Kaiser & Bénet Post, to appear in Proceedings of WCCFL 41 (2023)

Names and Pronouns with and without Gender Features: A Production Study of Singular *They*

Elsi Kaiser & Claire Benét Post

1. Introduction

The English pronominal system, in particular *they*, has received considerable attention recently. Although singular *they* is sometimes presented in the media as a new phenomenon, *they* has a centuries-long history of occurring with singular antecedents when these are non-specific or quantified (ex.1, e.g. Balhorn 2004). However, as noted by Bjorkman (2017), what is a newer phenomenon is use of *they* "even with an antecedent that is singular, definite, and specific, referring to an individual whose binary gender is known to both speaker and hearer" (p.2). This is illustrated in (2). Although (2a-b) are not acceptable for all English speakers,[1] there exists an innovative variety whose speakers accept this use.

1. *Examples and judgments from Bjorkman 2017*
a. Everyone should know their own phone number.
b. *(Seeing an unidentified distant figure:)* They're waving at us.

2. *Examples and judgments from Bjorkman 2017*
a. % The professor$_i$ said they$_i$ cancelled the exam.
b. % I'll let my cousin$_i$ introduce themselves$_i$.

However, Bjorkman reports that even those who speak the variety in (2) may not accept singular *they* in examples like (3), where the antecedent is realized with a gendered first name or a gender-specific noun. But Konnelly & Cowper (2020) provide evidence for a more innovative variety, whose speakers do accept sentences like (3) (also Conrod 2019, 2022). In this variety, "singular *they* can be used to refer to definite, singular individuals of any gender (binary or non-binary)" (Konnelly & Cowper 2020:4).

3. *Examples from Bjorkman 2017, who marks them as ungrammatical*
a. Janet$_i$ said they$_i$ cancelled the exam.
b. I'll let my sister$_i$/father$_i$/aunt$_i$ introduce themselves$_i$.

Thus, there is ongoing change and individual variation in singular *they* usage. Borrowing Conrod's (2022:231) paraphrasing of Konnelly & Cowper's analysis, at Stage 1 "speakers use and accept singular *they* with quantified, generic, or indefinite antecedents." At Stage 2, they do so "with definite or specific antecedents so long as the antecedents lack lexical gender specification," and at Stage 3, they do so "with any antecedent, regardless of lexical gender specification." This variation has inspired various accounts, in particular Bjorkman (2017), Konnelly & Cowper (2020) and Conrod (2022). Unfortunately we cannot do justice to the nuances of these accounts here given space constraints. Broadly speaking, Konnelly & Cowper's account of Stage 2 puts the burden of capturing speaker differences largely on whether the antecedent has gender features: for Stage 2 speakers who accept examples like those in (2), Konnelly & Cowper argue that some nouns (e.g. *professor, cousin*) lack lexically-specified gender features while other nouns (e.g. *mother, father*, certain proper names) still have contrastive gender features, and the

* Elsi Kaiser, University of Southern California, emkaiser@usc.edu. Claire Benét Post, cpost@usc.edu, University of Southern California. A big thank you to audience members at WCCFL for helpful feedback and comments. We also gratefully acknowledge support from USC's Summer Undergraduate Research Fund (SURF).

[1] Other work (e.g. Conrod 2019, Camilliere et al. 2021, Han & Moulton 2022) provides experimental evidence that antecedent type matters, and that bound-variable *they* (1a) is more acceptable for many than referential *they* (2).

pronoun system still has contrastive gender features. Thus, nouns whose gender is unspecified can be referred to with *they*, but nouns with gender features must be referred to with *she/he*. At Stage 3, gender features become optional modifiers across the board, and are no longer contrastive, even on *he* and *she*. As a consequence, "*they* can be used to refer to any non-inanimate singular individual, whether or not the antecedent is linguistically present, and whether or not the antecedent is a nominal bearing semantic gender features" (Konnelly & Cowper 2020:10).

In contrast, Conrod's account derives the observed patterns from variation in whether the referential pronouns themselves have a gender feature. On their account, all English lexical nouns lack gender features, and thus the explanatory burden lies on the featural properties of pronouns. What is crucial in their account is that Stage 3 speakers (and Stage 2 speakers) who accept examples like (2) and (3), where singular *they* refers to referential antecedents, lack a uGender feature in the pronominal DP complex (specifically, on the referential D head). It is also worth noting that, on Conrod's view, what determines whether a pronoun matches an antecedent is not syntax proper. (The uGender feature, if is present, just needs some value, any value, for the derivation to go through). Instead, Conrod posits that a crucial role is played by another process, which they suggest can be formalized with use-conditioned semantics à la Gutzmann & McCready (2014), that evaluates the sociopragmatic appropriateness of the pronoun for the referent. On this view, pronoun gender features are treated in terms of appropriateness, not truth conditions: an 'ungrammatical' pronoun is sociopragmatically inappropriate, not morphosyntactically ill-formed. This allows Conrod to capture why speakers who accept the examples in (2) and (3) "will occasionally avoid or reject singular *they*"; this happens "based not on grammatical constraints but on social relational knowledge about what pronoun is appropriate to the context" (p.237).

1.1. Aims of this work: Using a production task with first and last names

The present paper does not seek to directly pit the competing theoretical accounts each other. Rather, we seek (i) to complement existing work, which is largely based on acceptability judgments, by reporting an experiment that uses a language production paradigm and (ii) to broaden the empirical basis of inquiry by comparing gender-specific first name antecedents (e.g. *Sophia, Daniel*) to last-name-only antecedents (e.g. *Chapman, Fields, Hughes*), including different types of last names.

Antecedents realized with first and last names are referential and specific, but differ in how 'reliably' their gender can be inferred. Many first names are gender-specific, but (in English) last names are not usually gender-specific. However, referring to a person with only their last name (e.g. *Fields was promoted*) is associated with a male bias, at least in the U.S. context (e.g. McConnell-Ginet 2003, Atir & Ferguson 2018, Kaiser et al. 2022). In a variety of contexts (e.g. politics, academia, sports, casual conversation), last-name-only format is used more often for men than for women (e.g. Atir & Ferguson 2018). However, last-name-only can be used for women as well (e.g. *Curie won the Nobel Prize twice*).

Thus, last names allow us to test use of singular *they* in a probabilistic situation where gender is signaled by a (violable) social convention, but not explicitly encoded on the noun. What assumptions do people make about the referent of a (proper) noun with a 'socially implied' male bias, and (how) does this impact use of singular *they*? This has not been systematically tested in prior work on singular *they*.

In addition to comparing singular *they* usage with first and last names, we explore a more speculative line of research and take initial steps to see whether lexical/semantic properties of last name types modulate the strength of the male bias. Do names with male semantic associations (e.g. *Knight, Carpenter*) or male-derived endings (e.g. *Henderson, Hoffman*) exhibit a stronger male bias than names lacking such clear associations (e.g. *Saunders, Hughes*)? Does this impact use of singular *they*?

In the rest of this section, we describe our motivation for using a production task, review existing work on last-name-only, as well as last-name categorization systems. Section 2 describes the language-production study, and the results are presented in Section 3. Section 4 concludes the paper.

1.2. Methodology: Production task

A potential limitation of prior work on *they* is the use of grammaticality/acceptability/naturalness judgments. While such judgments are widely used, highly informative, and have yielded crucial insights into *they*, it would be helpful if they could be supported by converging evidence from another method. This is because participants could perhaps be rating as acceptable things that they themselves do not produce (see Kaschak & Glenberg 2004 on the distinction between comprehension and production of a

new construction). If someone rates singular *they* in (2) or (3) acceptable, but does not produce singular *they* in such contexts, what would this mean for the status of singular *they* in that speaker's grammar? A related concern that can arise with experiments is satiation: could exposure to multiple occurrences of singular *they* in a study increase the likelihood of people rating them as acceptable? (See e.g. Snyder 2000, Sprouse 2009, Goodall 2011, Chaves & Dery 2014, Do & Kaiser 2017 on syntactic satiation: after sufficient exposure, people perceive some ungrammatical constructions as less ungrammatical.)

Thus, our aim is to complement existing work with a production-based task, where we test whether participants opt to produce singular *they*. We do this in a context where the critical trials do not provide participants with a pronoun to judge; instead, participants are shown a person's name (e.g. a last name such as *Carpenter, Olson* or a gender-specific first name such as *Sophia, Ethan*) and when writing a paragraph about that person, can freely choose what form to produce (*she, he, they,* something else). With this open-ended task, we aim to see if patterns similar to those found in earlier work emerge even when participants are not explicitly presented with singular *they*. This can help address potential concerns about discrepancies in acceptability vs. production, i.e., whether people might be willing to accept innovative sentences that they would not actively produce.

1.3. A closer look: Use of singular 'they' with first- vs. last-name antecedents?

Given that speakers of the less innovative variety of English are less likely to accept singular *they* when the antecedent is specific and definite compared to non-specific or quantified ((1) vs. (2)) and even less likely to accept singular *they* when the antecedent's gender is specified (3), what happens when the antecedent is realized with only a last name – when it is specific and definite but its gender is not rigidly specified? Last names are especially interesting because although there is no requirement for them to be used for male referents, they do exhibit a strong male bias (which could vary from speaker to speaker). In what follows, we consider three possible outcomes for use of singular *they* (vs. *he* and *she*).

At the one extreme, we may find speakers who produce *they* with both last-name and first-name antecedents ('across-the-board *they* users').[2] These innovative speakers would presumably be in what Konnelly & Cowper (2020) call Stage 3 (i.e. using singular *they* with all antecedents, regardless of the antecedent's gender specification), as indicated by their use of *they* with gender-specific first names.

We may also find speakers who never produce *they* with first- *or* last-name antecedents ('non-*they* users'). These speakers are likely to be in what Konnelly & Cowper (2020) call Stage 1, only using *they* with indefinite or quantified antecedents, not with specific/definite antecedents. More speculatively, if it's the case that for these non-*they* users, last-name-only format is so extremely male biased that it reaches the status of being lexically specified as male, these participants could even be in Stage 2, where the grammar allows singular *they* with definite antecedents but not if these are lexically specified for gender. In this situation, the grammar would block *they* for last-name antecedents due to their lexical specification. However, our study does not distinguish between these two situations for non-*they* users.

Between the extremes of non-*they* users and across-the-board *they* users, we may find speakers who produce *they* in one configuration but not in the other, namely 'last-name *they* users': Stage 2 speakers who produce *they* with last-name antecedents (assuming these are not lexically specified for gender) but not gender-specific first-name antecedents (assuming these are lexically specified for gender).

Before continuing, it's worth noting that we can also follow Conrod's approach and treat the linguistic representation of English lexical nouns as lacking an obligatory formal gender feature, and still predict the existence of 'last-name *they* users' – because it is still the case that speakers' world knowledge about referents includes information about gender identity (see also Conrod 2022:227 on the denotations of lexical nouns being gendered). Building on Conrod's discussion of gender-ambiguous names like Taylor, it could be the case that 'last-name *they* users' are essentially guided by a constraint like *Maximize Presupposition!* (Heim 1991): For these speakers, with antecedents realized as a gender-specific first-name, use of *they* is blocked by the availability of a more specific pronoun (*she*/*he*) – due to *she/he* being (for these speakers) a pragmatically appropriate pronoun for the antecedent (or due to featurally matching the antecedent's gender, under an account where lexical nouns have gender features).

On the other hand, with antecedents realized with last-name-only, assuming that such antecedents are not lexically specified for gender, then – for these speakers – use of *they* becomes possible when the

[2] As Conrod (2022) notes, people who can use *they* for gender-specific names do not always do so, e.g. for reasons of sociopragmatic appropriateness. For use of *they* to be informative, it need not occur 100% in a particular context.

more specific option (*she* or *he*) is not available. Unavailability of *she/he* may be due to its use conditions not being pragmatically appropriate for reference to a particular last-name-only referent, or due to its gender feature not matching the antecedent (depending on one's analysis).

If speakers' production of *they* is modulated in principled ways by differences in the antecedent-pronoun relation (whether we conceptualize these in terms of first/last names differing in the presence/absence of gender features or different pronouns diverging in their sociopragmatic appropriateness for first/last name referents), we do *not* expect to find speakers who use singular *they* with last-names without also using it with first-names. In other words, specifically 'first-name *they* users' are predicted not to exist. This is because last-name-only format is less informative regarding gender than gender-specific first names, and thus if the use conditions for *they* are met with a gender-specific first-name antecedent, they should also be met with a last-name antecedent (other things being equal). In featural terms, if a speaker's grammar allows for *they* with gender-specific antecedents (first names), it should also allow it with antecedents whose gender is not lexically specified (last names).

1.4. Comparing different kinds of last names

In addition to testing production of singular *they* with first- vs. last-name antecedents, we also test different types of last names. Hereditary last names (surnames, family names) have a variety of historical origins and have emerged through different routes and at different points in history in different societies (see e.g. Hanks 2003, Hanks & Parkin 2016). As our study mostly used last names with historical roots in Europe (mostly in English, German and Scandinavian),[3] this overview is limited to the European context. However, we emphasize the necessity of using names from more diverse sources in future work.

In Europe, hereditary last names[4] started to emerge in medieval times (Hanks 2003). In England, hereditary last names evolved from nonhereditary bynames describing "some aspect or feature of their bearer, distinguishing him (or her) from other people by reference to occupation, geographical location or origin, relationship to another person, or some physical or behavioural characteristic" (Hanks & Parkin 2016: 216). Typically, English last names fall into a number of core categories, including (i) locative names (including topographic, landscape-related names like *Hill*, *Marsh*, and toponymic names based on pre-existing town/farm/etc names, e.g. *Copplestone, Burford*); (ii) occupation names (including profession-based names such as *Baker, Smith, Potter* and status-based names like *Knight, Squire*), (iii) nicknames (e.g. referring to appearance, clothing or personality, such as *Cape, Fox*) and (iv) patronomic names based on the personal name of an ancestor, often with a patronymic ending such as *-son* (e.g. *Williamson*) (e.g. McKinley 1990, Hanks 2003, Hanks & Parkin 2016, Parkin 2013). In comparison, *-man* is more complex (Hanks 2003): sometimes *-man* denoted a servant, e.g. *Bateman* 'servant of Bartholomew,' but sometimes it is suffixed to an occupation, e.g. *Millman* (Hanks 2003).

We explore whether different types of last names differ in the strength of their male bias as a function of their lexical/semantic properties. For example, does a name like *Knight*, with the same form as the common noun *knight* whose lexical semantics are associated with male referents, have a stronger male bias than a name like *Saunders* which lacks such a clearly male semantics? On the one hand, given the literature on semantic priming with common nouns, one might wonder whether last names that can be interpreted as having components typically linked to male referents exhibit a stronger male bias. On the other hand, in light of the literature arguing for the Millian view that proper names have reference but no meaning (see e.g. Searle 1969, Kripke 1972, Lyons 1977), one might expect that proper names, even if they resemble/are identical to common nouns, would *not* activate semantic representations of this type.

We tested seven different kinds of last names, in addition to female and male first names (see Table 1). We speculated that four of these categories might show an especially strong male bias: (i) last names ending in *-son*, (ii) last names ending in *-man*, (iii) last names that are professions which have historically been stereotypically male, and (iv) last names based on words for male nobles. We also tested (v) names referring to natural features/locations and properties of landscapes as well as (vi) neutral last names that do not have strong semantic links to common nouns. We speculated that the male bias of categories (v-vi) might be weaker than (i-iv). Furthermore, we tested (vii) last names with components that could be

[3] This restriction stems from our aim of testing *-man*/*-son,* and names that present-day participants perceive as linked to professions (e.g. *Farmer, Fisher*) and nature (e.g. *Fields*).

[4] 'Hereditary' is not the same as patronymic. In a patronymic system, a person is named based on their father. E.g., Petter *Lars*son can be the son of Lars *Hans*son. Patronymic names are not hereditary across multiple generations.

construed as having stereotypically female associations (e.g. flower terms, as in *Rosewood*). If participants are sensitive to some or all of these lexical cues, some last names could have a stronger male bias than others – conversely, some last names may have a weaker male bias, and participants might use *she* or singular *they* more when referring to a character with such a last name.

Neutral	Nature	*-man*	*-son*	Male nobility	Male-biased profession	Stereotypically female
Saunders	*Fields*	*Hoffman*	*Atkinson*	*King*	*Fisher*	*Rosewood*

Table 1: The seven last-name types tested, with an example of each

2. Experiment

Participants saw a first/last name and wrote a paragraph about the person, based on five gender-neutral bullet points. Participants' texts were analyzed for what kinds of pronouns, if any, they produced. This preliminary work aims to provide a foundation for future studies using more names and participants.

2.1. Participants, materials and design

Out of 64 US-born native English speakers who participated, 61 were included in data analysis. Two were excluded for failing to provide data for over half of the items, and one for reporting a hearing impairment (which could impact language exposure). Participants (ages 20-67, mean 36.5; recruited via Prolific) did the study on Qualtrics. (The present paper does not report age-based analyses, although they are planned for future work. See Conrod 2019 for discussion of age effects.)

On each trial – the study included 27 trials in total – people saw five gender-neutral pieces of information, presented in bullet point/fragment format, about a person whose name was at the top (Figure 1) and were asked to write a natural-sounding biographical description of each person. The information included topics such as place of birth, where someone went to school, number of siblings and so on. We created multiple versions of the study to ensure variance in which names were paired with which sets of bullet points. The bullet points provided no information about gender; any inferences about gender could only be made based on the name. Participants were asked to write the description in their own words. They could freely choose what pronouns to use, if any. Examples are in (4).

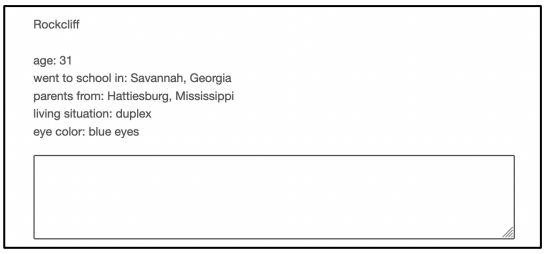

Figure 1. Sample screenshot of a trial. Each item was presented on a separate screen.

4. *Example paragraphs produced by participants*

a. Coded as *she:* Chapman's parents are from Oshkosh, Wisconsin, but she was born in Twin Falls, Idaho and raised with one brother. At the age of 41, she currently resides in Little Rock, Arkansas.

b. Coded as *he*: Fields was born in Orlando, Florida, the only child of parents from Belleville, Illinois. Growing up, he lived in a condo while he attended school in Fort Lauderdale, Florida.

c. Coded as singular *they*: Fields is 67 years old and has brown eyes. They were born in Henderson, Nevada and attended school in Lawrence, Kansas. They currently live in Great Falls, Montana.

d. Coded as *no pronoun*: Blue-eyed Knight studied in Savannah, Georgia, and is 31 years old. Knight's parents come from Hattiesburg, Mississippi. Knight currently lives in a duplex.

We manipulated name type, comparing gender-specific female and male first names (e.g. *Lucy,*

Greg) and subtypes of last names. As shown in Table 1, in addition to first names, we tested seven types of last names. We tested last names with male-associated suffixes (*-man*, *-son*), stereotypically male associations (male nobility and male-biased professions), stereotypically female associations (e.g. linked to flowers), names linked to natural/geographical features and semantically 'meaningless' neutral names.

Participants' texts were analyzed for how they referred to the named referent (with *she*, *he*, *they*, another form, or no use of anaphors). Participants did not mix pronoun types within paragraphs, so it was possibility to analyze each paragraph as belonging to a single category.

A significant caveat is that we only included three names in each of the nine categories. This is because the experiment was already long (median duration 40-50 min) and we wanted to avoid making it even longer. In future work, to assess the generalizability and robustness of our findings, larger sets of names need to be tested. Furthermore, the names selected for the different categories should be normed; our name choices in this preliminary study were based on prior work and our own judgments.

No information was presented about whether the named characters identify as (non)binary. As Konnelly & Cowper (2020) note, *they* is the personal pronoun used by "many non-binary individuals, those whose gender identity is not, or is not exclusively, masculine or feminine" (p.1-2). Because no information about gender identity was provided in the experiment, the results do not speak to use of *they* for binary and non-binary individuals. (See Arnold et al. 2021 for a comprehension study on how information about referents' gender identity can influence interpretation of *they* as singular vs. plural.)

2.2. Procedure

Participants did the study over the internet, using Qualtrics (Provo, UT). On each trial, they saw a name and five bullet points, and were asked to write a natural-sounding biographic paragraph about the person. At the end of the study, to tap into potential differences in gender attitudes, participants rated whether they (dis)agree with eleven statements about men's and women's roles (mostly adapted from Prasad & Baron 1996, Yoon et al. 2015). Analyses of these data are on-going and are not reported here.

3. Results

3.1. Pronouns produced for different name types

Before looking at individual differences, we collapse the data from all participants to see what kinds of pronouns are used for first vs. last names, and different kinds of last names. When people wrote about referents with female and male first names, there is a clear bias to use *she* and *he* respectively (>79%, two leftmost bars in Figure 2). Singular *they* is used but does not exceed 9%.

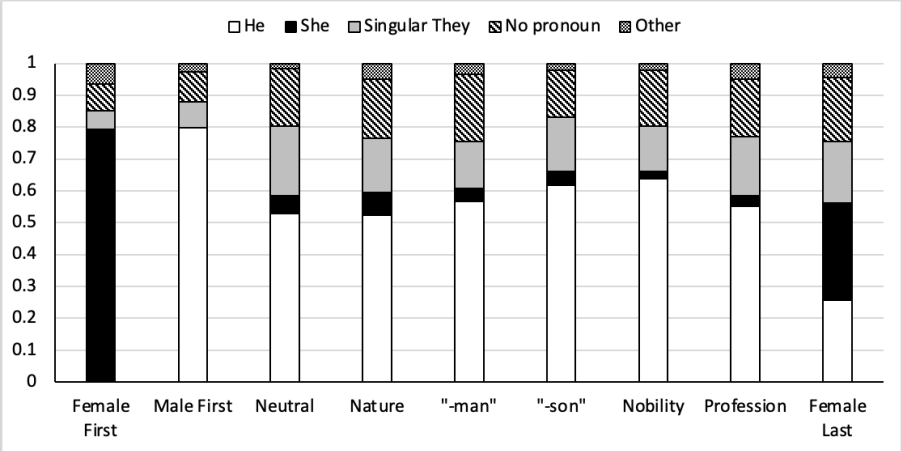

Figure 2. Proportion of trials where people used different pronominal forms ("No pronoun": the paragraph did not contain a pronominal. "Other": participant used another form, e.g. '*he/she*')

The last-name conditions yield different patterns. The third through sixth bars show a clear male bias (big white segments; more uses of *he* than other forms), although the *he* rates in the *neutral, nature, -man, -son, nobility* and *profession* conditions are significantly lower than in the male first-name

condition (z's > |4|, p's < .001).[5] Thus, last names do not trigger as strong a male bias as male first names. Nevertheless, last names in the *neutral, nature, -man, -son, nobility* and *profession* conditions do elicit more male than female pronouns (>50% *he*, <10% *she*), replicating the male bias found in prior work.

Strikingly, in these last-name conditions, singular *they* is produced more frequently (overall, it is used on 17% of trials in these last-name conditions) than with male first names. Singular *they* is only produced on 8% of trials in the male first-name condition.[6] Thus, although these last names have a male bias, people are still relatively more willing to use *they* with last-names than male first names.

What about subtypes of last names? Based on pairwise comparisons (with emmeans and Bonferroni correction), the rates of producing *he* do not differ significantly between the *neutral*, *nature*, *-man*, *-son*, *nobility* and *profession* conditions.[7] These conditions also do not differ in how often singular *they* or *she* are produced. Thus, we fail to find strong evidence for the idea that the endings of names or potential semantic associations involving the corresponding common nouns modulate the strength of the names' male bias. This result is compatible with the view, traditionally attributed to John Stuart Mill, that proper names have reference but no meaning (Section 1.4).

What about last names with stereotypically female associations (rightmost bar)? These elicit singular *they* rates comparable to other last names and, strikingly, rates of *he* (26%) and *she* (31%) that are similar to each other. Thus, we no longer see clear evidence for a strong male bias. Indeed, this condition yields significantly lower rates of *he* than the other last-name conditions (pairwise comparisons using emmeans with Bonferroni correction, z's > |6|, p's < .0001), and higher rates of *she* than the other last-name conditions (z's > |6|, p's < .0001), although rates of singular *they* do not differ reliably from other last-name conditions. While more work is needed to see if this preliminary result replicates and generalizes, this is the first hint that last names may not all be equally male-biased; perhaps gender assumptions triggered by last names can be sensitive to the name components' stereotypical semantic associations. This finding merits further investigation, especially in light of on-going philosophical and semantic debates regarding the nature of proper names.

3.2. Individual differences

We now turn to individual differences in *they* use. Figure 3 shows how often first names elicit singular *they* from each participant, and how often last names elicit *they* and *he* from each participant.

The black bars show the rate of singular *they* production on last-name conditions, in other words, how frequently each person produced singular *they* when presented with a last name. (Here, all different last name types are collapsed/combined.) Absence of a black bar signals that that a participant did not use singular *they* on any last-name trials. As can be seen in the figure, 30 people out of 61 have black bars (of varying heights), indicating they produced singular *they* with last names at least once. As the bar heights indicate, some participants when presented with a last name, produce singular *they* at least 80% of the time (e.g. participants 1,2,4 and 8[8]), whereas others only produce *they* occasionally on last-name trials (e.g. participants 21, 30, 35, 39 and 44 produce *they* on 10% or less of last-name trials).

What about how often participants produce *he* on last-name conditions? This is shown by the white bars, which represent the *he* rate on last-name conditions: In other words, how frequently did each person produce *he* when presented with a last name? (*She* productions are not shown.) Thus, the white bars and the black bars are inversely related. As Figure 3 shows, many participants have white bars: many participants produce *he* on last-name conditions – in line with the overall finding that last-name-only has a male bias. Furthermore, the fact that these white bars are often tall shows that many people who produce *he* on last-name conditions do so frequently.

The bars with diagonal lines show the rate of singular *they* production on *first*-name conditions (how

---

[5] Data was analyzed with R Statistical Software (R Core Team). We fit logistic mixed effects regression models to our data (with proportions of *he* uses, *she* uses and *they* uses as the dependent variables in different analyses as appropriate) and used the emmeans package v.1.8.2 to obtain Bonferroni-corrected pairwise comparisons.

[6] In recent work, Gardner & Brown-Schmid (2019) used a different production task with last names (but did not manipulate name types). They elicited very few uses of singular *they*, for reasons that are not clear.

[7] Although it looks like the *nobility* condition elicits more *he* and less *she* and *they* than the *neutral, nature, -man, -son and profession* conditions (especially *neutral* and *nature*), the differences do not reach significance under Bonferroni. Perhaps differences could be detected with a higher-powered study; we leave this for future work.

[8] Participant ID numbers in Figure 3 are based on reverse rank order by frequency of *he* usage with last names.

frequently each person produced singular *they* when presented with a gendered first name). Indeed, only eight of out 61 participants have bars with diagonal lines (of varying heights). Most participants (53 people) did not produce singular *they* at all on first-name conditions.

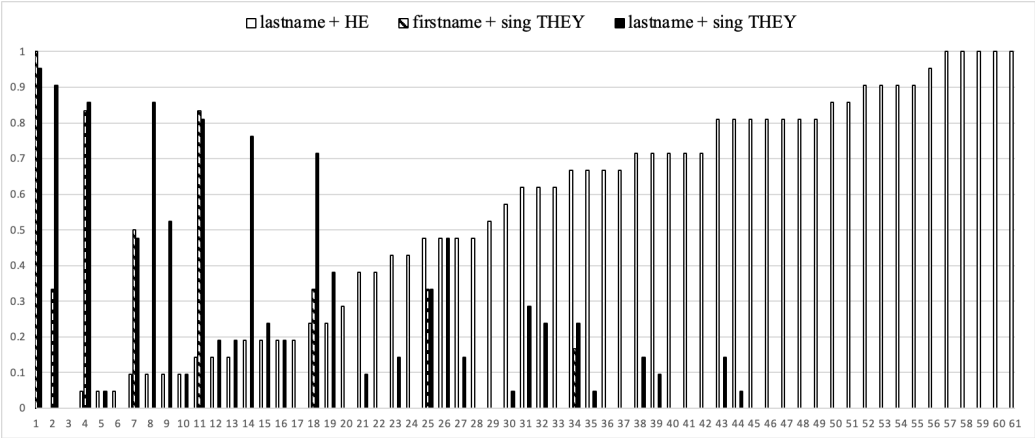

Figure 3. Data for each individual participant (proportion of singular *they* on first- and last-name conditions, proportion of *he* on last-name conditions; *she* not plotted; different last names collapsed)

Crucially, the patterns in Figure 3 show that participants fall into the three groups that we predicted. First, 31 people are 'non-*they* users' and do not use singular *they* at all, in any condition (participants with only white bars, indicating use of *he*). Second, 22 people are 'last-name *they* users' who use singular *they* for last-name but not first-name antecedents (participants with white and black bars.) Thus, these people use singular *they* when the gender of the referent is unclear, even if male-biased. Third, there are eight 'across-the-board *they* users' who use singular *they* both for first-name and last-name referents (participants with three bars). They are the most grammatically innovative group, as they use singular *they* with male-biased last names as well as first names that have clear gender associations.

Moreover, no one uses singular *they* with first names without also using it with last names; there are no specific 'first-name *they* users', as we predicted. This population-level pattern is echoed within-participants: the majority of those who use *they* with first names do so at rates comparable to last-name referents (i.e., their black and striped bars are similar heights); no one uses *they* at clearly higher rates for first than last names (no person's striped bar is clearly higher than their black bar).

4. Discussion

This paper reports a preliminary production study using both first-name and last-name antecedents that extends prior judgment-based research on singular *they* and broadens its empirical scope. On the one hand, our results provide strong evidence that last-name-only style (e.g. *Fisher walked in*) has a strong male bias and show that this bias arises with a variety of last name types. Even when the bullet points provided no information about gender, participants tended to assume that last-name-only antecedents are male, as shown by the predominance of using *he* relative to *she*.

On the other hand, participants' pronoun production patterns also show that forms which are merely male-biased, like *Atkinson* and *Knight*, pattern differently from forms that are (more rigidly) gender-specific, like *Ethan* and *Greg*. Thus, although last-name-only triggers strong assumptions about gender, use of the gender-neutral form *they* reveals the limits of these biases: many participants use singular *they* with male-biased (last) names but not with male-specific (first) names. We suggest that these individual differences largely fit with the stages proposed in prior work, but more in-depth analysis is needed. In future work, we also plan to look at whether people's attitudes about gender modulate their pronoun production choices with last names.

We also find interesting hints, whose validity should be assessed in future studies using a larger number of names, that the male bias may not be equally strong for all last names and may be modulated by the semantic associations of name components – findings which are potentially unexpected in light of Millian analyses treating proper names as having only reference and lacking meaning.

References

Arnold, Jennifer E., Mayo, Heather C. & Dong, Lisa. (2021) My pronouns are they/them: Talking about pronouns changes how pronouns are understood. *Psychonomic Bulletin & Review* 28, 1688-1697.

Atir, Stav, & Ferguson, Melissa J. (2018). How gender determines the way we speak about professionals. *Proceedings of the National Academy of Sciences*, 115(28), 7278-7283.

Balhorn, Mark. (2004). The rise of epicene *they*. *Journal of English Linguistics* 32(2), 79-104.

Bjorkman, Bronwyn M. (2017). Singular 'they' and the syntactic representation of gender in English. *Glossa: a journal of general linguistics*, 2(1).

Camilliere, Sadie, Izes, Amanda, Levanthal, Olivia, & Grodner Daniel. (2021). They is Changing: Pragmatic and Grammatical Factors that License Singular they. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 43, no. 43.

Chaves Rui P. & Dery, Jeruen, E. (2014). Which subject islands will the acceptability of improve with repeated exposure? In Robert E. Santana-LaBarge (Ed.), *Proceedings of the 31st West Coast Conference on Formal Linguistics*. Somerville, MA: Cascadilla Proceedings Project, 96-106.

Conrod, Kirby. (2022). Abolishing Gender on D. *Canadian Journal of Linguistics/Revue canadienne de linguistique*, 67(3), 216-241.

Conrod, Kirby. (2019). *Pronouns raising and emerging*. Doctoral dissertation, University of Washington.

Do, Monica L., & Kaiser, Elsi. (2017). The Relationship between Syntactic Satiation and Syntactic Priming: A First Look. *Frontiers in Psychology* 8, 1851.

Gardner, Bethany & Brown-Schmidt, Sarah. (2019). "Call me by my last name": Gender bias through production about and memory for names. Poster presented at the *60th Annual Meeting of the Psychonomic Society,* November 14-17, 2019, Montréal, Canada.

Goodall, Grant. (2011). Syntactic satiation and the inversion effect in English and Spanish *Wh*-questions. *Syntax* 14, 29-47.

Gutzmann, Daniel & McCready, Elin. (2014). Using descriptions. *Empirical issues in Syntax and Semantics*, 10, 55-72.

Han, Chung-Hye H. & Moulton, Keir. (2022). Processing bound-variable singular 'they.' *Canadian Journal of Linguistics/Revue canadienne de linguistique*, 67(3), 267-301.

Hanks, Patrick. (2003). *Dictionary of American Family Names*. Oxford: Oxford University Press.

Hanks, Patrick & Parkin, Harry. (2016). Family Names. In Carole Hough (Ed.), *The Oxford Handbook of Names and Naming*. Oxford Academic.

Heim, Irene. (1991). Artikel und Definitheit. In Arnim von Stechow and Dieter Wunderlich (Eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung,* 487-534. Berlin: de Gruyter Mouton.

Kaiser, Elsi, Ho, Deborah, Hsu, Haley, Post, Claire & Rouse, Madeline. (2022). Referring to someone using only their last name: Insights from gender-marked pronouns. Talk at *Linguistic Evidence 2022*, Paris (Université de Paris - CNRS, France), October 7, 2022.

Kaschak Michael P. & Glenberg Arthur M. (2004). This construction needs learned. *Journal of Experimental Psychology: General* 133, 450-467.

Konnelly, Lex, & Cowper, Elizabeth. (2020). Gender diversity and morphosyntax: An account of singular they. *Glossa: a journal of general linguistics*, 5(1).

Kripke, Saul. (1972). Naming and necessity. In Donald Davidson & Gilbert Harman (eds.), *Semantics of natural language* (pp. 253-255). Dordrecht: D. Reidel.

Lyons, John. (1977), *Semantics*. Vols. 1 and 2. Cambridge: Cambridge University Press.

McKinley, Richard. (1990). *A History of British Surnames*. London: Longman.

McConnell-Ginet, Sally. (2003). "What's in a name?" Social labeling and gender practices. In Janet Holmes & Miriam Meyerhoff (eds.), *The handbook of language and gender*, 69-97. Blackwell.

Parkin, Harry. (2013). Surname Typology and the Problem of Inconsistent Classification. *Names*, 61:4, 200-211.

Prasad, P., & Baron, J. (1996). *Measurement of gender-role attitudes, beliefs, and principles*. Philadelphia, PA: University of Pennsylvania.

Searle, John R. (1969). *Speech Acts. An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.

Snyder, William. (2000). An experimental investigation of syntactic satiation effects. *Linguistic Inquiry* 31, 575-582.

Sprouse, Jon. (2009). Revisiting Satiation: evidence for an equalization response strategy. *Linguistic Inquiry* 40, 329-341.

Yoon, Eunju, Adams, Kristen, Hogge, Ingrid, Bruner, John P., Surya, Shruti & Bryant, Fred B. (2015). Development and validation of the Patriarchal Beliefs Scale. *Journal of Counseling Psychology* 62, 264.