The Role of Feedback in Aligning Perspectives in Referential Communication

Timothy M. Gann (tgann@ucmerced.edu)

Sierra Nevada Research Institute and Cognitive and Information Sciences, University of California, Merced

5200 North Lake Rd. Merced, CA 95343 USA

Abstract

Successful dialogue frequently requires that interlocutors construct and align their conceptualizations of referents. This study presents data from a referential communication experiment the manipulates contextual factors such as the availability of feedback and role constancy in order to investigate how conversational partners reconcile their perspectives in the face of mutual uncertainty about what constitutes common ground. The results show that speakers tend to incorporate information about the addressee's perspective, and that this information tends to come through direct feedback rather than through indirect channels such as turn-taking.

Keywords: Psycholinguistics; language production; audience design; perspective taking.

Introduction

When deciding how to describe an object, people must first categorize it in a way that is useful to themselves them and ultimately for their partners. Sometimes, the taxonomic, established categorization is not well suited to the task and interlocutors may be forced to develop an ad hoc, or socially emergent category that is more suited to achieving their current conversational goals (Barr & Kronmüller, 2007; Barsalou, 1991). For example, when two people are discussing where they wish to go out to eat, they may form a category of acceptable restaurants on the spot. This category emerges over the conversation and is neither pre-existing in memory nor particularly enduring (unless they frequently go out together). This socially emergent conceptualization then becomes the foundation for future exchanges. What this means is that early on in a conversation, taxonomic categories are going to be activated more often and attract more attention. However, as time goes by and speakers learn to put more emphasis on shared information, and other types of information are found to be in conflict with the goal of successfully referring to an object, the socially constructed categories should be used more often. These socially emergent categories would probably demonstrate some degree of being person-specific (Horton & Gerrig, 2005), but that would not preclude their activation and propagation to new conversational partners if they were useful (Garrod & Doherty, 1994).

The development of new conceptualizations during conversation should also lead to changes in how attention is allocated. Rehder and Hoffman (2005) conducted an experiment in which participants had to do a category learning task while having their eyes tracked. At the beginning of the experiment, participants would overtly attend to each of the features in the display and consider them all before making a decision. By the end of the experiment, participants learned to overtly attend only to those features that were necessary for the task of choosing which category was represented by the display.

Similarly, in a referential communication context involving perspective taking, as time passes and socially emergent categories become more dominant, attention should shift away from privileged competitors. By the end of a conversation they should be attended to about as little as an object unrelated to the target. This shift should also happen in situations in which the speaker does not initially have access to what the listener knows. As evidence about listener's knowledge, or lack of knowledge, becomes available they should gradually begin to focus on the stimuli that are most consistent with the listener's perspective.

The question then becomes, what factors influence the rate of learning during conversation? One possibility is the potential for feedback. Closely related to feedback are the roles of the interlocutors: Are they in a didactic situation in which the roles are fixed, or is it a fully interactive dialogue in which there is turn taking? Unrestricted feedback can come in many forms. First, speakers can receive pragmatic feedback in the form of knowledge about the success or failure of the listener in establishing joint attention to a target object. It is implicitly accepted that pragmatic feedback is necessary for learning to happen, with positive feedback leading to a reinforcement of the speaker's current conceptualization, and negative feedback encouraging re-conceptualization. Second, the listener can give signals to the speaker, who can make inferences about the listeners confidence about an interpretation. For example, Barr (2003) found that listeners will use speakers' confidence as a cue that the current referent is either a typical or atypical member of a category. In principle, speakers should also similarly use the listener's confidence in their selection of a target as a cue to how well their message was received.

Finally, listeners can give direct verbal feedback to the speaker. Schober and Clark (1989) conducted an experiment in which a participant acted as a non-interactive observer to a conversation. This participant attempted to do the same task the actual listener was doing, but without the ability to give feedback to the speaker.They found that the observers never aligned as closely to the speaker because they were unable to provide feedback and have the speaker adjust to or correct their interpretations. The lack of verbal feedback can also make the speakers prone to give more information as a hedge against possible misunderstandings (Gann & Barr, 2012). For example, Krauss and Weinheimer (1966) found that when listeners were unable to give feedback, speakers were not inclined to simplify their utterances.Instead they kept utterances long so they could minimize the possibility of giving too little information. While it is possible for alignment to happen without it, direct verbal feedback likely serves to speed up the

process. Without it, speakers may not be able to easily settle on a shared perspective and fully acquire a socially emergent category that corresponds to the intersection of their knowledge with their partner's.

Role switching may serve a similar function. In a natural conversation, both participants are speakers and listeners. These participants both offer their respective conceptualizations of the referents they are bringing to each others attention, and both are able to interactively sculpt each others understanding. Even simply priming from ones partners utterances may serve to make dimensions more salient than the other might be from an individuals egocentric perspective (Pickering & Garrod, 2004). In addition, in situations where one person is the primary speaker, there might be less incentive for that individual to align, and thus, the speaker might rely on the listener to align to align to speaker knowledge.

In many referential communication experiments that manipulate common ground, knowledge is asymmetric in one direction. Specifically, the speaker knows more than the listener, or the listener knows more than the speaker. This experiment focuses on a case in which the speaker and listener both must learn to account for this difference in knowledge. The major question is whether or not conversational partners in such a situation tend to reconceptualize the stimuli in a way that is consistent with a *union* of their perspectives, thus taking into account both their sets of privileged information, or an *intersection* of their perspectives, focusing only on that information that is mutually shared. If in a particular circumstance, speakers tend towards reconceptualizing the referent in terms of an intersection between their knowledge and their partner's, will speakers then begin to ignore competitors that are not relevant to that overlap? Additionally, is their tendency to adopt one scheme or another influenced by the form of feedback they experience: either explicit feedback, or implicit feedback through role switching? Direct feedback should allow addressees to clearly indicate to the speaker what their informational needs are, and should thus be associated with rapid alignment between the interlocutors. Learning about a partner's perspective in the role switching condition presumably relies on more indirect learning, and may not have the same impact as direct feedback.

Method

Participants

The participants were all native English speakers, drawn from a pool of undergraduates at the University of California, Riverside. A total of 32 students participated in the study, forming 16 dyads. For half of the dyads, the participants were assigned to either the role of the director or the matcher. The other half of the dyads did not have fixed roles, but rather switched roles after each trial of the experiment.

Apparatus

An ISCAN ETL-400 table-mounted eye-tracker, sampling at a rate of 60 Hz, was used to track the director's gaze through-

out the experiment. The experimental stimuli were displayed for the participants on two LCD monitors. Each participant's screen was only visible to themselves. The director was given a set of headphones through which instructions could be given without the knowledge of the matcher.

Materials

The experimental stimuli consisted of 32 sets of five items. Each stimulus item was a 300 x 300 pixel colored bitmap of an object placed on a black background. The sets were constructed such that there was a target object, three competitor objects that are typically referred to by the same name, and an unrelated filler item. The competitor objects contrasted with the target object along the dimensions of size, openness, and material (or color). The target was always consistently open and larger; each of those two dimensions were visible to only one of the participants. Material differed from object to object; this dimension was always shared between the participants. So for example, one participant may have seen a target that was a large open trashcan, a smaller identical competitor, a plastic trashcan, and an unrelated fourth object. The other participant would see the same target, the plastic trashcan, the same unrelated competitor, but would have a closed competitor that is otherwise as physically similar to the target as possible. See Figure 1 for an example.

Procedure

Participants were assigned to either the role of director or matcher by having one of the participants choose a face down card at random that specified a role. If the dyad was participating in the condition in which their roles would switch throughout the experiment, it was explained to them that they were picking an initial role. For the switching participants, the participant who was eye-tracked was always the initial director.

The experiment was described to the participants as a simple communication game that they would be playing together in which the director would be describing a target object such that the matcher could pick out which object on their screen. They were shown a set of example stimuli demonstrating the perspective of both partners. The sample stimuli did not include competitors that varied along the privileged dimensions, only one that contrasted along the shared dimension. The participants were instructed that the target would be indicated to the current director with a red border. The matcher would listen to the director's description and select the object being referring to on their screen by selecting it with their mouse.

Design

The experimental variables manipulated in the design of the experiment included whether the partners switched roles or not, and level of feedback (2x2 between-subjects design). The dependent variables are accuracy (whether the addressee chooses the right target), speech onset time, use of an adjective corresponding to the speaker's privileged perspec-

Figure 1: Each of the displays above corresponds to the perspective of one participant. In this example, the target would be the closed metal trashcan in the upper left corners. The display on the left contains a competitor based on openness. The display on the right shows a competitor based primarily on material/color. Each display has a competitor in common based on size.

tive/dimension, use of an adjective corresponding to the addressee's privileged perspective, use of an adjective corresponding to the shared perspective, and the proportion of time the speaker spent gazing at the shared and privileged competitors prior to speech onset.

Analysis

The director's speech was recorded, transcribed, and coded for accuracy, speech onset, and the use of modifiers matching each of the three dimensions on which the competitors contrasted with the target. The eye-tracking data were coded as number of frames (sampled at 60 hz) that the director spent looking at each object prior to speech onset. Eye-tracking data for trials in which the participants switched roles was removed from the data set due to the inability to make an apples to apples comparison across the condition due to the current director only contributing eye-data for half the trials in the switching condition.

The analyses were conducted using linear mixed effects models with random effects included for subjects and items. Models were fit within R using the lmer function within the package lme4 version 0.999375-39 (Bates & Maechler, 2010). Appropriate link functions were chosen depending on the distribution of the outcome variable of interest. The significance of the fixed effects were assessed using a χ^2 model comparisons approach in which a model without the variable of interest is compared with the full model (Barr, Levy, Scheepers, & Tily, 2013).

Results

Thirteen of 512 (2.5%) trials were dropped from the analyses because of experimenter and procedural errors. Across all conditions the listener correctly identified the referent in 72.55% of trials (Figure 2 shows accuracy plotted by condition). There were significant effects of feedback ($\chi^2_{(1)} = 30.431, p < 0.01$) and trial order ($\chi^2_{(1)} = 11.709, p < 0.01$), with the presence of feedback and increasing trial order being

associated with greater accuracy. However, a three-way interaction of these factors with role switching ($\chi^2_{(1)} = 5.214, p < 0.05$) showed that while the feedback effect was stable in both switch conditions, order only had an effect when there was no role switching and it interacted positively with feedback when there was role switching ($\chi^2_{(1)} = 6.034, p < 0.05$).

Speech Analysis

Potentially, speakers could refer to three dimensions in their descriptions of the referent corresponding to the contrasts between the referent and the shared competitor, the speaker's privileged competitor, and the listener's privileged competitor. At the beginning of the experiment speakers are much more likely to use the first two by virtue of their lack of access to the third, but as the experiment progresses and evidence for that dimension becomes available it should be seen to rise. The use of adjectives associated with these dimensions speaks to the underlying scheme the speaker has about what information is necessary to identify the target. Because the common competitor is always necessary (and indeed, is referred to in 88% of trials) its use is a useful check when it comes to assessing whether a speaker understands and is faithfully attempting to do the task. However, because the underlying hypotheses are focused on how speakers reconcile their privileged knowledge with that of their listeners, how they used modifiers related to the two privileged contrasts is more theoretically interesting.

Use of a modifier relating to the speaker's privileged competitor was influenced by whether or not the participants were switching roles ($\chi^2_{(1)} = 6.500, p < 0.05$; see Figure 3). For the non-switchers, use of the speaker's privileged competitor stayed relatively close to ceiling over the course of the experiment whereas its use dropped over time for partners who switched roles. This may be because, for a speaker in the no-switch condition, their privileged competitor was always salient to them, whereas in the switching condition speakers experienced trials in which its presence was not linked to pro-
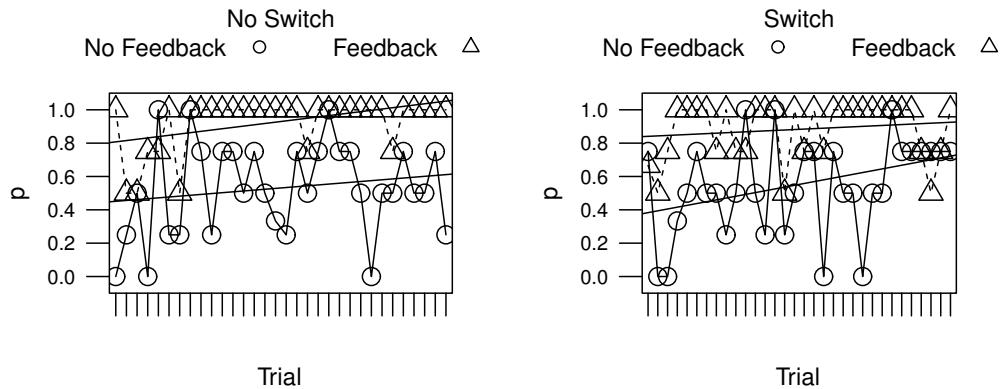
Figure 2: Proportion of targets correctly selected by the addressee.

duction. Because the use of the privileged competitor was not strongly associated with accuracy ($\bar{r}_\phi = -0.028$, p = 0.595), there was not necessarily a strong reason to reduce its salience to speakers in the non-switching condition in the absence of another role.

On the other hand, the speaker had more substantial motivation to learn and use the listener's privileged competitor because its use was strongly associated with success in the task ($\bar{r}_\phi = 0.310, p < 0.01$). In their use of modifiers that matched the listener's privileged competitor, speakers were significantly influenced by trial order (using more over time; $\chi^2_{(1)} = 26.866, p < 0.01$) and an interaction between the feedback and switching conditions ($\chi^2_{(1)} = 6.899, p < 0.01$), such that the effect of feedback was larger when there was no role switching. It is interesting that there doesn't appear to be an additive effect between the two routes to getting partner feedback associated with the two factors, with the greatest degree of use of the listener's privileged competitor being when there was no role switching (see Figure 4). It's possible that the switching condition in this case reduced the pressure on the speaker to unilaterally integrate the perspective of their partners that is granted primarily through the channel of direct feedback. Indirect feedback about their partner's knowledge may not figure into their subsequent production in the switching conditions. This is because some of the information is coming during trials in which they do not have to integrate this information into an utterance themselves, and for which it is unnecessary from their perspective, thus reducing its salience.

Speech onset offers a window into feedback's anticipatory effect in regards to how much planning speakers have to do when they are aware of the possibility that their addressee can make a clarification request. Planning (time to onset of speech) was influenced by the opportunity for feedback ($\chi^2_{(1)} = 10.235, p < 0.01$) with planning time decreasing with the presence of feedback. This is likely due to reduced pressure on the speaker to carefully craft their utterances due to

the lack of a chance for correction. Time needed for planning showed a significant decreasing trend over the course of an experimental session ($\chi^2_{(1)} = 34.045, p < 0.01$), which is consistent with a general increase in familiarity with the task.

Gaze Analysis

One of the primary claims of the hypotheses motivating this experiment is that patterns of gaze during planning may reflect the underlying categorical structure a speaker is considering. Rehder and Hoffman (2005) found that in category discrimination tasks participants learned to only attend to the features in a display that are required for discriminating category membership for the pattern as a whole. Likewise, it was hypothesized that gaze in a referential communication task would demonstrate the shifting category membership status of the objects under consideration by the speaker. For these analyses, due to the absence of eye-tracking data for the second participant in the switching group, and the inclusion of trial order as a factor in the analyses, only the first sixteen trials of the non-switching group were analyzed and compared with the performance of the eye-tracked partner in the switching group.

Attention directed to the privileged competitor (measured here as the proportion of pre-onset fixation time) does not appear to be influenced by the same factors that influence the use of the privileged competitor. In this case, the only factor of influence is feedback ($\chi^2_{(1)} = 4.820, p < 0.05$), with feedback being associated with increased attention being directed to the privileged competitor. Theoretically, attention to the shared competitor should be relatively insensitive to condition due to its constant use regardless of circumstance, but it appears to be more heavily influenced by experimental factors than attention to the privileged competitor. This is demonstrated by a three-way interaction between switching, feedback, and trial order ($\chi^2_{(1)} = 4.777, p < 0.05$), in which feedback led to more attention to the shared competitor at the beginning of the experiment, but less at the end. This may be due to feedback initially bringing a greater focus to the
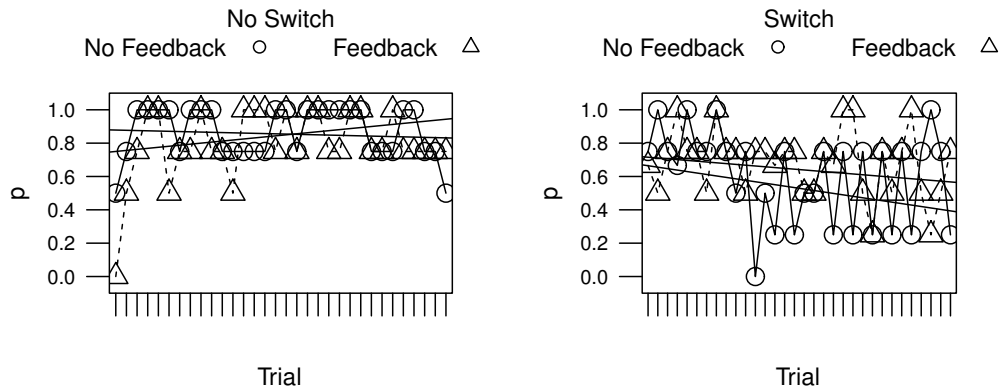
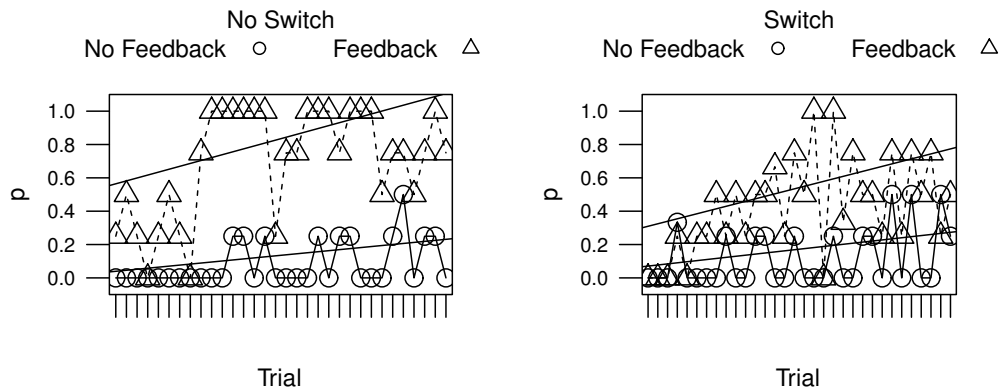Figure 3: Proportion of utterances that refer to the dimension that is seen only from the speaker's perspective.

Figure 4: Proportion of utterances that refer to the dimension that is seen only from the addressee's perspective.

shared competitor as parters in the role switching condition acclimated to the task.

Discussion

Speakers and their addressees tend to align their conceptual and semantic representations over the course of a conversation. However, the final form that these representations takes may depend on the characteristics and constraints under which the interaction is undertaken. In this experiment, each participant had information that was privileged and the goal was to determine what effect conversational constraints, such as role constancy and the opportunity (and actuality) of feedback, had on the nature of the interlocutors' apparent representations as evidenced by the form of their utterances. These changed representations are supposed to reflect the creation of socially emergent categories that arise out of the alignment process.

Interestingly, the most successful descriptions included the listener's privileged competitor. This is why performance was lowest by far in the most feedback-impoverished condition. This is notable because the task was, in principle,

solvable if the addressee aligned to the speaker by realizing that they could not see the addressee-privileged competitor. Just as the speaker had the opportunity to learn what the listener's perspective was through feedback, the listener could infer the speaker's perspective through their speech. However, the speaker's use of their own privileged perspective had relatively little bearing on the outcome of an individual trial.

Speakers seemed to acquire information about the addressee's perspective most effectively through direct, corrective feedback. Despite the opportunity to model their partner's knowledge, or at least benefit from priming, when they occupied the directing role in the role switching condition, speakers appeared to primarily acquire this information through direct feedback such as the listener asking about their privileged dimension. However, the significant role switching by feedback interaction speaks to a boosted effect of feedback when partners did not switch roles. This is interesting because it suggests that perhaps it is more beneficial to have role constancy and consistent feedback than to experience both roles and be in their partner's shoes. A participant who is consistently the director may feel it is part of their

role to adapt to their partners needs, whereas partners who switched roles may have been less sure about which partner should be aligned to. On the other hand, a speaker's use of their own privileged competitor seemed to be sensitive only to role switching, with use staying at a constant, high rate when partners have stable roles, but declined steadily when they switch roles. One possible explanation is that addressee feedback is more likely to be about communicating the dimension that is salient to them, and is not used to elicit clarifications about the superfluous (from their perspective) dimension used by the speaker. An impoverished description is potentially more damaging, than over-specification, but is amenable to feedback correction (Gann & Barr, 2012). Direct feedback could put an upward pressure on the use of the addressee's perspective, but be neutral in regards to the speaker's. Thus salience from the speakers perspective is a sufficient reason for a speaker to include information in a given description, even if it has no apparent communicative value to the addressee.

Thus, when speakers are role switching and have feedback, they are likely to align in a way that favors aligning to the addressee through the speaker's increasing use of the listener's privileged dimension. This suggests these participants are moving toward a conceptualization of the shared perspective that is a "union" of their perspectives: The speaker makes use of adjectives related to both partners privileged dimensions. The adaption that results in this union is driven by the structure of the task, and the quality of the feedback experienced. Additionally, as seen in previous experiments such as Wardlow Lane, Groisman, and Ferreira (2006), anything that raises the salience of a particular competitor is likely to increase mention of that competitor, which might explain why the privileged competitor continues to be used despite it only being relevant from the speaker's perspective. Because its use doesn't seem to have a negative effect, there is little pressure to reduce its use absent another factor.

The online measures paint a more mixed picture. The effect of speech onset in regards to feedback seemed to mirror that which was seen in the prior experiment. When speakers were not permitted the opportunity for direct advance feedback, they took a little longer to plan their utterances. The eye-tracking data were less clear. At the outset it was anticipated that the attention paid to the competitors would help reveal the underlying conceptual structure the speaker was consulting in order to formulate her utterances (Rehder & Hoffman, 2005). However, because the listener's competitor is never visible to the speaker, there was no way to match up eye-movements with its consideration, except for perhaps if the speaker preferentially fixates nothing when considering their partner's knowledge (which is one potentially interesting future analytical direction). In addition, the relative salience of the competitors on the screen was under different selective pressures than the features in the categorization task of Rehder and Hoffman (2005), as additional decisions had to be made about how to describe the contrast rather than

making a simple judgment of category membership.

References

Barr, D. J. (2003). Paralinguistic correlates of conceptual structure. *Psychonomic Bulletin and Review*, *10*(462-467).

Barr, D. J., & Kronmüller, E. (2007). Conversation as site of category learning and category use. In B. H. Ross & A. Markman (Eds.), *The Psychology of Learning and Motivation* (Vol. 47). Academic Press.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure in mixed-effects models: Keep it maximal. *Journal of Memory and Language*.

Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation: Advances in Research and Theory* (Vol. 27, pp. 1–64). San Diego, CA: Academic Press.

Bates, D., & Maechler, M. (2010). lme4: Linear mixed-effects models using S4 classes [Computer software manual].

Gann, T. M. (2012). *Toward a comprehensive theory of audience design*. Unpublished doctoral dissertation, University of California, Riverside.

Gann, T. M., & Barr, D. J. (2012). Speaking From Experience: Audience Design as Expert Performance. *Language and Cognitive Processes*.

Garrod, S., & Doherty, G. (1994). Conversation, coordination and convention: an empirical investigation of how groups establish linguistic conventions. *Cognition*, *53*, 181–215.

Horton, W. S., & Gerrig, R. J. (2005). Conversational common ground and memory processes in language production. *Discourse Processes*, *40*(1), 1–35.

Krauss, R. M., & Weinheimer, S. (1966). Concurrent feedback, confirmation, and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, *4*, 343–346.

Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, *27*, 1–22.

Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*, 1–41.

Schober, M., & Clark, H. H. (1989). Understanding by addressees and overhearers. *Cognitive Psychology*, *21*(2), 211–232.

Wardlow Lane, L., Groisman, M., & Ferreira, V. S. (2006). Don't talk about pink elephants! Speakers' control over leaking private information during language production. *Psychological Science*, *17*(4), 273–277.