Experimental Evidence for Speakers' Sensitivity to Common vs. Privileged Ground in the Production of Names

Daphna Heller (daphna.heller@utoronto.ca)
Department of Linguistics, 130 St. George Street
Toronto, ON M5S 3H1 CANADA

Kristen Skovbroten (kskovbroten@bcs.rochester.edu)
Department of Brain and Cognitive Sciences, University of Rochester
Rochester, NY 14627 USA

Michael K. Tanenhaus (mtan@bcs.rochester.edu)
Department of Brain and Cognitive Sciences, University of Rochester
Rochester, NY 14627 USA

Abstract

The distinction between shared and privileged information is important for the production of referring expressions: in order for a referring expression to be felicitous, it has to be based on shared information. But determining what information is shared and what is privileged requires gathering information from multiple sources, and constantly coordinating and updating them, which might be computationally too complex to affect production in real time. Previous work has found that speakers produce over-informative referring expressions which violate Grice's Maxim of Quantity, concluding that this is because they do not track the distinction between shared and privileged. We demonstrate that speakers are in fact quite effective in tracking this distinction and they mark it in the form of their utterances: under certain circumstances speakers choose to over-specify.

Keywords: common ground; perspective taking; referring expressions, names.

Introduction

It is standardly assumed that speakers obey Grice's Maxim of Quantity: 'be only as informative as required'. When applied to the production of referring expressions, and to names in particular, this means that a cooperative speaker should use a name if and only if this name can be assumed to be known to the addressee, i.e. only if the name is in common ground. Using a name that is known to the speaker alone could cause confusion on the part of the addressee, and may result in the addressee not succeeding in identifying the intended referent. Technically, such use would violate the Maxim of Quantity by being over-informative.

It is an open question whether speakers actually exhibit this pattern of optimal behavior, and there are reasons to believe this might not be the case. First, this pattern assumes that speakers are strictly Gricean, but it is well known that speakers tend to under-specify or over-specify under certain circumstances. More importantly, this behavior requires that speakers distinguish shared and privileged information, but given that computing this distinction is complex, interlocutors might resort to a simpler heuristic. Indeed, while there isn't much evidence concerning perspective taking in production, existing studies seem to suggest that speakers do not track and use the shared versus privileged distinction. Wardlow Lane, Groisman & Ferreira (2006) found that speakers produced a referring expression that was informationally-appropriate from their own perspective, but over-informative to their addressee. Wu & Keysar (2007b) who examined perspective taking behavior in a more complex context propose that speakers did not simply use their own perspective, but rather use a global strategy that relies on information overlap to determine how much information is shared. That is, existing psycholinguistic evidence seems to suggest that, despite its relevance to referring expressions, speakers do not actually use the distinction between shared and privileged information in production.

Our original goal was to examine whether certain factors, such as structuring information, would improve the ability of speakers to use the distinction between shared and privileged information in a setup similar to Wu & Keysar's. During coding, we discovered – to our surprise – that speakers were in fact quite effective in tracking and marking this distinction, and that this has been masked by the assumption that speakers obey the Maxim of Quantity in their name use. That is, under the assumption that speakers are always Gricean, Wu & Keysar concluded that they do not track the shared versus privileged status of names. In this paper we demonstrate that the opposite is the case: speakers are quite effective in tracking shared versus privileged and they mark this distinction in their utterances, but they do not strictly conform to the Maxim of Quantity as they choose to over-specify privileged names.

Perspective Taking in Comprehension and Production

Determining what information is shared and what is privileged requires interlocutors to gather information from multiple sources, such as the physical environment, the linguistic context, and more general information about the other interlocutor, known as 'community membership'

(Clark 1996). Coordinating the different types of information and constantly updating them during conversation,might be computationally intensive, and thus too slow or burdensome to affect language use in real time. For comprehension, there is a growing body of psycholinguistic evidence that listeners can effectively track the distinction between shared and privileged information and use this distinction from the earliest moments of processing (Nadig & Sedivy 2002; Hanna, Tanenhaus & Trueswell 2003; Wu & Keysar 2007a; Brown-Schmidt, Gunlogson & Tanenhaus 2008; Heller, Grodner & Tanenhaus 2008; but see Keysar et al. 2000; Keysar, Lin & Barr 2003). However, using this distinction in production might be a harder task. For one, listeners may get additional cues about the speaker's perspective from her speech, but this source of information is not available to speakers.

Note, also, that production contrasts with comprehension in that there does not seem to be an obvious heuristic that speakers can adopt to avoid computing the distinction between shared and privileged information. In particular, listeners might be able to focus on shared information (or common ground), blocking out privileged information altogether (but see Brown-Schmidt et al. 2008 and Heller et al. 2008 for evidence that this is not what listeners do). This heuristic, however, would not work for speakers: while the heuristic might be appropriate for referring expressions that require reference to information in common ground, assertions are generally required to contribute information that is not already in common ground (Stalnaker 1978). Thus utterance planning for an assertion would require speakers to consult their privileged information.

Wu & Keysar (2007b) propose that speakers use a different heuristic in production. In their study, speakers learned artificial names for novel shapes, some together with their addressee and some alone. When participants performed a referential communication task, speakers used privileged names that were not known to their addressees. Assuming that speakers would not intentionally violate Grice's Maxim of Quantity, Wu & Keysar concluded that speakers were not able to keep track of the shared versus privileged status of names. Instead, since they found that speakers used significantly more names when they shared more information with their addressee, they proposed that speakers rely on a global heuristic which depends on an estimate of the overall overlap in information. In particular, their Information Overlap Heuristic states that "when overlap in information between two people is extensive, using one's own information should work just fine because it is most likely to be shared." (p. 4).

Our original goal was to examine whether changing the nature of the stimuli will allow speakers to track the distinction between shared and privileged. In particular, given that information in real conversation is not arbitrarily categorized as shared or privileged, but rather follows some classification (e.g. of 'community membership'), we tested whether providing speakers with structured stimuli would allow them to better track the shared versus privileged status of names. To our surprise, we noticed during coding that even in the unstructured conditions speakers consistently distinguished shared and privileged names. While our results replicate Wu & Keysar's when using their coding schemes which assume that speakers respect Grice's Maxims, we found that speakers are often over-informative when talking about privileged shapes and utter their names in addition to a description. The structure manipulation did not have any significant effects on speakers' behavior; speakers reliably distinguished between shared and privileged names, regardless of structure; we therefore collapsed over the structure manipulation in all our analyses.

Methods

Participants

Forty pairs of naïve participants are included in the analysis. Participants arrived together; they were all native speakers of English recruited from the University of Rochester community. One pair was excluded because they were unable to complete the training and another because of computer failure. Six additional pairs were excluded from analysis because they did not achieve the required 90% accuracy on the referential communication task.

Materials

The novel shapes and their names were adapted from Wu & Keysar (2007b). Twenty four experimental shapes had novel names, which were slightly modified from Wu & Keysar's names such that six of the twenty-four names shared the onset /fl/. Thirty additional novel shapes that were not named were used during testing.

Procedure

Training. Participants sat together across from the experimenter, who had index cards with the twenty-four named shapes. Participants learned the names of the shapes in four blocks of six. On each trial, the experimenter presented the card, articulated the name, and waited for the participants to repeat the name before proceeding to the next card. After going through the six shapes in the block once, the experimenter presented the card and waited for the participants to name the shape; the experimenter articulated the name or corrected any errors if the participants could not name the shape correctly. The experimenter repeated this procedure for the block until both participants could name all six shapes flawlessly, and then moved on to the next block.

Common ground was manipulated by having the (randomly selected) addressee learn only a subset of the names learned by the speaker. Participants first learned some names together; then the speaker continued to learn more names alone, and the addressee played a non-linguistic computer game while listening to music over headphones (the addressee stayed in the same room).

Two factors were manipulated in training, creating a 2x2 between-subjects design: Overlap (High vs. Low) x Category (Category vs. No-category). Overlap manipulated the relative amount of information in common versus privileged ground. In the High Overlap conditions, participants learned eighteen names together, whereas Low Overlap meant participants learned only six names together. Category manipulated the whether the six shapes in the last block of training shared the onset /fl/ (Category conditions) or had no properties in common (No-category condition).

Testing: the referential communication task. Participants sat in front of different computers and could freely communicate over a network. The speaker was presented with one shape (shared, privileged, new), and had to instruct the addressee, who saw three shapes, to click on the target shape "as quickly and accurately as possible". Trials were advanced when the addressee clicked on a shape, even if it was the wrong one (an error sound was heard). The referential communication task had two practice trials followed by 18 experimental trials, six for each shape type. The shared shapes were always drawn from the first block of training, and the privileged shapes from the last block of training – see Figure 1. Therefore, testing did not differ between the different training conditions (order of presentation also did not change).

Coding and Results

On the referential communication task, addressees performed at 97.9% (pairs who performed at less than 90% were excluded from analysis).

There were no significant effects of the Category manipulation so we collapsed the levels of the Category factor and compared only High Overlap and Low Overlap. We observed during coding that speakers used shared and privileged names in different ways. Therefore, we performed analyses to quantify these differences.

Analyses I: Wu & Keysar's Coding Schemes

The Wu & Keysar coding scheme assumes a pattern of optimal behavior where speakers are being only as informative as necessary, in accordance with Grice's Maxims of Quantity. In particular, a speaker who obeys the Maxim should refer to a shape that has a shared name by using its name; using a description in this case would be under-informative. If, however, a name is only known to the speaker, this name should not be used, because using a privileged name will be over-informative and is likely to cause confusion on the part of the addressee; under these circumstances, a description should be used. This behavior would give rise to a one-to-one correlation between name use and shared names, so counting utterances that contain names will provide a measure of the speaker's assumptions about which names are shared.

We first followed Wu & Keysar's analysis and counted all utterances containing names ("all-names" analysis). Proportions were logit-transformed and then submitted to a 2 (Overlap) x 3 (Shape) ANOVA. There were main effects of Shape ($F(2,76)=99.85$, $p<.001$) and of Overlap ($F(1,38)=7.68$, $p<.01$). High Overlap speakers were more likely to use names than Low Overlap speakers: shared .80 vs. .65, privileged .31 vs. .16, new .05 vs. 0.

We also followed Wu & Keysar in performing a more conservative analysis, looking at just those utterances where the name occurred before any description ("name-first" analysis). The goal was to exclude cases "in which speakers first described the object in order to identify it and then named it in order to inform their addressee about the name" (Wu & Keysar 2007b, p.7). As was the case in Wu & Keysar (2007b), this analysis gave rise to the same pattern as the "all-names" analysis. There were main effects of both Shape ($F(2,76)=93.63$, $p<.001$) and Overlap ($F(1,38)=10.64$, $p<.01$). High Overlap speakers were significantly more likely to use names than their Low Overlap counterparts both for shared shapes ($F(1,38)$-4.91,$p<.05$), and for privileged shapes ($F(1,38)=4.28$, $p<.05$) – see Figure 1.

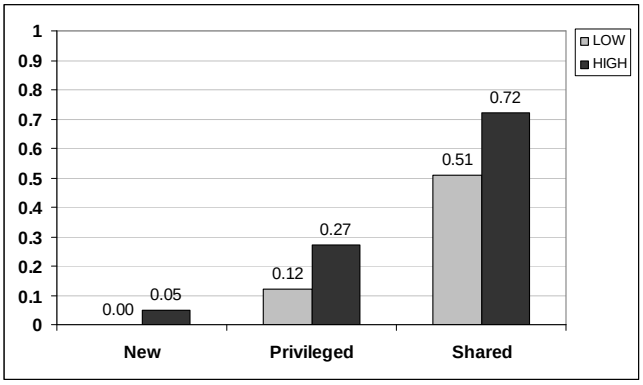

Figure 1: The proportion of "name-first" utterances for new, privileged and shared shapes in High and Low Overlap conditions.

In sum, more names were used by High Overlap speakers than by Low Overlap speakers both for shared shapes and, crucially, for privileged shapes. Given that this pattern was found for the same twelve shapes (six shared and six privileged) in both conditions, Wu & Keysar argue that speakers do not track the shared versus privileged status of individual names, but rather rely on the overall information overlap in determining whether to use names. Our findings replicate their pattern, potentially serving as further evidence for their proposal.

Analysis II: Listeners' Judgments

Wu & Keysar's conclusions crucially depend on the assumption that speakers' behavior obeys the Maxim of Quantity, and that speakers only use names when they assume that they are shared. During coding we observed a clear divergence between shared and privileged names – see Table 1. This may suggest that while speakers indeed used privileged names in violation of Grice's Maxims, they were nonetheless aware that these names were privileged. If this

is indeed the case, then Analysis I classifies too many names as shared. The crucial question is, of course, whether shared uses of names can be reliably distinguished from other uses independent of the status of shapes..

Names for shared shapes
• uhm **cortlog**.
• ah, **banpar**, your favorite.
• That's an **intra**?
• um, **abypit** I think it's called, I forget.
• It's like another, it's like an **abypit**. It's just a … yeah.
Name for privileged shapes
• oh you don't know this, ok it's square with arrows coming out of it… it's called **floogle** if you were interested.
• ah, **inta**, you haven't seen it, it's four arrows.
• ah, this is called **molget**, it's like a triangle and a rectangle.
• it's called **flazap**, it looks like a flag.
• **cortlog**, it's a guy kicking.

Table 1: Example of utterances containing shared and privileged names.

Three naïve coders judge the status of names (they were blind to the conditions of training and to the status of shapes). Name uses were classified as "assume shared" if the coder thought that the speaker expected the addressee to know the name, and as "assume privileged" if the coder thought that the speaker did not expect the addressee to know the name. If speakers systematically mark the shared versus privileged status of information (and if listeners are sensitive to these cues), then shapes with shared names should be judged as "assumed shared". If these judgments do not actually reflect something real, these labels should be evenly distributed between shared and privileged shapes. Table 1 shows that speakers sometimes explicitly mentioned their assumptions about the status of the names; it should be noted that such explicit comments were found for only 3% of trials.

There was agreement among all three coders for 83% of trials, and among at least two coders for 99% of the trials. We found that "assume shared" judgments correlated with shapes with shared names. Proportions were logit-transformed and then submitted to a 2 (Overlap) x 3 (Shape) ANOVA. There was a main effect of Shape ($F(2,76)=173.07$, $p<.001$) because there were more "assume shared" names for shared shapes than for privileged shapes ($F(1,39)=263.58$, $p<.001$) – see Figure 2a. In fact, "assumed shared" judgments were not more likely for privileged shapes than for new shapes ($F<1$). Thus, once we focus on those names that sound shared, they were usually used for shared shapes and only rarely used for privileged shapes. This contrasts with Analysis I where we found significantly more names for privileged shapes than for new (=unnamed) shapes (the same comparison was significant for "name-first" trials: $F(1,39)=13.34$, $p=.001$). This result suggests that speakers are in fact quite effective in identifying shared shapes. Note that the main effect of Overlap ($F(1,38)=12.62$, $p<.001$) persists under the current analysis.

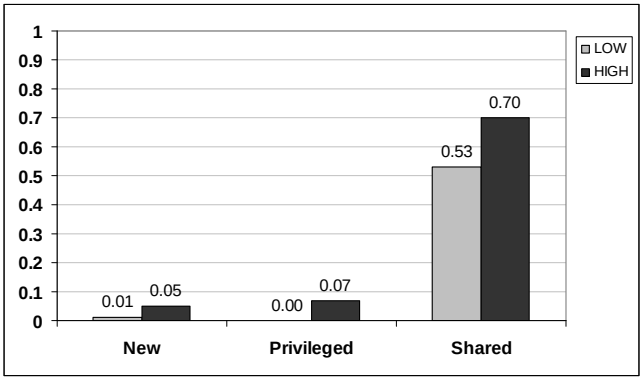

Figure 2a: The proportion of "assume shared" trials for new, privileged and shared shapes in High and Low Overlap conditions.

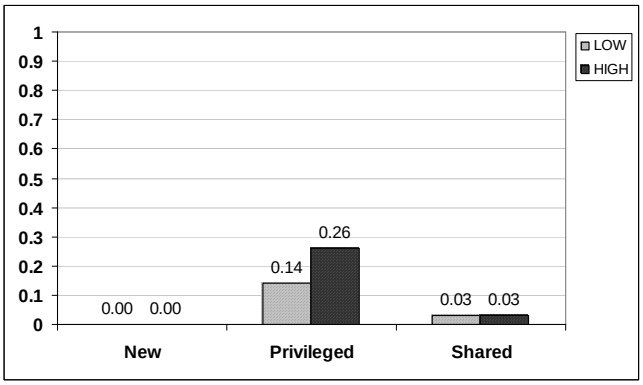

Figure 2b: The proportions of "assume privileged" trials for new, privileged and shared shapes in High and Low Overlap conditions.

This conclusion is further supported by the distribution of "assume privileged" judgments. These were mostly found with privileged shapes, and rarely with shared or new shapes. Proportions were logit-transformed and submitted to a 2 (Overlap) x 3 (Shape) ANOVA, revealing a main effect of Shape ($F(2,76)=26.21$, $p<.001$) (there was no effect of Overlap in this case) – see Figure 2b.

Analysis III: Name Form

The patterns observed in Analysis II suggest that speakers systematically mark the distinction between shared and privileged names. Our second analysis focused on how speakers mark names as shared or privileged. We focused on "name-first" trials, distinguishing trials that included just names ("name-alone") and trials in which the name was followed by a description ("name-then-description"). We hypothesized that "name-alone" uses are likely to be found for those names speakers assume to be shared, because on

these trials there was no additional information the addressee could use to pick out the referent.

As in previous analyses, proportions were logit-transformed and submitted to a 2 (Overlap) x 3 (Shape) ANOVA. For "name-alone" utterances, there was a main effect of Shape ($F(2,76)=269.40$, $p<.001$). Specifically, this strategy was more often used with shared shapes than with privileged shapes ($F(1,39)=463.97$, $p<.001$). Privileged and new shapes did not differ ($F<1$). That is, the "name-alone" strategy was used primarily for shared shapes – see Figure 3a.

For "name-then-description" utterances, there was again a main effect of Shape ($F(2,76)=71.51$, $p<.001$), but this strategy is used more with privileged shapes than with shared shapes ($F(1,39)=9.46$, $p<.01$). It was rarely used for shared or new shapes, indicating that this is a strategy speakers chose for privileged names – see Figure 3b.

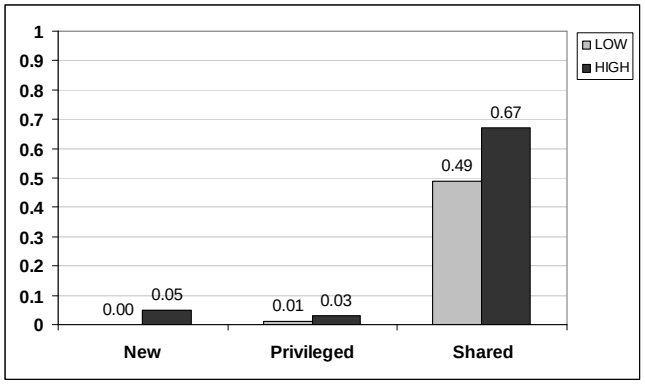

Figure 3a: The proportion of "name-alone" trials for new, privileged and shared shapes in High and Low Overlap conditions.

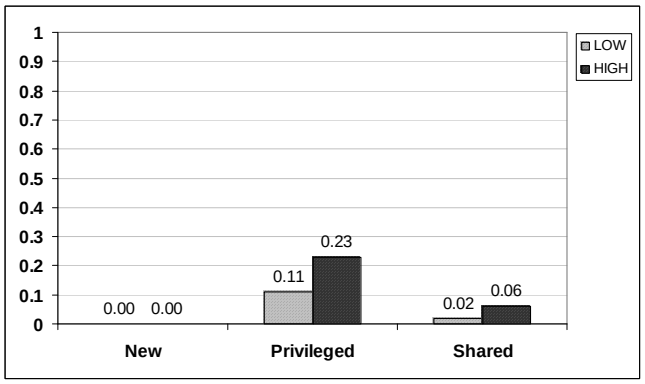

Figure 3b: The proportion of "name-then-description" trials for new, privileged and shared shapes in High and Low Overlap conditions.

Note that for both types of trials there was also a significant main effect of Overlap ("name-alone": $F(1,38)=6.65$,$p<.05$, "name-then-description": $F(1,38)=6.11$, $p<.05$).

It seems likely, then, that the two different forms ("name-alone" vs. "name-then-description") were the cues used by our coders in judging name uses as shared or privileged. We found that out of the "name-alone" trials, 97% were judged as "assume shared" and only 1% were judged as "assume privileged". By contrast, only 14% of "name-then-description" trials were judged as "assume shared" and 86% were judged as "assume privileged". This pattern suggests that, for the most part, coders used the form of the utterance as the cue for judging the status of names as shared or privileged.

Discussion

In Analyses II and III we have quantified the difference between two kinds of name uses in two different ways and have demonstrated that each is correlated with shared or privileged shapes. This closer examination of the data reveals that interlocutors are, in fact, extremely effective in keeping track of the status of individual names as shared or privileged, and in marking this status in the form of their utterances. The Wu & Keysar results, which we nonetheless replicated in Analysis I, were thus not due to the fact that speakers are unable to track shared versus privileged information, but rather because speakers did not strictly obey Grice's Maxim of Quantity and chose to over-specify the referring expression used for privileged shapes, including a privileged name in addition to the expected description.

Two questions are left open. The first is why did speakers choose to use names for privileged shapes. The second question is why were speakers more likely to use names in High Overlap condition compared to the Low Overlap condition. One possibility is that speakers were trying to teach the names to their addressees, hoping that it would make communication more effective if the shape occurred again during testing. Every shape occurred only once, but speakers did not know this before this phase ended. In fact during debriefing, a number of our participants volunteered that they used this teaching strategy. Teaching the addressee the names for the privileged shapes would make more sense when the addressee already knows most of the other names.

Analysis IV: "name-then-description" as a repair strategy?

The "name-then-description" strategy speakers used for privileged names raises the question whether this form was planned from the earliest moments, or whether the description is added later as an afterthought, once the speaker had realized that the name is privileged and thus uninformative for the addressee. Our data contained clear cases of repair, some where the description was added after the addressee did not respond to the name, and some where

the speaker interrupted herself in the middle of the name and produced a description – see Table 2.

• **flanzo** [break] ok, it looks like Ganzo.
• um, **flu**- it's like, it looks like a person sort of.
• oh **flan**- you don't know this one, it looks like a bunny, rabbit.

Table 2: Example of utterances containing repair.

Importantly, "name-then-description" trials did not involve a break or an interrupted name. We hypothesized that if speakers were planning a "name-then-description" utterance from the earliest moments, this should be reflected in the way they pronounce names, and those would differ from "names-alone" trials. To this end, we had naïve listeners listen to names drawn from both types of trials that were truncated at the end of the name, and judge on a seven-point scale whether they expected a continuation (we tested utterances from the two participants who produced the most trials in the relevant forms).

We found that listeners assigned a significantly different rating to names that had continuations and those that did not. This indicates that speakers pronounced names differently depending on whether or not there was a subsequent description, suggesting that the description had been planned early and was not a late repair.

Conclusions

While speakers are more likely to use names when more information is shared, this is not because of an increased level of confusion about whether a certain name is shared or privileged. Speakers clearly distinguish shared and privileged names in the form of their utterances: this distinction is one that our listeners were sensitive to, even those who listened to just the beginning of the utterances, suggesting that this difference exists in the earliest moments of production. This shows that, like listeners, speakers are effective in distinguishing shared and privileged information.

In addition, these results complement previous studies showing that speakers do not always strictly follow Grice's Maxim of Quantity, but sometimes chose to produce over-specified referring expression to meet other conversational goals (Issacs & Clark 1987; Engelhardt, Bailey & Ferreira 2006).

Acknowledgments

We are extremely grateful to Dana Subik for recruiting, training and testing all of the participants.

This research was partially supported by NIH grant HD-27026 to MKT and an NSF Graduate Research Fellowship to KS.

References

Brown-Schmidt, S., Gunlogson, C., & Tanenhaus, M. K. (2008). Addressees distinguish shared from private information when interpreting questions during interactive conversation. *Cognition*, *107*, 1122-1134.

Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.

Engelhardt, P., Bailey, K.G.B., & Ferreira, F. (2006). Do speakers and listeners obey the Gricean Maxim of Quantity? *Journal of Memory and Language*, *54*, 554-573

Hanna, J. E., Tanenhaus, M. K. & Trueswell, J. C. (2003). The effects of common ground and perspective on domains of referential interpretation. *Journal of Memory and Language*, *49*, 43–61.

Heller, D. Grodner, D. & Tanenhaus, M. K. (2008). The Role of Perspective in Identifying Domains of Reference. *Cognition*, *108*, 831-836.

Isaacs, E. A., & Clark, H. H. (1987). References in conversations between experts and novices. *Journal of Experimental Psychology: General* , 116, 26-37.

Keysar, B., Barr, D.J., Balin, J.A. & Brauner, J.S. (2000). Taking perspective in conversation: The role of mutual knowledge in comprehension. *Psychological Science*, *11*, 32-37.

Keysar, B., Lin, S. & Barr, D. J. (2003). Limits on theory of mind use in adults. *Cognition*, *89*, 25-41.

Nadig, A. S., & Sedivy, J. C. (2002). Evidence of perspective-taking constraints in children's on-line reference resolution. *Psychological Science*, *13*, 329–336.

Stalnaker, R. (1978). Assertion. In P. Cole (ed.), Syntax and Semantics 9: Pragmatics. Academic Press.

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*, 273-277.

Wu, S., & Keysar, B. (2007a). The effect of culture on perspective taking. *Psychological Science*, *18*, 600–606.

Wu S. & Keysar B. (2007b). The Effect of Information Overlap on Communication Effectiveness. *Cognitive Science*, *31*, 1–13.