

Language Users as Problem Solvers: Just What Ambiguity Problem Do They Solve?

Boaz Keysar

University of Chicago

Everyone who uses language knows that it is ambiguous. Consider this excerpt from a recommendation letter: "Anybody who will get this fellow to work for him would indeed be fortunate." The writer might be either congratulating or warning future employers. Typically, the inherent ambiguity of language gets in the way of successful communication, or at least, it gets in the way of those who attempt to explain how communication succeeds even though language has such high potential for ambiguity.

Ambiguity abounds in all levels of language, from the pragmatic down to the phonological level. As Nusbaum and Henly (in press) observed, there is no one-to-one mapping between the linguistic stimulus and the intended meaning at any level of linguistic analysis. Consequently, theories of language must explain how language users overcome this problem. How do they typically arrive at a unique interpretation of a potentially multiple-meaning linguistic stimulus? In this chapter I briefly describe the major traditional theoretical solution to this problem at the semantic-pragmatic level and illustrate the theoretical and empirical problems associated with this solution.

I talk about "language users" in general but mainly focus on the problem that addressees, readers, and overhearers solve in interpreting utterances. Speakers face a similar problem when they attempt to construct utterances to be interpreted unambiguously by their addressees, but although the problem is similar, it is not obvious that the solution is the same. I discuss speakers' solutions briefly at the end of this chapter.

My main argument is that language users actually attempt to solve a problem different from the one that has traditionally been assumed. Therefore, the traditional approach is wrong in its assumption about the ambiguity problem that addressees solve. The traditional approach assumes that addressees solve the following problem: In view of the semantics of a sentence, or the literal meaning of an utterance, what is a speaker's intention? Of the different speech acts that a sentence can convey, which one does a speaker actually intend to

convey? I argue that addressees do not attempt to solve this problem. Instead, they attempt to solve the following problem: In view of the meaning of this utterance from my perspective, or given the intention this utterance conveys when interpreted from my perspective, what did the speaker actually intend?

THE PERSPECTIVE-FREE APPROACH

Linguistic theory and philosophy of language have provided a clean, elegant solution to the ambiguity problem in language by proposing that people use *literal meaning* as an anchor. Literal meaning has received bad press in psycholinguistics during the last 10 years, but the basic idea is both highly attractive theoretically and intuitively appealing. The traditional theory assumes that sentences have a meaning that is context independent, a meaning that any competent language user can compose from knowledge of the meaning of words and rules of composition (e.g., Katz & Fodor, 1963). One way to "identify" literal meaning is with the "anonymous letter criterion"—a sentence's meaning that a person arrives at when given the sentence with no context whatsoever (Katz, 1972). This notion is attractive because it sounds intuitively right. People feel as if they know the literal meaning of "the cat is on the mat." More important, this idea simplifies theories, and ties them to a long tradition in philosophy of mind, a tradition that developed a logic to represent semantic relations giving rise to literal meaning (Ayer, 1936; Carnap, 1956). This notion also provides a simple starting point for computer simulations of language use, and it draws a clear line between meaning and actual language use. If meaning is anchored in literal meanings, then all that is left to explain is how people use meaning to convey intentions.

What is the Problem that Language Users Solve? Language users must solve the many-to-many mapping between utterances and intentions (Nusbaum & Henly, in press). In the literal-meaning tradition, the anchor of linguistic computation is the literal meaning. Therefore, the problem of an addressee in this tradition is construed as "In view of the literal meaning of an utterance, what did the speaker intend?" In Grice's (1968) terms, the problem is to get from what a speaker said to what that speaker intended, or to the "speaker meaning." The traditional solution is Grice's important insight: Because linguistic communication is a cooperative behavior, language users observe norms of cooperation. When a sentence meaning seems to violate these norms, the speaker's meaning is different from the sentence meaning (Grice, 1975). In an analogous solution proposed by Searle (1975) as part of Speech Act theory, literal meaning is the basis for computation of intention. When a literal meaning conflicts with a context, an indirect speech act is computed as the speaker's intention.

Evaluation of the Traditional Solution

The traditional solution has been extremely influential and is taken for granted across the board in cognitive science disciplines. Yet in the last 15 years, research in psycholinguistics has tended to undermine some assumptions of this solution. Gibbs (1984, 1994) argues strongly against the very concept of literal meaning; he suggests that literal meaning is an ill-defined concept and that many utterances do not have a unique, identifiable literal meaning. Gibbs also challenges the context-free aspect of literal meaning, which he finds both conceptually and psychologically inadequate concept.

Further, the assumptions of the traditional view did not hold under empirical investigation. The stage model implied by Searle and certain interpretations of Grice's theories assume that a nonliteral meaning is computed only when a literal meaning is rejected. This assumption, which predicts that nonliteral meaning should take longer to comprehend than does literal meaning, has been refuted by several studies demonstrating that, with appropriate context, figurative language such as metaphors and idioms as well as indirect speech acts do not take longer to understand than does literal meaning (e.g., Gerrig, 1989; Gibbs, 1979, 1983, 1984; Ortony, Schallert, Reynolds, & Antos, 1978). The only exception to these results is the finding that a metaphorical reference takes longer to determine than does a literal reference (Gibbs, 1990), but so far it is unclear why referential metaphors show delayed comprehension when comparable nominative metaphors do not (Onishi & Murphy, 1993). Empirical research has also refuted the traditional assumption that comprehension of nonliteral language requires the literal meaning to be anomalous. Glucksberg, Gildea, and Bookin (1982) demonstrated that familiar metaphors are understood automatically, and Keysar (1989) showed that this understanding occurs whether the sentence has a semantically adequate or anomalous literal meaning.

The traditional approach assumes that language users make inferences when they attribute intentions to speakers, but such inferences are supposed to take place only when a meaning diverges from the literal meaning. Language users supposedly attribute a literal meaning to a speaker if the literal meaning makes sense; but when they attribute a nonliteral meaning to a speaker they make an inference about the status of the literal meaning. In contrast, experiments that I have conducted suggest that inferences about a speaker's intended meaning occur in a much more symmetrical manner than has been previously assumed (Keysar, 1994a; Keysar & Glucksberg, 1993). The experiments demonstrate that under some conditions readers follow the inference assumed by traditional theories: Readers infer a metaphorical interpretation because they reject a potential literal interpretation. Yet the experiments also demonstrate that under the same conditions readers do the same thing when they arrive at a *literal* interpretation:

Readers infer a literal interpretation because they reject a potential metaphorical interpretation.

This symmetry of inference is embodied in an example from Kosinski's 1971 book *Being There*. At a certain point in the book, the President of the United States deplored the economic situation and then turned to a Mr. Gardiner for his opinion. Mr. Gardiner responded: "In a garden [. . .] growth has its season. There are spring and summer, but there are also fall and winter. And then spring and summer again. As long as the roots are not severed, all is well and all will be well" (p. 45). The President was supposed to follow the traditional inference and probably inferred that Mr. Gardiner must have been speaking metaphorically. In contrast, readers who know that the speaker is a gardener who knows about nothing else make the mirror image inference: They infer that Mr. Gardiner was speaking literally.

The symmetry between the comprehension of literal and nonliteral utterances goes even further. Glucksberg and Keysar (1990) demonstrated that a metaphorical sentence of the form "An A is a B," which has been traditionally thought of as an implicit comparison ("An A is (like) a B"), is actually understood as a true class-inclusion statement. So when people understand "A cigarette is a time bomb," they understand it as categorizing a cigarette in the superordinate category that borrows the name "time bomb." Moreover, Gernsbacher and Keysar (1995) demonstrated that the same comprehension mechanisms underlying comprehension in general and language comprehension in particular operate in exactly the same way in the comprehension of metaphorical sentences.

Despite the lack of empirical support for the traditional view, the issue is by no means resolved and is still under lively debate. The assumptions of literal meaning still implicitly and sometimes explicitly pervade theories of language use (e.g., Dascal, 1987). In light of this perplexing situation, researchers have hypothesized the existence of "anchors," which might serve as alternatives to literal meaning. One such suggestion has been extensively developed by Raymond Gibbs in his 1994 book *The Poetics of Mind*. In contrast to the traditional assumption about the primacy of literal meaning, Gibbs suggests that the human mind is based on nonliteral principles. He argues that human thought is metaphorical and metonymic and that language reflects these basic conceptual relations (cf. Lakoff, 1987; Lakoff & Johnson, 1980). Gibbs supports his argument in favor of a conceptual metaphor-based approach with a variety of empirical findings (Gibbs, 1992; Gibbs & Nayak, 1991; Gibbs & O'Brien, 1990). As a whole, Gibbs' ideas have the potential to challenge basic principles in cognitive science (Keysar, 1996), but it is still too soon to determine the extent to which this approach can account for actual language use (Keysar & Bly, in press). Specifically, before accepting this conceptual metaphor approach, researchers must answer strong objections raised by Gregory Murphy, who

argues that the current version of a conceptual metaphor theory is not theoretically coherent and that the supporting data can be better accounted for by an alternative view (Murphy, 1996).

PROPOSING AN ALTERNATIVE: A PERSPECTIVE-DEPENDENT APPROACH

I propose that the crux of what is wrong with the traditional approach is the way it construes the problem that an addressee attempts to solve in interpreting an utterance. The traditional assumption is that the addressee solves the problem of a speaker's intention in view of the literal meaning of an utterance. This solution assumes an anchor that is context-free and perspective-free. In other words, the literal meaning of an utterance does not depend on a speaker's or addressee's perspective but is "given" in the language. Research in psycholinguistics has provided a compelling case against the assumption of context-free computation. In fact, experimental evidence has demonstrated that addressees use context incrementally to create temporary interpretations as utterances unfold (Tanenhaus, Spivey-Knowlton, Eberhard & Sedivy, 1995). Because addressees use context to constrain potential interpretation, what interpretation problem do they solve? I propose a new explanation for the interpretation problem.

What is the Problem that Language Users Solve?

I propose that the interpretation problem is a function of perspective. The notion of perspective has been extensively investigated in the social sciences and consequently, like the notion of meaning, has too many different meanings (e.g., Anderson & Pichert, 1978; Cox, 1977; Graumann, 1989; Graumann & Herrmann, 1988; Krauss & Fussell, 1988; Long & Andrews, 1990; Mead, 1934; Piaget & Inhelder, 1956; Rommetveit, 1974; Schober, 1993; Traxler & Gernsbacher, 1993; Yaniv & Shatz, 1990). Krauss and Fussell (1996) provide a comprehensive review of the literature on perspective-taking in the context of models of communication. As Krauss and Fussell (1996) point out, most research on perspective-taking in communication focuses on the way that speakers attempt to take the perspective of addressees. Consequently, little is known about the way addressees, readers, and overhearers take the perspectives of speakers and writers. My proposal is targeted at understanding the interpretation problem that addressees solve in terms of perspective.

I use the term *perspective* to refer to access to information. An addressee's perspective is the information that the addressee has access to. Similarly, a speaker's perspective defines the information to which the speaker has access. The interaction between different perspectives in discourse can be considered in

the following way: An addressee has a perspective, a speaker has a perspective, and these perspectives might overlap. To the extent that the perspective overlap is mutually known to speaker and addressee, this overlap constitutes mutual knowledge (Clark & Marshall, 1981) or part of the common ground between them (Clark, Schreuder, & Buttrick, 1983). In other words, if both speaker and addressee know that they mutually share information, this information is their mutual perspective. When addressees interpret utterances, they can in principle use one of the three perspectives: (a) They can use their own perspectives to arrive at an interpretation that uses egocentric information; (b) they can use a speaker's perspective; or (c) they can use the common perspective, their mutual knowledge.

Together with students in my laboratory (Keysar, Barr, Balin, & Paek, 1997), I have proposed that addressees follow the first interpretation strategy: They interpret utterances in a context-sensitive manner, but the context is their own, given by their own perspective and insensitive to what they know about the speaker's perspective. This interpretation strategy sometimes leads them astray, especially when a speaker's meaning depends on the mutuality of perspective. With this model in mind, I propose that the problem that an addressee solves is the following: In view of the interpretation from my perspective, what did the speaker intend? In other words, did the speaker intend the (egocentrically) computed intention, or could the speaker have had a different intention in mind? To solve this problem, our model assumes an adjusting correction mechanism that attempts to take into account the common perspective whenever the interpretation from an addressee's own perspective fails. Hence the model is named the *Perspective Adjustment Model*.

The proposed model (Keysar et al., 1997) makes the following assumptions about the two different processes. The process that generates an interpretation from an addressee's perspective is what the "comprehension system" knows how to do best; it is quick, efficient and relatively automatic. In contrast, the adjustment process that incorporates a speaker's perspective is relatively slow and dependent on available resources. This process attempts to implement what addressees know they "should" do. Addressees believe that they should take the speaker's perspective, that they should use the mutual perspective, and so on. This intuitive theory that addressees hold about communication is high-level knowledge that prescribes behavior. By using this intuitive theory, addressees attempt to be in line with normative rules of cooperative behavior.

The model does not make a strong assumption about serial versus parallel processing. There is no reason to assume that the "check" for violation of mutual perspective must be initiated only after an interpretation process is completed. Therefore, the two processes might be operating in parallel. We do not make an assumption of serial processing and it is possible that the time course of the two processes is different. Consequently, the faster, "egocentric"

process is more likely to quickly constrain the interpretation and to result in a need for a correction whenever a check for violation of mutual perspective detects that addressees used information they "should" not have.

How is this model different from the traditional model? Apart from the obvious differences, the model assumes a different building block for interpretation, a contextualized, computed intention instead of a "context free" literal meaning. Also, instead of a perspective-free basic meaning, it assumes that even the basic computation is perspective-bound. An addressee's perspective determines the context used for the computation. The resulting interpretation from an addressee's perspective, then, is an intention: a particular speech act. By using their own perspectives as the basis for interpretation, addressees face a simpler ambiguity problem because their perspectives put constraints on the possible meanings that each utterance can convey. Therefore, the problem of the inherent ambiguity of language is sharply reduced: Even though each utterance can in principle have a large number of interpretations, addressees are not typically faced with all these potential interpretations, only with the interpretations specific to their perspective.

The general approach and the model I have described give rise to a variety of interesting predictions. In the rest of this chapter, I will review empirical studies that support implications of the perspective-bound approach.

ARE LANGUAGE USERS PERSPECTIVE-BOUND?

Illusion of Transparency

In general, the perspective-bound approach implies that language users perceive language as relatively unambiguous because they do not perceive utterances but perceive perspective-bound interpreted utterances. That is, they perceive utterances as "carrying" a *particular*, and often unique intention, even when the same utterance can convey other intentions. When Terry Gross interviewed Peter O'Toole on National Public Radio, the actor explained that he had been spending much of his time on writing his memoirs.

Terry: So how much are you acting now?
 Peter: I am not acting at all [pause] . . . I am just sitting here
 with the headphones on trying to answer your questions.

Both interlocutors seem to have had a particular interpretation for the ambiguous utterance "How much are you acting now?" The miscommunication resulted because both considered only their own interpretations and did not take into

account alternative interpretations of the ambiguous utterance.¹ Each seems to have thought the meaning to be relatively transparent or as adequate to the utterance. In my model, interpretations or perceived intended meanings generally *seem* relatively transparent to language users: Once an ambiguous utterance is interpreted in a unique way, the interpretation seems more adequate than possible alternative interpretations. Therefore, to the extent that utterances are ambiguous, their potential ambiguity should be underestimated, and the extent to which a particular utterance seems to carry an intention should be overestimated. This implication derives from the model's assumption that a language user's own perspective is the basis for disambiguating utterances. I review several experiments that reveal different facets of such illusion of transparency in language.

Overestimation of the Transparency of Idiomatic Meaning. Idioms are by definition expressions whose meanings are not a compositional function of their words. In the prototypically opaque idiom "to kick the bucket," the connection of the expression to its meaning (to die) is not at all transparent. Yet many idioms seem relatively transparent to native speakers of a language. For example, native English speakers think that "to go against the wind" is transparent: It apparently makes sense to them that the idiom means to do something the hard way. For instance, walking against the direction of the wind is harder than walking in the same direction because of the wind's resistance. Together with Bridget Bly, I investigated this subjective feeling of transparency. We hypothesized that native speakers of a language perceive such idioms as more transparent than they actually are (Keysar & Bly, 1995, in press).

Our argument was that once language users know the actual meaning of the idiom and once it makes sense to them, they discount the possibility that alternative meanings could make just as much sense. They also discount the possibility that the actual meaning would not make as much sense as it does had they not learned it as the idiomatic meaning of the expression. This argument is counterintuitive: If "to go against the wind" meant to do something with greater ease, then this hypothetical meaning would make a lot of sense and its opposite meaning (to do something the hard way) would be perceived as less transparent.

To investigate our hypothesis, we used archaic English idioms that are no longer in common use, such as "the goose hangs high." We taught participants that the idiom meant either that the future is highly promising (the original meaning) or that the future is very gloomy. We demonstrated that once people learn a meaning for the idiom, they perceive this meaning as relatively transparent—a reaction that did not depend on whether they learned the real, original meaning of the idiom or its conceptual opposite.

¹It is also possible that O'Toole was joking and only pretended to misunderstand Terry Gross.




FIG. 8.1. Ratings of the extent to which idioms make sense as a measure of perceived transparency. Both the original and the reversed meanings were perceived as more transparent than their opposite meaning when they were learned as the "true" meaning of the idiom (From "Intuitions of the Transparency of Idioms," by B. Keysar & B. Bly, 1995, *Journal of Memory and Language*, 34, p.101.) Copyright 1995. Reprinted with permission).

The illusion of transparency was revealed in two ways. First, we reasoned that if participants think that one meaning of the idiom makes more sense than another, they should believe that an uninformed person would agree on the idiom's meaning. Indeed, once they learned a particular meaning they did think that an uninformed individual would be more likely to think that the idiom meant what they believed it to mean than they would the opposite meaning. We also directly tested how transparent a learned meaning becomes by asking participants to rate how much sense each meaning makes, and a learned meaning was perceived as much more sensible than its opposite (Figure 8.1). Moreover, the more they used the idiom with the learned meaning in mind, the less sense the alternative meaning made to them. These findings suggest that our participants thought the meaning they learned was "carried" by the expression and was more transparent than the alternative meaning. Furthermore, the feeling of transparency is to a large extent a function of the belief that the idiom has a particular meaning and of the effort involved in making sense of that meaning.

Perhaps our results could not possibly be generalized from unused, archaic idioms to those that are used in current language; perhaps the

example of "to go against the wind" is strong counterevidence. No native English speakers in their right minds would believe that had the idiom meant doing something with relative ease, this meaning would have made sense. Indeed, any idiom dictionary reveals numerous idioms of the form "to (verb) against (noun)," as in "to go against the grain," and all these idioms refer to an activity done the hard way. Such an objection, however, reflects the presence of the intuition of inflated transparency that we are arguing about. As an exercise, try to imagine a situation in which going against the wind is actually easier than going with it. Pilots know that a head wind facilitates takeoff and the same is true for trying to fly a kite. After this exercise, the alternative meaning makes a bit more sense. A hard-core skeptic might still argue that if this meaning makes sense, then there would be examples of such idioms in the language. In fact, there is at least one such example, in which an expression of the form "to (verb) against (noun)" is used to convey ease of movement. Commuters want to travel against traffic; the idea makes sense.

Transparency of Intentions. Our findings with idiom transparency reflect the way human minds attempt to make sense of linguistic expressions that result in overestimation of transparency of meaning. But this overestimation should not affect communication among native speakers of a language because the meanings of idioms are highly conventional. Most native speakers of English know the meaning of "to go against the wind," and consequently they share the meaning's sense of transparency. But when a native speaker speaking to a non-native speaker of English chooses to use an idiom, the speaker might expect that the other person understands the idiom because of the false sense of transparency. Yet the issue of overestimated transparency might be much more pervasive than in the relatively constrained case of idiomatic meaning. Pragmatic intentions are generally expected to be perceived as relatively transparent. Because every utterance can convey more than one speech act, the illusion of transparency of intention should be fairly pervasive, if we are correct.

In a study demonstrating that young children behave as if speakers' intentions are relatively transparent, Olson (1991; Olson & Torrance, 1987) described a situation in which experimenters told 5- to 8-year-old children a story about Lucy. Lucy has two pairs of red shoes, one old and one new. She wanted Linus to bring her the new shoes but asked him for her "red shoes." Even though the reference is ambiguous, the children tended to believe that the expression actually referred to the intended pair of shoes, as if they took the expression to be transparent vis-à-vis the actual *intention* of the speaker.

Of course, such a result can also be explained in other ways. The experiment might reveal nothing about overestimation of the transparency of intentions; instead it might reveal young children's confusion about who knows what.

They might have difficulty keeping track of what character in a story has access to what information, as much research on the development of theory of mind has suggested (e.g., Astington, Harris, & Olson, 1988; Perner, Leekam, & Wimmer, 1987; Wellman, 1990). Although this explanation is possible, I argue that this finding actually points to a fundamental aspect of language use, which extends beyond young children and holds for adults as well (Keysar, 1993). I argue that adults and children alike tend to perceive intentions as relatively transparent and by doing so underestimate the ambiguity of utterances. Adults and children differ, of course, but the main difference is that adults have acquired strategies to compensate for this basic tendency. Part of becoming a competent user of a language involves learning strategies to adjust from a relatively egocentric perspective and to correct for the perspective of others. I describe adults' basic illusion of transparency and then turn to the hypothesized adjustment processes.

What would it mean for adults to overestimate the transparency of intentions? Imagine the following situation: Jane and John go to a movie that is supposed to be fairly good. Both hate the movie because they think it pretentious and boring. As they leave the theater they spot a mutual friend, Mary, who is about to go to the next showing. John then comments to Mary, with a straight face and flat intonation, "You are just going to love this movie." John is being sarcastic, but, as the utterance was ambiguous, does Mary perceive the intended sarcasm? From Mary's perspective, because she does not know whether Jane and John actually liked the movie, both sarcastic and nonsarcastic interpretations are possible. Jane knows that Mary does not know John's true opinion of the movie. This situation is analogous to the young child's position, when the child knows that Linus does not know which shoes Lucy is asking for. My prediction is that Jane would underestimate the ambiguity of John's utterance because she knows what he intended. In this case, Jane knows that John was being sarcastic; thus she would tend to overestimate the extent to which Mary would be able to perceive his intended meaning.

I have documented this phenomenon in a variety of studies (Keysar, 1994b, 1995). In a typical experiment participants read scenarios such as this (for more examples see the appendix in Keysar, 1994b):

Mark asked his office mate, June, to recommend a restaurant; his parents were in town, and he wanted to take them to a good place. "I strongly recommend this new Italian place, called Venezia. I just had dinner there last night and it was marvelous. Let me know how you all enjoy it." That evening, Mark and his parents ate there; the food was unimpressive and the service was mediocre. The next morning, Mark said to her: "You wanted to know about the restaurant, well, marvelous, just marvelous."

In this scenario, the participants knew that Mark was being sarcastic because they knew that he did not enjoy the dinner. In a comparison scenario, the information was identical, except that Mark enjoyed the dinner. In this case, participants typically thought that Mark was not being sarcastic.

The participants' main task was to evaluate the addressee's perspective. In this case, they decided whether June would perceive Mark's message as sarcastic or sincere. In principle, if participants evaluate the utterance and whatever contextual information June might have, they should provide the same answer in both cases. Their evaluation of June's interpretation should not be affected by their privileged information about the valence of Mark's experience (positive or negative) or about what Mark actually intended. In contrast, if they take the utterance as less ambiguous after they know what Mark intended, then they should conclude that June will be more likely to perceive sarcasm when Mark actually intended it than when he did not. This is precisely what happened. In several experiments, participants tended to perceive the intended meaning as relatively transparent and to believe that addressees would be more likely to perceive it than an alternative meaning (Figure 8.2). The measure of transparency is the *difference* between the two conditions, not the absolute number: The phenomenon is demonstrated by the fact that with negative event information participants attribute *more* perception of sarcasm to an addressee than they do with positive event information.

There is one caveat here. I could explain the results differently by assuming that participants inferred that Mark conveyed the sarcasm with his intonation. If he did, then the results should depend on the fact that Mark said the utterance. To test this possibility we compared the same two conditions with an additional change: Instead of "telling" June that the dinner was marvelous, Mark left her a note. The results were identical. The phenomenon (i.e., the difference between the negative and positive information) was the same regardless of the modality of Mark's utterance. Therefore it is safe to conclude that these results genuinely reflect an illusion of transparency of intentions.

This conclusion led us to explore an even more radical possibility. Perhaps the illusion of transparency of intentions originates from a speaker's communicative intentions, or what Grice (1968) termed *m-intention*: an intention which is supposed to have its effect on an addressee by virtue of the addressee's realizing that the speaker wanted him or her to perceive this intention. For example, if I would like my 3-year-old to eat an apple, I could tell him, "Please eat the apple." My *m-intention* is fairly clear; I want him to realize that I would like him to eat the apple. Because he probably would not eat the apple under these circumstances, I might get him to do so by saying instead, "This apple is for me." In both cases I had an intention to get him to eat the apple, but only in the first case was the intention also an *m-intention*. In the experiment, Mark's sarcasm was a communicative intention. It is possible,

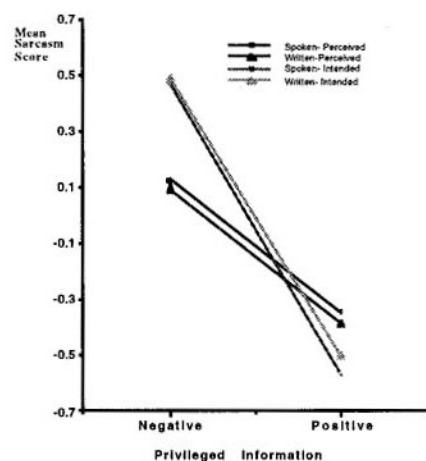


FIG. 8.2. Ratings of intended and perceived sarcasm. The intended meaning ratings indicate that speakers were perceived as more sarcastic when an event was negative than when it was positive. The perceived meaning ratings indicate that participants attributed to uninformed addressees the perception of sarcasm more when the event was negative than when it was positive—both when the message was spoken and when it was written. (From “The Illusory Transparency of Intention,” by B. Keysar [1994], *Cognitive Psychology*, 26, p. 183. Copyright 1994. Reprinted with permission.)

then, that participants perceived Mark’s sarcastic intention as transparent because it was an m-intention.

To test this hypothesis we separated the intended meaning from the information itself (Keysar, 1995). Recall the restaurant story. In addition to the two cases where the event was either positive or negative, we added a condition in which the event was negative: Mark hated the restaurant, but “Mark did not want June to feel bad.” Mark had a bad experience, but his utterance, “Marvelous, just marvelous,” was no longer sarcastic. Instead of a sarcastic m-intention, he was telling a white lie. When Mark attempted to conceal the way he really felt, his m-intention conveyed a sincere gratitude. As we expected, when we eliminated the sarcastic m-intention, there was no longer a difference in participants’ attribution of perceived intention to the addressee (Figure 8.3). Even though the privileged information was negative, participants did not use it when taking the addressee’s perspective. Instead, they relied on the inferred

communicative intention. Because the m-intention was the same with the positive event as with the white lie, they attributed the same perceived intention to the addressee. Perhaps, then, Gricean m-intentions are at the core of the illusion of transparency of intentions.

In a pilot study, we also demonstrated that the phenomenon generalizes to other speech acts. For example, in one scenario a speaker said, “I haven’t seen Donald for ages.” The speaker was either making an informative statement or an indirect request to be invited to a dinner with Donald. Overall, participants tended to attribute to the addressee the perception of the speaker’s actual intention, which was privileged to them.

But, a skeptic might say, perhaps the participants are simply confused about who knows what. If this is the case, then there is no illusion of transparency of intention, but instead participants have difficulty keeping track of who has access to what information. I propose that this is an unlikely explanation for the findings in the studies I mentioned. There are two main reasons to doubt it. First, Graesser, Bowers, Bayen, and Hu (in press) found that readers are very good at keeping track of who knows what and rarely confuse protagonists’ accessibility to information. Their findings suggest that people have the tools to identify who has access to what information and are able to keep this information in memory. What our findings suggest is that people do not use this information when they evaluate the interpretation of an utterance from the perspective of a protagonist. They do not discount the information privileged to

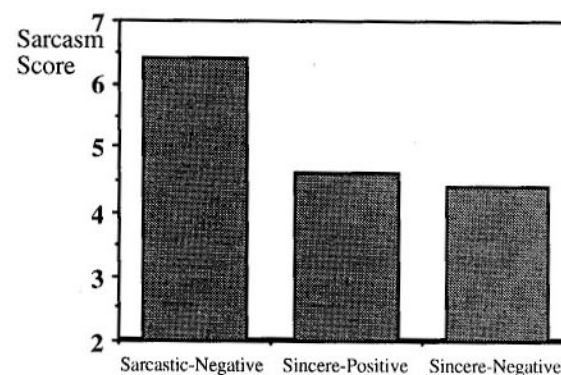


FIG. 8.3. Ratings of perceived sarcasm by addressee. The higher the number, the more participants thought that the addressee perceives sarcasm. Attribution of perception of sarcasm was a function of the speaker’s m-intention (sarcasm or sincerity), not a function of the event’s valence.

them, information that they know the protagonist does not know. Why? The reason, I suggest, is that they use their privileged information to disambiguate the speaker's utterance; afterward, the intention seems relatively transparent to them. The other reason to doubt this alternative explanation is that the phenomenon seems to be tightly related to inferred *m*-intentions, not simply to the privileged information about facts (e.g., good or bad experience at the restaurant). As Figure 8.3 illustrates, when Mark attempts to conceal the facts, participants no longer attribute the perception of sarcasm to the addressee. If they were simply confused about who knows what, their knowledge about the negative experience at the restaurant should have caused the same effect regardless of Mark's *m*-intention; in other words, the white lie condition and the negative information condition should have yielded similar results.

These findings support one aspect of the proposed model: They demonstrate that there is a general tendency to underestimate the ambiguity of utterances once the intended meaning is computed. Consequently, language users experience an illusion of transparency of intention.

Perspective Adjustments

Language users do know about perspective differences. They have metaknowledge about knowledge. Such metaknowledge includes knowledge about who knows what. They also know that they should not use information that is inaccessible to an addressee when taking the perspective of such addressee. In the model that we have proposed (Keysar, et al., 1997), we assume that language users make allowances for others' perspectives. In the case of addressees, whenever their perspectives diverge from a speaker's, then egocentric interpretation is adjusted. Overhearers, as in our sarcasm experiments, adjust to approximate the perspective of the addressee to evaluate the intention that the addressee would perceive. In general, we hypothesized an adjustment process that uses metaknowledge to compensate for perspective differences. In this sense, language users arrive at a mutual perspective only via a corrective process. The correction is a function of language users' application of normative knowledge about how communication takes place.

If language users make adjustments for perspective, why then do we find a relatively egocentric attribution of intended meaning in the sarcasm experiments? If they know that June does not know how much Mark enjoyed the dinner, the adjustment to June's perspective should lead them to the same answer when he liked the dinner and when he did not. The only way the model can explain the fact that we did find a difference between the two cases is by assuming that the adjustment was incomplete. Such incomplete adjustment processes have been demonstrated for a variety of judgment tasks (Tversky & Kahneman, 1974). In general, people's estimate of a particular value could be

anchored even in relatively irrelevant value information. In our case, we expected a person's own perspective to act as an effective anchor and we evaluated this possibility with the following assumptions: The adjustment process is relatively slow and cognitively-taxing whereas an egocentric interpretation of a speaker's intent is relatively quick and easy. If this assumption is correct, perspective-taking under pressure should affect the adjustment process and result in greater underadjustment than when participants perform the task without pressure.

To test this hypothesis we replicated the same experiment with the scenarios that had either negative or positive event information—Mark had either a good or a bad dinner experience. The crucial addition to the design was a manipulation of pressure to respond. In one case, participants were instructed to take their time and ponder their answer. In the other case they were required to answer as quickly as they could. As the model predicts, when participants were under time pressure to respond, the effect was much larger than when they responded at their leisure (Keysar, 1995). We interpreted this pattern to mean that under pressure, a participant is more likely to fall back on his or her own interpretation and attribute the perception of that interpretation to the addressee. This idea suggests that the perspective of the addressee was indeed considered only as part of a correction process. The time course of the speeded participants provided another piece of evidence. In the case when Mark hated the restaurant and left a sarcastic note, a participant might answer that June would perceive sarcasm or that she would not. According to the model, only the second answer (no sarcasm would be perceived) is the result of an adjustment, and so this answer should take longer than a response that does not require an adjustment. Indeed, responses that the model assumes are the result of adjustment took longer than responses that are less likely to be the result of such an adjustment process.

Do Addressees Adjust? Once we discovered that readers and observers follow the Perspective Adjustment model we were intrigued by the possibility that the model is a general model of linguistic perspective-taking. If so, addressees should follow the same course: They should construct interpretations that are perspective-bound, but should also use their metaknowledge to adjust to a speaker's or their mutual perspective.

Consider the way addressees disambiguate referential expressions: John welcomes his daughter's friend Mary into the house, and Mary says, "Is she ready?" How does John identify a referent for the pronoun *she*? In virtually any theory of pronoun resolution (e.g., Gernsbacher, 1989, 1990; McKoon, Ward, & Ratcliff, 1993), he searches his memory for potential female referents and selects his daughter, whom he knows is preparing to go out. But suppose that he also knows that his wife is getting ready to go out with him, and he knows that Mary does not know this. Clearly, Mary could not have referred to John's wife because she does not have access to the information that the wife is getting ready

at the same time. If John searches for potential referents of the speaker, his wife should not be in the referent set. In contrast, the Perspective Adjustment model makes the counterintuitive prediction that he might actually consider his wife as a potential referent of the pronoun *she*. In this case, the expression "Is she ready?" is ambiguous from John's perspective because it could refer either to his wife or to his daughter. The model also predicts that he will make allowances for Mary's perspective and eventually arrive at the intended referent, his daughter, via a correction process. If so, then the very presence of the information about his wife should result in interference.

These results occurred in an experiment reported in Keysar, et al. (1997). We modeled an experiment after the situation with John's wife and daughter and demonstrated that when privileged information provides a potential referent for a pronoun, addressees are slower to respond correctly and they make more errors than when the privileged information does not include a potential referent for the pronoun. This finding strongly suggests that addressees interpret utterances without regard to a speaker's perspective, but make adjustments when they are needed.

When People Speak. Research on language production has reported that speakers tailor their utterances to addressees (e.g., Clark & Marshall, 1981; Clark & Murphy, 1982; Krauss & Fussell, 1991; Levelt, 1989). For example, Krauss and Fussell (1991) discuss several experiments showing that speakers produce different descriptions depending on their intended addressees, and that intended addressees understand these utterances better than do nonintended addressees. Although current research demonstrates that speakers can take their addressees' perspective into account, it has not shown how they do it. In general, process models of production (e.g., Bock, 1995; Fromkin, 1971; Garrett, 1988; Levelt, 1989) have identified a variety of processes. Researchers have agreed that some processes are involved in conceptually designing an utterance, some in putting it in linguistic form, and some in monitoring and revising utterance plans. Speakers could use their addressees' perspectives or mutual perspectives when they design their utterances. In contrast, it is possible that speakers incorporate their addressees' perspective only as part of a correction mechanism. Because utterance production need not be a mirror image of utterance comprehension, both are viable options. Horton and Keysar (1996) hypothesized that when speakers plan their utterances, these plans are not audience-sensitive; they are not designed for specific addressees in a way that takes into account their mutual perspective. We also argued that as part of utterance monitoring, speakers detect and correct utterances that violate principles of audience design (Clark & Murphy, 1982). This "monitoring and adjustment" model, then, assumes that speakers take into account their addressees' perspectives only as part of the monitoring process.

To investigate the predictions of the model, Horton and Keysar (1996) asked people to describe simple shapes for addressees. The shapes appeared in a context that was either privileged to the speakers or shared with the addressees. For example, people described a circle in the context of a larger circle, and the context circle was either seen by both speaker and addressee or only by the speaker. We measured the extent to which the speakers relied on context by looking at the adjectives they used. In this case, referring to the circle as "small" indicated use of context. Overall, speakers relied on shared context more than they relied on privileged context. This fact is not surprising because they knew that the privileged context was inaccessible to their addressees. The question is, how did speakers incorporate this perspective information?

On the basis of our model, the difference between the shared and privileged context conditions is due to the monitoring function. Therefore, the model predicts that under pressure to initiate an utterance, speakers would be less able to monitor and that consequently they would not adjust for their addressees' perspectives. Indeed, when we asked speakers to initiate their utterances within 1.5 seconds the differences between context conditions disappeared. Speakers relied on shared and privileged contexts to the same degree (See Figure 8.4). We interpreted this finding to mean that under pressure, speakers produce an utterance that is a closer reflection of the utterance plan. This utterance plan does not take into account another's perspective. Only when speakers monitor do they detect cases that require adjustment and correct for perspective discrepancies. We presented evidence for pre-articulatory monitoring and correction, but it is possible that overt repairs function similarly.

Perspective Adjustment and Collaboration. The study of conversation has given rise to what Krauss and Fussell (1996) call "dialogic models" (e.g., Clark & Schaefer, 1989; Clark & Wilkes-Gibbs, 1986; Wilkes-Gibbs & Clark, 1992). Such models focus on the collaborative nature of a conversational interchange. By analyzing conversations, these researchers have demonstrated the ways that participants collaborate to arrive at a successful interaction. For example, Clark and Schaefer (1989) showed that conversations are composed of units of presentation and acceptance, where one party presents an utterance to the other and the conversation progresses only after the other party signals that he or she understands and therefore accepts the presentation. Wilkes-Gibbs and Clark (1992) further demonstrated that interlocutors are sensitive to their party's knowledge of referents: They provide longer descriptions of referents when their addressee is new to the conversation than when they had already conversed with the addressee about the topic.

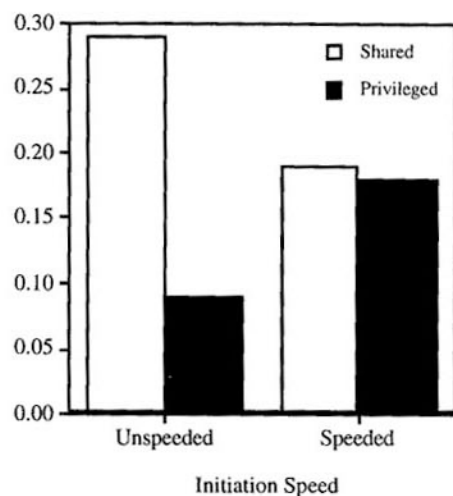


FIG. 8.4. Mean ratio of context-related adjectives to the total number of adjectives plus nouns per description as an index of reliance on context. Although speakers relied on shared context more than on privileged context when not speeded, they relied on shared and privileged context to the same degree when speeded in initiating their utterances. (From "When do Speakers Take Into Account Common Ground," by W. S. Horton & B. Keysar [1996], *Cognition*, 59, p. 108. Copyright 1996. Reprinted with permission.)

These findings or this approach may seem inconsistent with the model I have proposed because conversation appears to be geared toward shared meaning and interlocutors seem to be sensitive to the knowledge of their conversational partners. Such an interpretation is misguided. The fact that interlocutors collaborate is completely consistent with the Perspective Adjustment model. When an addressee interprets an utterance, the model assumes an egocentric interpretation, but also assumes a check vis-à-vis the knowledge of a speaker. If such a check either detects a violation of common ground or the check is uncertain about the validity of the interpretation, then the addressee pursues the path described by the Collaboration model—he or she asks questions of clarification. Similarly with speakers: They may monitor an utterance and realize that they are not sure about an addressee's ability to comprehend it. Then they may inquire, as the Collaboration model suggests. The main difference between the Perspective Adjustment model and the Collaboration model is in the level of

analysis. The Perspective Adjustment model describes cognitive processes whereas the Collaboration model describes higher level, interpersonal moves.

Evidence presented in support of the Collaboration model is also consistent with the Perspective Adjustment model, for instance, Wilkes-Gibbs and Clark's (1992) report that interlocutors are sensitive to another's perspective. Our model can explain the results as a function of speakers' monitoring. As in Horton and Keysar's (1996) study, although speakers showed sensitivity to their addressees' perspectives, they did so only as part of monitoring, not as part of their utterance planning. They used their intuitive theory of conversation to guide them in evaluating their utterances and in correcting them when needed. This result could have occurred in the Wilkes-Gibbs and Clark (1992) study, although this conclusion is uncertain because their study focused on final products, on the conversation itself. If we can investigate the cognitive processes underlying the conversation in the Wilkes-Gibbs and Clark experiment, we might be able to determine the extent to which sensitivity to addressees guides a speaker's very production or acts only as a corrective measure.

Dialogic models bring up an important issue that I do not address in this paper, namely, the role of feedback and the functioning of comprehension mechanisms in an ongoing dynamic conversation. To date, this question is interesting, yet open. How do feedback and the possibility of feedback affect comprehension processes? Would addressees engage in adjustment processes, as we found in our experiments, if they are part of an ongoing conversation? Would they anchor on egocentric interpretations as we found in our studies? It would be interesting to investigate the extent to which comprehension is sensitive to situational demands such as the possibility of feedback. We know of only one piece of evidence related to this issue, but it is about production, not comprehension. Schober (1993) found that speakers provide more egocentric spatial descriptions when they speak to a *real* addressee than to an imaginary one. Perhaps speakers might "relax" their laborious monitoring when they have a real addressee. Real addressees can give them feedback when they do not understand; therefore, speakers need not be too vigilant about monitoring for common ground. Consequently, their utterances are more likely to reflect their inherently egocentric plans.

CONCLUSIONS

The experiments reviewed here suggest that two major processes are involved in language comprehension. Addressees, as well as observers and readers, quickly interpret utterances in a way that is contextualized but bound by the perspective of the comprehender. They also adjust to another's perspective by using metaknowledge about what is mutually known. An analogous model holds for

speakers as well: They only take into account their addressees' perspectives when they monitor utterances, and correct them if needed.

This model could be refuted and counterevidence could be provided by showing that mutual perspective guides comprehension so that no egocentric interpretation takes place. To date, there is little direct evidence for the claim that comprehenders use mutual knowledge to interpret utterances, because studies investigating common ground and mutual knowledge *confound* mutual knowledge with self-knowledge. A typical experiment has attempted to demonstrate that an addressee used common ground by showing that the addressee used mutually known information. For example, participants were presented with a picture of flowers with a prominent flower at the center. The experimenters then asked an ambiguous question, "How would you describe the color of this flower?" (Clark, Schreuder, & Buttrick, 1983). Although a definite reference requires a unique referent, participants had no problem identifying "the flower" uniquely as the most salient flower. These results could be interpreted in terms of mutual perspective: Perhaps because participants knew that the flower is mutually salient, they were able to understand the question as referring to this flower. Although this interpretation is appealing, it is unwarranted by the data (Keysar, in press); a simpler theory that does not assume the use of mutual perspective can account for the data as well. Participants possibly interpreted the utterance egocentrically; because the flower was salient to them, they arrived at a unique interpretation. In other words, it is not sufficient to demonstrate that a comprehender used mutual knowledge to argue that mutual perspective guided the interpretation or that the information was used because it was mutual. Because every mutually known fact is also by definition known to an individual addressee, the addressee might have used the information because it was known to him or her, not because it was mutual. So this inherent conceptual confound must first be resolved to provide evidence relevant to the issue of common ground or common perspective.

There is a second reason why demonstrating that mutual knowledge is insufficient to refute the proposed model. Even if mutual knowledge is used because it is mutual, it is still consistent with our Perspective Adjustment model. The model does not assume that interlocutors do not use this metaknowledge; instead, it describes the processes underlying such use and the conditions under which these processes occur. The model assumes that an addressee uses mutual perspective but only as a corrective measure. To refute the model, the mutual perspective must be shown to play any other role, for example, to guide or restrict the very interpretation so that no correction is needed.

This account does not describe how language users arrive at the interpretation that I claim to be the basic building block. True, our model does

not attempt to answer the question, "How do addressees create interpretations from their own perspective?" Such an account is not very easy to construct. One possibility is Sperber and Wilson's theory of relevance (Sperber & Wilson, 1982, 1986), but the exact predictions of relevance theory are still to be spelled out and tested in this context (see the debate in Sperber & Wilson, 1987). My current attempt is simply to suggest that perceived intentions are a basic building block of language use. This notion is a particular interpretation of Clark's (1992) assumption that "in language use, speaker's meaning is primary, and word or sentence meaning are derivative" (p. xv). We assume that computing a speaker's intention is what the comprehension system does well and fast. Although the system computes intentions, this computation does not involve considerations of mutual knowledge: Instead, it does the computation from the perspective of an addressee. The system also does not, as the traditional approach suggests, compute decontextualized meanings, literal meanings, or otherwise perspective-free meanings.

I propose that such quick and efficient computation of intention is one important way that our comprehension system tackles the inherent problem of ambiguity. People are basically insensitive to the extent of ambiguity because they comprehend utterances from their own perspectives. Consequently, intentions seem relatively transparent and utterances seem less ambiguous than they really are. The illusion of transparency that I have described is a natural outcome of the way the comprehension system works to reduce ambiguity. Using a personal perspective does not always eliminate ambiguity completely, but it reduces it drastically and makes the interpretation process feasible.

Language users could have reduced ambiguity differently; they could have arrived at intentions by using other perspectives. For example, addressees might use speakers' perspectives, their mutual perspectives, or mutual knowledge to interpret utterances. The Perspective Adjustment model and our data (Keysar et al., 1997) suggest that addressees do not. Comprehenders do not use speakers' or mutual perspectives to restrict ambiguity and arrive at intended meaning. The model assumes that they do use mutual perspective, but only when there is a need. They use it as a corrective measure.

Perceiving intentions as transparent seems to be part of acquiring language, as Olson and Torrance's (1987) young participants demonstrated. Adults might be perceiving Lucy's request for her "red shoes" as relatively unambiguous as well, because they know her intent. But, because they know that Linus does not know her actual intention, they know that from his perspective the utterance is actually ambiguous, and they make allowances for the perspective discrepancy. Therefore, I propose that language users solve a different problem from the one that the traditional theory assumes. Language users do not solve the problem, "How does literal meaning relate to speaker's meaning?" Instead, they solve the problem, "How does the contextualized computed intention relate to the

speaker's actual intent?" Typically, the "transparent" intention is identical to the speaker's actual intent; sometimes it is not. When an interpretation is erroneous and comprehenders detect the error, they adjust.

ACKNOWLEDGMENTS

The preparation of this paper was supported by Public Health Service, grant R29 MH49685, to the University of Chicago, Boaz Keysar, Principal Investigator. I am grateful to Public Health Service for this support. I am indebted to Linda Ginzel and William Horton for thoughtful comments on earlier drafts and to my graduate students and research assistants who helped me think about these issues.

REFERENCES

- Anderson, R. C., & Pichert, J. W. (1978). Recall of previously unrecalled information following a shift in perspective. *Journal of Verbal Learning and Verbal Behavior*, 17, 1-12.
- Astington, J., Harris, P., & Olson, D. (Eds.) (1988). *Developing theories of mind*. New York: Cambridge University Press.
- Ayer, A. (1936). *Language, truth and logic*. London: Gollancz.
- Bock, K. (1995). Sentence production: From mind to mouth. In J. Miller & P. Eimas (Eds.), *Handbook of perception and cognition: Speech, language, and communication*. San Diego, CA: Academic Press.
- Carnap, R. (1956). *Meaning and necessity* (2nd ed.). Chicago: Chicago University Press.
- Clark, H. H. (1992). *Arenas of language use*. Chicago: University of Chicago Press.
- Clark, H. H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. H. Joshi, B. Webber, & I. A. Sag (Eds.), *Elements of discourse understanding*. Cambridge, England: Cambridge University Press.
- Clark, H. H., & Murphy, G. L. (1982). Audience design in meaning and reference. In J. F. Le Ny & W. Kintsch (Eds.), *Language and comprehension*. Amsterdam, Netherlands: North-Holland.
- Clark, H. H., & Schaefer, E. F. (1989). Contributing to discourse. *Cognitive Science*, 13, 259-294.
- Clark, H. H., Schreuder, R., & Buttrick, S. (1983). Common ground and the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behavior*, 22, 245-258.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22, 1-39.
- Cox, M. V. (1977). Perspective ability: The relative difficulty of the other observer's viewpoints. *Journal of Experimental Child Psychology*, 24, 254-259.
- Dascal, M. (1987). Defending literal meaning. *Cognitive Science*, 11, 259-281.
- Fromkin, V. (1971). The nonanomalous nature of anomalous utterances. *Language*, 47, 27-52.
- Garrett, M. F. (1988). Processes in language production. In F. J. Newmeyer (Ed.), *Linguistics: The Cambridge survey, Vol. 3, Language: Psychological and biological aspects*. Cambridge, England: Cambridge University Press.
- Gernsbacher, M. A. (1989). Mechanisms that improve referential access. *Cognition*, 32, 99-156.
- Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gernsbacher, M. A., & Keysar, B. (1995, November). *The role of suppression in metaphor interpretation*. Paper presented at the 36th annual meeting of the Psychonomic Society, Los Angeles, CA.
- Gerrig, R. J. (1989). Empirical constraints on computational theories of metaphor: Comments on Indurkha. *Cognitive Science*, 13, 235-241.
- Gibbs, R. (1979). Contextual effects in understanding indirect requests. *Discourse Processes*, 2, 1-10.
- Gibbs, R. W. (1983). Do people always process the literal meanings of indirect requests? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 9, 524-533.
- Gibbs, R. W. (1984). Literal meaning and psychological theory. *Cognitive Science*, 8, 275-304.
- Gibbs, R. W. (1990). Comprehending figurative referential descriptions. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16, 56-66.
- Gibbs, R. W. (1992). What do idioms really mean? *Journal of Memory and Language*, 31, 485-506.
- Gibbs, R. W. (1994). *The poetics of mind*. Cambridge, England: Cambridge University Press.
- Gibbs, R. W., & Nayak, N. P. (1991). Why idioms mean what they do. *Journal of Experimental Psychology: General*, 120, 93-95.
- Gibbs, R. W., & O'Brien, J. E. (1990). Idioms and mental imagery: The metaphorical motivation for idiomatic meaning. *Cognition*, 36, 35-68.
- Glucksberg, S., Gildea, P., & Bookin, H. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior*, 21, 85-98.
- Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*, 97, 3-18.
- Graesser, A. C., Bowers, C., Bayen, U. J., & Hu, X. (in press). Who said what? Who knows what? Tracking speakers and knowledge in narrative. In W. van Peer, E. Andrica, D. Schram, & E. Tan (Eds.), *Narrative perspective: Cognition and emotion*.
- Graumann, C. F. (1989). Perspective setting and taking in verbal interaction. In R. Dietrich & C. F. Graumann (Eds.), *Language processing in social context*. Amsterdam, Netherlands: Elsevier Science.
- Graumann, C. F., & Herrmann, T. (1988). Other-relatedness in language processing. *Journal of Language and Social Psychology*, 7, 159-168.
- Grice, H. P. (1968). Utterer's meaning, sentence meaning, and word meaning. *Foundations of Language*, 4, 225-242.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics: Vol. 3, Speech acts*. New York: Academic Press.
- Horton, W. S., & Keysar, B. (1996). When do speakers take into account common ground? *Cognition*, 59, 91-117.
- Katz, J. (1972). *Semantic theory*. New York: Harper & Row.

- Katz, J., & Fodor, J. (1963). The structure of semantic theory. *Language*, 36, 170-210.
- Keysar, B. (1989). On the functional equivalence of literal and metaphorical interpretations in discourse. *Journal of Memory and Language*, 28, 375-385.
- Keysar, B. (1993). Common sense and adult theory of communication. *Behavioral and Brain Sciences*, 16, 54.
- Keysar, B. (1994a). Discourse context effects: Metaphorical and literal interpretations. *Discourse Processes*, 18, 247-269.
- Keysar, B. (1994b). The illusory transparency of intention: Linguistic perspective taking in text. *Cognitive Psychology*, 26, 165-208.
- Keysar, B. (1995). *Exploring linguistic perspective taking*. Unpublished manuscript, University of Chicago.
- Keysar, B. (1996). From figures of speech to figments of thought. [Review of the book: *The poetics of mind: Figurative thought, language, and understanding*.] *Contemporary Psychology*, 41, 222-223.
- Keysar, B. (in press). Unconfounding common ground. *Discourse Processes*.
- Keysar, B., Barr, D. J., Balin, J. A., and Paek, T. (1997). *Definite reference and mutual knowledge: A processing model of common ground in comprehension*. Manuscript submitted for publication.
- Keysar, B., & Bly, B. (1995). Intuitions of the transparency of idioms: Can one keep a secret by spilling the beans? *Journal of Memory and Language*, 34, 89-109.
- Keysar, B., & Bly, B. (in press). Swimming against the current: Do idioms reflect conceptual structure? *Journal of Pragmatics*.
- Keysar, B., & Glucksberg, S. (1993). Metaphor and communication. *Poetics Today*, 13, 633-658.
- Kosinski, J. (1971). *Being there*. New York: Bantam Books.
- Krauss, R. M., & Fussell, S. R. (1988). Other-relatedness in language processing: Discussion and comments. *Journal of Language and Social Psychology*, 7, 263-279.
- Krauss, R. M., & Fussell, S. R. (1991). Perspective-taking in communication: Representations of others' knowledge in reference. *Social Cognition*, 9, 2-24.
- Krauss, R. M., & Fussell, S. R. (1996). Social psychological models of interpersonal communication. In E. T. Higgins & A. Kruglanski (Eds.), *Social psychology: Handbook of basic principles*. (pp. 655-701). New York: Guilford Press.
- Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago: University of Chicago Press.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
- Long, E. C. J., & Andrews, D. W. (1990). Perspective taking as a predictor of marital adjustment. *Journal of Personality and Social Psychology*, 59, 126-131.
- McKoon, G., Ward, G., & Ratcliff, R. (1993). Morphosyntactic and pragmatic factors affecting the accessibility of discourse entities. *Journal of Memory and Language*, 32, 56-75.
- Mead, G. H. (1934). *Mind, self and society*. Chicago: University of Chicago Press.
- Murphy, G. L. (in press). On metaphoric representations. *Cognition*, 60, 173-204.

- Nusbaum, H. C., & Henly, A. S. (in press). Understanding speech perception from the perspective of cognitive psychology. In J. Charles-Luce, P. A. Luce, & J. R. Sawusch (Eds.), *Spoken language processing*. Ablex.
- Olson, D. R. (1991). Children's understanding of interpretation and the autonomy of written texts. *Text*, 11, 3-23.
- Olson, D. R., & Torrance, N. (1987). Language, literacy, and mental states. *Discourse Processes*, 10, 157-168.
- Onishi, K. H., & Murphy, G. L. (1993). Metaphoric reference: When metaphors are not understood as easily as literal expressions. *Memory & Cognition*, 21, 763-772.
- Ortony, A., Schallert, D. L., Reynolds, R. E., & Antos, S. J. (1978). Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior*, 17, 465-477.
- Perner, J., Leekam, S., & Wimmer, H. (1987). Three-year-olds' difficulty with false belief: The case for a conceptual deficit. *British Journal of Developmental Psychology*, 5, 125-137.
- Piaget, J., & Inhelder, B. (1956). *The child's conception of space*. New York: Norton.
- Rommetveit, R. (1974). *On message structure: A framework for the study of language and communication*. New York: Wiley.
- Schober, M. F. (1993). Spatial perspective-taking in conversation. *Cognition*, 47, 1-24.
- Searle, J. R. (1975). Indirect speech acts. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3, Speech acts*. New York: Academic Press.
- Sperber, D., & Wilson, D. (1982). Mutual knowledge and relevance in theories of comprehension. In N. Smith (Ed.), *Mutual knowledge*. London: Academic Press.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
- Sperber, D., & Wilson, D. (1987). *Precis of relevance: Communication and cognition*. *Behavioral and Brain Sciences*, 10, 697-754.
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268, 1632-1634.
- Traxler, M. J., & Gernsbacher, M. A. (1993). Improving written communication through perspective-taking. *Language and Cognitive Processes*, 8, 311-334.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.
- Wellman, H. M. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.
- Wilkes-Gibbs, D., & Clark, H. H. (1992). Coordinating beliefs in conversation. *Journal of Memory and Language*, 31, 183-194.
- Yaniv, I., & Shatz, M. (1990). Heuristics of reasoning and analogy in children's visual perspective taking. *Child Development*, 61, 1491-1501.