Interpretations of plurality in personal pronouns differ across person categories: An experimental study

Mora Maldonado[1] and Carmen Saldana[2,3]

[1] Laboratoire de Linguistique de Nantes, CNRS/Nantes Université
mora.maldonado@univ-nantes.fr
[2] Department of Comparative Language Science, University of Zurich
[3] Center for the Interdisciplinary Study of Language Evolution, University of Zurich
carmen.saldanagascon@uzh.ch

1 Introduction

Personal pronouns refer to individuals defined by the communicative context, in which there are at least a speaker and an addressee, and potentially other individuals who play no active role in the conversation. The person category of a given personal pronoun specifies the main referent of the pronoun, (a.k.a. its *focal referent*). For example, singular first person pronouns (1SG) refer to the speaker in the context, singular second person pronouns (2SG) to the addressee and singular third person pronouns (3SG) to someone who is not actively participating in the conversation. When it comes to the plural forms of these pronouns, an interesting observation emerges, since plural pronouns do not seem to be interpreted in exactly the same way across person categories.

First and second plural pronouns (1PL and 2PL), such as English *we* or *you* (or *youse* or *y'all*), are not the plural of their singular alternatives like *dogs* is the plural of *dog*: Unlike these plural nouns, 1PL and 2PL pronouns do not necessarily refer to a *collection* of speakers or addressees (henceforth, *additive meaning*). Rather, they refer to a group of people that *includes* the speaker or addressee (henceforth, *associative meaning*) [4, 13, 5, 7]. For example, the pronoun *we* does not refer to multiple people speaking in unison (except in some rare cases, i.e., *chorus we*), but it is used to refer to groups whereof the speaker is part, together with other people; and the plural *you* might sometimes refer to a group containing multiple addressees (i.e., an audience) and sometimes to a single addressee plus associates. That is, both 1PL and 2PL pronouns seem to be compatible with both additive and associative meanings.

Things are different with plural third person pronouns (3PL). For a plural pronoun to be morphologically third person, it generally needs to be the case that *no* member of the group is actively participating in the conversation. Just like standard nominal plurals[1], 3PL pronouns normally refer to a homogeneous group, a collection of individuals who are neither speakers nor addressees. Thus, 3PL pronouns seem to typically have additive meanings.[2]

Altogether, there seems to be a tight relation between the interpretation of plurality in pronouns and their person category. However, there is hitherto no experimental evidence to

[1] We treat standard nominal plurality as having an additive interpretation, however, associative nominal plurality also exists for some categories of nouns in a rather large subset of the world's languages [4, 13, for an in-depth discussion]. Around half of these languages even have different markers for associative and additive meanings [8]. This is not the case for pronouns: No known language has a specific form to refer to a homogeneous group of speakers or addressees, different from a form denoting the heterogeneous, associative meaning [9, 14, 5].

[2] It has been argued that 3PL pronouns can also have associative readings, referring to a heterogeneous group of one focal non-participant plus associates [7, 18]. Teasing homogeneous and heterogeneous readings apart in this case requires a way of distinguishing between non-participants and associates, which is not entirely obvious in most uses.

confirm it. Here, we provide the first experimental investigation on speakers' interpretation of plurality in personal pronouns. We do this via an online experiment where we ask participants to choose between associative and additive meanings of plural person pronouns.

Before we move onto presenting our experiment, and in order to better frame our experimental hypotheses, we first develop further assumptions in semantic theory regarding the interpretations of plurality in personal pronouns conditioned on their person category.

1.1 The semantics of personal pronouns

Standard semantic accounts of personal pronouns [16, 11] derive these observations by assuming that 1PL and 2PL pronouns refer to a group with one focal referent plus some underspecified, additional associate(s), whereas 3PL refer to an homogeneous group of non-participants. More concretely, these proposals are built on two main assumptions. First, following common assumptions in morphological theory [14, 1], personal pronouns are defined by their reference to just two conversational roles: speaker and addressee. Second, personal pronouns only require the *inclusion* of communicative roles in the person referent: A first person pronoun (either singular or plural) requires the speaker(s) in the context to be *part of* the individual denoted, but not to be *the* individual itself. The second person imposes the same requirement but with respect to the addressee(s). A simplified denotation for each person category is given in (1), following the implementation in [12], where person and number requirements are presuppositional.

(1) a. $[\![1]\!]^c = \lambda x_e : x$ includes the speaker(s) in c. x
b. $[\![2]\!]^c = \lambda x_e : x$ includes the addressee(s) in c. x
c. $[\![3]\!]^c = \lambda x_e : x$ excludes the speaker(s) and the addressee(s) in c. x

(2) a. $[\![\text{SG}]\!]^c = \lambda x_e : |x| = 1.\ x$
b. $[\![\text{PL}]\!]^c = \lambda x_e : |x| > 1.\ x$

These assumptions, together with an intuitive semantics for number marking (as in 2), give the desirable results for first and second person pronouns. In the singular, the only way for an individual to have a cardinality of one and to include the speaker/addressee is *to be* the speaker/addressee. For 1PL and 2PL pronouns, the requirement is weaker: the group denoted by the pronoun should simply include the speaker(s) or addressee(s). Associative interpretations arise when there is a single speaker/addressee in the context, whereas additive meanings, such as the *chorus we*, arise whenever there are multiple speakers/addressees.

For third person pronouns, the only requirement is to *not* refer to speaker and addressee, given that only two roles (speaker and addressee) are assumed—but see [15]. As a result, 3PL are necessary homogeneous with respect to person roles, and end up looking more similar to plural nouns which only have additive interpretations.

By relying on the semantics of person, the account above generates both associative and additive meanings for 1PL and 2PL pronouns. However, this theory does not say why associative interpretations of 1PL and 2PL pronouns are intuitively preferred over additive interpretations, and even in many cases the only available reading (in particular for 1PL pronouns).[3]

[3]It has been alternatively proposed that plurality has a different meaning when it applies to pronouns than when it applies to nouns [6, 18]. Different approaches vary in their implementation, but the general idea is that the semantics of plural marking when applied to pronouns is inherently associative (i.e., it doesn't mean 'more than one' but rather 'plus associates'). These approaches differ substantially from the person-based proposal above. However, the asymmetry in the interpretation *across* plural personal pronouns is left partly unexplained in the same way, and we thus will not discuss these theories here.

To account for the asymmetries between the different interpretations of plurality across person categories, some authors have constrained the possible meanings of 1PL and 2PL pronouns by ruling out the existence of multiple speakers and addressees in the ontology [10]. Since there can only be one speaker/addressee in the context, the semantics in (1) can be kept as it is and additive readings will be fully blocked. A more widespread view is that preference for associative meanings is something that pertains to world knowledge: e.g., the *chorus we* reading is dispreferred because it's extremely infrequent to encounter situations with multiple speakers in the world. Interestingly, leaving interpretative preferences to world knowledge allows to account for an additional intuition about the interpretation of 1PL and 2PL pronouns, namely that the later gives rise to more additive readings than the former [19, 7, for discussion].

In order to understand how functional pressures, such as world knowledge, may interact with morphosemantic knowledge, we need a better estimation of the relation between the interpretation of plurality in pronouns and their person category. This study aims to tackle this question by providing the first quantitative investigation on speakers' interpretation of plurality in personal pronouns. We assess the hypothesis that the likelihood of interpreting a plural pronoun additively depends on the person category in a way that follows the hierarchy: 1PL $<$ 2PL $<$ 3PL. That is, the likelihood of interpreting a 1PL as meaning 'multiple speakers' is lower than the likelihood of interpreting a 2PL as meaning 'multiple addressees', which is in turn lower than the likelihood of interpreting a 3PL as 'multiple non-participants'.

2 Methods

We test the aforementioned hypothesis via two online experiments where we ask Spanish-speaking participants to choose between associative-like and additive-like meanings of Spanish personal pronouns. Experiment 1 (Exp 1) compares the interpretation of 1PL and 2PL pronouns, and Experiment 2 (Exp 2) compares the interpretation of 2PL and 3PL pronouns. Both experiments use an analogous methodology. All experimental materials, data, and analysis reported here are available at osf.io/fbnc7. The preregistered design and analysis plan is accessible at osf.io/sdtb8. The study was approved by the Ethics Committee of the School of Philosophy at the University of Zurich (Authorisation Nr. 22.6.11).

2.1 Participants

118 Spanish-speaking adult participants (59 per experiment) were recruited on Prolific. All participants were located in Spain. They were paid £1 for a five-minute experimental session.

2.2 Materials

Participants were told they would see themselves at a party (depicted as a cartoon character) chatting to acquaintances (also depicted as cartoon characters of different colours). Participants in the experiment were assigned an avatar for themselves (e.g., the white one). In each trial, participants were presented with a GIF depicting the situation at the party, and were provided with a sentence that some of the other characters uttered to them.

The GIFs displays three "focal" characters which remain still: the participant' character (which given the scenario acts as a *focal addressee*), a character standing close to it and directing its gaze towards it (*focal speaker*), and a character further away from these two, apparently not participating in the interaction (*focal other*). Each of these stationary characters was joined by a moving partner character who stood next to them for 80ms and could then be interpreted

as a non-focal, additional speaker, addressee or other. There were two variants of the GIF, differing on the distribution of the coloured characters in the scene. An example of a GIF is available here.

All sentences involved overt subject personal pronouns, for a subset of person-number combinations in Spanish. Exp 1 involved 1PL and 2PL pronouns (*nosotros* vs. *vosotros*), and Exp 2 involved 2PL and 3PL pronouns (*vosotros* vs. *ellos*). Half of the trials involving these sentences were our critical trials, the other half were used as fillers (see Section 2.3). Singular pronouns (*yo*, *tú*, *él*) were uniquely used as filler trials in both experiments. Sentences for first person pronouns were declarative statements (e.g., 'I like Rosalía'), whereas sentences for second and third person pronouns were interrogative statements (e.g., 'Do you like Rosalía?'). The predicates involved in these sentences always denote habitual activities, which in the context of a party could easily refer to characters present or absent in the scene.

2.3 Procedure

The procedure was minimally different between experiments. In Exp 1, after being presented with a GIF-sentence pair, participants had to decide: (1) who uttered the sentence (Q-SPEAKER), (2) who was the sentence addressed to (Q-ADDRESSEE), and (3) what character or group of characters the pronoun in the sentence referred to (Q-REFERENCE). All these questions involved a forced-choice amongst two alternatives, and were presented consecutively.

In Q-SPEAKER and Q-ADDRESSEE, participants had to decide whether there was a single speaker/addressee (focal referent alone) or multiple speakers/addressees by selecting amongst two descriptions. Descriptions indicated the characters by naming their color. For example, given the scenario depicted in Fig. 1, participants would have to indicate whether the sentence was uttered by the red character alone or by the red and the blue characters together.[4]

The Q-REFERENCE depended on the specific type of trial (critical or filler). In critical trials, participants had to choose between two images: a *homogeneous* picture, which displayed two characters who could arguably be assigned the same person role and a *heterogeneous* picture, which displayed one character present in the context and one character absent in the context. An illustration of these two alternatives for Q-REFERENCE is provided in Fig. 1.

The experiment involved four critical trials, two per plural pronoun. We also included 12 filler trials including singular pronouns, and four filler trials including plural pronouns but with different alternatives in Q-REFERENCE than those in the critical trials. In filler singular trials, alternatives for Q-REFERENCE involved choosing between pictures with a single character, which could be focal or not. In filler plural trials, participants were asked to choose between one heterogeneous or homogeneous picture, and a *combined* picture, which displayed all the relevant characters together (focal, non-focal and absent in the scene) together. The description of options for fillers can be found at osf.io/2bzvm. Trials for each pronoun involved one of four different predicates; the same pronoun would thus appear in four different sentences. The same four predicates were used with the two plural pronouns, and other four different predicates were used for the trials involving singular pronouns.

The combination of answers to Q-SPEAKER, Q-ADDRESSEE and Q-REFERENCE in critical trials allow us to discern between additive and associative readings. 1PL pronouns are considered to be interpreted additively if participants first consider that there are multiple speakers (in Q-SPEAKER) and then choose an homogeneous picture as the reference for the pronoun (in Q-

[4]To prevent automatic responses, where participants always select e.g., the focal referent alone, in filler trials the single speaker/addressee description could involve the non-focal referent alone (in the example, the blue character alone) rather than the focal one.

Reference). Likewise, 2pl pronouns are interpreted additively whenever participants select multiple addressees in Q-Addressee and later a homogeneous picture in Q-Reference.

Exp 2 was identical to Exp 1 except that participants were not asked the Q-Addressee. The reasoning behind this modification was that the Q-Addressee could boost additive readings of 2pl pronouns by making them available to participants (we will discuss this further in Section 3). We thus excluded this question to be able to compare directly 2pl and 3pl pronouns. Table 1 summarises the alternative responses per experiment and question type.

3 Results

Fig. 2 shows the proportion of responses compatible with an additive interpretation for each experiment and person category. For Exp 1, a response is considered to be compatible with an additive interpretation when participants first select multiple speakers or multiple addressees for 1pl (in Q-Speaker) and 2pl pronouns (in Q-Addressee) respectively, and later select homogeneous pictures (in Q-Reference).[5] In Exp 2, responses are considered to be compatible with an additive interpretation when participants select homogeneous pictures in Q-Reference.

We use *R*'s *brms* package [2] as an interface to *Stan* [3] to run Bayesian binomial mixed-effects regression models predicting the additive-like responses of plurality by person. We run two separate models, one for Exp 1 and another for Exp 2, with the same model structure. Our dependent variable is participants' responses compatible with an additive interpretation as explained above (coded as 1 if compatible, and 0 if not). We apply treatment coding to our categorical predictor of person. In Exp 1, we compare 2pl to 1pl trials; in Exp 2, we compare 3pl to 2pl. We include random intercepts for participants as well as by-participant random slopes for the effect of person. We set the same student-t prior on our fixed effect as well as on the intercept ($DF = 6, \mu = 0, \sigma = 1.5$); for the random effects, we set a half-Cauchy prior with scale parameter 10.

Results from our models suggest that, in accordance with our hypothesis, participants are more likely to choose additive-like responses for 2pl pronouns than for 1pl in Exp 1 ($\hat{\beta} = 2.574$, 90%CI $= [1.705, 3.617]$, $SE = 0.589$, $P(\hat{\beta} > 0) = 1$) and for 3pl pronouns than for 2pl pronouns in Exp 2 ($\hat{\beta} = 1.848$, 90%CI $= [0.911, 2.939]$, $SE = 0.625$, $P(\hat{\beta} > 0) = 0.99$).

In the results in Fig. 2 we nonetheless observe that the proportion of additive responses for 2pl notably varies between experiments. We believe this is due to a sort of *priming* effect, whereby in Exp1, the Q-Addressee influences subsequent Q-Reference responses in 2pl trials, boosting additive-like interpretations. Recall that this was the reason why we excluded these trials in Exp 2, as it would not have allowed us to compare 2pl and 3pl.

In order to explore this effect, we look at the proportion of multiple addressee choices in Q-Addressee questions across 1pl and 2pl in Exp 1. The idea is that participants might be more likely to conceive multiple addressees when they are specifically asked about the addressee(s) of 2pl sentences, possibly because they find it hard to distinguish between the addressee(s) and the pronominal referents. Indeed, in Fig. 3 we observe that participants choose multiple addressees with a grand mean probability of $P = 0.44$ in 1pl trials, but for 2pl trials, they select them with $P = 0.78$. After the selection of multiple addressees in 2pl trials, the only possible compatible choice in Q-Reference for critical trials is the homogeneous image: the alternative heterogeneous image does not contain both of the chosen addressees and selecting it would thus be incongruous. Altogether, this suggest that the use of 2pl pronouns in sentences is

[5] We exclude trials in Exp 1 where participants select multiple speakers or addressees in Q-Speaker or Q-Addressee and later do not select pictures that include these referents in Q-Reference. These are considered to be incongruous (errors). We excluded the responses of 5 (out of 118) 1pl and 7 (out of 118) 2pl trials.

indeed priming participants to select multiple addresses more than they would do for sentences containing another pronoun (i.e., 1PL), and it is inevitably boosting the choice of homogenous images in Q-REFERENCE, which it is not the case in Exp 2.

In our confirmatory results (Fig. 2) we also observe that, despite its rarity, participants are equally likely to select multiple speakers in 1PL trials than single speakers (which is the norm). This can be also due to the difficulty of teasing apart speakers from pronominal referents in 1PL trials, the same way we observe it is hard to distinguish addressees and referents in 2PL trials. Fig. 4A shows that participants choose multiple speakers in Q-SPEAKER notably more often for 1PL than for 2PL trials, where they select singular speakers almost exclusively (as in 2PL and 3PL trials in Exp 2 as well, see Fig. 4B).

4 Discussion

Our study provides the first experimental evidence confirming that the likelihood of additive interpretations of plural pronouns depends on the person category. Additive readings are more likely in 3PL than in 2PL pronouns (Exp 2), and in 2PL pronouns than in 1PL pronouns (Exp 1). This is consistent with our hypothesised hierarchy: $1 < 2 < 3$.

Although intuitive, the difference in interpretation between 1PL and 2PL pronouns is not accounted for by most theories of person—but see [19]. Our results provide evidence in favour of an account whereby the source of this asymmetry is not due to a difference in the semantics of these pronouns alone, but rather to the combined influence of world knowledge and semantics: the existence of multiple speakers is less likely than the existence of multiple addressees (independently of language). Exp 1 moreover suggests that additive interpretations are *possible* for both 1PL and 2PL pronouns, indicating that the person ontology should allow for multiple speakers and addressees—*contra* [10]. Interestingly, research on the cross-linguistic distribution of person systems has argued that some morphological patterns (e.g., the likelihood of person-number suppletion) also seem to obey a $1 < 2 < 3$ hierarchy [17], suggesting that the interaction of morphosemantic knowledge and functional pressures not only influences the preference of a certain interpretation but also finds its way into the grammar.

Note, however, that even though the results from Exp 2 confirm our hypothesis, we cannot establish an additive interpretation for what we have been calling 'additive-like' responses in Q-REFERENCE. Unlike for Exp 1, here we do not know whether participants conceive multiple or single addressees when they select homogeneous images, and thus we are inflating additive interpretations by only taking into account Q-REFERENCE responses. However, even with this conservative approach, we still observe the predicted $2 < 3$ asymmetry.

Surprisingly, we also found comparable results for 2PL in Exp 1 and 3PL in Exp 2 (see Fig. 2). This could suggest a hierarchy $1 < 2/3$ (rather than $1 < 2 < 3$) with regards to plural interpretation. However, we also observe a notable difference between additive responses for 2PL across experiments, which we have argued results form the priming effect the Q-ADDRESSEE has on the selection of multiple addressee responses in communicative contexts where the sentence uttered contains a 2PL pronoun. We showed that participants are indeed more likely to select homogeneous images in Q-REFERENCE in Exp 1 than in Exp 2, indicating that a direct comparison between Exp 1 and 2 is hard to make.

References

[1] Jonathan David Bobaljik. Missing persons: A case study in morphological universals. *The Linguistic Review*, 25(1–2):203–230, 2008.

[2] Paul-Christian Bürkner. Advanced Bayesian multilevel modeling with the R package brms. *The R Journal*, 10(1):395–411, 2018.

[3] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1), 2017.

[4] Greville G Corbett. *Number*. Cambridge University Press, 2000.

[5] Michael Cysouw. *The Paradigmatic Structure of Person Marking*. OUP Oxford, 2003.

[6] Michael Cysouw. The expression of person and number: A typologist's perspective. *Morphology*, 21(2):419–443, 2011.

[7] Michael Daniel. Associative plural as indexical category. *Language Sciences*, 81:101256, 2020.

[8] Michael Daniel and Edith Moravcsik. The associative plural. In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig, 2013.

[9] Joseph H Greenberg. The first person inclusive dual as an ambiguous category. *Studies in Language. International Journal sponsored by the Foundation "Foundations of Language"*, 12(1):1–18, 1988.

[10] Daniel Harbour. *Impossible Persons*. Linguistic Inquiry Monographs. MIT Press, 2016.

[11] Irene Heim. Features on bound pronouns. In Daniel Harbour, David Adger, and Susana Béjar, editors, *Phi theory: phi-features across modules and interfaces*, chapter 2. Oxford University Press, Oxford, 2008.

[12] Angelika Kratzer and Irene Heim. *Semantics in generative grammar*, volume 1185. Blackwell Oxford, 1998.

[13] Edith A Moravcsik. A semantic analysis of associative plurals. *Studies in Language. International Journal sponsored by the Foundation "Foundations of Language"*, 27(3):469–503, 2003.

[14] Robert Rolf Noyer. *Features, positions and affixes in autonomous morphological structure*. PhD thesis, Massachusetts Institute of Technology, 1992.

[15] Uli Sauerland. On the semantic markedness of phi-features. In Daniel Harbour, David Adger, and Susana Béjar, editors, *Phi theory: phi-features across modules and interfaces*, chapter 3. Oxford University Press, Oxford, 2008.

[16] Philippe Schlenker. Person and binding (a partial survey). *Italian Journal of Linguistics/Rivista di Linguistica*, 16(1):155–218, 2004.

[17] Anna Siewierska and Dik Bakker. Suppletion in person forms: The role of iconicity and frequency. *Languages across boundaries: Studies in memory of anne siewierska*, pages 347–396, 2013.

[18] Jolijn Sonnaert. *The atoms of person in pronominal paradigms*. Netherlands Graduate School of Linguistics, 2018.

[19] Stephen Wechsler. What 'you' and 'I' mean to each other: Person indexicals, self-ascription, and theory of mind. *Language*, 86(2):332–365, 2010.

	Exp 1	Exp 2
Q-Speaker	(a) focal speaker (*single*) (b) focal + non-focal speakers (*multiple*)	(a) focal speaker (*single*) (b) focal + non-focal speakers (*multiple*)
Q-Addressee	(a) focal addressee (*single*) (b) focal + non-focal addressees (*multiple*)	
Q-Reference	(a) focal + non-focal (*homogeneous*) (b) focal + associates (*heterogeneous*)	(a) focal + non-focal (*homogeneous*) (b) focal + associates (*heterogeneous*)

Table 1: Alternative answers per question and experiment. Labels are indicated in italics.

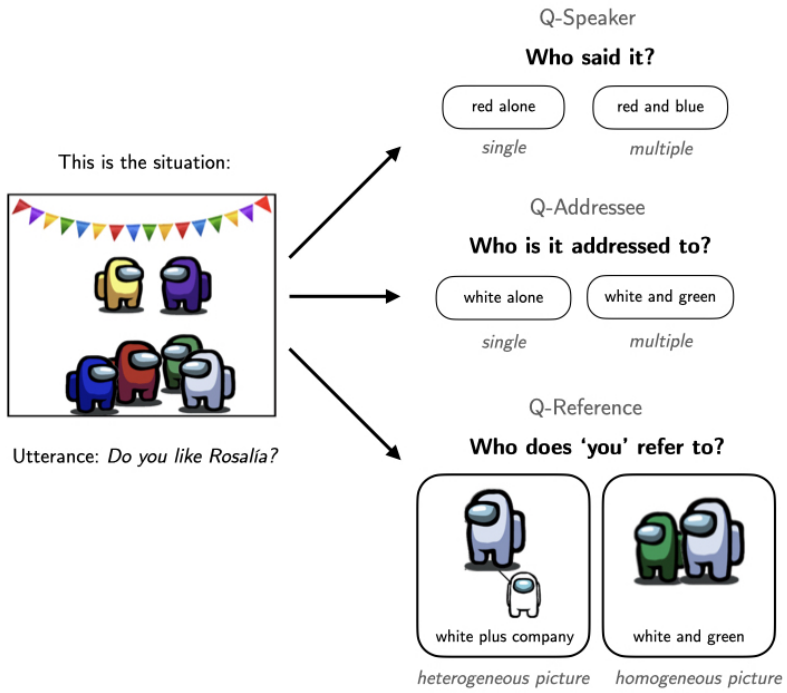

Figure 1: Schematic trial (Exp 1: 2PL pronoun). Each question was presented with the GIF-sentence pair. Original materials in Spanish can be found in osf.io/fbnc7. Labels in italics were not part of the trial.

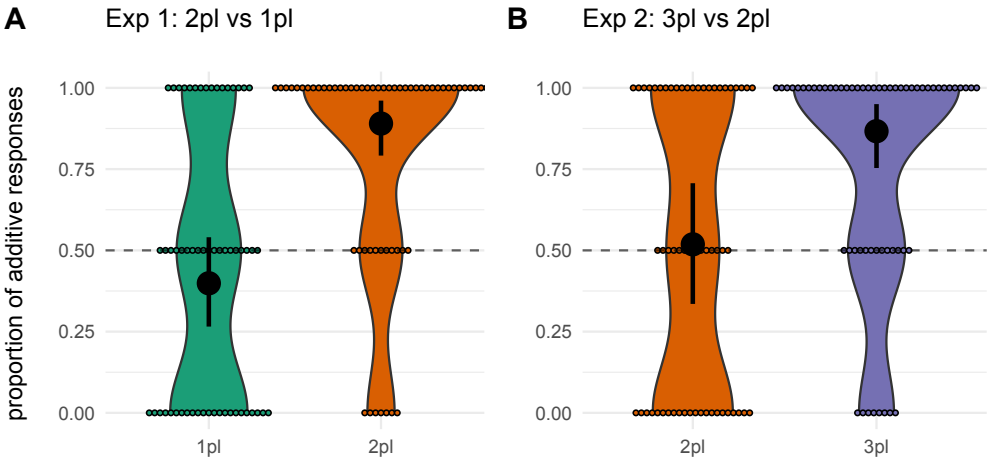

Figure 2: Proportion of additive-like responses in Exp 1 (A) and Exp 2 (B). We show the model's estimate means (black dot) and 90%CI, as well as kernel density plots indicating the peaks in the data distribution. The small coloured dots illustrate participants' individual scores. The dash line indicates the chance level.

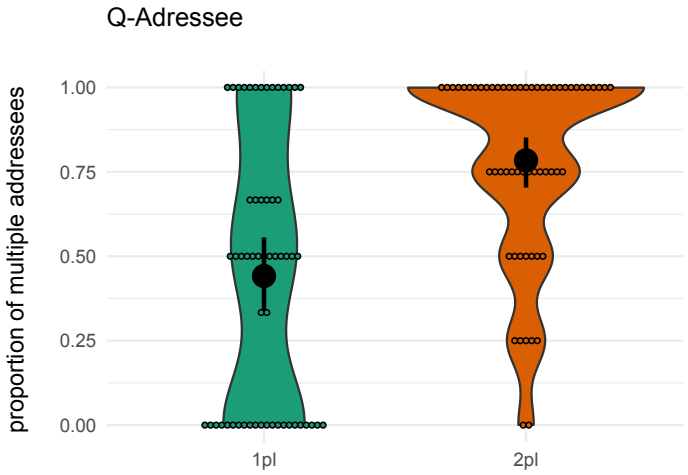

Figure 3: Proportion of selection of multiple addresses in Q-ADDRESSEE for 1PL and 2PL trials in Exp 1. We show the means and standard errors, as well as kernel density plots indicating the peaks in the data distribution. The small coloured dots illustrate participants' individual scores. Note that we include the responses of all Q-ADDRESSEE plural trials, regardless of whether they are followed by a critical Q-REFERENCE trial or not.

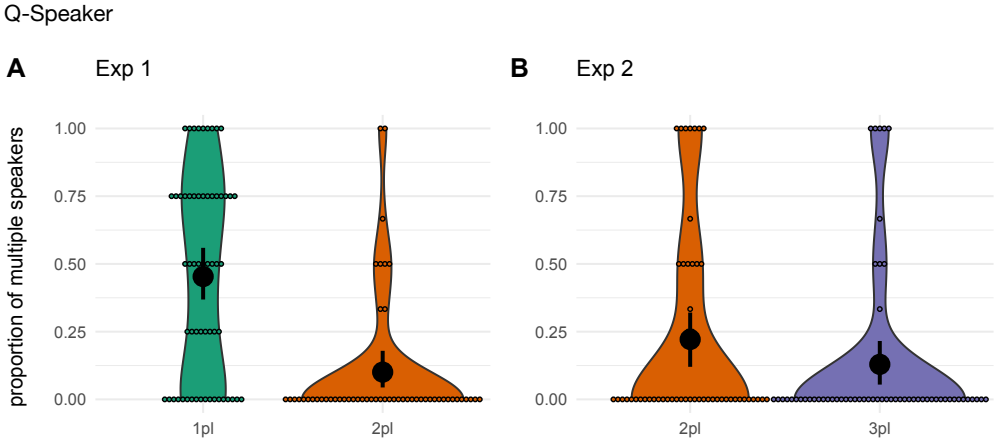

Figure 4: Proportion of selection of multiple speakers in Q-SPEAKER for 1PL and 2PL trials in Exp 1 (A), and for 2PL and 3PL trials in Exp 2 (B). We show the means and standard errors, as well as kernel density plots indicating the peaks in the data distribution. The small coloured dots illustrate participants' individual scores. Note that here we also include the responses of all Q-SPEAKER plural trials, regardless of whether they are followed by a critical Q-REFERENCE trial or not.