Available online at www.sciencedirect.com
SCIENCE DIRECT®
COGNITION
ELSEVIER
Cognition 96 (2005) 127–142
www.elsevier.com/locate/COGNIT

The impact of memory demands on audience design during language production

William S. Horton*, Richard J. Gerrig

State University of New York at Stony Brook, Stony Brook, NY, USA

Received 2 October 2003; revised 12 March 2004; accepted 21 July 2004

Abstract

Speakers often tailor their utterances to the needs of particular addressees—a process called *audience design*. We argue that important aspects of audience design can be understood as emergent features of ordinary memory processes. This perspective contrasts with earlier views that presume special processes or representations. To support our account, we present a study in which Directors engaged in a referential communication task with two independent Matchers. Over several rounds, the Directors instructed the Matchers how to arrange a set of picture cards. For half the triads, the Directors' card categories were initially distributed orthogonally by Matcher (e.g. Directors described birds and dogs with one Matcher and fish and frogs with the other). For the other triads, the Directors' card categories initially overlapped across Matchers (e.g. Directors described two members of each category with each Matcher). We predicted that the orthogonal configuration would more readily allow Directors to encode associations between particular cards and particular Matchers—and thus allow those Directors to provide more evidence for audience design. Content analyses of Directors' utterances from two final rounds supported our prediction. We suggest that audience design depends on the memory representations to which speakers have ready access given the time constraints of routine conversation.

Keywords: Language production; Memory; Audience design; Communication

* Corresponding author. Address: Georgia Institute of Technology, School of Psychology, 654 Cherry Street, Atlanta, GA 30332-0170, USA. Tel.: +1 404 385 0008; fax: +1 404 894 8905.
E-mail address: william.horton@psych.gatech.edu (W.S. Horton).

0022-2860/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2004.07.001

Consider this pair of expressions referring to a picture of a bird, produced on separate occasions by the same speaker for two different addressees:

(1) a. *the black bird with uh spots of white and he's looking down to the right*
b. *the firecracker bird*

Based on these descriptions, each addressee was able to identify the correct bird. Even so, the speaker clearly chose to provide different information in the two cases, despite having already described the picture several times previously in each instance. We suggest that the speaker modified these descriptions to suit the (perceived) communicative needs of each addressee. In the first case, it was the first time that the picture was being described for that particular addressee, and so the speaker provided information that was generically descriptive. In the second case, however, the speaker had used *firecracker bird* to refer to the picture with that addressee before. This general phenomenon—speakers taking their current conversational partner into account when formulating utterances—is called *audience design* (Clark, 1996; Clark & Carlson, 1982; Clark & Murphy, 1982).

Because it is possible to find evidence of audience design in virtually every instance of language use, there is little controversy as to its importance as a feature of conversation. What has remained controversial, however, is the exact nature of the mechanisms that enable speakers to engage in audience design (Brown & Dell, 1987; Horton & Gerrig, 2002; Horton & Keysar, 1996; Lockridge & Brennan, 2002; Schober & Brennan, 2003). Understanding audience design requires understanding how language users take into account the information they believe they share with their conversational partners—their common ground.

Previous theoretical approaches to the issue of common ground, however, have provided only partial descriptions of the cognitive mechanisms that mediate the ability of language users to produce and understand language with respect to others. For example, in the influential account of common ground and definite reference proposed by Clark and Marshall (1981), the use of common ground information is characterized as being dependent upon specialized memory structures—*reference diaries*—that encode exclusively for partner-related information. Under this view, individuals are presumed to routinely search their diary entries for records of past events that provide evidence for whether particular information can be taken as shared with others. Although Clark and Marshall's notion of reference diaries recognizes the importance of memory for conversational interaction, we believe that it does not provide a general solution to the problem of common ground (Horton & Gerrig, in press). For example, it is not clear at what level of specificity mutually-experienced events are to be encoded for the purposes of establishing commonality. That is, what range of circumstances would require individuals to *infer* that particular knowledge can be taken as shared rather than having this information directly available in the reference diary? Also, what are the boundaries on events such that they can be encoded as such in a diary entry? For a given episode like "the birthday party," would evidence about commonality have to be encoded separately for different moments within the event like "the opening of the gift" and "the presentation of the birthday cake?" This problem becomes even more apparent for multiparty experiences, which would appear to require information be encoded separately with respect to each participant.

An alternative approach to the issue of common ground in language use comes from Keysar and colleagues (Barr & Keysar, 2002; Horton & Keysar, 1996; Keysar, Barr, Balin, & Brauner, 2000), who suggest that the bulk of language processing occurs without regard to beliefs about common ground. In this view, language use is anchored in the information that is most immediately accessible to individuals, which necessarily includes their private knowledge. As a result, language use is postulated to have an egocentric basis, with partner-related adjustments carried out only as a consequence of monitoring the products of initial processing (Keysar, Barr, & Horton, 1998). An advantage of this approach is that it makes testable predictions concerning the time course with which partner-specific information has an impact upon language processing. However, this position has limitations as well. Not only is common ground still viewed as a specialized resource of partner-specific knowledge, but the mechanisms responsible for "checking" the products of initial processing with respect to common ground remain unclear.

Not surprisingly, these conflicting accounts have led to debates concerning the ubiquity and time course of common ground effects in language use (e.g. Barr & Keysar, 2002; Hanna, Tanenhaus, & Trueswell, 2003; Metzing & Brennan, 2003). These debates have been fostered in part by the fact that existing views have tended to treat common ground as somehow exceptional with respect to ordinary psycholinguistic processing. In contrast, the account of audience design motivating the current research departs from these prior views by eliminating the need to presume either special representations or special processes. In general, we wish to "normalize" audience design as a straightforward consequence of how partner-specific information is encoded and retrieved during routine utterance planning. The goal of this article, therefore, is to provide data in favor of an account of audience design that can be integrated with other aspects of language production, and with cognitive psychology more generally.

1. Audience design and language production

We conceptualize audience design as one of a number of constraints that potentially have an impact on the surface realization of utterances. In a general sense, theories of language production characterize the types of information that are brought to bear on utterances before they reach threshold for articulation (Bock & Levelt, 1994; Levelt, 1989). As such, many of the features of utterances can be predicted from the time course with which information impinges on production processes (e.g. Bock, 1987; Ferreira & Dell, 2000). On our view, audience design is subject to similar constraints (Horton & Gerrig, in press). We propose that people will perform utterances that show evidence of audience design whenever the memory representations that permit audience design become accessible with the appropriate time course. Given the time constraints of spontaneous conversation, this suggests that whether a given utterance reflects partner-specific considerations may often depend on the immediacy with which suitable memory representations are available as input for ongoing production processes (cf. Bard & Aylett, 2000; similar proposals in the context of language comprehension have been offered by Hanna et al. (2003) and Metzing and Brennan (2003)).

On this account, a speaker's ability to produce evidence of audience design depends critically on the accessibility of suitable memory representations. We propose that the kind of experiences that speakers have with different interlocutors is an important factor that affects the relative accessibility of particular memory representations. As an example of what we mean, consider a situation in which a speaker, named Alex, has talked about birds on numerous occasions with her friend Oliver, but has never discussed birds with her friend Ned. If Alex wishes to talk about birds with Oliver again, conversational precedents should be relatively accessible in memory, especially for idiosyncratic terms like *firecracker bird*. If she wishes to talk about birds with Ned, however, there should be no (or few) such precedents available. (More formally, Oliver should act as a cue to make bird-related memories more accessible whereas Ned should not.) Suppose, however, that Alex has frequent conversations about movies with both Oliver and Ned. In these circumstances, because both Oliver and Ned serve as cues for movie-related memories, there are more likely to be instances in which Alex is unable to ascertain—particularly by the time she wants or needs to perform an utterance—which memories are associated with each individual.

An advantage of our proposal is that it incorporates important features of the prior views of common ground described above. Like the view proposed by Clark and colleagues, it suggests that, given the right sort of cues, common ground can have an early and immediate impact upon audience design. And like the view espoused by Keysar and colleagues, it sees language use as opportunistic, using whatever information is most accessible within the time course necessary for utterance planning. We explicitly assume, however, that the representations taken as input into the language production system are domain-general episodic traces, and that the processes that act upon these representations work in a similar manner as other sorts of memory-dependent processes, such as source memory (Johnson, Hashtroudi, & Lindsey, 1993) and context-dependent memory (Smith, 1994). The essence of our claim is that conversational partners act as contextual cues for the retrieval of associated information. When the situation is such that particular partners serve as a highly-effective cues for the retrieval of relevant knowledge, utterances should more likely reflect partner-specific considerations.

2. The current experiment

To provide evidence for this claim, we carried out an experiment in which participants, playing the role of Director in a referential communication task, repeatedly instructed two other participants, serving as Matchers, how to arrange arrays of picture cards. Crucially, Directors in this task were given one of two possible experiences as they initially described cards for each Matcher. Some Directors had the experience of describing different categories of cards with different Matchers. We refer to this situation as the *Orthogonal* card condition. Other Directors, however, described the same card categories with both Matchers. We refer to this situation as the *Overlapping* card condition. We expected these two situations to place different memory demands upon Directors. In the Orthogonal case, Directors had the opportunity to use the higher-level card categories as a metric for associating particular types of cards with specific Matchers. That is, Directors in this

condition merely needed to encode, for example, that Matcher A has the birds and dogs and Matcher B has the fish and frogs. In the Overlapping case, however, the equivalent memory encoding should have been more difficult. Here, category information would be relatively useless by itself for keeping track of who has what. Instead, Directors in this condition were called upon to associate specific members of each card category with a particular Matcher (e.g. Matcher A has the big shaggy dog with a tail while Matcher B has the smaller dog with a collar).

To test whether these different memory demands have an impact upon speakers' ability to adjust their utterances for particular partners, we then had Directors carry out a final round with each Matcher in which they described card arrays that included complete sets of all four critical card categories. Thus, from the perspective of the Matchers, some of the cards in this last round were cards that they had matched previously, while other cards were cards that they were dealing with for the first time. With respect to audience design, then, we were interested in whether there would be any differences in how Directors described "old" versus "new" cards in these final two rounds. If the distribution of experiences in the initial communicative situation makes it more difficult to encode particular categories with respect to particular conversational partners, then the Directors' utterances in the Overlapping card condition should be relatively less likely to exhibit evidence of audience design.

To investigate the impact of our experimental manipulations, we considered two sorts of evidence. First, we examined the time course with which Directors initiated their descriptions for each card. Through this measure, we assessed whether Directors chose to prolong the amount of time they spent planning their descriptions when placed in more demanding conversational circumstances. Second, we also examined whether Directors made adjustments to the content of their descriptions that reflected the history of interaction they shared with particular addressees. If such adjustments are more apparent in circumstances in which card categories were initially clearly differentiated by partner, this will provide evidence that memory accessibility has an influence upon audience design.

3. Method

3.1. Participants

Seventy-two undergraduate students (24 groups of three) from the State University of New York at Stony Brook volunteered for this study as partial fulfillment for a psychology course requirement. All were native speakers of English.

3.2. Materials

We created two identical sets of picture cards, one for the Director and one for the Matchers, by attaching color pictures of living things onto 4-in.×6-in. index cards. Each card set contained 16 experimental and 16 filler items. The experimental items consisted of four dogs, four birds, four fish and four frogs. The filler items consisted of four cats,

four flowers, four lizards, and four turtles. We chose the individual members of each card category such that there was a reasonable amount of cross-item similarity, reducing the possibility that participants would be able to use unique descriptive features to identify a particular card.

From the Directors' set of 16 experimental cards, we generated 12 different item lists. For the Orthogonal condition, we created lists by assigning different card categories to different Matchers. For example, the cards to be matched with Matcher A might include all four birds and all four dogs, and the cards to be matched with Matcher B might include all four fish and all four frogs. Pairing all possible combinations of card categories produced six lists. For the Overlapping card condition, we created lists in which card categories were split between Matchers. Thus, the Directors' experimental cards were subdivided into two sets, one for each Matcher, with each set containing two dogs, two birds, two fish, and two frogs. Again, pairing different combinations of card categories produced six lists. Table 1 presents an example of how the experimental cards were distributed across all 10 rounds in both the Orthogonal and Overlapping card conditions. The 16 filler cards were similarly subdivided into either orthogonal or overlapping card category combinations. All participants saw the same orthogonal or overlapping distributions of filler cards, which served to fill out the complete array of 16 cards during the initial rounds of the experiment.

3.3. Apparatus

The Director sat on one side of a long table in front of a vertical barrier that blocked his or her view of the Matcher. We used colored poster board and black tape to subdivide the Director's side of the table top into two regions—a central section containing a 4×4 grid (with each cell of the grid sized to the dimensions of a single card) and a peripheral section above and to the sides of the central grid, large enough to contain the 16 cards *not* being described in each round. The Matchers' side of the table also contained a 4×4 grid, used to place successfully matched cards during the course of each round. The participants'

Table 1
Sample orthogonal and overlapping distributions of cards matched by participant pairs in each round

	Current Matcher	Orthogonal Card Distribution		Overlapping Card Distribution	
Rounds 1–4	Matcher A	Bird 1 2 3 4		Bird 1 2	Fish 1 2
		Dog 1 2 3 4		Dog 1 2	Frog 1 2
		plus eight filler cards		*plus eight filler cards*	
Rounds 5–8	Matcher B		Fish 1 2 3 4	Bird 3 4	Fish 3 4
			Frog 1 2 3 4	Dog 3 4	Dog 3 4
		plus eight filler cards		*plus eight filler cards*	
Round 9	Matcher A[a]	Bird 1 2 3 4	Fish **1 2 3 4**	Bird 1 2 **3 4**	Fish 1 2 **3 4**
		Dog 1 2 3 4	Frog **1 2 3 4**	Dog 1 2 **3 4**	Frog 1 2 **3 4**
Round 10	Matcher B[a]	Bird **1 2 3 4**	Fish 1 2 3 4	Bird **1 2** 3 4	Fish **1 2** 3 4
		Dog **1 2 3 4**	Frog 1 2 3 4	Dog **1 2** 3 4	Frog **1 2** 3 4

Note. Each number following the card categories refers to a different category exemplar. The numbers marked in bold in Rounds 9 and 10 represent the cards that were coded as "new" for this configuration of cards.

[a] For Rounds 9 and 10, the order of the Matchers was counterbalanced across triads.

conversations were recorded with a stereo microphone connected to a Sony Professional Walkman, Model No. WM-D3. The microphone was attached to the top of the vertical barrier, allowing clear recording of both the Directors' and Matchers' speech.

3.4. Design

For 10 rounds, Directors repeatedly instructed Matchers how to arrange arrays of 16 picture cards. In Rounds 1–8, each Matcher participated in the task separately for four consecutive rounds: one Matcher was present for Rounds 1–4 while the other Matcher was present for Rounds 5–8. The distribution of experimental cards (Orthogonal or Overlapping) in these eight rounds was manipulated between groups, as shown in Table 1.

For Rounds 9 and 10, the Directors matched the complete set of 16 experimental items with each Matcher individually. In these two rounds, half of the cards were "old" cards that each Matcher had matched previously, while the other half were "new" cards that each Matcher had not matched in the earlier rounds. For half the triads, the Matcher originally present for Rounds 1–4 came back for Round 9, while for the remaining triads, the Matcher present for Rounds 5–8 continued with the task for Round 9. The other Matcher then returned for Round 10. This counterbalancing of Matcher order equalized the distance, in rounds, between the Matchers' initial rounds (1–4 or 5–8) and final rounds (9 or 10), across triads of participants.

Thus, card distribution (orthogonal vs. overlapping) and Matcher order (A-B-A-B vs. A-B-B-A) were between-participants factors, while round (Round 9 vs. Round 10) and card status (old vs. new) were within-participant factors. In initial four-way ANOVAs carried out for both the description onset latencies and the proportions of reconceptualizations, there were no reliable effects involving either Round or Matcher order. Thus, our primary analyses for both measures collapse over these two factors.

3.5. Procedure

From each triad of participants, we randomly selected one person to play the role of the Director. The other two individuals became Matcher A and Matcher B. We instructed the participants that their goal in each round was to have the Matcher accurately reproduce the arrangement of cards in the Director's target array, and that they should interact as much as necessary to accomplish this goal. We also informed them that they would be carrying out the card-matching task multiple times, and that one Matcher would be present for Rounds 1 through 4 but would then switch off with the other Matcher for Rounds 5 through 8. In the initial instructions, no mention was made of Rounds 9 and 10. The participants took part in a brief practice round in which the Director and one Matcher arranged a set of four cards. Then, depending on the particular Matcher order designated for each group, the appropriate Matcher was asked to step outside the room for the first four rounds.

Prior to each experimental session, we divided the Director's card set into two subsets according to the particular card distribution and item list designated for that group. At the beginning of each round, the experimenter shuffled the appropriate subset of cards and placed them in the central 4×4 grid on the Director's side of the visual barrier. The Director was told that this was the target arrangement for that round, and that

he or she should describe the cards in order from top left to bottom right. The remaining 16 cards from the Director's set were also shuffled and distributed evenly in the peripheral region surrounding the central grid. It was explained that this would allow the Director to see the complete set of cards that the Matcher was selecting from. At the same time, the Matcher received the complete set of 32 picture cards and sorted them any way he or she saw fit before each round.

Following Round 4, we switched Matchers. At the same time, we exchanged the two subsets of the Director's cards between the central and peripheral areas of the display. Then, after four more rounds with the second Matcher, the experimenter informed the Director that he or she would now be asked to complete one last round with each Matcher. At this point, the 16 filler items were removed from both the Director's and Matchers' card sets, leaving only the 16 experimental items. For Rounds 9 and 10, the two Matchers then took turns matching the complete set of 16 experimental cards.

During each round, the experimenter wrote down the target order of the Director's cards for that round. We gave feedback about any mismatches to the participants between rounds by pointing out the cards that were misplaced. Mismatches occurred less than 1% of the time. The entire session lasted about an hour.

3.6. Analysis and coding

Description onset. To measure how long it took for Directors to initiate their descriptions for each card, we converted the audiotapes of the experimental sessions from analog to digital format by sending the output from a tape player into the sound card of a personal computer. Using sound editing software at a 16 kHz sampling rate, we converted each description from Rounds 9 and 10 into a digital sound file. We then auditorily and spectrographically analyzed each file for two pieces of information: (1) the offset of the Matcher's *okay* indicating a match for the previous card and (2) the onset of the first content word that was part of the description for the current card, most often the first noun or adjective. In the following example, these two points are marked with arrows:

(2) D: second one is a dog with a black collar
M: okay ↑
D: next one is the ↑ bird that's half blue and half yellow
M: okay

We measured latencies to the point at which the Directors began their actual descriptions, rather than the point at which they began speaking, because Directors typically prefaced descriptions with introductory phrases like "And the next card is…," as in the example above. By Rounds 9 and 10, these preambles were highly routinized, and frequently were also drawn out as Directors considered what to say next (e.g. "Aaand the next card is uhhhh…"). As a result, we felt that the point at which Directors simply began speaking would not reliably indicate when they had settled on a conceptualization for each card (see Ferreira & Swets, 2002).

There were, however, a number of instances in which we were unable to cleanly measure onset latencies. On some occasions, at the end of a round, Directors would simply list the last few remaining cards (e.g. "and then there's the last dog and the last fish"),

making it difficult to measure the initiation times for these final descriptions. Additionally, because Directors could formulate a description for the first card in the array before being told to begin, we did not measure the initiation time for the first description in each round. These considerations removed 7.4% of the descriptions in Rounds 9 and 10 from the analysis of onset latencies.

Reconceptualizations. The first author and two trained research assistants transcribed verbatim all audiotaped sessions. To assess the extent to which Directors engaged in audience design, we restricted our analyses of these transcripts to the Directors' utterances produced in advance of feedback from the Matchers. That is, we were particularly interested in the adjustments Directors made as they produced their *initial* descriptions for each item—our aim was to assess the degree to which Directors' descriptions showed evidence of audience design prior to any feedback about the success or failure of individual descriptions.

Researchers have often assessed the consistency with which speakers describe the same referent on multiple occasions by measuring *lexical entrainment* (Bortfeld & Brennan, 1997; Brennan & Clark, 1996). Entrainment occurs when speakers choose the same manner of referring to a particular item from one occasion to the next. For example, consider the following moments taken from the current experiment, in which the same Director is describing the same picture of a fish over several rounds:

(3) Round 8: Fish with the red black and green tail, shiny and blue
Round 9: Alright the first picture is of a fish with a red black and green tail, it's shiny and blue
Round 10: Next one is the fish that looks shiny and blue and its tail has red green and black

Despite minor variations in framing, these three descriptions all contain the same descriptive content, and therefore represent instances of lexical entrainment. In general, entrainment provides evidence that the speaker is maintaining the same conceptualization for a given referent. To the extent that this perspective had been mutually established with a specific partner, reusing this perspective should help that same partner readily identify the intended referent (Brennan & Clark, 1996).

In the current experiment, however, our focus is upon the adjustments that speakers make as they interact with different conversational partners. For this reason, we were mainly interested in identifying descriptions that contained modifications of previously established referring expressions (Bortfeld & Brennan, 1997; Horton & Gerrig, 2002). Here are two examples of such modified descriptions, each from a single Director describing a bird to one Matcher in Round 8 and the same bird to a different Matcher in Round 9:

(4) Round 8: The next is a perched bluejay
Round 9: The next is a bluejay that's perched without the legs

(5) Round 8: Next one is the bird with the yellowish underbelly
Round 9: Next one over is the skinny long bird with the yellowish spotted underbelly

These Round 9 descriptions represent what we will call *reconceptualizations*. By examining such reconceptualizations, we can obtain an indication of how frequently speakers are willing to depart from a previously established description when referring to specific objects with either the same or different addressee.

However, as we argued in Horton and Gerrig (2002), reconceptualizations by themselves do not provide unambiguous evidence with respect to audience design. Consider example (4), which constitutes a reconceptualization by virtue of the information "without the legs" added in Round 9. We cannot be certain whether the added clause was planned in advance of an expectation of feedback (in which case we would want to count it as audience design), or whether it arose interactively because the Matcher did not provide immediate feedback following the first part of this utterance. To avoid this ambiguity, we focus upon a subset of reconceptualizations that provide a clearer indication of having been planned that way. Namely, we limit our analyses to reconceptualizations in which new material is introduced prior to or internal to the entrained expression, as in example (5) above. The fact that "skinny long" and "spotted" appear amid components of the previous description allow the stronger claim that these changes were incorporated as part of the speaker's initial utterance plan. We interpret these *utterance-medial reconceptualizations* as comprising the set of cases in which we can be relatively certain that the Directors were planning from the start to modify an earlier description.

In our coding, we compared each Director's descriptions from Rounds 9 and 10 to the same Director's earlier descriptions of the same items from the last round in which each item had appeared previously, which was either Round 4 or Round 8 (the particular comparison round depended upon the configuration of cards in the initial phase of the experiment, as shown in Table 1).[1] When the Director described the card for the same Matcher in Round 9 or 10 as in the comparison round, the particular item could be classified as "old." When the Director described the card for different Matchers across the two rounds, the item was classified as "new." For each comparison, we counted as an utterance-medial reconceptualization each instance in which the Director modified an earlier description by adding new information that either preceded or was interleaved within the old material. All comparisons were independently coded by the two authors, who were blind to the experimental condition of each description. Any differences were resolved through discussion (inter-coder reliability: $K = .93$).

4. Results

4.1. Description onset

Our main purpose for examining description initiation latencies was to determine whether Directors adjusted the time course with which they performed their utterances to

[1] Recall that we counterbalanced the order of Matchers in the final two rounds (A then B or B then A). Across triads, then, the distances involved in comparing old and new cards to the appropriate comparison round were equivalent.

Table 2
Mean onset latencies (and standard errors) for the first content word of the Directors' descriptions from Rounds 9 and 10, in milliseconds

	Card Distribution			
	Orthogonal		Overlapping	
Card Status	*M*	SE	*M*	SE
New	2350	99	2420	77
Old	2258	91	2185	71

reflect the relative difficulty of the Orthogonal versus Overlapping card conditions. The mean onset latencies for the critical portions of the Directors' descriptions are shown in Table 2. The average latencies for descriptions in the Orthogonal ($M = 2305$ ms) and Overlapping ($M = 2301$ ms) card conditions were nearly identical. Analyses of variance for this and all other effects reported here were conducted with both participants ($F1$) and items ($F2$) as random variables. Analyses of the onset latencies revealed no effect of card distribution (both Fs < 1). This result provides an important context for the interpretation of performance differences. Specifically, these initiation times suggest that the Directors were not generally allowing themselves the extra time that may have been necessary to retrieve appropriate memory representations for the Overlapping condition. We return to this observation in the general discussion.

By contrast, Directors were significantly slower to initiate descriptions for cards that were new for the current Matcher ($M = 2384$ ms) compared to descriptions for cards that the current Matcher had dealt with before ($M = 2221$ ms; $F1[1,22] = 8.45$, MSe = 37,210, $P < .01$; $F2[1,15] = 8.58$, MSe = 52,131, $P < .01$). Thus, the Directors' descriptions involved somewhat more planning time when they referred to a card that was unfamiliar to the current addressee. Finally, the pattern of means presented in Table 2 suggests an interaction between card distribution and card status, such that Directors in the Overlapping condition appear to have taken longer to initiate their descriptions for new versus old cards compared to Directors in the Orthogonal condition. This would support the possibility that the preparation involved in describing the new cards was that much harder in the Overlapping case. However, this interaction was not borne out statistically ($F1[1,22] = 3.15$, MSe = 37,210, $P = .09$; $F2[1,15] = 1.47$, MSe = 60,599, $P = .24$). Even so, the fact that Directors generally took longer to initiate descriptions for new cards provides a useful context for the analysis of reconceptualizations: It suggests that Directors often recognized instances in which extra effort was warranted.[2]

[2] An examination of the actual time to begin speaking revealed a similar pattern. As with the description onset latencies, the mean speech onsets were nearly identical in the Overlapping ($M = 999$ ms) and Orthogonal ($M = 998$ ms) conditions. Also, the cell means for this measure (Orthogonal-old: 988 ms; Orthogonal-new: 1012 ms; Overlapping-old: 990 ms; Overlapping-new: 1008 ms) indicate that the Directors took somewhat longer to begin speaking when describing new cards. Analyses carried out on these data, however, revealed no reliable effects (all Fs < 1.7). Thus, significant differences in the timing of the Director's utterances emerged only at the onset of the first content word, and then only in the comparison between cards classified as old versus new. This is consistent with the possibility that Directors often began speaking before conceptualizing their descriptions.

Table 3
Mean proportion (and standard errors) of Directors' initial descriptions in Rounds 9 and 10 containing utterance-medial reconceptualizations (compared to Round 4 or Round 8)

	Card Distribution			
	Orthogonal		Overlapping	
Card Status	*M*	SE	*M*	SE
New	0.57	0.04	0.44	0.04
Old	0.35	0.03	0.36	0.03

4.2. Reconceptualizations

Next we turn to our reconceptualization analysis, which allows us to test our main prediction that Directors' descriptions in the Orthogonal card condition would show relatively more evidence of audience design. We are interested in the extent to which Directors adjusted their descriptions based on the communicative needs of their current partner. As such, our analyses focus on how frequently the Directors' initial descriptions in Rounds 9 and 10 showed evidence for utterance-medial reconceptualizations.[3] These data are shown in Table 3. Because this measure involves examining proportions, we first applied an arcsine transformation to our data before carrying out analyses of variance. For clarity of presentation, we present means in the form of raw proportions.

If Directors are being attentive to their shared experiences with their Matchers, they should reconceptualize their descriptions most often when a card is "new" for the current partner. In fact, out of the total number of initial descriptions in Rounds 9 and 10, there was a higher proportion of utterance-medial reconceptualizations for New items ($M = 0.50$) than for Old items ($M = 0.35$; $F1[1,22] = 22.64$, MSe = .0567, $P < .001$; $F2[1,15] = 42.09$, MSe = .0381, $P < .001$). We carried out comparisons for new versus old items in each of the Orthogonal and Overlapping card conditions. That difference was reliable for both the Orthogonal condition ($t1[11] = 4.07$, $P < .002$; $t2[15] = 5.97$, $P < .001$) and the Overlapping condition ($t1[11] = 2.47$, $P < .04$; $t2[15] = 2.22$, $P < .05$). These results provide strong evidence for the Directors' general predilection toward audience design.

Beyond this fundamental demonstration, of course, we wished to demonstrate that Directors would provide more evidence of audience design when the memory burdens of the task allowed them to do so. Specifically, we predicted that Directors' descriptions would show a stronger pattern of reconceptualizations in the Orthogonal versus Overlapping card condition. The data support this prediction. As shown in Table 3, for the Orthogonal condition Directors offered 22% more utterance-medial reconceptualizations for new versus old cards. This difference was only 8% for the Overlapping condition. This interaction between card distribution and card status was reliable ($F1[1,22] = 4.95$, MSe = .0567, $P < .04$; $F2[1,15] = 6.62$, MSe = .0555, $P < .03$). We suggest that this result reflects the differential accessibility of memory representations. In the Orthogonal card

[3] As previously noted, we focus our results upon utterance-medial reconceptualizations because they provide a more conservative measure of the Directors' willingness to modify earlier descriptions. The general pattern of results is similar if we include all reconceptualizations in our analyses.

condition—within the span of time Directors gave themselves to plan their utterances—it was relatively more possible for the Directors to produce an appropriate reconceptualization.

5. General discussion

The goal of this project was to demonstrate that the memory requirements of particular circumstances of language use have an impact on speakers' ability to provide evidence of audience design. In our study, Directors provided descriptions of pictures, drawn from four categories, to two independent Matchers. In the Orthogonal card condition, the categories were distributed so that—in the initial four rounds of descriptions with each Matcher—the Director discussed different categories of cards with each Matcher. In the Overlapping card condition, cards from all four categories were discussed initially with both Matchers. In both conditions, however, the Director then carried out two final rounds in which he or she described the full set of card categories once more to each Matcher. If Directors are being sensitive to audience design, they should use their previous conceptualizations for *old* cards (i.e. those cards for which they already share experience with the current Matcher) but provide reconceptualizations for *new* cards in these final rounds. In fact, we found that Directors consistently provided more reconceptualizations for new cards. However, as we had predicted, that pattern was enhanced for Directors in the Orthogonal card condition.

On our analysis, Directors' utterances in the Orthogonal condition were able to exhibit more evidence of audience design because the memory representations required for audience design were more readily accessible during the period of time they allowed themselves to plan their utterances. The results for the initiation time measure (i.e. the amount of time Directors waited before they spoke the first content word of their descriptions) showed that Directors took equivalent amounts of time in the Orthogonal and Overlapping card conditions. This didn't need to be the case: We put no pressure on the Directors to begin to speak with any particular time course. Their timing may have reflected the social situation (with a Matcher waiting across the table) and the Directors' sense of how much their first utterances mattered.

Had the Directors taken more time for utterance planning in the Overlapping condition (and also, presumably, the Orthogonal condition), it is very likely that they could have done better with respect to audience design—where by "better" we mean that they would have offered more reconceptualizations for new cards (and fewer for old cards). We think better performance is possible because the Directors had four rounds of experience with one of the two Matchers on each card. When they saw the cards in Rounds 9 and 10, it was immediately apparent that we had united the card sets. We suspect that the Directors could have gone card-by-card and explicitly determined which was old or new for each Matcher. In fact, it is always possible for speakers to slow down language production to accumulate evidence for whether particular information is shared with an addressee. Consider this example from a portion of a telephone conversation (Kingsbury, Strassel, McLemore, & McIntyre, 1997; see Horton & Gerrig, in press):

(6) Yeah, I've got another buddy who, uh, is a Marine pilot. I'm trying to think if you had ever met this guy. I don't think so.

We can see the process of memory at work here, as the speaker does an explicit search for evidence of common ground. Similarly, it would have been possible for Directors always to seek evidence for common ground in the current experiment. Clearly, they chose not to—at least not for every card. Our results suggest, in fact, that Directors mostly used information to the extent that it was readily accessible to them.

Directors presumably chose not to query their memory card-by-card because they had no strong motivation to do so. Because we were most interested in the impact of the accessibility of memory representations on audience design, our analyses focused on the Directors' first utterances. However, the Directors were engaged in interactive dialogues with the Matchers and did not necessarily need to get their descriptions exactly right the first time out (Fussell & Krauss, 1989). It was possible for the Directors to say what they were prepared to say and then let Matchers play a role giving feedback. Against that background, we still believe that it is important that Directors showed a consistent pattern of performing the "right" utterances with each Matcher.

We intended this study to demonstrate the importance of memory considerations in evaluating speakers' ability to provide evidence of audience design. Our results are valuable not only because they attest to a role for memory in audience design, but also because, as discussed in the introduction, there has sometimes been an impulse within research on pragmatics to make general statements—with respect to all circumstances of language use—about the extent to which speakers develop models of the common ground they share with their addressees (for a review, see Schober & Brennan, 2003). We believe, quite simply, that audience design, as a process, cannot be turned on or off or otherwise modularized with respect to other constraints on language production.

In recent years, a number of theorists have suggested that psycholinguistic models need to move beyond being able to describe how people process words and sentences in isolation and toward a consideration of the broader range of phenomena involved in how people use language in actual situations (e.g. Clark, 1996; Trueswell & Tanenhaus, 2004). We believe that progress toward this goal will be well-served by a better understanding of how fundamental cognitive processes like memory encoding and retrieval constrain moment-by-moment aspects of language processing. In this article, we provide a demonstration of how such memory considerations can shape higher-level utterance planning. As stated earlier, our goal is to situate audience design more firmly within the domain of "ordinary" cognitive processes (Horton & Gerrig, in press). As with other factors that affect language production, audience design relies upon the nature of the memory representations that are available as input during utterance planning. We recognize that this claim has precedents in other theories of language processing. However, we believe that much can be gained by recognizing the centrality of memory processes in language use. At the very least, theories in pragmatics and theories concerned with other aspects of language processing should not be treated as existing in separate domains.

In sum, the results that we have presented here are intended to show that aspects of language use like audience design cannot be divorced from considerations of more basic

aspects of cognitive processing. Furthermore, by proposing a specific role for memory associations in making partner-specific information accessible during language use, we hope to flesh out existing accounts concerning the role of common ground in language production. In general, speakers will show evidence of audience design to the extent that suitable memory representations become accessible with an appropriate time course. In that sense, it should be relatively easy to create circumstances in which speakers perform better or worse with respect to some external metric for the "amount" of audience design (see also Bard & Aylett, 2000). The distinction between the Orthogonal and Overlapping card distributions in the current study was intended to provide such a contrast in circumstances. We believe that the impulse to engage in audience design was equivalent across the two conditions. What differed was the accessibility of the memory representations that allowed Directors to follow through on that impulse.

Acknowledgements

We would like to thank Susan Brennan, Zenzi Griffin, Joy Hanna, Sarah Haywood and three anonymous reviewers for their helpful comments on this article, and Kimberly Hines and Jennifer Puig for their invaluable assistance with data collection and coding. This material is based upon work supported by the National Science Foundation under Grants No. 0082602 and 9980013.

References

Bard, E. G., & Aylett, M. P. (2000). Accessibility, duration, and modeling of the listener in spoken dialogue. *Göteborg 2000, fourth workshop on the semantics and pragmatics of dialogue*. Göteborg, Sweden: Göteborg University.

Barr, D. J., & Keysar, B. (2002). Anchoring comprehension in linguistic precedents. *Journal of Memory and Language*, *46*, 391–418.

Bock, J. K. (1987). An effect of the accessibility of word forms on sentence structures. *Journal of Memory and Language*, *26*, 119–137.

Bock, J. K., & Levelt, W. J. M. (1994). Language production: Grammatical encoding. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 945–984). San Diego: Academic Press.

Bortfeld, H., & Brennan, S. E. (1997). Use and acquisition of idiomatic expressions in referring by native and non-native speakers. *Discourse Processes*, *23*, 119–147.

Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1482–1493.

Brown, P. M., & Dell, G. S. (1987). Adapting production to comprehension: The explicit mention of instruments. *Cognitive Psychology*, *19*, 441–472.

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

Clark, H. H., & Carlson, T. B. (1982). Hearers and speech acts. *Language*, *58*, 332–373.

Clark, H. H., & Marshall, C. M. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. L. Webber, & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 10–63). Cambridge: 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* (pp. 287–299). Amsterdam: North-Holland.

Ferreira, F., & Swets, B. (2002). How incremental is language production? Evidence from the production of utterances requiring the computation of arithmetic sums. *Journal of Memory and Language*, *46*, 57–84.

Ferreira, V. S., & Dell, G. S. (2000). Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive Psychology*, *40*, 296–340.

Fussell, S. R., & Krauss, R. M. (1989). The effects of intended audience on message production and comprehension: Reference in a common ground framework. *Journal of Experimental Social Psychology*, *25*, 203–219.

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.

Horton, W. S., & Gerrig, R. J. (2002). Speakers' experiences and audience design: Knowing when and knowing how to adjust utterances to addressees. *Journal of Memory and Language*, *47*, 589–606.

Horton, W. S., & Gerrig, R. J. (in press). Conversational common ground and memory processes in language production. *Discourse Processes*.

Horton, W. S., & Keysar, B. (1996). When do speakers take into account common ground? *Cognition*, *59*, 91–117.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, *114*, 3–28.

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*, 32–38.

Keysar, B., Barr, D. J., & Horton, W. S. (1998). The egocentric basis of language use: Insights from a processing approach. *Current Directions in Psychological Science*, *7*, 46–50.

Kingsbury, P., Strassel, S., McLemore, C., & McIntyre, R. (1997). CALLHOME American English transcripts, LDC97T14. Philadelphia: Linguistic Data Consortium.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.

Lockridge, C. B., & Brennan, S. E. (2002). Addressees' needs influence speakers' early syntactic choices. *Psychonomic Bulletin and Review*, *9*, 550–557.

Metzing, C., & Brennan, S. E. (2003). When conceptual pacts are broken: Partner-specific effects on the comprehension of referring expressions. *Journal of Memory and Language*, *49*, 201–213.

Schober, M. F., & Brennan, S. E. (2003). Processes of interactive spoken discourse: The role of the partner. In A. C. Graesser, M. A. Gernsbacher, & S. R. Goldman (Eds.), *Handbook of discourse processes* (pp. 123–164). Mahwah, NJ: Lawrence Erlbaum.

Smith, S. M. (1994). Theoretical principles of context-dependent memory. In P. E. Morris, & M. Gruneberg (Eds.), *Theoretical aspects of memory* (2nd ed., pp. 168–195). London: Routledge.

Trueswell, J. C., & Tanenhaus, M. K. (in press). *Approaches to studying world-situated language use: Bridging the language-as-product and language-as-action traditions*. Cambridge, MA: MIT Press.