

Gender representation in linguistic example sentences

Kotek, Hadas, Sarah Babinski, Rikker Dockum, & Christopher Geissler*

Abstract. Prior studies have shown that example sentences in syntax textbooks systematically under-represent women and perpetuate gender stereotypes (Macaulay & Brice 1994, 1997; Pabst et al. 2018). We examine the articles published over the past 20 years in *Language*, *Linguistic Inquiry*, and *Natural Language & Linguistic Theory*, and find striking similarities to this prior work. Among our findings, we show a stark imbalance of male (N=10807) to female (N=5019) arguments, and that male-gendered arguments are more likely to be subjects, and female arguments non-subjects. We show that female-gendered arguments are less likely to be referred to using pronouns and are more likely to be referred to using a kinship term, whereas male-gendered arguments are more likely to have occupations and to perpetrate violence. We show that this pattern has remained stable, with very little change, over the course of the twenty years that we examine, leading up to the present day. We conclude with a brief discussion of possible remedies and suggestions for improvement.

Keywords. gender representation; implicit bias; syntax; publications; academia;

1. Introduction. Example sentences are a central tool in teaching and research in the field of linguistics, often considered a key source of evidence to substantiate claims about a language or construction. Such examples often introduce the protagonists *John* and *Mary*, as well as gender-neutral names such as *Alex*, *Bobby*, *Chris*, and *Dana*, as advocated for in the Linguistic Society of America (LSA)’s Guidelines for non-sexist usage, adopted in 1997. That same year, Macaulay & Brice (1997) (henceforth M&B) published a paper in *Language*, the flagship journal of the LSA, which asked whether the use of gender-neutral names may suggest that example sentences are balanced and therefore do not exhibit gender bias. These authors presented a study of example sentences in eleven then-current syntax textbooks, and concluded that “the majority of constructed example sentences in syntax textbooks are biased toward male-gendered NPs, and [...] contain highly stereotyped representations of both genders.”

Twenty years later, Pabst et al. 2018 conducted a follow-up study of recently published textbooks in syntax to test whether the problem of gender bias has been rectified in the time since M&B’s ground-breaking work. In short, these authors find that the majority of problems identified by M&B in 1997 still plague syntax textbooks today: male arguments are over-represented in linguistic examples and are presented in a more positive light compared to female arguments.

* For questions and comments, we thank audiences at Yale University, Brandeis, and MIT, as well as the LSA 2020 Annual Meeting in New Orleans. Special thanks to the Yale undergraduates who participated in this study, listed in footnote 3. We additionally thank the Yale Women Faculty Forum and Claire Bown for providing funding which supported this work. The first author would like to additionally thank the member of the LSA Committee on the Status of Women in Linguistics and Katharina Pabst, Paola C  peda, and Kristen Syrett, for their support of this and related work. Finally, we thank Monica Macaulay, for comments on an almost-final draft of this paper, as well as for her ongoing support of this and related projects designed to draw attention to the representation of women in Linguistics and to improve it. Authors: Hadas Kotek, (hkotek@alum.mit.edu), Sarah Babinski, Yale University (sarah.babinski@yale.edu), Rikker Dockum, Swarthmore College (rikker@aya.yale.edu), Christopher Geissler, Yale University (christopher.geissler@yale.edu).

This paper goes beyond textbooks to test the distribution and role of arguments in example sentences in journal papers published in three major linguistics journals over the past twenty years — covering the period between the publication of Macaulay & Brice 1997 and that of Pabst et al. 2018. In short, we observe the same issues of imbalance in example sentences used in the research literature as in textbooks. Although we observe a positive trend, in its current pace the field will not reach parity in this regard for another fifty years. These findings affect all three journals we study to the same degree, suggesting that this is an entrenched way of thought in our field, which we argue must change. In what follows, we first present the main results of M&B 1997 and Pabst et al. 2018. We then present our current study and its results, and conclude by discussing broader implications for linguists and educators.

2. Preliminaries: Macaulay & Brice 1997 and Pabst et al. 2018. Macaulay & Brice 1997 present two studies: a careful study of a single textbook, published separately in Macaulay & Brice (1994), and a comparative study of 10 additional textbooks, whose goal is to ascertain whether the gender imbalance found in the first study generalizes across other textbooks. Here we concentrate on the latter study, whose design and results inspired both Pabst et al. 2018 and the study we present in section 3. We first present the results of this study, and then of the followup study in Pabst et al.

2.1. MACAULAY & BRICE 1997. M&B study 10 syntax textbooks published between 1969 and 1994. 7 textbooks had male authors, and 3 had female authors. 200 examples were randomly sampled from each textbook, and manually coded by for the following:

- (1) *M&Bs coding:*
 - a. Grammatical function (subject, direct object, indirect object, etc.)
 - b. Theta roles (agent, patient, experiencer, recipient, etc.)
 - c. Lexical choices (pronouns, proper names, violence, appearance, etc.)

Here we summarize some of the main findings. We refer the reader to M&B (1997) for a more comprehensive discussion, and for example sentences illustrating each one of the findings below. In short, M&B find that example sentences introduce male protagonists at higher rates than female ones, and that they perpetuate gender biases, as summarized in (2). When female-gendered arguments are over-represented in example sentences, this, too, is done in a way that perpetuates stereotypes, as outlined in (3).

- (2) *M&B's findings: male-gendered arguments*
 - a. Appear as arguments more often than female-gendered arguments
 - b. Are more likely to be subjects and agents
 - c. Are subjects and agents more often than other roles
 - d. Have pronouns mentioned more often than female-gendered arguments
 - e. Have names mentioned twice as often as female-gendered arguments
 - f. Are engaged in 'intellectual activities' such as book reading/handling, and working with cars more often than female-gendered arguments
 - g. Are described as having occupations more often than female-gendered arguments, and in a wide range of occupations
 - h. Perpetrate violence more often than female-gendered arguments

(3) *M&B's findings: female-gendered arguments*

- a. Often lack names, and are referred to with kinship terms (X's wife, mother) more often than male-gendered arguments are
- b. Have their appearance described more often than male-gendered arguments

These findings are shown in the selected examples below, representing only a small portion of those shown by M&B's to illustrate their observations:

(4) *M&B's findings: selected examples*

- a. Harry watches the fights and his wife the soap operas.
- b. Bill is proud of his father and tired of his mother.
- c. Every painting of Maja and photograph of Debbie pleased Ben.
- d. The man is hitting the woman with a stick.
- e. The man who shot her believed there was someone else who was seeing Helen.

In addition, the syntax textbooks studied by M&B commonly used examples that contained explicit and suggestive language:

(5) *M&B's findings: explicit and suggestive language in example sentences*

- a. Max doesn't beat his wife because he loves her.
- b. After Rambo as a lover, she was exhausted.
- c. She's fond of John naked.
- d. I can't imagine you in kinky boots.
- e. John's turned on by Mary in tight trousers.
- f. She'll soon tire of her sexploits.
- g. What a nice pear Mary's got!
- h. John forced Mary to be kissed by Bill.

Finally, the gender of the textbook author played an important role: male authors were on average much more likely to choose biased examples, whereas female authors tended toward a more balanced sample.

M&B thus conclude that "Our results clearly illustrate the need for such scrutiny: females are simply not significant actors in the world constructed by sample sentences" (p. 816).

2.2. PABST ET AL. 2018. The majority of M&B's findings are replicated in the study of textbooks published between the years 2005–2017 presented in Pabst et al. 2018. Results for male-gendered arguments are given in (6). The only finding in M&B (2) that was not apparent in Pabst et al. 2018 (6) that men often worked with cars.

(6) *Pabst et al.'s (2018) findings: male-gendered arguments*

- a. Appear as arguments more often than female-gendered arguments
- b. Are more likely to be subjects and agents
- c. Are subjects and agents more often than other roles
- d. Have pronouns mentioned more often than female-gendered arguments

- e. Have names mentioned twice as often as female-gendered arguments
 - f. Are engaged in ‘intellectual activities’ such as book reading/handling more often than female-gendered arguments
 - g. Are described as having occupations more often than female-gendered arguments, and in a wide range of occupations
 - h. Perpetrate violence more often than female-gendered arguments
- (7) *Pabst et al.’s (2018) findings: female-gendered arguments*
- a. Often lack names, and are referred to with kinship terms (X’s wife, mother) more often than male-gendered arguments are
 - b. Exhibit more negative emotions than male-gendered arguments

Some findings about female-gendered arguments are replicated as well, summarized in (7). Unlike M&B (1997), however, Pabst et al. no longer find that female-gendered arguments have their appearance described more often than male ones — in fact, they find a very small number of examples that describe physical appearances in general. On the other hand, they find that female-gendered arguments exhibit a greater proportion of negative emotions in their sample, a finding that was not investigated in the original M&B study.

In general, Pabst et al. find very little sexually explicit or suggestive language, but they find that the majority of stereotypes about both male- and female-gendered arguments is maintained in recent textbooks. Some examples are shown below, and the reader is referred to the original text for more data.

- (8) *Pabst et al.’s findings: selected examples*
- a. She snarled at the students who hadn’t read the book.
 - b. Bruce loved and Kelly hated phonology class.
 - c. Joan believes he is a genius even more fervently than Bob’s mother does.
 - d. Mary entertained the men during each other’s vacation.
 - e. He drove her hard, he stole her fame or would have if he could have.
 - f. Mohammed buys a house.
 - g. The woman bought rice for the children.
 - h. Slavko left his wife.
 - i. Mary may wonder if John cheats on her.

Finally, Pabst et al. find some differences between male and female authors, such that female authors tended to use a higher proportion of non-gendered arguments than male authors. One female author used a more balanced proportion of gendered arguments, but the other two used similar proportions to the men. Overall there were no statistically significant differences between the proportions of male to female arguments based on the textbook authors’ genders.

To summarize, the vast majority of problems that afflicted example sentences in syntax textbooks 20 years ago are still present today. The main change has affected explicitly suggestive and stereotypical examples. While Pabst et al. find a small number of such examples, they are no longer as blatant or numerous. Instead, the discrepancies are made more difficult to detect, although they remain present: the skew now requires a broader lens to observe. Once this

is done, however, Pabst et al. find a great range of ways in which implicit gender biases are present in their sample, such that men are over-represented and presented more favorably, and women are suppressed. They additionally point out that non-binary identities are non-existent in their sample, and that there is a Western bias in choice of names and contexts.

3. Study. On the heels of these results, and following a recent study of French example sentences (Richy & Burnett 2019), in this study we examine all papers from 3 journals: *Linguistic Inquiry*, *Natural Language & Linguistic Theory*, *Language* between the years 1997–2018. We find similar findings to those reported in the above studies.

3.1. METHODS AND DESIGN. Unlike previous studies, instead of sampling, we automatically extracted *all* example sentences from all papers published during the time period under investigation. To do so, we relied on the fairly standard ways in which example sentences are formatted in linguistic journals: marked by a number enclosed in parentheses, removed from the margin, perhaps including a gloss and translation. We used regular expressions to capture all such sentences, which we then manually explored.^{1,2}

Twenty-four Yale University undergraduate students were hired to clean up the examples and provide the coding.³ As a first step, these students were instructed to remove all data points which were not example sentences, but rather were mistakenly captured or are generalizations, rules, descriptions of data, or any other text that appeared in numbered examples. Next, examples that did not contain third person arguments were removed. Finally, for each third person argument in each of the remaining example sentences, our coders provided information along the same lines as the studies above:

(9) *Argument coding in this study:*

- a. Grammatical function (subject, direct object, indirect object)
- b. Theta roles (agent, patient, experiencer, recipient)
- c. Lexical choices (pronouns, proper names, violence, appearances, etc)

We decided to focus on major syntactic and semantic roles, as spelled out in (9a–b). Unlike prior work, we did not code minor theta roles such as *beneficiary* or *instrument*, anticipating the data to be too sparse for analysis, as was the case in those prior studies.

Coders were instructed to flag any example that they considered stereotypical in any way. In addition, they were given a list of specific properties to look out for. Following prior work, we were mainly interested in violence, appearances, emotions, “intellectual activities” such as writing/reading, and cars. Data for kinship terms was extracted automatically using regular expressions, as the set of such terms is sufficiently limited to easily allow such a search.

¹ This means that we did not explore any examples that appeared in the text of a paper, outside of numbered examples. Furthermore, we recognize that it’s possible that our automated means of extracting examples may have caused some to be missed. If such data is missing, we have no reason to believe that it’s a large amount, as we manually checked that large page ranges are not missing from the data; we also have no reason to believe that the missing data should be in any way skewed. Therefore, we believe that our results are representative and valid.

² We chose to exclude Macaulay & Brice 1997, published in *Language*, to avoid skewing the results.

³ We would like to thank our coders: Aarohi Srivastava, Abigail Fortier, Amelia Lake, Calvin Kaleel, Faren Roth, Georgia Michelman, Joe Class, Joshua Celli, Justin Yamamura, Karina Di Franco, Lena Venkatraman, Michael Gancz, Nanyan Wu, Nico Kidd, Oliver Shoulson, Prastik Mohanraj, Ronnie Rodriguez, Serena Puang, Shayley Martin, Slater Smith, Stella Fitzgerald, Stella Xu, and Zhiliang Fang. We thank the Yale Women Faculty Forum and Claire Bower for providing funds to support this part of the project.

In the case of non-English examples, our coders were instructed to rely on the English translation, and in addition to inspect the glosses to ensure that the translation was a good match. We were specifically concerned with possible cases of mismatches between gloss and translation, such as a non-gendered third-person pronoun being translated into English as “he.” Further, in cases of gendered languages, we did not include in our data any inanimate objects, even if they bore masculine or feminine agreement. Only human arguments were included in our analysis.

Coders received hands-on training in conducting the coding described here. Their work was inspected on an ongoing basis by one of the paper’s authors, and feedback was provided as needed to improve their work. We additionally conducted some post-hoc tests to verify the accuracy of the coding, such as the use of regular expressions to find examples with relevant traits and confirm that they were flagged by the coders, and spot-checking of random individual examples.

3.2. RESULTS. Based on these considerations, we analyzed a total of 873 papers from our 3 journals of interest, resulting in 15,826 third person human/animate arguments for our analysis. We first show results for each of the properties listed in (9), and then discuss a meta-analysis over time and considering language of example in the next section. All results are statistically significant, as confirmed by t-tests.

3.2.1. DISTRIBUTION OF ARGUMENTS. We begin by considering the overall ratio of male to female arguments in our sample. Of the 15,826 arguments we coded, 32% were female, for a ratio of 2.1 male arguments for every 1 female argument.




Figure 1: Gender distribution of arguments in all example sentences in this study

These ratios are consistent across the three journals that we studied:




Figure 2: Gender distribution of arguments in all example sentences by journal

Specifically, we find that the ratio of female-gendered arguments in the data to the total number of arguments is 31% for papers published in *Natural Language & Linguistic Theory*, 32% for paper published in *Linguistic Inquiry*, and likewise 32% for papers published in *Language*. This will allow us to show combined results for all three journals in the data that follows

3.2.2. GRAMMATICAL FUNCTION AND THEMATIC ROLES. Next, we consider the distribution of arguments with respect to syntactic and semantic roles. We find that 74% of male arguments, while only 66% of female arguments are subjects. That is, female-gendered arguments are over-represented in non-subject roles (here, direct and indirect object) compared to their male counterparts. Similarly, we observe a skew in theta roles: Female arguments represent 30% of agents and experiencers, 34% of patients, and 43% of recipients. These arguments are significantly over-prevalent as patients and recipients, when we take into account their already diminished distribution overall, as presented in Figure 1.

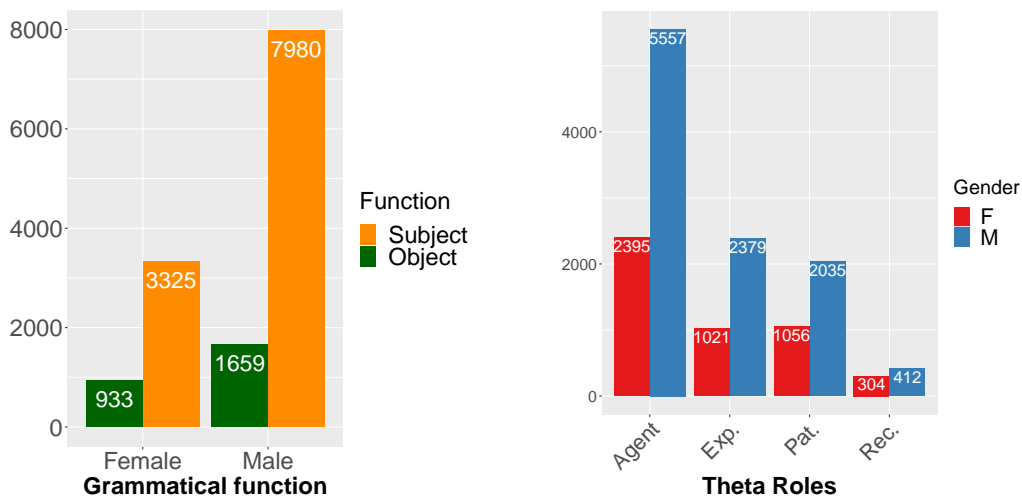


Figure 3: Distribution of arguments by grammatical function (left) and thematic role (right)

3.2.3. PRONOUNS AND PROPER NAMES. Considering the distribution of pronouns and proper names, we find that male-gendered pronouns are significantly more common than female ones, at 32% vs 25%, respectively. We tentatively suggest that this may be a side-effect of the overall prevalence of male-gendered arguments, especially subjects, in the data. If a pronoun were to refer back to an argument, that argument is more likely to be male than female.⁴ Proper names, on the other hand, are equally distributed, at 53% vs 54%, respectively.

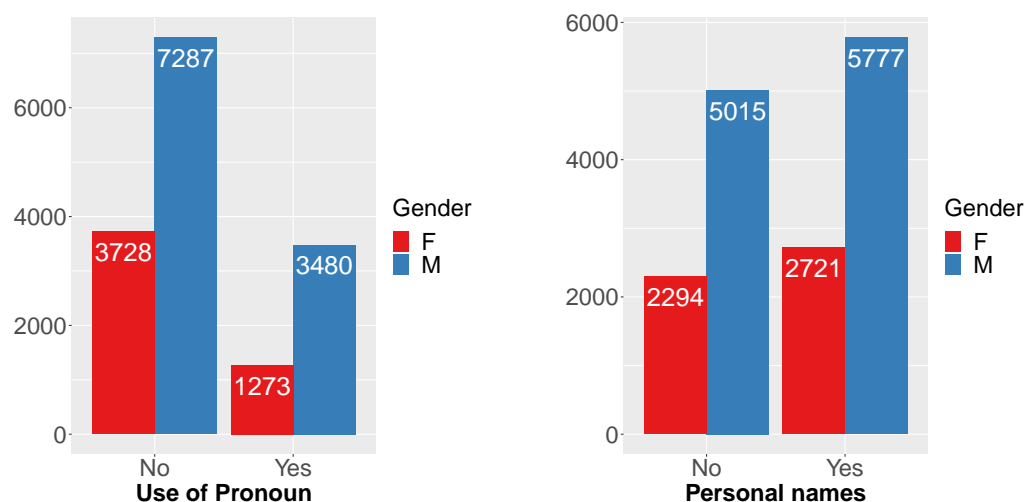


Figure 4: Pronouns (left) and proper names (right) in the data

Concentrating more closely at names, we observe that 37% of all male names are *John*. Concomitantly, 37% of all female names are *Mary*. In fact, 3 of the top 5 male names are “John” variants: *John*, *Juan*, and *Ivan*, whereas 4 of the top 5 female names are “Mary” variants: *Mary*, *Maria*, *Marie*, and *Mari*.




Figure 5: Top five most common male (left) and female (right) names

⁴We thank Susan Fisher (p.c.) for this suggestion.

3.2.4. LEXICAL CHOICES. Finally, we consider various lexical choices about the example sentences in our study. We start with examples that refer to an argument’s occupation. We find that male-gendered arguments are over-represented in such examples, even given their 2:1 overall distribution in the data: they are described as having occupations in 74% of all such examples.

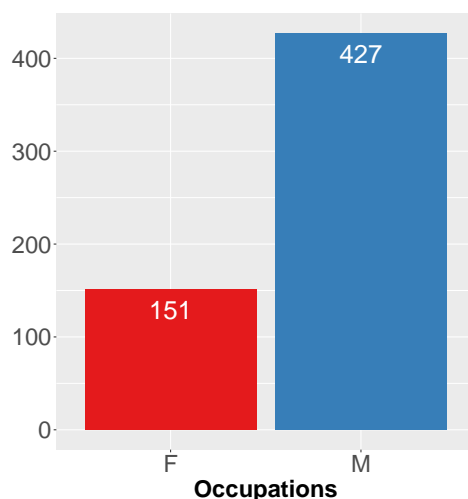


Figure 6: Arguments whose occupation is discussed in the example

Next we concentrate on examples which involve violent events of some kind. Observing first the overall distribution of arguments, we find that male-gendered arguments are massively over-represented in such examples: 85% of the arguments in these examples are male. When we next observe the distribution of arguments as subjects vs non-subjects within each gender category, however, we find no differences: male-gendered subjects comprise 72% of all male arguments, while female-gendered arguments comprise 70% of all female arguments.




Figure 7: Violent events by gender (left) and by gender×syntactic position (right)

Turning to examples involving the expression of romantic or sexually suggestive content, we now find that the gender distribution in examples is remarkably different. Here, female-gendered arguments are over-represented, comprising 49% of all arguments. Considering the distribution of subject and non-subject positions, female-gendered arguments are again over-represented as *objects* in such sentences: only 57% of female-gendered arguments are subjects, whereas 73% of male-gendered arguments are subjects.




Figure 8: “Romantic” events by gender (left) and by gender×syntactic position (right)

Finally, female-gendered arguments are massively over-represented among those referred to by kinship terms: 57% of all such arguments are female-gendered. Considering the overall 2:1 male skew in the data, this is a remarkable finding.




Figure 9: Use of kinship terms by gender

3.2.5. SOME ILLUSTRATIVE EXAMPLES. Like in Pabst et al. 2018 and unlike in the earlier Macaulay & Brice 1997 study, we do not find many sexually explicit or suggestive language. Nonetheless, we find that stereotypes of both genders are commonly used in linguistic example sentences. In the following example we provide a representative sample of such sentences:

- (10)
- a. John ate the meal and Mary cleaned the dishes
 - b. John didn't eat the meal because he would have to clean the dishes
 - c. * Kelly broke again tonight when she did the dishes
 - d. John (not Peter) washed cars well
 - e. Tomas replaced Ricardo as the captain
 - f. John thinks that he himself is a war hero
 - g. John told Bill that Mary began to cry without any reason
 - h. The boys had thrown no rock at the cars
 - i. Maria reviewed the novel, she didn't write it
 - j. The senator killed herself
 - k. Iraqi father drowned his 17 year old
 - l. Mary, being dumb, needs to sit down
 - m. Which Nobel prize winning author came in his car?
 - n. I called for a policeman, not a policewoman
 - o. Every male student doesn't fear tigers
 - p. At least one student of every professor_i is horrified at his_i grading procedure

3.3. META-ANALYSIS. We comment on two additional aspects of the data we have collected. First, we note that whether the example was in English or not did not affect the results we observe. We find a total of 33% of female-gendered arguments in English examples, and 31% female-gendered arguments in non-English examples.




Figure 10: Gender distribution of arguments in all example sentences in this study

We believe that this, again, is suggestive of a broad issue in our field. The bias is not introduced by non-English examples, where some constraints on data collection may affect the sentences tested by field linguists in various ways. Instead, English and non-English examples show similar bias, indicating that access to data or speakers, or familiarity with the language, do not significantly affect the gender distribution of arguments.

Finally, considering the distribution of male vs female arguments in example sentences over time, we notice a positive trend over the past 20 years. However, at no point — in no year over the 20 years of this study — were the genders at parity. On average, we have moved 5% closer to parity over the course of two decades. At this pace, all things being equal, we would not expect to reach equality in the use of male and female arguments for at least 50 more years.




Figure 11: Gender distribution of arguments over time

3.4. SUMMARY. To summarize our main findings, male-gendered arguments are over-represented in the sample overall and in particular as subjects and agents. They are also over-represented in some stereotypical examples. These results are consistent across journals and time, and replicate findings for textbooks in prior work.

- (11) *Main findings: male-gendered arguments*
- a. Appear twice often as *arguments* in total
 - b. Appear more often as *subjects and agents*
 - c. Are referred to with pronouns more often
 - d. Engage in significantly more *violence*
 - e. Have significantly more *occupations*

- (12) *Main findings: female-gendered arguments*
 - a. Are over-represented as *non-subjects*, especially as *recipients*
 - b. Are over-represented in sentences involving *romantic/sexual language*
 - c. Are massively over-referred to using *kinship* terms
- (13) *Main findings: other trends in the data*
 - a. No effect of language of example (English vs. non-English) on the results
 - b. A small trend of improvement in gender ratios: from low-30% to mid-30%
 - c. Few or no overtly suggestive or explicit examples
 - d. But gendered stereotypes are very much evident

4. Discussion. Constructed examples sentences are one of the main sources of data in theoretical linguistics. Many examples are cited again and again, often divorced from their original source, and some become ‘canonical’ examples of particular phenomena. As scientists we think of these examples as a tool in our arsenal, used in furtherance of a better understanding of the language faculty in all its aspects. However, we often ignore the social aspects that these examples occur in and that they exemplify. We have demonstrated here that example sentences used in the literature may encode implicit biases (even at a very subtle level). These then get handed down to new generations of linguists, perpetuating the cycle.

These trends do not occur in isolation: rather, they reflect biases that are prevalent across our society at large, and the field of linguistics more broadly. For example, studies have found similar biases in English language textbooks for various dialects of English (e.g. Bergvall 1996; Lee & Collins 2010; Lewandowski 2014; Tarrayo 2014), and recently in example sentences in French journals (Richy & Burnett 2019). The issue has garnered some attention in recent LSA presentations (Boyce et al. 2019; Haugen & Margaris 2020, as well as the 2018 panel *Our Linguistics Community: Addressing bias, power dynamics, harrassment*). It has featured in discussions of the representation of gender more broadly, including in reference to language acquisition (e.g. Eckert & McConnell-Ginet 2013; McConnell-Ginet 2014; Leslie et al. 2015; Meyer et al. 2015; Bian et al. 2017, among others).

Nonetheless, as scientists, we believe that we *could* and *should* strive to avoid perpetuating bias, even if implicit, in our work. We argue here that better constructed example sentences, using inclusive language, can send an important message to the field: inclusive language encourages participation from underrepresented groups, leading to a better community and therefore to better science, at the cost of just a little more thoughtfulness.

To be clear, we list here several suggestions that could lead to better constructed examples. First, be mindful of the composition of your example sentences and the distribution of arguments in them: ensure that male arguments aren’t over-represented, and avoid all stereotypes — about both men and women. Second, avoid using sexually explicit, demeaning or biased language. Do not perpetuate bias in the examples you cite or construct — if needed, paraphrase a flawed original rather than reproducing it verbatim. Likewise, avoid the use of gendered lexical items such as *-man* and *he* where unnecessary. Adopt and encourage instead the use of singular *they*, as a more inclusive pronoun when referring to (singular) nouns whose gender is unknown. Consider using singular *they* even when the argument’s gender is known, but is irrelevant to the example. Use inclusive nouns such as *Congressperson* and *humankind*.

Keep the LSA's Guidelines for Inclusive Language, published in 2017, in mind when writing, as well as when reviewing and editing papers.

Consider going beyond the common “*John, Mary, Bill, and Sue,*” to embrace more diverse names and representations of the protagonists in your examples, including non-Western and those used more frequently in non-binary, and non-heterosexual communities. More generally, think past the first names that come to mind when you ask yourself who to feature in your examples, who to cite in your papers, who to teach in your syllabi and courses, and even who to invite to your organized events. In sum, remember that as a teacher and simply as a citizen of the field, you are in a position of authority and can have a positive influence on new members of the field and the community as a whole. Be sensitive to how you portray all individuals in your examples and what implicit messages you may inadvertently send through your examples.

5. Conclusion. This paper presents evidence in support of the conclusion that constructed example sentences exhibit a variety of gender biases and stereotypes. We showed that this has been the case in three leading journals in the field of linguistics, virtually unchanged over the course of the past twenty years, and that it is consistent across different languages used in the sentences. We additionally discussed several ways of mitigating these effects.

Before concluding, we will point out here several additional avenues of research which we leave to future work. First, we have not yet investigated the possible effects of properties of the authors themselves on the makeup of example sentences they use: including their gender, their PhD institution, their institution at the time of publication, their geographic location, their age, their race, etc. Prior work has shown that author gender, at least, may have an effect on these results. We have additionally not yet investigated the specific lexical items used in example sentences. For example, we know from prior work that male-gendered arguments have not only *more* occupations, but also more diversity in the *types* of occupations they are portrayed as having. Similarly, the degree of violence that men engage in in example sentences is often more severe than that perpetrated by women. Finally, we leave to future work sentiment analysis, to determine the effects of gender in examples involving positive and negative emotions and connotations.

An additional avenue for future research is a comparison of constructed example sentences with examples drawn from field elicitations and other corpora. In such cases, authors must consider both the factors discussed in this paper and the representation of the voices presented. We have already begun pursuing this work, although we do not include it here.

References

- Bergvall, Victoria. 1996. Humpty Dumpty does syntax: Through the looking glass, and what Alice found there. *Natural Language & Linguistic Theory* 14(2). 433–443.
- Bian, Lin, Sarah-Jane Leslie & Andrei Cimpian. 2017. Gender stereotypes about intellectual ability emerge early and influence children's interests. *Science* 355(6323). 389–391. <https://doi.org/10.1126/science.aah6524>.
- Boyce, Veronica, Titus von der Malsburg, Till Poppels & Roger Levy. 2019. Female gender is consistently under-expressed in pronoun production and under-inferred in comprehension. Poster presentaion. *The 93rd Annual Meeting of the Linguistic Society of America*. New York, NY.

- Eckert, Penny & Sally McConnell-Ginet. 2013. *Language and gender*. Cambridge: Cambridge University Press.
- Haugen, Jason D. & Amy V. Margaris. 2020. Faculty placements into linguistics PhD programs across the US and Canada: Market share and gender distribution. Presentation in the 94th Annual Meeting of the Linguistic Society of America.
- Lee, Jackie & Peter Collins. 2010. Construction of gender: A comparison of Australian and Hong Kong English language textbooks. *Journal of Gender Studies* 19. 121–137.
- Leslie, Sarah-Jane, Andrei Cimpian, Meredith Meyer & Edward Freeland. 2015. Expectations of brilliance underlie gender distributions across academic disciplines. *Science* 347(6219). 262–265. <https://doi.org/10.1126/science.1261375>.
- Lewandowski, Marcin. 2014. Gender stereotyping in EFL grammar textbooks. A diachronic approach. *Linguistik Online* 68(6).
- Macaulay, Monica & Colleen Brice. 1994. Gentlemen prefer blondes: A study of gender bias in example sentences. *Cultural performances: Proceedings of the Third Berkeley Women and Language Conference* 449–461.
- Macaulay, Monica & Colleen Brice. 1997. Don't touch my projectile: gender bias and stereotyping in syntactic examples. *Language* 73(4). 798–825.
- McConnell-Ginet, Sally. 2014. Meaning-making and ideologies of gender and sexuality. In Susan Ehrlich, Miriam Meyerhoff & Janet Holmes (eds.), *The handbook of language, gender, and sexuality*. Chapter 16. Hoboken, NJ: John Wiley & Sons.
- Meyer, Meredith, Andrei Cimpian & Sarah-Jane Leslie. 2015. Women are underrepresented in fields where success is believed to require brilliance. *Frontiers in Psychology* 6(235). <https://doi.org/https://doi.org/10.3389/fpsyg.2015.00235>.
- Pabst, Katharina, Paola Cépeda, Hadas Kotek & Kristen Syrett. 2018. Gender bias in linguistics textbooks: Has anything changed since Macaulay & Brice (1997)? Manuscript.
- Richy, Célia & Heather Burnett. 2019. Jean does the dishes while Marie fixes the car: a qualitative and quantitative study of social gender in French syntax articles. *Journal of French Language Studies*.
- Tarrayo, Veronico Nogales. 2014. Gendered word (or world): Sexism in Philippine preschool English language textbook. *Journal on English Language Teaching* 4(2). 25–32.