RESEARCH ARTICLE

The P600 for singular “they”: How the brain reacts when John decides to treat themselves to sushi.

Grusha Prasad^a and Joanna Morris^b

^aJohns Hopkins University ^bHampshire College, Rhode Island School of Design

ARTICLE HISTORY

Compiled January 24, 2020

ABSTRACT

There has been increased awareness that individuals need not have a binary gender identity (i.e., male or female), but rather, gender identities exist on a spectrum. With this increased awareness, there has also been an increase in the use of *they* as a singular pronoun when referring to individuals with a non-binary gender identity. Has the processing of singular *they* changed along with a change in its usage? Previous studies have demonstrated that sentences in which *they* is co-indexed with singular antecedents, are judged acceptable and are easy to process, but only if the antecedents are non-referential and/or have ambiguous gender; co-indexing *they* with referential antecedents with unambiguous gender (e.g., *Mary*) results in lower acceptability ratings and greater processing effort. We investigated whether participants who frequently interacted with individuals with a non-binary gender identity and/or identified as having a non-binary gender themselves would process sentences in which *themselves* was co-indexed with singular antecedents similarly. We found a significant P600 effect for sentences in which *themselves* was co-indexed with singular referential antecedents with unambiguous gender, but failed to find a P600 effect when the antecedents were non-referential and/or had an ambiguous gender. This pattern of results is consistent with behavioural results from previous studies, suggesting that the change in the usage of singular *they* has not resulted in a corresponding change in the way in which this pronoun is processed.

KEYWORDS

singular they; P600; sentence processing; non-binary gender;

1. Introduction

There has been increased awareness that individuals' gender identity does not have to be binary (i.e., male or female), but rather, gender identities exist on a spectrum. With this increased awareness, there has also been an increase in the use of *they* as a singular pronoun when referring to individuals with a non-binary gender identity. In a recent study, Camilliere et al. (2019) found that participants (especially those who were more accepting of non-binary individuals) rated sentences like (1) as being fairly "natural" (with a mean score of 5.2 on a 7-point likert scale).

- (1) My nurse_i was able to explain how my medication would affect me even though they_i had no say in prescribing it.

In contrast, two decades ago, Foertsch and Gernsbacher (1997) found that participants read sentences in which singular antecedents were co-indexed with *they* (like (1)) more slowly than minimally different sentences in which singular antecedents were co-indexed with *she* or *he* (like (2)). This suggests that participants found it difficult process sentences in which singular antecedents like *my nurse* were co-indexed with *they*.

- (2) My nurse_i was able to explain how my medication would affect me even though she_i had no say in prescribing it.

If the change in processing of *they* changed along with its usage, then participants who rate sentences like (1) as acceptable, are also likely to find such sentences easy to process. On the other hand, if the processing of *they* did not change along with its usage, then sentences like (1) are likely to be more difficult to process than sentences like (2), despite the fact that both (1) and (2) are rated as being acceptable.

Based on previous theoretical analyses of singular *they* and analyses of how gender is cognitively and linguistically encoded, we hypothesized that participants would find sentences like (1) easy to process if they were able to dissociate gendered antecedents *nurse* from their associated gender. We further hypothesized that there were two groups of participants who were most likely to be able to dissociate the gender as-

sociated with nouns like *nurse* or names like *Mary*: individuals with a non-binary gender identity (*NBIs*) and individuals who frequently interact frequently with NBIs. We tested these hypotheses in an ERP experiment, using the P600 amplitude as a measure of processing effort.

To foreshadow our results, we observed a P600 effect when participants were listening to sentences in which *themselves* was co-indexed to referential antecedents with unambiguous gender (e.g., *Mary*). We failed to find such a P600 effect when participants were listening to sentences in which *themselves* was co-indexed with antecedents that were non-referential (e.g., *some woman*), had an ambiguous gender (e.g., *the participant*) or were both non-referential and had an ambiguous gender (e.g., *someone*). These results are consistent with studies looking at the processing of singular *they* prior to the rising popularity of singular *they* (Foertsch and Gernsbacher, 1997), suggesting that the change in the usage of singular *they* has not resulted in a corresponding change in the way in which this pronoun is processed.

1.1. Analysis of singular *they*

Even prior to the rising use of *they* to refer to NBIs, *they* was used to refer to singular antecedents that were unmarked for gender. This was true in cases when the antecedents were non-referential and did not pick out a specific individual (e.g., (3) and (4)) as well as in cases when the antecedents were referential and picked out a specific individual (e.g., (5)).

- (3) Everyone_i takes their_i time finding a seat. (Balhorn, 2004)
- (4) If a student_i is getting a low grade, they might want to go and talk to their_i teacher. (Balhorn, 2004)
- (5) A runner I knew_i always ate lots of pasta the night before a race even when they_i would rather have steak. (Foertsch and Gernsbacher, 1997)

Balhorn (2004) accounts for the use of *they* in these sentences by arguing that *they* has a greater “semantic justification” than either *he* or *she* — if a speaker decides that the gender of the antecedent does not need to be made salient in the discourse, then

the speaker might use *they* since it is unmarked for gender, whereas *he* and *she* are marked for gender. This account cannot explain the use of *they* in sentences like (6), however. In (6), the antecedent is unambiguously masculine, and therefore there is a greater “semantic justification” to use *he* than to use *they*.

- (6) There is not a man_i I meet but doth salute me. As if I were their_i well-acquainted friend. (Bjorkman, 2017)

Bjorkman (2017) accounts for the use of *they* in such examples by proposing that although antecedents are usually marked with [*Masc*] or [*Fem*], the gender of antecedents can remain unmarked in some circumstances i.e. ($\emptyset$). Specifically, antecedents are unmarked for gender if they are not semantically associated with masculine or feminine features (e.g., *a runner I knew*) or if they are semantically associated with masculine or feminine features but are non-referential (e.g., *a man*). Non-referential antecedents *could* be marked for gender, but it is not necessary for them to be. Therefore, according to Bjorkman, *a runner* is unmarked for gender, *a man* is optionally marked for gender and *the man* is obligatorily marked for gender. Consequently, her account predicts that since *they* is unmarked for gender, antecedents like *a runner* and *a man* which are also unmarked for gender can be co-indexed with *they*, whereas antecedents like *the man* which are marked for gender cannot be co-indexed with *they*.

Results from Foertsch and Gernsbacher (1997) provide empirical support for Bjorkman’s account. As discussed earlier, participants read sentences like (1) more slowly than sentences like (2) (repeated below as (7a) and (7b)) suggesting that participants incurred a processing cost when a referential antecedent with unambiguous gender was co-indexed with *themselves*.¹ However, the reading times for (8a) and (8b) were equivalent as were the reading times for (9b) and (9a), suggesting that participants did not incur a processing cost when non-referential antecedents or antecedents with unambiguous gender were co-indexed with *themselves*.

- (7) *Referential antecedent with unambiguous gender*

¹*My nurse* does not technically have unambiguous gender — male nurses exist. However, due to gender stereotypes and frequency effects, participants often expect nurse to be female.

- a. My nurse_i was able to explain how my medication would affect me even though they_i had no say in prescribing it.
 - b. My nurse_i was able to explain how my medication would affect me even though she_i had no say in prescribing it.
- (8) *Non-referential antecedent with unambiguous gender*
- a. A nurse_i should have an understanding of how a medication works even if they_i will not have any say in prescribing it.
 - b. A nurse_i should have an understanding of how a medication works even if she_i will not have any say in prescribing it.
- (9) *Referential antecedent with ambiguous gender*
- a. A runner I knew_i always ate lots of pasta the night before a race even when she/he_i would rather have steak.
 - b. A runner I knew_i always ate lots of pasta the night before a race even when they_i would rather have steak.

Additionally, Bjorkman found that most people judged sentences like (10) as ungrammatical when the gender of the antecedent was known, despite the fact that the noun *professor* is usually unmarked for gender. However, a certain subset of speakers (who Bjorkman calls *creative users*), judged sentences like (10) as grammatical, suggesting that these creative users did not mark the antecedent with a gender even when the gender of the antecedent was known.

- (10) The professor_i said they_i cancelled the exam.

Yet, even these *creative users* judged sentences like (11) and (12) as being ungrammatical, suggesting that these speakers did mark antecedents like *actress* and *Janet* with a gender ([*FEM*]).

- (11) I just saw the lead actress_i, and I really love their_i costume.
- (12) Janet_i left their_i sweater here.

Based on these set of results, Bjorkman argues that even when creative users know the gender of some referential antecedent, they do not obligatorily mark the antecedent with the known gender. However, even for such creative users, there are certain noun classes for which the gender obligatorily marked (e.g., kinship relations, stereotypical proper names and nouns with morphological gender marking). This is because these noun classes are uniformly associated with masculine or feminine features, which results in the gender being overtly marked on the lexical item.

This account predicts that if a subset of speakers did not uniformly associate masculine or feminine features with kinship relations and stereotypical proper names (*super-creative users*), then gender would not be overtly marked on the lexical items like *sister* or *Mary*; consequently, gender marking for antecedents like *my sister* or *Mary* for such users would not be obligatory. Therefore, these super-creative users would be able to co-index *they* with antecedents like *Mary*, even if *Mary* is a stereotypically feminine name, much like the creative users in Bjorkman’s study were able to co-index *they* with *the professor*, even when the gender of the professor was known. In the following section, we review a framework proposed by Ackerman (2019) on how gender is cognitively and linguistically represented. Using this framework, we propose that NBIs and individuals who interact with *NBIs* are most likely candidates for super-creative users (if such users exist).

1.2. Cognitive and linguistic encoding of gender

Ackerman proposes a three-tier model of how speakers cognitively and linguistically encode gender. In the *exemplar tier*, speakers cluster the individuals they encounter and interact with based on the individuals’ perceived gender expression — i.e., all the individuals that the speaker perceives as being feminine cluster together on one end of a spectrum, and all the individuals that the speaker perceives as being masculine cluster together on another end of the spectrum. While occasionally an individual might have a gender expression that causes the speaker to place the the individual in the middle of the spectrum, Ackerman postulates that the distribution of individuals for most speakers is bimodal along the ends of the spectrum.

In the *category tier*, speakers form semantic categories of *masculine* and *feminine*

based on features shared by clusters in the exemplar tier. For example, if many individuals with a masculine gender expression shared a certain feature (e.g., being interested in cars), then that feature would be a part of the semantic category of *masculine*. If a certain feature was shared both by individuals with a masculine gender expression and by individuals with a feminine gender expression (e.g., liking chocolate), then that feature would be a part of both semantic categories (or a part of neither).

In the *feature tier*, speakers assign the grammatical feature $[MASC]$ to lexical items associated with the masculine semantic category (e.g., *car*) and the grammatical feature $[FEM]$ to lexical items associated with the feminine semantic category (e.g., *pink*). Lexical items associated with both semantic categories or neither (e.g., *chocolate*) are assigned the $\emptyset$ feature.

According to this framework, since a name like *Mary* is shared by many individuals with a feminine gender expression, it would be a part of the feminine semantic category and would therefore be marked with the $[FEM]$ feature. Consequently individuals would find it easier to process a sentence in which an antecedent like *Mary* is co-indexed with a gender matching anaphor like (13a) than a sentence in which the antecedent is co-indexed with a gender mismatching anaphor like (13b) or (13c).

- (13) a. $Mary_i$ decided to treat herself_i to sushi.
 b. $Mary_i$ decided to treat himself_i to sushi.
 c. $Mary_i$ decided to treat themselves_i to sushi.

This framework also predicts that if a speaker's exemplar tier were organized such that if the name *Mary* was shared by individuals with different gender expressions, then *Mary* would be marked with $\emptyset$. If we assume that antecedents marked with $\emptyset$ can be co-indexed with anaphors with any gender marking, we would expect an equivalent processing effort for (13a), (13b) and (13c).

There are two (not mutually exclusive) reasons why a speaker's exemplar tier could be organized such that *Mary* was shared by individuals with different gender expressions. First, a speaker encounters roughly equal number of individuals with masculine gender expression and individuals with a feminine gender expression who are named *Mary*. This is unlikely to happen because currently the number of individuals with

a feminine gender expression named *Mary* largely outnumbers the number of individuals with a masculine gender expression named *Mary*.² Second, the speaker might have non-standard ways of perceiving gender and therefore might not perceive typically feminine presenting individuals (including individuals named *Mary*) as having a feminine gender expression. Such a speaker’s exemplar tier is likely to have a different distribution than most other speakers’ exemplar tiers. For example, instead of being bimodal along the ends of the spectrum, the exemplar tier could have a uniform distribution or a unimodal distribution with a peak in the middle of the spectrum.

There are at least two groups of speakers who are likely to have non-standard ways of perceiving gender: NBIs and speakers who frequently interact with NBIs. Standard ways of perceiving gender can often lead to miscategorizing the gender of NBIs. Therefore, in order to avoid the intellectual, social and emotional cost that is associated with the miscategorization, both of these groups of speakers are likely to have actively tried to alter their perception of gender. If either of these groups succeeded in altering their perception of gender such that individuals named *Mary* are not clustered together on end of the gender spectrum, then such a group might be able to mark *Mary* with $\emptyset$, successfully dissociating the gender from *Mary*. In such a scenario, we would expect equivalent processing effort for sentences like (13a), (13b) and (13c). If on the other hand, neither of these groups succeeded in marking *Mary* with $\emptyset$ — i.e., could not dissociate the gender associated with *Mary* — then, we would expect sentences like (13b) and (13c) to be more difficult to process than sentences like (13a). In either of these cases, as was observed in previous studies, we would expect participants to have no difficulty processing sentences like (14), (15) and (16), where *themselves* is co-indexed with singular antecedents that were non-referential and/or have ambiguous gender.

(14) Someone_i decided to treat themselves_i to sushi.

(15) Some woman_i decided to treat themselves_i to sushi.

(16) The participant_i decided to treat themselves_i to sushi.

²Although, in the future it is possible for *Mary* to become a gender neutral name like *Charlie* did.

1.3. Using event-related potentials to study the processing of sentences with gender and/or number mismatch

Event-related potentials (*ERPs*) measure changes in electrical activity at the scalp in response to specific ‘events’ (e.g., words, sounds, images etc). This technique has been widely used in psycholinguistics to study participants’ process of phonology (Praamstra and Stegeman, 1993), morphology (Morris and Holcomb, 2005), syntax (Osterhout and Holcomb, 1992), semantics (Kutas and Federmeier, 2011) and pragmatics (Nieuwland and Van Berkum, 2008).

Relevant to this paper, Osterhout and Mobley (1995) used ERPs to study the processing of sentences in which either the gender marking on the antecedent (i.e., [*Masc*], [*Fem*] or $\emptyset$) mismatched with the gender marking on the anaphor (as in (17a)) or the number marking on the antecedent (i.e. [*Singular*], [*Plural*]) mismatched with the number marking on the anaphor (as in (18a)).

- (17) a. The successful woman_i congratulated himself_i on the promotion.
b. The successful woman_i congratulated herself_i on the promotion.
- (18) a. The hungry guests_i helped himself_i to the food.
b. The hungry guests_i helped themselves_i to the food.

When the gender and/or number marking on the anaphor mismatched with that of the antecedent, they found the so called *P600 effect* — i.e. in the posterior electrodes, the ERPs were more positive around 600 milliseconds after the anaphor was presented for sentences like (17a) and (18a), relative to sentences without a mismatch like (17b) and (18b).

Based on these results, when sampling from the general population, we would expect to find a P600 for sentences like (19a) and (20a), relative to the baselines (19b) and (20b) respectively.

- (19) a. Mary_i decided to treat himself_i to sushi.
b. Mary_i decided to treat herself_i to sushi.
- (20) a. Mary_i decided to treat themselves_i to sushi.

- b. Mary and John_i decided to treat themselves_i to sushi.

However, if we were able to sample from a population of super-creative users who were able to dissociate the gender from names like *Mary*, we would not expect to find a P600 for (19a) or (20a).

Irrespective of whether or not our participants were sampled from a population of super-creative users, we would expect to find a P600 for sentences like (21a) relative to the baseline (21b) because we do not expect even super-creative to be able to dissociate the gender associated with antecedents like *woman* and *man*.

- (21) a. Some woman_i decided to treat herself_i to sushi.
- b. Some woman_i decided to treat himself_i to sushi.

For either populations of participants, we would not expect to find a P600 for sentences like (22a), (23a) and (24b) relative to the baselines (22b), (23b) and (24b) — the antecedents of these sentences are non-referential and/or are ambiguous with respect to gender and Foertsch and Gernsbacher (1997) demonstrated that participants did not find it difficult to process sentences in which such antecedents were co-indexed with *they*.³

Non-referential antecedents

- (22) a. Someone_i decided to treat themselves_i to sushi.
- b. Some people_i decided to treat themselves_i to sushi.
- (23) a. Some woman_i decided to treat themselves_i to sushi.
- b. Some women_i decided to treat themselves_i to sushi.

Antecedents with ambiguous gender

- (24) a. The participant_i decided to treat themselves_i to sushi.
- b. The participants_i decided to treat themselves_i to sushi.

We would also not expect to find a P600 for (25a) and (26a) relative to the baselines (25b) and (26b) — the antecedents in these sentences have ambiguous gender and

³We used *themselves* instead of *they* in our experiment to avoid any ambiguity in co-reference.

Critical sentence	Baseline	Super-creative uses	Normal population
Someone decided to treat herself to sushi	Someone decided to treat himself to sushi	No P600	No P600
Some woman decided to treat herself to sushi	Some woman decided to treat himself to sushi	P600	P600
The participant decided to treat herself to sushi	The participant decided to treat himself to sushi	No P600	No P600
Mary decided to treat herself to sushi	Mary decided to treat himself to sushi	No P600	P600
Someone decided to treat themselves to sushi	Some people decided to treat themselves to sushi	No P600	No P600
Some woman decided to treat themselves to sushi	Some women decided to treat themselves to sushi	No P600	No P600
The participant decided to treat themselves to sushi	The participants decided to treat themselves to sushi	No P600	No P600
Mary decided to treat themselves to sushi	Mary and John decided to treat themselves to sushi	No P600	P600

Table 1. Predicted P600 patterns for sentences in which the anaphor is either *himself* or *herself* (top half) or the anaphor is *themselves* (bottom half).

therefore both *himself* and *herself* would be equally acceptable.

- (25) a. Someone_i decided to treat herself_i to sushi.
b. Someone_i decided to treat himself_i to sushi.
- (26) a. The participant decided to treat herself_i to sushi.
b. The participant decided to treat himself_i to sushi.

In the following sections we describe an ERP study where we test these predictions (summarized in Table 1).

2. Methods

2.1. *Participants*

We surveyed 256 volunteers from Hampshire College and the Pioneer Valley (Massachusetts) about their gender identity, knowledge of gender neutral pronouns and the frequency with which they interacted with individuals with a non-binary gender identity (questionnaire included in the Supplementary Materials). Each of these factors was associated with a numerical score, which were then summed to get a score for the survey takers' degree of familiarity with individuals with non-binary gender identity; the maximum possible score was 13 and the minimum possible score was 3. We recruited 45 participants with a score of 10 or higher to participate in the ERP experiment. All participants were right handed monolingual native speakers of English. There were 23 participants with a non-binary gender identity and the other 22 identified as either male or female (15 male, 7 female). These two groups of participants did not differ from each other on their familiarity scores, linguistic proficiency or linguistic background (details in Supplementary Materials). Data was collected over two sessions of roughly 3 hours each and participants were compensated with \$50 for both of these sessions combined.

2.2. *Materials*

All of our sentences contained one of four types of antecedents: non-referential ambiguous (*someone*), non-referential unambiguous (*every woman*), referential ambiguous (*the stranger*) and referential unambiguous (*John*).⁴ For each antecedent type, we constructed 50 sentence frames, resulting in 200 sentence frames in total. Each frame was filled with four different combinations of antecedent and reflexive pairing: a singular antecedent co-indexed with either *himself*, *herself* or *themselves* and a plural antecedent co-indexed with *themselves*. Therefore in total, participants were presented with $50 \times 4 \times 4 = 800$ critical items. Table 2 contains example sentences for each of the 16 conditions.

⁴We picked 50 frequent male and female names in the US, such that each sentence frame for a given antecedent type contained a unique name.

	SING himself	SING herself	SING themselves	PLU themselves
Non-referential ambiguous	Someone in the group needs to pull himself together.	Someone in the group needs to pull herself together.	Someone in the group needs to pull themselves together.	Some people in the group need to pull themselves together.
Non-referential unambiguous	Every woman must learn to stand up for himself .	Every woman must learn to stand up for herself .	Every woman must learn to stand up for themselves .	All women must learn to stand up for themselves .
Referential ambiguous	The stranger poured himself a cup of coffee.	The stranger poured herself a cup of coffee.	The stranger poured themselves a cup of coffee.	The strangers poured themselves a cup of coffee.
Referential unambiguous	John decided to treat himself to sushi.	John decided to treat herself to sushi.	John decided to treat themselves to sushi.	John and Mary decided to treat themselves to sushi.

Table 2. Example sentences

In addition to the 800 critical sentences, we constructed 1200 filler sentences. 800 of these filler sentences were the stimuli for a related experiment that studied the use of singular *they* when it is embedded in an idiomatic context. The remaining 400 sentences did not contain any pronouns or anaphors. These sentences were repeated twice.

In total, participants were presented with 2400 sentences. We presented these sentences in two sessions with a gap of between one and two weeks between the sessions, since presenting 2400 sentences in one session was not feasible. In each session, participants were presented with 1200 sentences — 400 critical items and 800 fillers (400 idiomatic with pronouns and 400 without pronouns or anaphors). Each session was divided into 10 blocks of 120 sentences each. Sentences were pseudo-randomly assigned to blocks and sessions, ensuring that the four types of antecedents were evenly distributed across blocks and sessions. For each sentence frame, all four versions of the sentence with different antecedent-anaphor pairing occurred in the same session, but never in the same block. The order of sentences within each block was randomized across participants. Each session was between one and two weeks apart.

In order to avoid straining the eyes, which could result in excessive blink related artefacts, all sentences were presented auditorily. The sentences were recorded by a native

speaker of American English over several sessions. The recordings were normalized to have the same volume level across multiple sessions. In order to send the time-locking trigger, critical items were split into two files at the /s/ of the anaphor (the point at which the anaphor becomes uniquely identifiable). The time-locking trigger was sent in a four millisecond gap between the two split files.⁵

2.3. Procedure

2.3.1. Experimental procedure

Participants sat at a comfortable distance from the computer screen and listened to the sentences through headphones. In order to make sure they were paying attention, they answered comprehension questions for a tenth of the sentences. The correct responses were equally distributed between ‘yes’ and ‘no’. The buttons corresponding to ‘yes’ and ‘no’ were switched across the sessions, with the order counterbalanced. Participants were asked to keep their eyes focused on the spot underneath a triangular sticker, that was placed right above the center of the screen. All text in the experiment appeared at the center of the screen. The background of the screen was black and the font colour for text was white. The background was chosen to be black to minimize straining participants’ eyes.

At the end of every sentence, there was a 500 millisecond gap. There was an additional 1000 millisecond gap if the anaphor occurred sentence finally. On trials with comprehension questions, the question was presented for 1500 milliseconds after this gap, and participants were given 1000 milliseconds to respond. Participants were asked to try not to blink during trials, however, after every third trial, there was a blink break of 500 milliseconds, where the characters ‘(- -)’ were presented. This was followed by a blank screen for another 500 milliseconds. Participants were given an opportunity to take a self-timed break after every block. The sequence of events for trials is schematized in Figure 1.

⁵All the stimuli and the split audio files can be found on OSF: <https://osf.io/2vjyp/>

Figure 1. The sequence of events for trials in the experiment.

2.3.2. ERP data acquisition procedure

EEG activity was measured from 30 scalp locations using tin electrodes attached to an elastic cap. The montage is displayed in Figure 2. Apart from these 30 scalp electrodes, we used an additional five electrodes: two on the left and right mastoids, two on the outer canthus of the left and right eye and one underneath the left eye. Electrical activity was re-referenced to the right mastoid during recording and referenced to the average of both mastoids during preprocessing. Horizontal eye movements were monitored using the eye electrodes on the outer canthi of the eyes and vertical eye movements (and blinks) were monitored using the electrode underneath the left eye.

The data were collected at a 1000 Hz and downsampled to 200 Hz during pre-processing. The data were filtered online for frequencies between 0.1 to 100 Hz and were also filtered offline using a bandpass filter of 0.1 to 30Hz. A prestimulus baseline correction (200 ms before the stimulus) was run on the epoched data.⁶

⁶The scripts for pre-processing the data along with the scripts for all statistical analyses can be found in the OSF repository.

Figure 2. The montage of electrodes. Channels used in the data analyses are highlighted in blue.

3. Results

3.1. Data exclusion

We excluded trials with artifacts like eye blinks and excessive muscle movements that occurred before 800 ms after the stimulus onset. Since data from eye-electrodes was very noisy for most participants, we used the FP2 electrode to detect eye blinks. We excluded any participant with greater than 35% trials rejected (both sessions combined), resulting in the exclusion of three participants. All the artifact-free-trials were averaged together to form the ERP sets for each participant. This contained activity from 200 ms before the pronoun was presented until 1000 ms after. Due to experimenter error, several sentences were repeated - not more than 5 per condition. The second occurrence of these sentences were also excluded when calculating the average ERPs. Further details about the number of trials excluded per condition can be found in the Supplementary Materials.

Figure 3. P600 effect for sentences in which antecedents were co-indexed with *himself* or *herself* in electrode PZ. The box highlights the time-window for which we analyzed our data. A similar plot for all the other electrodes can be found in the Supplementary Materials.

3.2. The P600 effect when antecedents are co-indexed with *himself* or *herself*

For each of our eight conditions in which antecedents were co-indexed with *himself* or *herself* (Referentiality $\times$ Ambiguity $\times$ anaphor type; $2 \times 2 \times 2$) we calculated the mean ERP amplitude between 500–800 ms after stimulus onset for each posterior electrode by averaging across all the trials in that condition. For antecedents with unambiguous gender (*some woman*, *Mary*), we calculated the P600 effect by subtracting the ERP amplitude averaged across all trials with a congruent anaphor (e.g., *Mary* co-indexed with *herself*) from the ERP amplitude averaged across all trials with an incongruent anaphor (e.g., *Mary* co-indexed with *himself*). For antecedents with ambiguous gender (*someone*, *the participant*), we calculated the P600 effect by subtracting the ERP amplitude averaged across all trials with the anaphor co-indexed with *himself* from the ERP amplitude averaged across all trials with the anaphor co-indexed with *herself*.

Based on results from Foertsch and Gernsbacher (1997), we predicted that the P600 effect for antecedents with unambiguous gender would be greater than the P600 effect for antecedents with ambiguous gender. Additionally, if participants marked traditionally gendered names like *Mary* with $\emptyset$, then we predicted that the P600 for such antecedents would be smaller than the P600 for non-referential antecedents with unambiguous gender like *some woman*. It is also possible that only participants with non-binary gender identity would be able to mark names like *Mary* with $\emptyset$, in which case we would expect a difference in P600 between referential and non-referential antecedents with unambiguous gender only for individuals with non-binary gender identity, but not for other participants.

To test these predictions, we ran a repeated measures ANOVA looking at the interaction of referentiality, ambiguity and gender-identity.⁷ We found that as predicted, the P600 effect for antecedents with unambiguous gender was greater than the P600 effect for antecedents with ambiguous gender ($\mu_U = 1.08$, $\mu_A = 0.15$; $F(1, 36) = 4.27$, $p < 0.05$). We also found that the P600 effect for non-referential antecedents was greater than the P600 effect for referential antecedents ($\mu_R = 1.77$, $\mu_{NR} = -0.54$; $F(1, 36) = 25.63$, $p < 0.001$). Crucially we found a significant interaction between

⁷The code and data for all the analyses in the reported in this paper can be found on the OSF repository

referentiality and ambiguity ($F(1, 36) = 14.05$, $p < 0.001$). Planned paired t-tests indicated that the difference in the P600 effect between antecedents with ambiguous and unambiguous gender for non-referential antecedents ($t(759) = -10.85$; $p \ll 0.01$) was greater than the difference in the P600 effect between antecedents with ambiguous and unambiguous gender for referential antecedents ($t(759) = 2.83$; $p \ll 0.01$). These effects can be seen in Figure 3. Since we did not find any significant effect of gender identity, we collapsed across all participants in Figure 3.⁸

These results are consistent with the hypothesis that our participants were sampled from a population of super-creative users who marked antecedents like *Mary* with $\emptyset$. Since we hypothesized that such super-creative users would be able to process sentences in which antecedents like *Mary* were co-indexed with *themselves* without any difficulty, in the following analysis we did not expect to find a P600 effect when antecedents like *Mary* were co-indexed with *themselves*.

3.3. The P600 effect for singular antecedents co-indexed with themselves

As in the previous analysis, for each of our eight conditions in which antecedents were co-indexed with *themselves* (Referentiality $\times$ Ambiguity $\times$ Antecedent number; $2 \times 2 \times 2$) we calculated the mean ERP amplitude between 500–800 ms after stimulus onset for each posterior electrode by averaging across all the trials in that condition. We calculated the P600 effect by subtracting the ERP amplitude averaged across all trials with plural antecedents from the ERP amplitude averaged across all trials with singular antecedents. We then ran a repeated measures ANOVA looking at the interaction of referentiality, ambiguity and gender-identity.

We found that the P600 for referential antecedents was greater than the P600 effect for non-referential antecedents ($\mu_R = 0.49$, $\mu_{NR} = 0.00$; $F(1, 36) = 5.23$, $p < 0.05$). Crucially, there was a significant interaction between referentiality and ambiguity ($F(1, 36) = 4.76$, $p < 0.05$). Planned paired t-tests indicated that the P600 effect for referential antecedents with unambiguous gender (e.g., *Mary*) was greater than

⁸Unexpectedly, we also found that the P600 effect for non-referential antecedents with ambiguous gender was greater than the P600 effect for referential antecedents with ambiguous gender. In other words, participants found it more difficult to process *herself* when co-indexed with antecedents like *someone* than when co-indexed with antecedents like *the participant*. One possible explanation for this result is that generics are more strongly associated with a default masculine than common nouns.

Figure 4. P600 effect for sentences in which singular and plural forms of the antecedents were co-indexed with *themselves* in electrode PZ. The box highlights the time-window for which we analyzed our data. A similar plot for all the other electrodes can be found in the Supplementary Materials.

Figure 5. Scalp maps for sentences in which antecedents were co-indexed with *himself* or *herself*. For sentences with antecedents with unambiguous gender, the difference waves were calculated by subtracting the ERP waveforms for sentences in which the antecedent was co-indexed with a congruent anaphor from sentences in which the antecedent was co-indexed with the incongruent anaphor. For sentences with antecedents with ambiguous gender, the difference waves were calculated by subtracting the ERP waveforms for sentences in which the antecedent was co-indexed with *himself* from sentences in which the antecedent was co-indexed with the *herself*

the P600 effect for referential antecedents with ambiguous gender ($t(759) = 6.75$, $p \ll 0.01$). It was also greater than the P600 effect for both non-referential antecedents with ambiguous gender ($t(759) = 5.99$, $p \ll 0.01$) and unambiguous gender ($t(759) = 9.73$, $p \ll 0.01$). These effects can be seen in Figure 4. As in the earlier analysis, since we did not find an effect of gender identity, we collapsed across all participants in Figure 4.

The results from both these analyses taken together are inconsistent with our hypothesis that super-creative users can process sentences in which antecedents like *Mary* are co-indexed with *themselves* without any processing difficulty. The first analysis suggests that we sampled participants from a population of super-creative users who marked antecedents like *Mary* with $\emptyset$. Yet, results of the second analysis suggests that these participants find it difficult to process sentences in which antecedents like *Mary* are co-indexed with *themselves*.

3.4. *Were our participants really sampled from a population of super-creative users?*

Visual examination of data indicated that there was a strong anterior negativity from 300-800 ms after stimulus onset for sentences in which referential antecedents with unambiguous gender were co-indexed with an anaphor with mismatching gender (see Figure 5). The latency and topographic distribution of this effect resemble an N400 effect, albeit one that lasts longer and is more anterior than the typical N400 (Kutas and Federmeier, 2011). If this negativity, like the N400, is a measure of surprise (Frank et al., 2013; Kutas and Hillyard, 1984), then we can infer that participants were more surprised when they encountered sentences in which antecedents like *Mary* co-indexed with *himself* than when they encountered sentences in which antecedents like *Mary* were co-indexed with *herself*. This suggests that participants associated these antecedents with the traditional gender and did not mark them with $\emptyset$. In the following section, we speculate about why referential antecedents with unambiguous gender (like *Mary*) elicit an anterior negativity when co-indexed with an incongruent anaphor, whereas non-referential antecedents with unambiguous gender (like *some woman*) elicit a posterior positivity.

4. Discussion

Previous studies have established that while participants have no difficulty in processing sentences in which *they* is co-indexed with singular non-referential antecedents and/ or singular antecedents with ambiguous gender, they do have difficulty processing sentences in which *they* is co-indexed with singular referential antecedents with unambiguous gender (Balhorn, 2004; Bjorkman, 2017; Foertsch and Gernsbacher, 1997). The goal of this study was to test whether this pattern of data would hold for participants who frequently interacted with individuals with a non-binary gender identity and/or identified as having a non-binary gender themselves.

Drawing on arguments from Bjorkman (2017), we hypothesized that participants who were able to dissociate the gender associated with names like *Mary* (*super-creative users*), would also be able to process sentences in which antecedents like *Mary* were

co-indexed with *themselves* without difficulty. Drawing on a framework proposed by Ackerman (2019), we hypothesized that individuals with a non-binary gender identity (NBIs) or individuals who frequently interacted with NBIs were most likely to behave like super-creative users. We recruited such participants and tested this prediction in an ERP experiment, using the P600 effect as a measure of processing difficulty.

Although participants did not exhibit a P600 when listening to sentences in which antecedents like *Mary* were co-indexed with anaphors mismatching in gender, they displayed an anterior negativity resembling the N400. Additionally, consistent with previous behavioural results, we found that participants exhibited a P600 when listening to sentences in which *themselves* was co-indexed with referential antecedents with unambiguous gender like *Mary*, but not when co-indexed with singular antecedents that were non-referential and/or had an ambiguous gender. Taken together, these results suggest that participants were unable to dissociate the gender associated with a name like *Mary*. In other words, it was not sufficient for participants to have a non-binary gender identity and/or frequently interact with NBIs in order for them to behave like super-creative users.

4.1. Why was there an anterior negativity and not P600 when antecedents like Mary were co-indexed with anaphors mismatching in gender?

For antecedents like *some woman*, a mismatch in the gender of the anaphor resulted in a P600. For antecedents like *Mary*, on the hand, such a mismatch resulted in an anterior negativity resembling an N400. One possible explanation for this difference is that people tend to conceptualize or visualize antecedents when they are referential, but not when they are non-referential. This conceptualization or visualization could make the semantic features more accessible for antecedents like *Mary* than for antecedents like *some woman*, resulting in a more semantic-like mismatch when co-indexed with an anaphor mismatching in gender. Existing empirical evidence suggests that this explanation is unlikely to be true — sentences in which referential feminine antecedents like *actress* and *beautician* were co-indexed with *himself* resulted in a P600 and not an N400 (Osterhout et al., 1997).

Another possible explanation is that the process of co-indexing proper names with anaphors might differ from the process of co-indexing common nouns, because the manner in which gender is marked on proper names and common nouns can differ. Van Langendonck (2008) argues that proper names do not have an “asserted (definitional) meaning” — i.e., any properties associated with the name are presuppositional but not necessary, and are therefore not explicitly marked on the name. For example, while one might presuppose that a person bearing the name *Mary* is female, there is no requirement for this to be true. Therefore, gender (and other properties) are not explicitly marked on proper names. Common names on the other hand do have an “asserted definitional meaning”. For example, it is required for someone who is referred to as a “woman” to be female.⁹ Since gender is explicitly marked on common nouns but not on proper names, sentences in which common nouns like *woman* are co-indexed with anaphors that mismatch in gender result in a P600, whereas, sentences in which proper names are co-indexed with anaphors that mismatch in gender results in an anterior negativity resembling the N400.

4.2. *Future work*

While our results suggest that our participants were sampled from the normal population, and not a population of super-creative users, they do not confirm this for two reasons. First, since the observed anterior negativity does not have the topography and latency of a typical N400, we cannot be certain that our participants were surprised by sentences in which antecedents like *Mary* were co-indexed with a mismatching anaphor. Second, it is possible that while our participants were unable to dissociate the gender from names, they would be able to dissociate the gender from other gendered nouns like *nurse* or *mother*.

Future studies can mitigate both of these concerns by replacing antecedents like *Mary* with common gendered nouns like *the nurse* or *my mother*. If NBIs and individuals who interact with NBIs do not behave like super-creative users, as this paper suggests, then we expect to find a P600 effect when such antecedents are co-indexed

⁹Though there can be other properties associated with “woman” that are only presuppositional and not definitional (e.g., long hair).

with *themselves* or *himself*. If on the other hand, NBIs or individuals who interact with NBIs do behave like super-creative users, then we would not expect to find either a P600 effect or the observed anterior negativity when such antecedents were co-indexed with *themselves* or *himself*. Comparing NBIs and individuals who interact with NBIs with participants who do not have any practice using gender neutral pronouns can also be fruitful.

Future studies can also look at whether participants can be taught to behave like super-creative users for antecedents with varying degrees of gender rigidity (e.g., *teacher*, *nurse*, *mother*, *Mary* and *woman*). Such studies can give us insight into the malleability of gender associations (or the lack thereof) for different types of antecedents, and whether this malleability depends on factors such as an individual's willingness to learn.

5. Conclusion

There has been an increase in the use of *they* as a singular pronoun when referring to individuals with a non-binary gender identity. The goal of this paper was to investigate whether the change in the usage of *they* has also resulted in a change in the processing of this pronoun — i.e., can at least a subset of the population process sentences in which *they* is co-indexed with referential antecedents with unambiguous gender (like *Mary*) without processing difficulty? Based on prior empirical and theoretical work, we hypothesized that participants who were able to dissociate the gender from commonly gendered names like *Mary*, would also be able to process sentences in which *Mary* was co-indexed with *themselves* without additional processing difficulty. We further hypothesized that individuals with a non-binary gender identity, or individuals to interacted with people with a non-binary gender identity were most likely to dissociate the gender associated with names like *Mary*. We recruited such participants in an ERP experiment and found that these participants were not able to dissociate the gender from antecedents *Mary* and incurred a processing cost when such antecedents were co-indexed with *themselves*, suggesting that the change in the usage of singular *they* has not resulted in the way in which this pronoun is processed.

Acknowledgement(s)

We would like to thank Toby Goodman for recording the stimuli and Kirsten Lydic and Anastasia Karavan for assistance with data collection. We would also like to thank Lauren Ackerman, Mara Breen, Mark Feinstein, Neil Stillings, Daniel Altshuler, Guili Dussias, Kyle Rawlins, Brian Dillon and the Experimental Linguistics Group at UMass for helpful discussions and valuable feedback on earlier versions of this work. Finally we would like to thank the Hampshire College community for all the insightful conversations that planted the seed for this project.

Funding

This study was funded by school of Cognitive Science at Hampshire College and The Foundation for Psychocultural Research Hampshire College Program in Culture, Brain, and Development.

Data availability statement

The grand average of the ERP amplitudes for all participants used to generate the plots, the pre-processed data used in the statistical analyses and the scripts used for pre-processing and statistical analyses can be found on OSF: <https://osf.io/2vjyp/>

References

- Ackerman, L. M. (2019). Syntactic and cognitive issues in investigating gendered coreference. PsyArXiv:doi.org/10.31234/osf.io/6cf57.
- Balhorn, M. (2004). The rise of epicene they. *Journal of English Linguistics*, 32(2):79–104.
- Bjorkman, B. M. (2017). Singular they and the syntactic representation of gender in english. *Glossa: A journal of general linguistics*, 2(1):1–13.
- Camilliere, S., Izes, A., Leventhal, O., and Grodner, D. (2019). Multiple grammars for singular they. In *CUNY Conference on Human Sentence Processing*.
- Foertsch, J. and Gernsbacher, M. A. (1997). In search of gender neutrality: Is singular “they” a cognitively efficient substitute for generic “he”? *Psychological Science*, 8(2):106–111.

- Frank, S. L., Otten, L. J., Galli, G., and Vigliocco, G. (2013). Word surprisal predicts n400 amplitude during reading. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 878–883, Sofia, Bulgaria. Association for Computational Linguistics.
- Kutas, M. and Federmeier, K. D. (2011). Thirty years and counting: finding meaning in the n400 component of the event-related brain potential (erp). *Annual review of psychology*, 62:621–647.
- Kutas, M. and Hillyard, S. A. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307(5947):161–163.
- Morris, J. and Holcomb, P. J. (2005). Event-related potentials to violations of inflectional verb morphology in english. *Cognitive Brain Research*, 25(3):963–981.
- Nieuwland, M. S. and Van Berkum, J. J. (2008). The interplay between semantic and referential aspects of anaphoric noun phrase resolution: Evidence from erps. *Brain and Language*, 106(2):119–131.
- Osterhout, L., Bersick, M., and McLaughlin, J. (1997). Brain potentials reflect violations of gender stereotypes. *Memory & Cognition*, 25(3):273–285.
- Osterhout, L. and Holcomb, P. J. (1992). Event-related brain potentials elicited by syntactic anomaly. *Journal of memory and language*, 31(6):785–806.
- Osterhout, L. and Mobley, L. A. (1995). Event-related brain potentials elicited by failure to agree. *Journal of Memory and language*, 34(6):739.
- Praamstra, P. and Stegeman, D. F. (1993). Phonological effects on the auditory n400 event-related brain potential. *Cognitive Brain Research*, 1(2):73–86.
- Van Langendonck, W. (2008). *Theory and typology of proper names*, volume 168. Walter de Gruyter.

Appendix A. Screening survey

[IN THE FOLLOWING PAGE]

Survey for an ERP study

Hi! Thanks for taking the time out to read this.

I am a fourth year student at Hampshire College and I am looking for participants for a psycholinguistic study that I want to run. This will be an ERP study - which means that it will involve getting a cap with electrodes on your head. It is non-invasive, safe and not painful. If you are interested in participating, you will need to come in to the ERP lab at Hampshire for two sessions (each around 3 to 3.5 hours each) where you will be listening to sentences and occasionally answering questions. You will be paid \$50 for your participation.

If this sounds like something you might want to do, please take five minutes to fill out this survey - which will determine if you are eligible to participate in the study. Please be as honest as you can and don't spend too much time thinking about the questions. If you are eligible, I will get in touch with you with further details about the study and scheduling sessions if you are still interested. Filling this form does not obligate you to participate in the study! You can change your mind later :)

If you have any questions feel free to email me at gsp13@hampshire.edu or grusha.prasad@gmail.com

Continue »

25% completed

Some sentences

In this page there are 20 sentences. For each sentence, please indicate if it sounds 'weird' to you. Don't think too much. Just go with your gut instinct

The boys proudly presented the art to their mother *

- ☐ Sounds weird
- ☐ Does not sound weird

The lions gets ready to attack the deer *

- ☐ Sounds weird
- ☐ Does not sound weird

Mary found himself in trouble again *

- ☐ Sounds weird
- ☐ Does not sound weird

The students were very excited *

- ☐ Sounds weird
- ☐ Does not sound weird

The children plays outside everyday *

- ☐ Sounds weird
- ☐ Does not sound weird

The boy was upset with herself *

- ☐ Sounds weird
- ☐ Does not sound weird

I would like an ice cream *

- ☐ Sounds weird
- ☐ Does not sound weird

The flies were very annoying *

- ☐ Sounds weird
- ☐ Does not sound weird

John made himself a sandwich for lunch *

- ☐ Sounds weird
- ☐ Does not sound weird

That butterfly chirp beautifully *

- ☐ Sounds weird
- ☐ Does not sound weird

Michael had way too many things to do that day *

- ☐ Sounds weird
- ☐ Does not sound weird

We is very excited for the road trip

- ☐ Sounds weird
- ☐ Does not sound weird

Pat loved themselves *

- ☐ Sounds weird
- ☐ Does not sound weird

Jane's obsession with cartoons is getting a little out of hand *

- ☐ Sounds weird
- ☐ Does not sound weird

Could you give me three breads please? *

- ☐ Sounds weird
- ☐ Does not sound weird

I am craving a slice of chocolate cake right now *

- ☐ Sounds weird
- ☐ Does not sound weird

This dish has too few salt *

- ☐ Sounds weird
- ☐ Does not sound weird

There was too many homework that day *

- ☐ Sounds weird
- ☐ Does not sound weird

Cory gave themselves a hard time for doing poorly on the test *

- ☐ Sounds weird
- ☐ Does not sound weird

My shoelaces keep getting undone *

- ☐ Sounds weird
- ☐ Does not sound weird

Other information

Did you learn any languages other than English before the age of 5? If yes, please list them *

What is your dominant hand? *

- ☐ Left
- ☐ Right
- ☐ Ambidextrous

Have you had any major head injuries? If yes, please describe *

Are you diagnosed with any learning disabilities? Please indicate which one(s) *

Are you on any psychotropic medication? If yes, please indicate which one(s) *

What are your preferred gender pronouns? *

Preferred gender pronoun is the pronoun that people choose to use for themselves. Please check everything that is relevant

- ☐ He/ Him/ His
- ☐ She/ Her/ Hers
- ☐ They/ Them/ Theirs
- ☐ Ze/ Hir
- ☐ Other:

Do you know what gender neutral pronouns are? If yes, when was the first time you heard about them?
What are the ones you are familiar with? *

On a scale of 1 to 5 how frequently do you use gender neutral pronouns in day to day conversations *

1 2 3 4 5

☐ ☐ ☐ ☐ ☐

How many people in your friend circle use gender neutral pronouns? *

- ☐ Almost none
- ☐ A few
- ☐ A lot
- ☐ Almost everyone

How many of the people you interact with on a weekly basis use gender neutral pronouns? *

- ☐ Almost none
- ☐ A few
- ☐ A lot
- ☐ Almost everyone

Appendix B. Language background

	Male/female	Non-binary
% Exposure to other languages (Mean; Median)	7.75%; 5%	6.56 %; 5%
Number of people speaking a second language	15	11
Age of acquisition of second language: years (Mean; Median)	11.73; 13	10.45; 11
Proficiency of second language on a scale of 1–10 (Mean; Median)	4.47; 4	3.16; 2
Number of people speaking a third language	4	6
Age of acquisition of third language: years (Mean; Median)	19; 18.5	14.67; 15.5
Proficiency of third language on a scale of 1–10 (Mean; Median)	2.75; 2.5	2.5; 2
Mean score on LEXTALE test	92.41	94.54
Mean score on MELICET practice test	36.20	36.75
Mean NBI friend score (scale 1–4)	2.75	2.94
Mean NBI interaction score (scale 1–4)	3.05	2.94
Mean singular <i>they</i> use score (scale 1–5)	4.65	4.56

Table B1. Summary of language background

Appendix C. Additional figures

Figure C1. Scalp maps for sentences in which singular and plural antecedents were co-indexed with *themselves*. Difference waves were calculated by subtracting the ERP waveforms for sentences with singular antecedents from the ERP waveforms for sentences with plural antecedents.

Figure C2. P600 effect for sentences where antecedents were co-indexed with *himself* or *herself*. All posterior electrodes used in the statistical analyses are displayed here. Positive is plotted up.

Figure C3. P600 effect for sentences where singular and plural antecedents were co-indexed with *themselves*. All posterior electrodes used in the statistical analyses are displayed here. Positive is plotted up.

Appendix D. Information about excluded trials

Antecedent	Number	Anaphor	Number excluded
Generic Ambiguous (e.g. <i>someone</i>)	Singular	himself	1
	Singular	herself	3
	Singular	themselves	5
	Plural	themselves	2
Generic Unambiguous (e.g., <i>some woman</i> , <i>some man</i>)	Singular	himself (congruent)	2
	Singular	himself (incongruent)	2
	Singular	herself (congruent)	3
	Singular	herself (incongruent)	3
	Singular	themselves	3
	Plural	themselves	0
Specific Ambiguous (e.g. <i>the participant</i>)	Singular	himself	4
	Singular	herself	4
	Singular	themselves	2
	Plural	themselves	4
Generic Unambiguous (e.g., <i>Mary</i> , <i>John</i>)	Singular	himself (congruent)	2
	Singular	himself (incongruent)	1
	Singular	herself (congruent)	3
	Singular	herself (incongruent)	1
	Singular	themselves	4
	Plural	themselves	1

Table D1. Number of trials excluded per condition because they were repeated due to experimenter error.

Antecedent	Condition	% excluded binary gender	% excluded non-binary gender
Generic Ambiguous (e.g. <i>someone</i>)	himself	13.48	12.39
	herself	13.53	13.39
	themselves (singular)	12.71	13.47
	themselves (plural)	13.24	14.07
Generic Unambiguous (e.g. <i>some woman</i> , <i>some man</i>)	congruent gender	13.38	13.56
	incongruent gender	11.74	14.66
	themselves (singular)	12.49	11.74
	themselves (plural)	11.43	13.90
Specific Ambiguous (e.g. <i>someone</i>)	himself	13.35	12.29
	herself	12.75	13.30
	themselves (singular)	14.81	14.09
	themselves (plural)	11.89	10.33
Generic Unambiguous (e.g. <i>Mary</i> , <i>John</i>)	congruent gender	15.13	14.63
	incongruent gender	11.76	13.54
	themselves (singular)	12.30	12.92
	themselves (plural)	15.22	14.20

Table D2. Percentage of trials excluded per condition during artifact rejection.