

# How In-Group Bias Influences the Level of Detail of Speaker-Specific Information Encoded in Novel Lexical Representations

Sara Iacozza

Max Planck Institute for Psycholinguistics, Nijmegen, the Netherlands, and International Max Planck Research School for Language Sciences, Nijmegen, the Netherlands

Antje S. Meyer

Max Planck Institute for Psycholinguistics, Nijmegen, the Netherlands, and Radboud University

Shiri Lev-Ari

Max Planck Institute for Psycholinguistics, Nijmegen, the Netherlands, and Royal Holloway, University of London

An important issue in theories of word learning is how abstract or context-specific representations of novel words are. One aspect of this broad issue is how well learners maintain information about the source of novel words. We investigated whether listeners' source memory was better for words learned from members of their in-group (students of their own university) than it is for words learned from members of an out-group (students from another institution). In the first session, participants saw 6 faces and learned which of the depicted students attended either their own or a different university. In the second session, they learned competing labels (e.g., *citrus-peller* and *citrus-schiller*; in English, *lemon peeler* and *lemon stripper*) for novel gadgets, produced by the in-group and out-group speakers. Participants were then tested for source memory of these labels and for the strength of their in-group bias, that is, for how much they preferentially process in-group over out-group information. Analyses of source memory accuracy demonstrated an interaction between speaker group membership status and participants' in-group bias: Stronger in-group bias was associated with less accurate source memory for out-group labels than in-group labels. These results add to the growing body of evidence on the importance of social variables for adult word learning.

**Keywords:** in-group bias, word learning, indexical information, social identity, lexical representations

Previous research has established that people can quickly make use of any available cue (e.g., gender, age) to infer whether their interlocutors belong to their in-group (e.g., Bargh, Schwader, Hailey, Dyer, & Boothby, 2012). Group membership status has been shown to influence how attentively people process information about and from others, resulting in more accurate memory for information related to in-group compared with out-group members (e.g., Hugenberg, Young, Bernstein, & Sacco, 2010; Judd & Park, 1988; Van Bavel & Cunningham, 2012). For instance, after seeing a series of unfamiliar faces, people showed superior memory for


those faces that they believed to be in-group members, as compared with out-group faces (e.g., Van Bavel & Cunningham, 2012). Similarly, people were better at recalling in-group biographic information than out-group information (e.g., Judd & Park, 1988).

Here, we test the proposal that in-group biases permeate language learning as well. Specifically, we ask whether the group membership of the speakers from whom new words are learned modulates the level of detail with which information about the speaker identity is encoded. We suggest that representations of novel words are more speaker-specific when learned from in-group members than when learned from out-group members. Furthermore, we predict that the difference in level of detail between in-group and out-group representations are modulated by the strength of learners' own in-group bias, namely, by how much learners preferentially process in-group over out-group related information.

To test our predictions, we designed a word-learning study in which participants learned novel words from in-group and out-group speakers. Participants then performed a task that measured their in-group bias. We then tested their source memory for the words learned from in-group versus out-group members and assessed how source memory was related to individual in-group biases.

Before turning to the current study, we review the relevant literature. First, we focus on evidence that shows that speaker

This article was published Online First October 17, 2019.

 Sara Iacozza, Max Planck Institute for Psycholinguistics, Nijmegen, the Netherlands, and International Max Planck Research School for Language Sciences, Nijmegen, the Netherlands; Antje S. Meyer, Max Planck Institute for Psycholinguistics, and Donders institute for Brain Cognition of Behaviour, Radboud University; Shiri Lev-Ari, Max Planck Institute for Psycholinguistics, and Department of Psychology, Royal Holloway, University of London.

We thank Caitlin Decuyper and Rene Terporten for help with programming. We are grateful to Atsuko Takashima, Christina Bergmann, Rebecca Frost, Laurel E. Brehm, and Ian Hancock for helpful discussions regarding this work.

Correspondence concerning this article should be addressed to Sara Iacozza, Max Planck Institute for Psycholinguistics, P.O. Box 310, 6500 AH Nijmegen, the Netherlands. E-mail: [sara.iacozza@mpi.nl](mailto:sara.iacozza@mpi.nl)

social identity plays a role in language processing. Second, we describe existing exemplar-based models of language processing that account for the effects of speaker social identity on language processing. At this stage, we also propose modifications for existing models. We argue that existing models should not assume that all speakers are treated equally but instead assume that information provided by in-group speakers is encoded more fully than information from out-group speakers. We end by providing support for our proposed modifications from nonlinguistic studies in social psychology that report group membership effects on memory and information processing.

### Speaker Social Identity Can Influence Language Processing

Previous research indicates that people rapidly extract information about others' social identity (see Bargh et al., 2012 for a review). Such information has been shown to activate expectations and attitudes (e.g., Greenwald & Banaji, 1995; Jones & Fazio, 2010). For instance, children who were given silent demonstrations of unfamiliar objects' different functions selectively endorsed the functions provided by native-accented rather than foreign-accented speakers (Kinzler, Corriveau, & Harris, 2011), showing a native-speaker preference even when the language was not involved in the task. The use of others' social identity to guide processing extends to language processing in a variety of ways. For instance, a speakers' social identity can shape listeners' expectations. Upon hearing a speakers' voice, people activate expectations about what will likely be said. If these expectations are not met, such as when a habit of drinking wine is reported in a child's voice, language processing becomes harder (Van Berkum et al., 2008, see also Martin, Garcia, Potter, Melinger, & Costa, 2016; Walker & Hay, 2011). Similarly, speaker social identity can affect listeners' word processing (e.g., Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; Johnson, Strand, & D'Imperio, 1999; Niedzielski, 1999). For example, the gender of a face displayed when listening to ambiguous vowels affected participants' responses on a vowel identification task, suggesting that information about the speaker affects speech perception (Johnson et al., 1999). Finally, listeners detected fewer word changes in a story after listening to non-native as compared with native speakers (Lev-Ari & Keysar, 2012). Importantly, these differences disappeared when participants were instructed to pay attention to the exact wording of the stories, rather than simply comprehending them (Lev-Ari & Keysar, 2012). These results confirmed that people were in principle able to process the language of native and nonnative speakers but did so less when listening for comprehension to the nonnatives (unless given explicit instructions).

The literature shows that people encode speakers' social identity and this information can influence speech processing. Based on this evidence, one might expect speakers' social identity to influence language learning as well. When learning new words, the novel linguistic information might be encoded and stored in representations that contain details about both its linguistic content (e.g., its word form, its meaning) and its source (i.e., who said it). Indeed, existing exemplar-based models of speech processing assume such encoding of social information.

### Exemplar-Based Models: Concurrent Encoding of Social and Linguistic Input

Exemplar-based models of speech processing posit that linguistic experiences are encoded as rich episodic memories (i.e., exemplars; e.g., Goldinger, 2007; Hay, Nolan, & Drager, 2006; Nielsen, 2011). Exemplars contain information that is both language-specific (e.g., includes phonetic, lexical, and syntactic details) and related to the speakers' social identity, spanning from more intrinsic indexical characteristics such as age and sex to broader social categories, such as professions (e.g., a doctor) and membership in a specific community (e.g., health care professionals; see Drager & Kirtley, 2016 for a review). Associating speaker-related information to linguistic input can be beneficial for language users. It allows listeners to preactivate those associations in the appropriate context and with the relevant interlocutors. This may exploit the socially conditioned variation in speech to speed up and/or ease comprehension (Münster & Knoeferle, 2018 for a recent account of socially situated language processing).

In a recent extension of exemplar models, Sumner, Kim, King, and McGowan (2014) proposed the socially weighted encoding of spoken words. Specifically, they suggested that social characteristics, such as perceived prestige, might affect linguistic encoding in the regard that prestigious forms would receive greater weight. In support of a socially weighted encoding account, Sumner and colleagues referred to empirical evidence from work that compares standard with nonstandard phonetic variations. Extending their theory, one could hypothesize that people may encode linguistic variations more strongly if they are associated with speakers to whom a special status is ascribed.

We propose that learning new words from speakers that are ascribed a special status, such as in-group members, might lead to lexical representations that are richer in contextual information (e.g., speaker-related) compared with representations of words learned from speakers without a special status. In the current study, we tested the hypothesis that representations of novel words learned by in-group speakers contain individual-specific detail, whereas new words learned from out-group speakers contain broader group-specific information, and that such a pattern is further modulated by learners' in-group bias. This finding would contribute to the open debate in psycholinguistics on how words are represented in the mental lexicon and call for an extension of existing exemplar-based theories that assume an encoding of contextual information that does not vary in relation to the identity of the speakers. Though the role of group membership in language processing has not yet been explored, there is evidence to suggest that it influences processing and learning.

### In-Group Biases

Well-established self-biases, such as improved memory and faster processing of stimuli related to the one's self versus stimuli related to others (see Symons & Johnson, 1997, for a meta-analysis), have been shown to extend to in-group members as well (e.g., Aron, Aron, Tudor, & Nelson, 1991; Sui & Humphreys, 2015; Symons & Johnson, 1997; see also Cadinu & Rothbart, 1996). Thus, memory is better for in-group faces than for out-group faces (e.g., Hugenberg et al., 2010; Van Bavel, Packer, & Cunningham, 2008) and for information delivered by in-group

than by out-group members (e.g., Frable & Bem, 1985; Wilder, 1990). People also learn and process new associations between previously neutral stimuli (e.g., geometrical shapes) and in-group membership (e.g., the logo of their favorite football club) more quickly than associations involving out-group membership (Enock, Sui, Hewstone, & Humphreys, 2018; Moradi, Sui, Hewstone, & Humphreys, 2015). Like self-biases, in-group biases may rely on the recruitment of additional attentional resources (Meissner, Brigham, & Butz, 2005; Van Bavel & Cunningham, 2012). These additional resources have been suggested to lead to representations with a higher level of detail for in-group than out-group representations. For example, in-group faces are easier to recognize than out-group faces because more individual-specific information is encoded during the exposure phase (e.g., Hugenberg et al., 2010).

The better processing of in-group-related information extends to the individuals providing the information. Greenstein, Franklin, and Klug (2016) showed that source memory was superior when participants believed the sources to be in-group members than when they believed them to be out-group members. Such findings are highly relevant for language learning and language processing because they suggest that when listeners comprehend speech, the amount of both speaker-specific detail and linguistic information that is encoded might depend on whether the speaker is an in-group or an out-group member of the listener.

Thus, we propose that (a) listeners encode the group membership status of the speakers (i.e., in-group vs. out-group) from whom they learn new words and (b) group membership status and individual in-group bias influence the level of detail with which speaker-specific information is encoded. In the current study, we tested this idea by probing participants' source memory for novel words. We predicted that words learned from in-group speakers, relative to words learned from out-group speakers, would contain a higher level of detail about who produced them. Thus, we suggest that in-group linguistic representations would be more speaker-specific than would be out-group representations. Crucially, this in-group advantage should be stronger for participants exhibiting greater in-group bias.

## Method

We conducted a 2-day experiment. On Day 1, speaker group membership was established via a speaker familiarization task in which participants listened to six fictitious speakers refer to facts and habits about their lives that implied either attendance at the participants' university ( $n = 3$ ; in-group speakers) or at a different university ( $n = 3$ ; out-group speakers).

On Day 2, participants performed the word-learning task, which was a communication task. In the exposure phase of this task, they learned novel labels for 24 unfamiliar gadgets by listening to the speakers from Day 1 refer to these gadgets. Crucially, target gadgets received one label from an in-group speaker and an alternative, competing, but equally fitting, label from an out-group speaker (e.g., *citrus-peller* vs. *citrus-schiller*; in English, *lemon peeler* vs. *lemon stripper*). Afterward, source memory for these words was tested. Participants were shown a photo of one speaker and one label at a time and asked if the speaker had produced the label in the previous phase (eliciting a forced choice yes or no). This was done with all speakers and labels from the exposure

phase. Last, participants' implicit in-group bias was measured by their performance in a perceptual matching task.

Following our hypothesis suggesting that the level of detail encoded depends on the speakers' group membership, we predicted that representations of words learned from in-group members would include more detailed information about the speaker than representations of words learned from out-group members. This would be evidenced by greater accuracy for in-group than for out-group label-speaker pairs. Additionally, we expected participants to spontaneously encode the speakers' group membership status (even though they were not explicitly instructed to apply this strategy). Consequently, when asked to recognize the source of the new words, we expected participants to struggle remembering the exact source for each word, and instead to rely on speaker group membership to inform their decisions. Therefore, they should be more likely to misattribute words to incorrect speakers within the same university than between different universities, that is, there should be source memory confusion. Because in-group representations should be more speaker-specific, the source memory confusion for in-group words should be of a smaller magnitude as compared with confusion for out-group words. Last, we predicted that these effects would be modulated by participants' own in-group bias, such that the stronger the in-group bias, the larger the difference between speaker-specific details in in-group versus out-group representations.

## Participants

Sixty native Dutch speakers (age range = 18–27 years;  $M = 21.98$ ,  $SD = 1.89$ ) participated in the study after providing their informed consent, as approved by the Ethics Committee of the Social Sciences Department of Radboud University (Project Code: ECSW2014-1003–196). All participants were students or recent graduates of Radboud University Nijmegen.<sup>1</sup> Eight additional participants were tested, but their data were excluded from the study because they reported during the debriefing to have strong connections with the out-group university (e.g., their sister or best friend studying there). All participants were female, as were the speakers from whom they learned. This was done to avoid having an additional social dimension (gender) of in-group status interact with the one we manipulated (i.e., same/different university).

## Materials

**Speaker familiarization task.** Six fictitious speakers were created by pairing female faces selected from the Chicago Face Database (Ma, Correll, & Wittenbrink, 2015), with the voices of native Dutch, female speakers recorded in our laboratory. Prior to the experiment, voices were matched for perceived typicality and attractiveness (paired  $t$  tests,  $ps > .05$ ) via an online norming in which 20 different participants participated. Each speaker was a unique combination of one face and one voice, consistent across participants.

After a norming procedure (see Appendix A for details), we selected 24 sentences about student life and recognizable land-

<sup>1</sup> One participant's accuracy was more than three standard deviations away from the population mean in the perceptual matching task. We excluded her data from the word-learning analyses.

marks or actions. Twelve sentences referred to Radboud University and 12 to Groningen University. This was achieved by creating 12 unique sentential frames in which key words explicitly related to either Radboud or Groningen University were embedded. All statements were audio-recorded by each of the female voices described in the preceding text.

**Word-learning task.** The same speakers from the speaker familiarization task were used. Images of 24 unfamiliar gadgets (e.g., lemon peeler) and their corresponding labels were selected via a norming study (see [Appendix A](#) for details). Half of the gadgets, hereinafter referred to as *target gadgets*, were presented with two competing labels, which were equated for goodness-of-fit and frequency. The other 12 gadgets were presented with a single label and served as fillers. All labels were audio-recorded by the speakers.

**Individual in-group bias task.** Black and white pictures of the logos of Radboud and Groningen Universities were used. Black shapes for triangle, square, and circle were used.

**Gadget familiarization task and prior familiarity task.** The same 24 images of gadgets used in the word-learning task were used in these tasks. However, for the purposes of the gadget familiarization task, an additional exemplar for each gadget was selected online.

## Procedure

**Day 1: Gadget familiarization task.** The goal of the gadget familiarization task was to familiarize participants with the gadgets to reduce demands in the subsequent word-learning task. This task was a picture-matching task. On each trial, participants saw two images and had to indicate whether they depicted the same gadget. In half of the trials, different items were shown, whereas in the remaining trials, two instances of the same gadget (only one of which was later used in the word-learning task) were displayed. The gadgets were shown in two exposure rounds with a randomized trial order. This led to 96 trials in total (48 in each round). Participants' responses were self-paced. Performance on the task was at ceiling, as expected, and are not discussed further.

**Day 1: Speaker familiarization task.** The goal of the speaker familiarization task was to familiarize the participants with the speakers and ensure that they knew their group membership status before they did the word-learning task on Day 2. Participants were told that this task was part of a joint project of Radboud University, their in-group university, and Groningen University, an out-group university, in which they would see students of those universities describing their student life. The information provided was supposed to guide inferences about the speakers' academic affiliation: all habits implied attendance at either the in-group or out-group university. To ensure that these affiliations had been correctly inferred, we had participant perform a surprise memory test. In each trial of this memory test, they saw a photo of a speaker and six alternative sentences written down. They selected which sentence among the given six was the one that the speaker had previously uttered. The memory test was performed both on Day 1, immediately after the exposure phase, and on Day 2, at the beginning of the experimental session. Performance on these indirect tests allowed us to ensure that participants indeed learned speakers' academic affiliations.

On each exposure trial, participants saw a fixation cross (500 ms) followed by a photo of a speaker (500 ms). Then the sentence was played while the photo was still displayed. Two exposure rounds of 24 trials each were administered with a randomized trial order. In each round, participants heard all six speakers, three per university, produce four statements each. Speaker affiliation was randomized across participants.

After exposure, a surprise source memory test was administered. On each test trial, participants saw a photo of a speaker and six written sentences. These sentences were as follows: the correct response (i.e., the one sentence that the speaker had previously uttered), two incorrect sentences produced by other speakers from the speakers' university, and three incorrect sentences produced by speakers from the other university. Participants were instructed to indicate which sentence had been previously uttered by the speaker by pressing the corresponding key on a keyboard. Responses were self-paced and immediate visual feedback about the accuracy of their choice was provided along with the correct option.

Responses were coded as (a) correct, (b) within-university error (selection of a sentence produced by a speaker from the speakers' university), or (c) between-university error (selection of a sentence produced by a speaker from the other university). If participants made a between-university error, the trial was repeated at the end. This was done until no between-university errors were made to ensure participants correctly learned the speakers' affiliation before moving to the word-learning task. We did not repeat trials that led to within-university errors, since the goal of the training was only to teach participants which speakers were members of their in-group and which belonged to the out-group. This was critical for the assessment of our hypothesis concerning the effect of in-group bias on novel word learning.

**Day 2: Delayed speaker test.** The test of the speaker familiarization task from Day 1 was repeated (without repeating the exposure phase) to make sure participants still remembered the speakers' academic affiliations before the word-learning task. Speakers and sentences were the same as on Day 1, and the procedure of repeating trials that yielded between-university errors used on Day 1 was used on Day 2 as well.

**Day 2: Word-learning task.** The word-learning task was presented as a communication task. Thus, participants were told that they would select pictures according to information provided by the speakers from the previous task. In the exposure phase, participants saw 24 gadgets, all named by both in-group and out-group speakers. Half were target gadgets, for which in-group and out-group speakers provided competing, but equally fitting, labels, whereas the other half were fillers, for which unique labels were provided. Fillers were included to minimize participants' awareness of the nature of the manipulation (i.e., in-group vs. out-group) in the experiment. Note that not all speakers referred to all the gadgets. In fact, each gadget was only labeled by two of the six speakers (one in-group and one out-group speaker). Speaker group membership, speaker-label pairing, and label-group membership status pairings were randomized per participant. On each trial, a fixation cross (500 ms) preceded a photo of a speaker (800 ms). Then, the recording of the gadget label was played while the photo was still displayed. Simultaneously, the written form of the label appeared superimposed upon the speakers' mouth (1,500 ms). Next, three gadgets appeared on the screen and participants selected the one that fit the audio and the written label (see [Figure](#)

1 for an example of the learning display).<sup>2</sup> If the response was wrong, the audio was repeated. Three exposure rounds were administered so that each display was repeated three times, once per round, in a randomized trial order.

Next, participants performed a surprise source memory recognition task on the target gadget labels only. In each trial, they saw a photo of a speaker and a written label (see Figure 2). Participants indicated whether the speaker had produced the label in the previous phase via key press (yes/no). Trials were self-paced. There were 148 test trials in total in which all possible speaker–label pairings were shown. This means that there were 24 matching trials (i.e., the speaker had indeed produced the label) and 120 mismatching trials (i.e., the label had not been used by the speaker). Of those mismatching trials, 48 were within-university mismatching trials (showing a label along with a wrong speaker from the same university as the correct one), 48 were between-university mismatching trials (showing a label with a wrong speaker from the other university). The remaining 24 trials showed a speaker with a label that competed with the one she used (e.g., the speaker that had used *citrus-schiller* instead of *citrus-peller*) and were included only so that all possible combinations of speaker–label pairings were available; these were not analyzed. Note that in all mismatching trials, the correct answer was that the pairing was incorrect as the speaker depicted in the photo had not used the displayed label in the word-learning task.

**Day 2: In-group bias task.** The participants' in-group bias was measured in a perceptual matching task (Moradi et al., 2015), which has been shown to provide results that are reliable within individuals and across different test sessions (Stolte, Humphreys, Yankouskaya, & Sui, 2017). Three geometric shapes (circle, square, triangle) were randomly paired with logos of three universities: the in-group university (Radboud University) and two other universities (Groningen University and Tilburg University). Each association was initially presented 10 times. Then, participant performed a practice block of 24 trials, followed by two blocks of



citrus-peller

Figure 2. Example of a memory test trial. Participants indicated if the speaker had produced the label in the exposure task. Facial image was reprinted and adapted from The Chicago Face Database: A Free Stimulus Set of Faces and Norming Data, by Ma, Correll, & Wittenbrink, 2015, *Behavior Research Methods*, 47, 1122–1135. Copyright by the University of Chicago, Center for Decision Research. This figure is used with permission. See the online article for the color version of this figure.

120 experimental trials each. In both practice and test trials, a fixation cross (500 ms) preceded a blank screen (between 1,000 ms and 2,000 ms) and the simultaneous presentation of logo and shape (600 ms), following the timings used in Moradi et al. (2015). Participants had 1,500 ms to judge the accuracy of the pairing. Feedback was given only during practice. In-group bias in this task is typically indexed by faster and more accurate responses for stimuli that are newly associated with in-group membership compared with stimuli associated with out-group membership (e.g., Moradi et al., 2015). We replicated these general patterns (see Appendix B) and extracted individual measures of in-group bias by calculating a per-participant measure of effect size, Cohen's *d*, from both accuracy and response times (RTs) over in-group versus out-group matching trials. The measure calculated over RTs was not a significant predictor in any of the models we ran; thus, we focus on the measure derived from accuracy (see Appendix B for the distribution of the individual measures).

**Day 2: Rating of prior familiarity, written picture naming, and preference.** After the perceptual matching task described in the preceding text, participants carried out three further tasks. First, to exclude gadgets that participants already had labels for before the experiment, we asked them to indicate whether they already knew any gadget and their names before the experiment. Based on these ratings, gadgets familiar to a participant were excluded from their data set (11.81% of data across all participants). Second, participants saw all the pictures of the gadgets again and were asked to write down the labels they had learned, as well as any additional labels they might know or come up with. Finally, they



Figure 1. Example of the learning display. Participants had to select the gadget that was referred to. In this case, they were expected to indicate the first image. Facial image was reprinted and adapted from The Chicago Face Database: A Free Stimulus Set of Faces and Norming Data, by Ma, Correll, & Wittenbrink, 2015, *Behavior Research Methods*, 47, 1122–1135. Copyright by the University of Chicago, Center for Decision Research. This figure is used with permission. See the online article for the color version of this figure.

<sup>2</sup> Because of copyright issues, none of the pictures of the gadgets in the example correspond to actual stimuli, but they provide a good approximation of the type of stimuli we used.

were asked to indicate which label among those they had just written down (regardless of whether they had seen them in the experiment) was the most appropriate for each gadget. This task was meant to assess the participants' preference for labels acquired from in-group versus out-group speakers. However, probably because of unclear instructions, many participants were very creative and came up with their own alternatives or rarely wrote down two labels for one object. Therefore, this analysis could not be conducted as a result of data scarcity.

## Results

All analyses were performed with mixed-effects modeling as implemented in the lme4 package (Bates, Maechler, Bolker, & Walker, 2014) in R (R Core Team, 2013). The random structures of the models were determined following Bates, Kliegl, Vasishth, and Baayen (2015).

### Speaker Familiarization Task

We tested whether participants correctly learned speaker group membership status. Overall, participants answered 78.65% of trials correctly ( $SD = 41$ ; 81.23% on Day 1 and 76.05% on Day 2). Their performance was above chance level (16.67% because there was one correct response out of six presented options), as confirmed by a one-sample  $t$  test ( $t = 82.47, p < .001$ ). The overall accuracy rate was calculated across all trials, including those in which participants made within-university errors. This type of errors indicate that participants could not remember which exact statement corresponded to the speaker, but that they did correctly infer speakers' affiliation. When recoding within-university errors as correct responses (i.e., correctly identifying the university the speaker was from), the average of correct responses increased to 96.51% ( $SD = 18.35$ ) for Day 1 and to 97.1% ( $SD = 16.79$ ) on Day 2. These analyses confirmed participants successfully learned which speakers belonged to which academic affiliation.

### Word-Learning Task

The main task in the current study was the word-learning task. After an exposure phase, participants were tested with a memory recognition test. In this test, participants were presented with trials in which the speaker-label pairings either matched or mismatched with what they had seen in the learning phase. That is, the label had or had not been produced by the displayed speaker. We carried out two main analyses over the memory data. We looked first at the matching trials and then at the mismatching trials.

**Matching trials.** To test whether source memory was better for words learned from in-group than from out-group speakers, we ran a logistic mixed effects model with accuracy per trial as the dependent measure and fixed effects for group membership (in-group [reference level] vs. out-group), in-group bias ( $z$ -scored continuous measure), and their interaction. We added per participant and per items random intercepts and by participant slope for group membership.

Overall, participants' accuracy in the matching trials was 68.11% ( $SD = 46.62$ ) and above chance level, as confirmed by a one-sample  $t$  test (i.e., 50%;  $t = 14.35, p < .0001$ ). Neither group membership (out-group:  $M = 70.38, SD = 45.69$ ; in-group:  $M =$

65.84,  $SD = 47.46$ ;  $\beta = 0.19, SE = 0.13, z = 1.45, p = .15$ ) nor in-group bias ( $\beta = -0.02, SE = 0.33, z = -0.07, p = .95$ ) significantly predicted accuracy at the reference level. The Group Membership  $\times$  In-Group Bias interaction, however, was marginally significant ( $\beta = -0.72, SE = 0.39, z = -1.84, p = .07$ ).<sup>3</sup> Follow-up simple effects analyses confirmed that in-group bias affected the accuracy of participant responses differently for the in-group and the out-group matching trials. Although the simple effect of in-group bias did not predict response accuracy for in-group matches ( $\beta = -0.02, SE = 0.33, z = -0.07, p = .95$ ), there was a significant negative predictor of participants' performance for the out-group matches ( $\beta = -0.75, SE = 0.27, z = -2.73, p = .006$ ). In agreement with our predictions, the larger participants' in-group bias, the less accurate they were at identifying out-group speakers (see Figure 3).

**Mismatching trials.** To test whether speaker group membership influenced the level of detail in the speaker-specific information encoded with the new words, we analyzed accuracy on mismatching trials. By looking at participants' performance on within-university mismatching trials, where labels were paired with incorrect speakers who belonged to the same university as the original speaker, we could test whether the source-related information for novel words was speaker-specific (participants should have rejected the wrong source) or group-specific (participants would have incorrectly accepted the wrong source). Our hypothesis was that people encode more speaker-specific information from in-group members. We therefore predicted greater confusion among out-group speakers than in-group speakers in the within-university mismatching trials. We also predicted that this difference in accuracy would depend on in-group bias, such that greater in-group bias participants exhibited, the greater difference they should show between their accuracy on in-group versus out-group trials. Conversely, in between-university mismatches (i.e., where an in-group label was shown with out-group members, and vice versa) no differences due to our manipulation were expected.

To test these hypotheses, we ran a logistic mixed model analysis with fixed effects for mismatch type (within-university [reference level] vs. between-university), group membership (in-group [reference level] vs. out-group), in-group bias ( $z$ -scored continuous measure), and their interaction terms. We added per participant and per item random intercepts and by participant slopes for group membership and mismatch type. Overall, participants' accuracy on mismatching trials was 57.97% ( $SD = 49.36$ ) and above chance level (i.e., 50%;  $t = 12.88, p < .0001$ ).

Generally, participants were less accurate in the between-university mismatches than in the within-university mismatches ( $\beta = -0.19, SE = 0.08, z = -2.27, p = .02$ ;  $M = 56.39\%, SD = 49.6$  and  $M = 60.08\%, SD = 48.98$ , respectively). Participants' performance was also predicted by a significant interaction of group membership with in-group bias ( $\beta = -0.82, SE = 0.28, z = -2.96, p = .003$ ), suggesting that participants with different

<sup>3</sup> Visual inspection of the figure might imply that participants with low in-group bias (and potentially out-group bias) exhibit higher accuracy with out-group members. An analysis that focuses only on participants in the bottom quartile of in-group bias (all showing no bias or reversed bias on the in-group bias measure) does not show any significant effect though, and an analysis that excludes the bottom quartile of participants is qualitatively similar to the reported analysis.




Figure 3. Accuracy as a function of group membership and In-group Bias (centered; the black, dotted line represents in-group bias values equal to zero). Error bars represent standard errors. See the online article for the color version of this figure.

strength of in-group bias were differently affected by speaker group membership. Crucially, both effects were further modulated by the expected three-way interaction between mismatch type, group membership, and in-group bias ( $\beta = 0.95$ ,  $SE = 0.37$ ,  $z = 2.61$ ,  $p = .009$ ). The remaining main effects and interactions were not significant ( $ps > .57$ ).

Before further evaluating the three-way interaction, we tested its reliability. We ran a series of split-half analyses by randomly selecting subsets of participants composed of half the number of the original participants ( $n = 30$ ) over which we ran the same mixed-effect model. The interaction pattern was replicated in 15 out of 15 subsets, with the effect being significant in 13 subsets and marginally significant in two subsets. This procedure suggests that the significance of the interaction was well supported by the data and that the analysis was sufficiently powered such that even random subsets of half the sample led to comparable results.

To unpack the interaction, we ran separate analyses for each mismatch type condition (see Figure 4). The analysis of the between-university trials included fixed effects for group membership, in-group bias, their interaction terms, and per item and per participant random slopes and by participant slopes for group membership. Results did not show any significant predictor ( $ps > .42$ ; see Figure 4, left panel).

The analysis of the within-university mismatches included group membership, in-group bias and their interaction as fixed effects, and per participant and per item random intercepts and by participant slopes for group membership. In agreement with our predictions, the results showed a significant Group Membership  $\times$  In-Group Bias interaction ( $\beta = -0.83$ ,  $SE = 0.28$ ,  $z = -3.00$ ,  $p = .003$ ). Post hoc analyses revealed a cross-over interaction, but nonsignificant simple effects. The stronger the in-group bias, the more participants showed a difference between their ability to distinguish in-group members compared with out-group members (see Figure 4, right panel). These patterns show greater receptiveness to individual in-group speakers over out-group speakers and are in line with our predictions and with previous research showing that in-group bias stems from greater individuation of in-group members while viewing out-group members as less distinct (e.g., Hugenberg et al., 2010).

Next, we used signal detection theory to capture detection sensitivity ( $d'$ ) and decision criterion ( $C$ ). These analyses allow us to test whether participants' sensitivity and criterion during decision making differed for in-group versus out-group related decisions. We calculated two  $d'$  values and two  $C$  values per participant for in-group and out-group trials separately. In order to generate values that reflected participants' decisions to purely in-group or out-group trials,  $d'$  and  $C$  values were calculated from participants' performance in matching trials (i.e., hit rates) and within-university mismatching trials (i.e., false alarm rates). Between-university mismatches were not considered for these analyses because they contained an element (either label or speaker) from each group and were therefore not purely in-group or out-group related. We ran two linear mixed-effect models with either  $d'$  or  $C$  values as the dependent variable and group membership (in-group [reference level] vs. out-group), in-group bias, and their interaction as fixed effects. The models included per participant random intercepts.

The model which explored the relationship between individual  $d'$  and the independent variables showed no significant main effects for group membership ( $\beta = 0.13$ ,  $SE = 0.1$ ,  $t = 1.42$ ,  $p = .16$ ) or in-group bias ( $\beta = -0.01$ ,  $SE = 0.30$ ,  $t = -0.04$ ,  $p = .97$ ). However, there was a significant interaction between these two factors ( $\beta = -0.63$ ,  $SE = 0.29$ ,  $t = -2.13$ ,  $p = .04$ ; see Figure 5). Simple effects analyses revealed that the larger the participants' in-group bias, the less sensitive they were to out-group speakers ( $\beta = -0.64$ ,  $SE = 0.30$ ,  $t = -2.12$ ,  $p = .04$ ), whereas sensitivity to in-group speakers did not change as a function of in-group bias ( $\beta = -0.01$ ,  $SE = 0.30$ ,  $t < 1$ ,  $p = .97$ ).

The model exploring  $C$  values did not lead to any significant results ( $p > .23$ ) suggesting that participants did not differ in decision criterion as a function of group membership or in-group bias. This means that participants' tendencies to produce positive or negative responses were not affected by their social biases. That is, they were not more likely to say "yes" when seeing in-group speakers than when seeing out-group speakers, or vice versa.

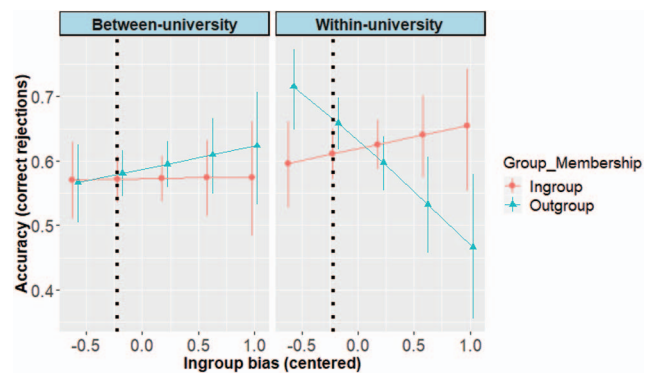


Figure 4. Accuracy for between-university mismatches (left panel) and within-university mismatches (right panel) as a function of in-group bias (centered; the black, dotted line represents in-group bias values equal to zero) and the group membership of the speaker from whom the word was originally. Error bars represent standard error. See the online article for the color version of this figure.

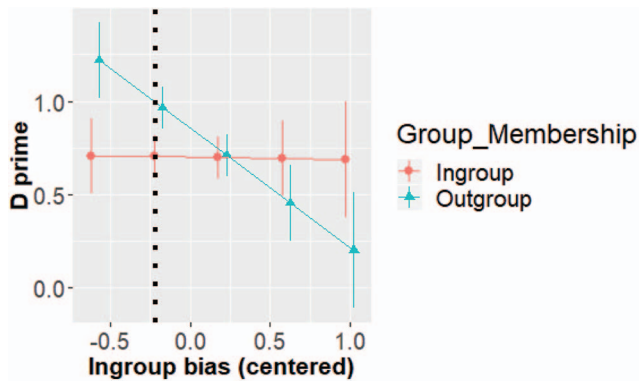


Figure 5.  $D'$  prime as a function of group membership and in-group bias (centered; the black dotted line represents in-group bias values equal to zero). Error bars represent standard errors. See the online article for the color version of this figure.

## Discussion

In a novel word-learning study, we tested whether listeners encoded speakers' social identity, such as their group membership status and their personal identity, and whether the level of detail of such information depended on the group membership status of the speakers and the individual in-group bias of the listeners.

In the word-learning task, participants had to identify matching speaker-label pairs in a forced-choice task (i.e., yes/no). Our results show that the level of detail with which participants encoded the identity of the speakers varied as a function of speakers' group membership status and participants' own in-group bias. Specifically, the analyses indicated that more in-group biased individuals exhibited greater confusion for out-group over in-group members. Our results are consistent with previous findings that showed greater individual-specific encoding for in-group than for out-group members (e.g., Frable & Bem, 1985; Greenstein et al., 2016; Hugenberg et al., 2010; Wilder, 1990) and extend them to the case of word learning.

To better understand participants' behavior, we also analyzed their decision criterion and detection sensitivity for in-group and out-group speakers. Again, these analyses showed results that were comparable to those obtained in the previous analyses: the larger the in-group bias, the greater participants' detection sensitivity to in-group as compared with out-group speakers. That is, across all analyses (for matching trials, for mismatching trials, and  $d'$ ), we found consistent results that demonstrate that speaker group membership and learners' individual in-group biases modulated the level of detail of speaker-related information contained in the representations of novel words.

One interesting aspect of our findings was that though we expected individual in-group biases to simultaneously increase source detection for the in-group and decrease it for the out-group, as Greenstein et al.'s (2016) results suggested, our data indicate mostly the latter phenomenon. The more in-group biased participants were, the less sensitive they were to individual out-group speakers, while their sensitivity to individual in-group speakers did not change as a function of in-group bias. There can be two potential explanations for these results. The most straightforward one relates to the possibility that in-group

individuation, that is, the encoding of speaker-specific detail for in-group members, might be the default mechanism, regardless of individual in-group bias, and it is whether people also encode specific information about out-group speakers that depends on learners' own in-group bias. In the study by Greenstein and colleagues (2016), participants were better at identifying in-group sources as compared with out-group sources. However, they did not include an individual measure of in-group bias in their analysis as a modulator of the effects of interest and, on the contrary, excluded participants who identified themselves with the in-group and out-group membership to a similar extent. Therefore, a direct comparison between our results and the results of their study is difficult to make.

An alternative explanation of this unexpected pattern is that the implicit measure we used to quantify in-group bias was better at capturing individual differences related to how well people attended to out-group members than differences related to the in-group. The distribution of accuracy rates in the perceptual matching task (see Appendix B) shows that participants' performance for the in-group condition was almost at ceiling. Therefore, the variation in the individual difference measure we used was mostly driven by differences in the out-group condition. Future research should use a more taxing task that limits ceiling effects to determine whether individual measures of in-group bias would then be seen to also modulate participants' performance in in-group within-university trials. Although the in-group bias did not directly modulate the level of detail encoded in representations of in-group labels, the patterns of results confirmed our predictions.

Taking a broader scope, our results address a primary question in the field of psycholinguistics, namely, how words are represented in the mental lexicon. A growing body of literature suggests that social information, such as speaker identity, is stored together with the lexical information (e.g., Goldinger, 2007; Hay, Nolan, & Drager, 2006; Nielsen, 2011; also see Drager & Kirtley, 2016, for a review). Importantly, all the existing models assume that the same information will be stored to the same degree regardless of who the speaker is. In contrast, we show that the type of social information that is encoded (group level or individual level) and its robustness depend on speakers' identity (e.g., in-group vs. out-group member) and the listeners' social biases. These findings, thus, call for the modification of existing models of language processing.

The nature of lexical representations is of great interest partly because it has implications for all aspects of language use—language processing, language production, and language learning. Our results therefore can have implications across the field of language use. We next illustrate that with a couple of examples. There is great interest in how learners learn from input and what should be its distributional characteristics to optimize learning. Studies and models on the topic often contrast intraindividual and interindividual variability. Our results, however, suggest that learners would be able to isolate intraindividual and interindividual variability for input provided by in-group speakers, but less so for input provided by out-group speakers. These differences could lead learners to draw different inferences from the same input, depending on who provides it.



Another hot topic in psycholinguistic is the role of prediction in language learning and processing. Listeners are assumed to constantly predict utterance continuations, and use the prediction error (the difference between the actual input and the input they predicted) to update their representations, and thus learn (e.g., Chang, Dell, & Bock, 2006; Garrod, Gambi, & Pickering, 2014; Kuperberg & Jaeger, 2016). Current research attempts to understand what information these predictions are based on and to which cases learned information is generalized (Kleinschmidt & Jaeger, 2015). Our results suggest that the identity of the speaker could influence the granularity of the predictions that listeners make, as well as how widely they would generalize the patterns they encounter. For example, listeners might be able to use speaker-specific information when predicting the speech of an in-group member, but only be able to use group-level information when predicting the speech of out-group members. Furthermore, poor individuation of out-group sources in people with high in-group bias may lead them to generalize information learned from a specific out-group member to all the out-group members. Imagine, for instance, what might happen if a politician from a disliked party labels an event in a particular way (say as “fake news” or “national triumph”). A member of an opposing party may overgeneralize the usage of this label (“fake news”) to all members of that disliked party, thereby, perhaps, feeding unfounded beliefs about the views of the members of the disliked party. Further research should explore the consequences of different speaker identities and listeners’ biases on language processing and learning as well as how in-group bias modulates the role of social factors in learning, processing, and transmitting language.

## References

- Aron, A., Aron, E. N., Tudor, M., & Nelson, G. (1991). Close relationships as including other in the self. *Journal of Personality and Social Psychology*, *60*, 241–253. <http://dx.doi.org/10.1037/0022-3514.60.2.241>
- Bargh, J. A., Schwader, K. L., Hailey, S. E., Dyer, R. L., & Boothby, E. J. (2012). Automaticity in social-cognitive processes. *Trends in Cognitive Sciences*, *16*, 593–605. <http://dx.doi.org/10.1016/j.tics.2012.10.002>
- Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). *Parsimonious mixed models*. Retrieved from <https://arxiv.org/pdf/1506.04967.pdf>
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). lme4: Linear mixed-effects models using Eigen and S4. *R package Version 1*(7), 1–23.
- Cadinu, M. R., & Rothbart, M. (1996). Self-anchoring and differentiation processes in the minimal group setting. *Journal of Personality and Social Psychology*, *70*, 661–677. <http://dx.doi.org/10.1037/0022-3514.70.4.661>
- Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, *113*, 234.
- Drager, K., & Kirtley, J. (2016). Awareness, salience, and stereotypes in exemplar-based models of speech production and perception. *Awareness and Control in Sociolinguistic Research*, *2016*, 1–24.
- Enock, F., Sui, J., Hewstone, M., & Humphreys, G. W. (2018). Self and team prioritisation effects in perceptual matching: Evidence for a shared representation. *Acta Psychologica*, *182*, 107–118. <http://dx.doi.org/10.1016/j.actpsy.2017.11.011>
- Frable, D. E., & Bem, S. L. (1985). If you are gender schematic, all members of the opposite sex look alike. *Journal of Personality and Social Psychology*, *49*, 459–468. <http://dx.doi.org/10.1037/0022-3514.49.2.459>
- Garrod, S., Gambi, C., & Pickering, M. J. (2014). Prediction at all levels: Forward model predictions can enhance comprehension. *Language, Cognition and Neuroscience*, *29*(1), 46–48.
- Goldinger, S. D. (2007). A complementary-systems approach to abstract and episodic speech perception. *Proceedings of the 16th International Congress of Phonetic Sciences* (pp. 49–54). Retrieved from <http://www.public.asu.edu/~sgolding/docs/pubs/other/GoldingerICPhS.pdf>
- Greenstein, M., Franklin, N., & Klug, J. (2016). In-group versus out-group source memory. *Experimental psychology*, *63*, 150–158. <http://dx.doi.org/10.1027/1618-3169/a000322>
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*, 4–27. <http://dx.doi.org/10.1037/0033-295X.102.1.4>
- Hay, J., Nolan, A., & Drager, K. (2006). From fush to feesh: Exemplar priming in speech perception. *Linguistic Review*, *23*, 351–379. <http://dx.doi.org/10.1515/TLR.2006.014>
- Hay, J., Warren, P., & Drager, K. (2006). Factors influencing speech perception in the context of a merger-in-progress. *Journal of Phonetics*, *34*, 458–484. <http://dx.doi.org/10.1016/j.wocn.2005.10.001>
- Hugenberg, K., Young, S. G., Bernstein, M. J., & Sacco, D. F. (2010). The categorization-individuation model: An integrative account of the other-race recognition deficit. *Psychological Review*, *117*, 1168–1187. <http://dx.doi.org/10.1037/a0020463>
- Johnson, K., Strand, E. A., & D’Imperio, M. (1999). Auditory–visual integration of talker gender in vowel perception. *Journal of Phonetics*, *27*, 359–384. <http://dx.doi.org/10.1006/jpho.1999.0100>
- Jones, C. R., & Fazio, R. H. (2010). Person categorization and automatic racial stereotyping effects on weapon identification. *Personality and Social Psychology Bulletin*, *36*, 1073–1085. <http://dx.doi.org/10.1177/0146167210375817>
- Judd, C. M., & Park, B. (1988). Out-group homogeneity: Judgments of variability at the individual and group levels. *Journal of Personality and Social Psychology*, *54*, 778–788. <http://dx.doi.org/10.1037/0022-3514.54.5.778>
- Keuleers, E., Brysbaert, M., & New, B. (2010). SUBTLEX-NL: A new measure for Dutch word frequency based on film subtitles. *Behavior Research Methods*, *42*, 643–650. <http://dx.doi.org/10.3758/BRM.42.3.643>
- Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children’s selective trust in native-accented speakers. *Developmental Science*, *14*, 106–111. <http://dx.doi.org/10.1111/j.1467-7687.2010.00965.x>
- Kleinschmidt, D. F., & Jaeger, T. F. (2015). Robust speech perception: Recognize the familiar, generalize to the similar, and adapt to the novel. *Psychological Review*, *122*, 148.
- Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience*, *31*(1), 32–59.
- Lev-Ari, S., & Keysar, B. (2012). Less-detailed representation of non-native language: Why non-native speakers’ stories seem more vague. *Discourse Processes*, *49*, 523–538. <http://dx.doi.org/10.1080/0163853X.2012.698493>
- LimeSurvey Project Team. (2015). *LimeSurvey: An open source survey tool*.
- Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, *47*, 1122–1135. <http://dx.doi.org/10.3758/s13428-014-0532-5>
- Martin, C. D., Garcia, X., Potter, D., Melinger, A., & Costa, A. (2016). Holiday or vacation? The processing of variation in vocabulary across dialects. *Language, Cognition and Neuroscience*, *31*, 375–390. <http://dx.doi.org/10.1080/23273798.2015.1100750>

- Meissner, C. A., Brigham, J. C., & Butz, D. A. (2005). Memory for own-and other-race faces: A dual-process approach. *Applied Cognitive Psychology, 19*, 545–567. <http://dx.doi.org/10.1002/acp.1097>
- Moradi, Z., Sui, J., Hewstone, M., & Humphreys, G. W. (2015). In-group modulation of perceptual matching. *Psychonomic Bulletin & Review, 22*, 1255–1277. <http://dx.doi.org/10.3758/s13423-014-0798-8>
- Münster, K., & Knoeferle, P. (2018). Extending situated language comprehension (accounts) with speaker and comprehender characteristics: Toward socially situated interpretation. *Frontiers in Psychology, 8*, 2267. <http://dx.doi.org/10.3389/fpsyg.2017.02267>
- Niedzielski, N. (1999). The effect of social information on the perception of sociolinguistic variables. *Journal of Language and Social Psychology, 18*, 62–85. <http://dx.doi.org/10.1177/0261927X99018001005>
- Nielsen, K. (2011). Specificity and abstractness of VOT imitation. *Journal of Phonetics, 39*, 132–142. <http://dx.doi.org/10.1016/j.wocn.2010.12.007>
- R Core Team. (2013). *R: A language and environment for statistical computing*. Retrieved from <http://www.R-project.org/>
- Stolte, M., Humphreys, G., Yankouskaya, A., & Sui, J. (2017). Dissociating biases towards the self and positive emotion. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 70*, 1011–1022. <http://dx.doi.org/10.1080/17470218.2015.1101477>
- Sui, J., & Humphreys, G. W. (2015). The integrative self: How self-reference integrates perception and memory. *Trends in Cognitive Sciences, 19*, 719–728. <http://dx.doi.org/10.1016/j.tics.2015.08.015>
- Sumner, M., Kim, S. K., King, E., & McGowan, K. B. (2014). The socially weighted encoding of spoken words: A dual-route approach to speech perception. *Frontiers in Psychology, 4*, 1015. <http://dx.doi.org/10.3389/fpsyg.2013.01015>
- Symons, C. S., & Johnson, B. T. (1997). The self-reference effect in memory: A meta-analysis. *Psychological Bulletin, 121*, 371–394. <http://dx.doi.org/10.1037/0033-2909.121.3.371>
- Van Bavel, J. J., & Cunningham, W. A. (2012). A social identity approach to person memory: Group membership, collective identification, and social role shape attention and memory. *Personality and Social Psychology Bulletin, 38*, 1566–1578. <http://dx.doi.org/10.1177/0146167212455829>
- Van Bavel, J. J., Packer, D. J., & Cunningham, W. A. (2008). The neural substrates of in-group bias: A functional magnetic resonance imaging investigation. *Psychological Science, 19*, 1131–1139. <http://dx.doi.org/10.1111/j.1467-9280.2008.02214.x>
- Van Berkum, J. J., van den Brink, D., Tesink, C. M., Kos, M., & Hagoort, P. (2008). The neural integration of speaker and message. *Journal of Cognitive Neuroscience, 20*, 580–591. <http://dx.doi.org/10.1162/jocn.2008.20054>
- Walker, A., & Hay, J. (2011). Congruence between ‘word age’ and ‘voice age’ facilitates lexical access. *Laboratory Phonology, 2*, 219–237. <http://dx.doi.org/10.1515/labphon.2011.007>
- Wilder, D. A. (1990). Some determinants of the persuasive power of in-groups and out-groups: Organization of information and attribution of independence. *Journal of Personality and Social Psychology, 59*, 1202–1213. <http://dx.doi.org/10.1037/0022-3514.59.6.1202>

(Appendices follow)

## Appendix A

### Material Norming

#### Norming of Statements About Student Life

We generated 32 pairs of sentences about student activities and student life such that each pair differed only in whether the location or name of the activity implied that it took place at Radboud or Groningen University (e.g., “I got to know most of my friends during the activities organized for the Radboud/ Groningen introductory week”). The 64 sentences were divided into two lists and participants were randomly assigned to one of the two lists. In this way, each participant only saw one sentence per pair. Twenty Dutch volunteers studying at Radboud University participated in an online survey. Their task was to indicate for each sentence how representative it was of student life at (a) Radboud University or (b) another university. Twelve pairs of sentences were selected, such that at least 70% of participants agreed for each of the sentences that it clearly represented either Radboud or another university.

#### Norming of Labels and Images for Word Learning

The aim of the norming study was to select the object images and labels for the word-learning task, such that target objects would have at least two equally good labels. To do so, a two-phase norming study was carried out. In each phase, a different group of Dutch university students completed an online survey hosted on LimeSurvey (LimeSurvey Project Team, 2015). Participants received either university-credit or monetary compensation.

#### First Phase: Written Picture Naming

Eighteen volunteers (12 women; age  $M = 20.89$ ,  $SD = 1.71$ ) participated. They saw 120 images of uncommon gadgets collected from the Internet (e.g., a corn peeler) and provided potential labels for them. After data collection, three Dutch native speakers corrected misspelt forms and excluded redundant forms, as well as regionalisms, non-Dutch forms, and inaccurate names (i.e., if the object was confused with something else). The majority of the

elicited labels were nominal compounds with a semantic head (on the right, as common in Dutch), which refers to the action or the purpose of the gadget (e.g., peeler), and another noun (or compound) referring to the object that receives the action (e.g., corn). To ensure competing labels were comparable, labels not following this structure were excluded at this stage. At the end of this process, 116 out of the initial 120 items were selected. Each gadget elicited between three and fourteen alternative labels.

#### Second Phase: Goodness-of-Fit Ratings

During the second phase, 40 university students (33 women; age  $M = 19.67$ ,  $SD = 1.96$ ) were recruited for a new online survey. The aim of the survey was to test how well the given labels described the gadgets. The 116 objects selected in the previous phase were pre-randomized and divided into three blocks. Six different versions of the survey were created in order to counter-balance the presentation order of the blocks. Each object was individually presented along with all the given labels. Participants had to rate how well each label described the item on a seven-point scale (1 = *inappropriate* [i.e., *ongeschikt*], 7 = *perfect name* [i.e., *perfecte benaming*]). Participants were instructed that several labels for a specific item could receive the same value if they were equally good.

Next, paired  $t$  tests were performed on the ratings of all possible pairwise combinations of labels for the same objects. Once we identified pairs which did not statistically differ in ratings ( $p > .05$ ), we compared the frequency of their nominal constituents using SUBTLEX-NL (Keuleers, Brysbaert, & New, 2010) to ensure they do not significantly differ in frequency either. For fillers, we selected items with one label that was rated as significantly better than all the other options (paired  $t$  tests,  $ps < .05$ ).

Forty-one target and 43 filler items fulfilled the requirements. To fulfill the requirements of another study that used the same stimuli, we also ensured that items' familiarity is not judged to depend on level of education. Twelve target items and 12 fillers met these criteria.

(Appendices continue)

## Appendix B

### Perceptual Matching Task Results

#### Analyses Over RTs

Prior to analyses, trials with incorrect responses or with RTs faster than 200 ms or slower than 2,100 ms were excluded. For these confirmatory analyses, we selected only matching trials (i.e., in which the logo of the university was displayed with the associated geometrical shape) which referred to the in-group university and the out-group university used in the study (i.e., Groningen University). We then performed an outlier removal procedure by removing trials with RTs 2.5 standard deviations or higher from the mean per condition

and per participant. The resulting dataset was analyzed using linear mixed-effect model in which  $\log(10)$ -transformed RTs were predicted by the fixed effect for group membership (in-group [reference level] vs. out-group; see Figure B1). We added per participant random intercept and by participant random slope for group membership. Results confirmed the usual trend for this task: Participants were faster at recognizing in-group-related associations ( $M = 672$  ms,  $SD = 150$  ms) than out-group-related associations ( $M = 735$  ms,  $SD = 176$  ms;  $\beta = 0.08$ ,  $SE = 0.01$ ,  $t = 7.37$ ,  $p < .0001$ ).




Figure B1. Density plot showing response time (RT) means per group membership. See the online article for the color version of this figure.

(Appendices continue)




Figure B2. Density plot showing accuracy rates per group membership. See the online article for the color version of this figure.

### Analyses Over Accuracy

Prior to analyses, we selected only matching trials (i.e., in which the logo of the university was displayed with the associated geometrical shape) which referred to the in-group university and the out-group university used in the study (i.e., Groningen University). Responses' accuracy on a single trial were analyzed using logistic mixed-effect model with a fixed effect for group membership (in-group [reference level] vs. out-group; see Figure B2). We added per participant random intercept and by participant random slope for group membership. Results confirmed the usual trend for this task: Participants were better at recognizing in-group-related associations ( $M = 96\%$ ,  $SD = 20\%$ ) than out-group-related associations ( $M = 89\%$ ,  $SD = 31$ ;  $\beta = -1.18$ ,  $SE = 0.29$ ,  $t = -4.73$ ,  $p < .0001$ ).

### Distribution of the Individual In-Group Bias

The in-group bias measure is  $z$  scored, therefore the values around 0 do not correspond to actual 0 (see Figure B3).




Figure B3. Distribution of the individual in-group bias measure.

Received December 19, 2018  
 Revision received August 2, 2019  
 Accepted August 13, 2019 ■