantic similarity between any pair of words by assessing the degree to which they can be found within the same word context. To reiterate, LSA is not merely a measure of the words' first-order co-occurrence (Dennis et al. 2003): *Her* and *florist* need not ever appear together in the same unit of discourse in order for LSA to deem them semantically similar. Additionally, the corpus on which our analysis was based is composed of fictional, in addition to nonfictional,

writings (TASA Inc. 2006). These methodological features render it unlikely that our findings only reflect real-world sex differences in occupational choice, rather than stereotypes per se. Still, even if there were correspondence to base-rate sex differences, this would not mean our results are unrelated to stereotypes, as stereotypes vary in the extent to which they are accurate (Judd and Park 1993). In other words, some stereotypes contain a 'kernel of truth.' For example, across several studies researchers have observed that the perceived degree of gender segregation in certain occupations correlates with the actual degree of gender segregation in those occupations (Beyer 1999). And so it is with LSA's 'perceptions': They do not comprise the accuracy criterion itself but, rather, could be compared to such.

Because language also shapes the way people perceive the world (Carroll 1956), language semantics may not merely reflect gender stereotypes, but may perpetuate them as well. Our findings thus point to the potentially intractable nature of gender stereotypes. According to the LSA model (Landauer and Dumais 1997), however, a word's meaning is never fixed, as it changes each time it is encountered in a new context. Thus, gender stereotypes could wane if the words that people use (and print) in the context of male and female category labels change. Future research, then, could examine the utility of LSA for the study of stereotype change over time. For example, repeated latent semantic analyses on a continuously updated corpus might yield an estimate of stereotype dynamism that is relatively free from biases (e.g., experimental demand).

As we mentioned in the introduction, the interpretation of LSA is controversial, and the nature of this controversy requires some attention. Many researchers do not accept the claim that a word's meaning can be deduced from its relationship to other words, for this is a logical impossibility. As Glenberg and Robertson (2000) put it, "To know the meaning of an abstract symbol such as an LSA vector or an English word, the symbol has to be grounded in something other than more abstract symbols" (p. 382). According to this framework, it is misleading to suggest—as we do—that our findings show that *woman*, for example, is closer in meaning to feminine than to masculine role-words. The meaning of these words, instead, is argued to be a mesh of an object's or event's affordances, the personal experiences a perceiver has had with the object or event (including cultural norms relevant to these experiences), and a perceiver's goals with respect to the object or event. Together, these aspects constrain 'meaning' and, thus, the array of actions available upon perception of the object or event (Glenberg 1997). In sum, according to this perspective, a word's meaning refers to this mesh, not to other words. Thus, our results may not speak to the grounded

meaning of gender referents but rather, more simply, to how these terms are represented in American English.

Indeed, LSA's proponents accept that it is not a wholly adequate model of human learning and cognition but, for them, the argument that LSA fails to ground word meaning does not bring down the house of cards. Landauer (2002) invokes Occam's razor when he suggests that the mechanisms underlying word–word associations should be no different than those underlying object/event-word associations. Thus, if one could input perceptual features and action tendencies into LSA in the same way as words, “the words ‘headache,’ ‘fireplace,’ ‘throw,’ and ‘kiss,’ for example, would surely have quite high cosines with their perceptual equivalents” (p. 64). And if LSA does not represent word meaning, then how—to provide just one example—can one account for the correspondence between human graders' assessment of written essays and LSA's assessment of those same essays (Landauer et al. 1997)? Add to this the fact that a great deal of human learning—especially of the formal variety—takes place via reading, then we are left to conclude that LSA does indeed extract word meaning to at least some extent. Put differently, *her* and *him* do not merely refer to other words, but to real-life exemplars, experiences, and action tendencies as well.

Returning to the results and their other potential implications, the findings are also consistent with a hypothesis we derived from social role theory (Eagly et al. 2000), which posits that beliefs about the gendered nature of traits stem from observations of the gendered nature of roles. Given that roles are primary in this framework, we anticipated that gender stereotyping and (differential) narrowness would be stronger among the role-words. Not only were the effects indeed stronger among the role-words, but there was no evidence for stereotype content and narrowness among the trait-words. This finding is surprising in light of research pointing to the ubiquity of each gender's association with particular traits (Blair and Banaji 1996; Blair et al. 2001). There may have been yet another difference between our role-words and trait-words however. To address this possibility, we reanalyzed a subsample of our words in which word-type differences in valence, frequency, length, and perceived gender could be controlled (those taken from Lenton et al., unpublished data). This reanalysis yielded a pattern of results identical to those reported. Of course there remain other differences between the role- and trait-words for which we could not control or account. For example, roles are more likely than traits to denote human involvement, perhaps because the latter are more likely to be polysemous than the former. As a consequence, there may be more fuzziness surrounding the similarity of trait-words (versus role-words) to the gender referents. Future research might investigate this idea more generally. In any case, a meta-analysis of more

traditional psychological measures would do well to examine the relative strength of role- versus trait-based gender stereotypes to further test our proposition.

Although they cannot account for our results of interest, let us briefly comment on some of the other comparisons one could make when inspecting Table 1. For example, masculine role-words are more strongly related to male referents than feminine role-words are to female referents. These particular results appear to be in line with research showing that male stereotypes are generally held more firmly than female stereotypes (Diekmann and Eagly 2000).

Table 1 also shows that the neutral role-words appear to be just as related to the female referent category as are feminine role-words. How can this finding be explained? Recall that our gender-related and gender-neutral words were not matched with respect to word frequency, valence, etc. When we subsequently examined the means for the subset of words in which these word categories are matched, the results are somewhat more in line with expectations: Feminine role-words were most strongly related to female referents ( $M=.14$ ,  $SE=.03$ ), followed by neutral role-words ( $M=.13$ ,  $SE=.04$ ) and, lastly, masculine role-words ( $M=.10$ ,  $SE=.03$ ). Still, while the difference between feminine and neutral role-words' semantic similarity to female referents was in the right direction, the difference was not great. We invite future researchers to replicate our results with other feminine, masculine, and gender-neutral words in order to examine their generalizability across different items.

Future LSA-based research also might examine other social category stereotypes. For example, words such as *slow* and *fragile* may share greater semantic meaning with *elderly* than with *young* (Hense et al. 1995). Similarly, LSA could be used to look at possible differences in stereotype narrowness (breadth) for these other categories. Perhaps normative groups are construed less broadly than non-normative groups more generally (Kahneman and Miller 1986). We encourage social psychologists to make use of LSA as a tool for understanding how people represent social information. We are aware of only one other social psychological application of LSA (Campbell and Pennebaker 2003).

#### Coda

The very meaning of the category referents *man/he/him* and *woman/she/her* is intricately tied to gender stereotypes. Our research shows that stereotypical roles share meaning with their matching category referent. Furthermore, the study suggests that while both the male and female categories are narrowly construed, the former effect was somewhat stronger. Our findings add to the literature on how the (American) English language reflects and reinforces existing gender stereotypes regarding men's and women's roles.



## Appendix

The feminine, neutral, and masculine role- and trait-words submitted to latent semantic analysis

	Roles	Traits
Feminine	Beautician	Affectionate
	Caregiver	Caring
	Cheerleader	Cheerful
	Dancer	Compassionate
	Decorator	Delicate
	Designer	Emotional
	Dietician	Flatterable
	Florist	Gentle
	Hairdresser	Gossipy
	Homemaker	Humble
	Housekeeper	Loyal
	Model	Moody
	Nanny	Nagging
	Nurse	Polite
	Receptionist	Sensitive
	Stylist	Shy
	Typist	Sympathetic
		Tender
		Understanding
		Warm
	Neutral	Assistant
Cashier		Candid
Clerk		Childlike
Doctor		Conceited
Editor		Conscientious
Lawyer		Conventional
Poet		Earnest
Reporter		Forward
Servant		Friendly
Worker		Gullible
		Happy
		Helpful
		Inefficient
		Irrational
		Jealous
		Likable
		Outspoken
		Reliable
		Ridiculous
		Secretive
		Sincere
	Solemn	
	Stubborn	
	Tactful	
	Theatrical	
	Truthful	
	Unpredictable	
	Unsystematic	
	Yielding	
Masculine	Architect	Aggressive
	Carpenter	Ambitious
	Coach	Analytical

Contractor	Arrogant
Detective	Assertive
Electrician	Athletic
Engineer	Authoritative
Farmer	Bold
Firefighter	Capable
Gambler	Charismatic
Inventor	Competitive
Machinist	Confident
Mechanic	Crude
Officer	Daring
Physicist	Decisive
Pilot	Dominant
Programmer	Forceful
Rancher	Independent
Sheriff	Individualistic
Soldier	Reckless
	Unyielding
	Vulgar

## References

- American National Corpus. (2007). Retrieved May 19, 2008, from <http://americannationalcorpus.org/SecondRelease/frequency2.html>.
- Anderson, C. A., & Sedikides, C. (1991). Contributions of a typological approach to associationistic and dimensional views of person perception. *Journal of Personality and Social Psychology, 60*, 203–217.
- Bem, S. L. (1974). The measurement of psychological androgyny. *Journal of Consulting and Clinical Psychology, 42*, 155–162.
- Berry, M. W., Dumais, S. T., & Letsche, T. A. (1995, December). Computational methods for intelligent information access. In Paper presented at the *Proceedings of Supercomputing*. San Diego, CA. Retrieved February 23, 2007 from <http://www.cs.utk.edu/~berry/sc95/sc95.html>.
- Beyer, S. (1999). The accuracy of academic gender stereotypes. *Sex Roles, 40*, 787–813.
- Biernat, M. (1991). Gender stereotypes and the relationship between masculinity and femininity: A developmental analysis. *Journal of Personality and Social Psychology, 61*, 351–365.
- Blair, I. V., & Banaji, M. R. (1996). Automatic and controlled processes in stereotype priming. *Journal of Personality and Social Psychology, 70*, 1142–1163.
- Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology, 81*, 828–841.
- Bosson, J. K., Prewitt-Freilino, J. L., & Taylor, J. N. (2005). Role rigidity: A problem of identity misclassification? *Journal of Personality and Social Psychology, 89*, 552–565.
- British National Corpus (2001). *The British National Corpus. version 2 (BNC World)*. Distributed by Oxford University Computing Services on behalf of the BNC Consortium. Retrieved February 21, 2007, from <http://view.byu.edu/>.
- Burn, S. M. (2000). Heterosexuals' use of "fag" and "queer" to deride one another: A contributor to heterosexism and stigma. *Journal of Homosexuality, 40*, 1–11.
- Campbell, R. S., & Pennebaker, J. W. (2003). The secret life of pronouns: Flexibility in writing style and physical health. *Psychological Science, 14*, 60–65.
- Carroll, J. B. (1956). Introduction. In J. B. Carroll (Ed.), *B. L. Whorf, Language, thought and reality: Selected writings of Benjamin Lee Whorf*. Cambridge, MA: MIT.

- Deerwester, S. C., Dumais, S. T., Landauer, T. K., Furnas, G. W., & Harshman, R. A. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science American Society for Information Science*, 41, 391–407.
- Dennis, S., Landauer, T., Kintsch, W., & Quesada, J. (2003). *Introduction to latent semantic analysis*. Slides from the tutorial given at the 25th Annual Meeting of the Cognitive Science Society, Boston. Retrieved February 27, 2007 from <http://lsa.colorado.edu/%7Esimon/LSATutorial.pdf>.
- Diekmann, A. B., & Eagly, A. H. (2000). Stereotypes as dynamic constructs: Women and men of the past, present, and future. *Personality and Social Psychology Bulletin*, 26, 1171–1181.
- Eagly, A. H., & Steffen, V. J. (1984). Gender stereotypes stem from the distribution of women and men into social roles. *Journal of Personality and Social Psychology*, 46, 735–754.
- Eagly, A. H., Wood, W., & Diekmann, A. B. (2000). Social role theory of sex differences and similarities: A current appraisal. In T. Eckes & H. M. Trautner (Eds.), *The developmental social psychology of gender* (pp. 123–174). Mahwah, NJ: Erlbaum.
- Ernst, S. B. (1995). Gender issues in books for children and young adults. In S. Lehr (Ed.), *Battling dragons: Issues and controversy in children's literature* (pp. 66–78). Portsmouth, NH: Heinemann.
- Gastil, J. (1990). Generic pronouns and sexist language: The oxymoronic character of masculine generics. *Sex Roles*, 23, 629–643.
- Glenberg, A. M. (1997). What memory is for. *The Behavioral and Brain Sciences*, 20, 1–55.
- Glenberg, A. M., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, 43, 379–401.
- Global Reach (2004). *Global Internet statistics: Sources and references*. Retrieved July 22, 2008 from <http://www.glreach.com/globstats/refs.php3>.
- Graddol, D. (2000). *The future of English? A guide to forecasting the popularity of the English language in the 21st century*. London, England: The British Council. Retrieved July 22, 2008 from <http://www.britishcouncil.org/de/learning-elt-future.pdf>.
- Helgeson, V. S. (1994). Prototypes and dimensions of masculinity and femininity. *Sex Roles*, 31, 653–682.
- Hense, R. L., Penner, L. A., & Nelson, D. A. (1995). Implicit memory for age stereotypes. *Social Cognition*, 13, 399–415.
- Judd, C. M., & Park, B. (1993). Definition and assessment of accuracy in social stereotypes. *Psychological Review*, 100, 109–128.
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93, 136–153.
- Kennison, S. M., & Trofe, J. L. (2003). Comprehending pronouns: A role for word-specific gender stereotype information. *Journal of Psycholinguistic Research*, 32, 355–378.
- Kontostathis, A., & Pottenger, W. M. (2002). *A mathematical view of latent semantic indexing: Tracing term co-occurrences*. Technical Report, LU-CSE-02-006, Department of Computer Science and Engineering, Lehigh University. Retrieved 16 September, 2007 from <http://www3.lehigh.edu/engineering/cse/research/reports/2002.asp>.
- Laham, D. (1997). Latent semantic analysis approaches to categorization. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the 19th annual meeting of the Cognitive Science Society* (p. 979). Mahwah, NJ: Erlbaum Retrieved March 1, 2007 from: <http://lsa.colorado.edu/>.
- Landauer, T. K. (2002). On the computational basis of learning and cognition: Arguments from LSA. In N. Ross (Ed.), *The psychology of learning and motivation*, 41 (pp. 43–84). New York: Academic.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent semantic analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–240.
- Landauer, T. K., Foltz, P. W., & Laham, D. (1998a). Introduction to latent semantic analysis. *Discourse Processes*, 25, 259–284.
- Landauer, T. K., Laham, D., & Derr, M. (2004). From paragraph to graph: Latent semantic analysis for information visualization. *Proceedings of the National Academy of Sciences of the United States of America*, 101, 5214–5219.
- Landauer, T. K., Laham, D., & Foltz, P. W. (1998b). Learning human-like knowledge by singular value decomposition: A progress report. In M. I. Jordan, M. J. Kearns, & S. A. Solla (Eds.), *Advances in neural information processing systems*, 10 (pp. 45–51). Cambridge: MIT.
- Landauer, T. K., Laham, D., Rehder, B., & Schreiner, M. E. (1997). How well can passage meaning be derived without using word order? A comparison of latent semantic analysis and humans. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the 19th annual meeting of the Cognitive Science Society* (pp. 412–417). Mahwah, NJ: Erlbaum.
- Lemaire, B., & Denhière, G. (2006). Effects of high-order co-occurrences on word semantic similarity. *Current Psychology Letters: Behaviour, Brain & Cognition*, 18.
- Lenton, A. P., & Webber, L. (2006). Cross-sex friendships: Who has more? *Sex Roles*, 54, 809–820.
- Lippa, R., & Connelly, S. (1990). Gender diagnosticity: A new Bayesian approach to gender-related individual differences. *Journal of Personality and Social Psychology*, 59, 1051–1065.
- Maass, A., Salvi, D., Arcuri, L., & Semin, G. (1989). Language use in intergroup contexts: The linguistic intergroup bias. *Journal of Personality and Social Psychology*, 57, 981–993.
- Maccoby, E. (1998). *The two sexes: Growing up apart, coming together*. Cambridge, MA: Harvard University Press.
- Prentice, D. A., & Carranza, E. (2002). What women and men should be, shouldn't be, are allowed to be, and don't have to be: The contents of prescriptive gender stereotypes. *Psychology of Women Quarterly*, 26, 269–281.
- Queller, S. (2002). Stereotype change in a recurrent network. *Personality and Social Psychology Review*, 6, 295–303.
- Queller, S., & Smith, E. R. (2002). Subtyping versus bookkeeping in stereotype learning and change: Connectionist simulations and empirical findings. *Journal of Personality and Social Psychology*, 82, 300–313.
- Rodler, C., Kirchler, E., & Holzl, E. (2001). Gender stereotypes of leaders: An analysis of the contents of obituaries from 1974 to 1998. *Sex Roles*, 45, 827–844.
- Schiffman, S. S., Reynolds, M. L., & Young, F. W. (1981). *Introduction to multidimensional scaling: Theory, methods, and applications*. San Diego, CA: Academic.
- Smith, E. R., & Conroy, F. R. (2007). Mental representations as states not things: Implications for implicit and explicit measurement. In B. Wittenbrink, & N. Schwarz (Eds.), *Implicit measures of attitudes* (pp. 247–264). New York, NY: Guilford.
- Smith, E. R., & DeCoster, J. (1998). Knowledge acquisition, accessibility, and use in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology*, 74, 21–35.
- Touchstone Applied Science & Associates (TASA) Inc. (2006). *Word frequency guide*. Retrieved January 22, 2008, from <http://www.tasaliteracy.com/wfg/wfg-main.html>.
- Watt, N., & Gow, D. (March 25, 2006). Chirac vows to fight the growing use of English. *Guardian News and Media Limited 2008*. Retrieved July 22, 2008 from <http://www.guardian.co.uk/world/2006/mar/25/france.eu>.
- Wolfe, M. B. W., & Goldman, S. R. (2003). Use of latent semantic analysis for predicting psychological phenomena: Two issues and proposed solutions. *Behavior Research Methods*, 35, 22–31.