

DRAFT: DO NOT CITE WITHOUT PERMISSION.

Running head: COMMON GROUND IN PRODUCTION

Common Ground in Production:

Effects of Mutual Knowledge on Word Duration

Michelle L. Gregory, Alice F. Healy and Daniel Jurafsky

University of Colorado at Boulder

Abstract

A key issue in models of lexical production is whether a model of hearer knowledge influences speakers' productions. In this study we investigate whether the durational shortening of repeated words is a consequence of mutual knowledge, or if it is a byproduct of speaker internal processes. In a study of spontaneously produced narrations, in Experiment 1 we demonstrated that repeated words are shorter than their first mentions. In Experiment 2, we found that words are shortened more in the presence of a repeated hearer compared to a new hearer. Results from Experiment 3 demonstrate that shortening continues across narrations for repeated hearers, but is blocked by new hearers. These results indicate that a speaker's model of hearer knowledge affects production.

Common Ground in Production: Effects of Mutual Knowledge on Word Duration

Speakers commonly change the way they speak in order to conform to their interlocutors and other aspects of the speech context. These effects are pervasive and obvious; speakers raise their pitch when speaking with babies, code-switch only with appropriate interlocutors, and only use pronouns with referents that are accessible to others in the speech context. As noted by Clark and Wilkes-Gibbs (1986), in order for language to be a communicative process, there must be some sort of shared knowledge and experience between interlocutors.

The knowledge that is shared between interlocutors is often referred to as *common ground*, "the background, the context, for everything the participants both do and say" (Clark, 1994, p. 990). While common ground is information that both speaker and hearer share, *mutual knowledge* is what the speaker, or hearer, knows (or assumes) about what her conversation partner knows (see Clark & Marshall, 1981). For example, if two people are talking, let us call them Cynthia and Jason, and Jason says to Cynthia "I was walking along today and the neighbors dog bit me," then at this point *the neighbors dog* is considered as part of the common ground between the speaker and hearer. In the next utterance, Jason continues, "It came right at me." Jason presumably uses the pronoun *it* to refer to the dog because he knows, or has reasonable ground to assume, that Cynthia shares mutual knowledge concerning the referent of the pronoun.

The intuitive notion that interlocutors rely on mutual knowledge in conversation is the foundation upon which many linguistic theories have been built. In fact, discourse pragmatic theories are largely grounded in the belief that speakers choose sentence types and forms of referring expressions based on assumptions concerning the hearer's knowledge:

If a speaker assumes, prior to uttering an intonation unit, that a concept is already active in the hearer's mind, he will verbalize that concept in an attenuated manner, giving it weak stress and probably pronominalizing it. (Chafe, 1987, p.48).

Because communication is a collaborative process, it is easy to assume that mutual knowledge is a factor in language processing. But to what extent do interlocutors make use of mutual knowledge during language processing? The debate over the use of mutual knowledge in processing concerns both the comprehension system and the production system. In this study we are concerned with common ground and the production system. Intuitively, it seems that both hearer design and speaker internal factors influence the shape and production of utterances. In this study we report on a psycholinguistic study that is designed to address whether durational shortening is influenced by hearer knowledge. We use a methodology that has been demonstrated to be sensitive to speaker internal sources of reduction. Previous research in this area has provided evidence for contributions of speaker based processes, but has not excluded contributions from the speaker's assumptions of hearer knowledge. Similarly, research demonstrating that assumptions of hearer knowledge shape productions has not excluded possible contributions from speaker internal processes. We review evidence for hearer-based models and speaker-based models of production in the following section. Finally, we report the results of three experiments designed to address the role of mutual knowledge in production.

Mutual Knowledge and Production

Human language production is the result of the collaboration and interaction of many layers of linguistic processes. A production begins with a thought, or a concept that the speaker wished to convey. From the initial concept, lexical items have to be chosen, syntactic structure is

constructed, phonological processes are consulted and then fed through the articulatory system, finally resulting in a speech stream (see Bock & Levelt, 1994; Garret, 1975; Levelt, 1989; among others). Traditionally, knowledge of the production system has come through the study of different types of speech errors (Cutler, 1982; Dell, 1984; Raymond, 2000; Stemmer, 1983; Shattuck-Hufnagel, 1979). More recently, researchers have begun to tap into the production system by investigating the effects of external factors on production, including aspects of common ground and mutual information.

Studies concerning common ground and the production system generally fall into two categories. The first involves aspects of mutual knowledge that are involved in off-line processes such as initial utterance design. Included in this category are the pragmatic factors that influence lexical selection (such as the choice to pronominalize or use a full referring expression). The second area of research is concerned with whether mutual knowledge is involved in fast, automatic processes such as articulation.

Common Ground and Initial Design

There have been a number of studies that demonstrate that *hearer knowledge* (defined as the speaker's assumptions of what the hearer knows) is a major factor in audience design (Clark & Brennan, 1991; Chafe, 1976; Gerrig et al., 2001; Gundel et al., 1993; Nadig & Sedivy, 2001; Prince, 1981). For example, in a corpus study of forms of referring expressions, Gundel et al. (1993) illustrated that the form of a referring expression that a speaker chooses is directly correlated with the assumptions speakers make concerning the cognitive status of the referring expressions in the mind of the hearer.

In a study of the use of the use of pronominal adjectives (e.g., *the big cup*), Sedivy (2002) found that adult speakers only use scalar and material adjectives in the presence of a contrast set

(where both a big and small cup were in the speaking environment). Nadig and Sedivy (2001) found that children as young as 5 or 6 are also sensitive to common ground in their productions. They found that children were more likely to use a scalar adjective when describing an object to a confederate when the confederate had access to a contrast set, but not when the contrast set was blocked from view of the confederate, even though the children still had access to the contrast set. These results demonstrate that children's utterance design is influenced by information perceived to be common ground.

There is evidence that speakers rely on more than mutual knowledge for lexical choice in conversations. Brennan and Clark (1996) argue that speakers and hearers engage in conceptual pacts in conversations. In a task where speakers had to give instructions to hearers on how to match cards with pictures of objects on them, they found that speakers propose a conceptualization of an object, which is then accepted or denied by the hearer. Once a conceptualization is agreed upon, both interlocutors tend to use that term for that object (the result of a process called *lexical entrainment*). Moreover, they found that the conceptualization pact is partner specific; the same speaker is more likely to refer to the pictured objects with the same lexical item when they are giving instructions to the same partner. Once they switch partners, they change their lexical choice based on the conceptualization pact they make with the new partner.

There have been a number of studies, however, that find that mutual knowledge is *not* a factor in initial utterance design. For example, Brown & Dell (1987) (see also Dell & Brown, 1991) conducted a series of experiments in which they tested if mutual knowledge between speakers and hearers played a role in the explicit mention of atypical instruments. Participants were asked to retell a story that they read. They found that participants were more likely to use t

atypical instruments than typical ones in their re-telling of the story (e.g., *ice pick* over *knife*). However, they found that the speaker's belief about the listener's knowledge of the instrument, as determined by the presence or absence of a picture of the scene, had little impact on whether the instrument was mentioned. They claim that the mechanisms behind what seems to be listener adaptation are "automatic products of the production system" (Dell & Brown, 1991, p.107).

Consistent with Dell and Brown (1991), Horton and Keysar (1996) also found that speakers do not take common ground into consideration during the initial stages of utterance planning. In an investigation of adjective use, they found that when under a time constraint, speakers do not use adjectives, like "the *big* ball," even when a contrast set was available to hearers. Horton and Keysar concluded that common ground does not play a primary role in utterance design, but that speakers do fall back on it during the monitoring and adjustment phase of production, which happens later in the production process (see Levelt, 1983).

Articulation and Mutual Knowledge

As with the debate concerning initial design, there is evidence for and against the notion that mutual knowledge is a factor in the fast, automatic processes that are involved in articulation. Perhaps the most well documented affects of mutual knowledge and articulation is the study of the duration and intelligibility of referring expressions.

Fowler and Housum (1987) found that second mentions of the referring expressions are durationally shorter and less intelligible than their first mentions in an investigation of 35 word pairs from a narrative. Speakers produced a longer referring expression the first time it was uttered, compared to subsequent productions of the same referring expression. The first use of a referring expression is also more intelligible than later uses. When the first use and second uses of a referring expression were excised and played to new participants, the participants were

better able to identify the referring expressions that were the first mention in the narration, compared to the second mentions. From these studies, Fowler and Housum concluded that speakers attenuate the production of old words. In a follow-up study, Fowler (1988) tested the same 35 pairs of referring expressions. In this study, she had speakers produce the referring expressions in a list, as opposed to a narrative context. She found that the second mention of the repeated referring expressions did not produce significantly shorter target forms than the first mentions when spoken in list format. From the results of both studies, Fowler observed that repeated referring expressions are only shortened in discourse contexts, concluding that speakers made use of a hearer model and shortened words that were predictable to hearers. In this study we refer to explanations such as these as *hearer-based* models.

On the other hand, researchers have found that durational shortening is not caused by hearer knowledge. In a study of the production of word pairs, Balota, Boland, and Shields (1989) found that word pairs that are semantically related to one another are said faster than word pairs that are not related, even in the absence of a hearer. For example, *dog-cat* presented together produced quicker articulations than the pair *dog-cup*. From these results, they concluded that durational shortening of referring expressions has to do with lexical access: Conceptual relationships influence activation levels during production, which affects articulation. While Brown and Dell did not manipulate hearer knowledge in their study, they did find evidence that speaker internal processes, viz., lexical access, are a source of durational shortening.

In a recent study, Bard et al. (2000) varied hearer knowledge to isolate possible articulatory adjustments resulting from mutual knowledge. In their first experiment, both a speaker and a hearer were given a map with landmarks such as *the park bench*, or *pelicans*. Speakers were instructed to give directions to a location on a map to two different hearers (using

the same map with each hearer). They compared the length and intelligibility of the first use of a referring expression in the first set of directions to the first use of the same referring expression in the second set of instructions (in both cases, the hearers only heard the directions once). Previous research has demonstrated that second mentions of words are shorter than their first mentions (see Fowler 1988; Fowler & Housum, 1987). Bard et al. (2000) predicted that if durational shortening and a decrease in intelligibility are affected by hearer knowledge, then those processes should be blocked in the presence of the new hearer in the second set of instructions. To the contrary, they found that the first use of a referring expression in the second set of instructions decreased in intelligibility, despite the presence of a new hearer. From these results, they conclude that speakers do not make use of a model of hearer knowledge to update and inform the articulation process: Articulations are only sensitive to speaker internal processes. However, unlike their results concerning intelligibility, their results concerning the durational shortening of the first use of a referring expression were inconclusive. Referring expression in the second narration were not significantly shorter, or longer, than the same referring expression in the first narration. We refer to models that claim that the articulatory system is not sensitive to hearer knowledge as *speaker-based* models.

In this review we have provided seemingly contradictory evidence regarding the use of mutual knowledge in production. On one hand, there is a body of research that suggests that it is not used in utterance design or during articulation. On the other hand, there seems to be a plethora of evidence that finds that mutual knowledge is used both in initial utterance design and during the fast, automatic processes involved in articulation. Why are there such conflicting results? We argue that by targeting either speaker-based effects *or* hearer-based effects, the studies reviewed here indeed find evidence for one or the other, but their results do not exclude

an additional influence from the alternate source. In the present study, we narrow our attention to issues concerning whether speakers' assumptions about hearer knowledge are available at the later stages of the production process, such as articulation. Thus, we narrow our focus to evidence from speaker-based models and hearer-based models as they relate to articulation, but not to utterance design or lexical selection.

The current experiments are designed to test the effects of mutual knowledge on articulatory processes affecting the duration of words. In the following experiments, we test (a) if referring expressions get shorter upon subsequent use, (b) if the presence of a repeated hearer causes more durational shortening in a second narration than the presence of a new hearer, and (c) if the attenuation of referring expressions is blocked by the presence of a new hearer. If in fact the presence of a repeated hearer yields shorter durations than the presence of a new hearer, it would indicate that speakers' assumptions concerning hearer knowledge do influence the articulation process. Note that although we refer to these as different experiments, there was only one set of data collection, materials, and participants for all of them. From this data collection, a corpus of narrations was created and the three experiments represent different comparisons within the corpus. The methodology for the corpus collection is described in the General Method, with the individual experiments following.

General Method

Participants

Groups consisting of two or three participants were employed. Participants were undergraduate students from the University of Colorado recruited as groups of friends. However, in three of the groups, the participants did not know each other prior to meeting at the

experiment. Participants either received course credit or 10 or 15 dollars each for their participation. Thirty-six participant groups were used.

Materials

We used two stories to develop the corpus used in these experiments. Although the names of the participants in the story were changed, the stories were based on news articles taken from *The Colorado Daily* newspaper, Sept. 15, 2000. Each story consists of three sentences (see the Appendix). The stories are matched for number of syllables and roughly matched for sentence length and number of words. Although the stories were taken out of the newspaper, a number of referring expressions (such as names of individuals and places in the stories, as well as posts that the individuals held) were changed to ensure that speakers did not have prior knowledge of the events described in the story. Additionally, both stories consist of only one episode to avoid any confounding effects of episode boundaries (see Fowler, Levy, & Brown, 1997).

Procedure

Participant groups were tested individually. For each participant group, one participant was chosen as the "speaker" and the other one or two participants were "hearers". Speakers were taken to a separate room from the hearers and were asked to read the short story. Reading the story was an untimed task. They could take as long as they needed to read the story, although the experimenter did ask that they not try to memorize the story verbatim. After indicating that they had finished the story, speakers completed a four question multiple choice comprehension exam. They no longer had access to the story during the narrations and exam. Speakers were then instructed to repeat the story to another participant from memory. The speakers were aware that their addressees were going to take a similar comprehension exam after hearing the story. Each

speaker was audio taped during the face-to-face interaction. The experimenter was not in the room during the narration of the story. The participants who were the hearers were instructed to listen to the story and they were informed that they would take a comprehension exam after hearing the story. Addressees were allowed to ask questions at any point during the narration, although they were instructed to never ask the speaker to repeat the story in its entirety. Each participant group belonged to either the *repeated hearer* group or the *new hearer* group. In the repeated hearer group, speakers told the story to the same addressee twice. In the new hearer group, speakers told the story to two different hearers.

Stimuli Analyses

All recorded materials were digitized at a sampling rate of 16 kHz using the Praat Speech analysis recorder (Boersma & Weenink, 1992). We used ESPS WAVES software to align the waveform of each narration to the words used in each narration. Each narration was then analyzed with Praat speech analysis software. For each of the two stories, six referring expressions were analyzed for duration and pitch accent. The referring expressions are listed in Table 1. We analyzed all six referring expressions for each narration of the story for every participant group.

[place Table 1 about here]

Dependant variable. Durations of each item were measured in milliseconds. The duration for each use of a referring expression was taken as a raw duration, the raw duration normalized by number of syllables (the duration of the referring expression divided by the number of syllables, resulting in a measure of the mean syllable duration), and raw duration normalized by number of phones. Because word duration is not a linear function, we employed the log duration for these measures for all analyses. All four measures produced nearly identical results, so we

only report analyses on log durations normalized by number of syllables. Note, however, that information in tables and charts is given in duration normalized by number of syllables.

Referring expressions. Not all speakers used all six referring expressions in each narration and not all uses of the referring expressions were included in analyses. For the multi-word referring expressions, a speaker had to use the full form in order for it to be included in analyses. For example, Speaker 23 used only the last name of the proper name *John Powers* (see Example 1). Thus, this token and similar ones were excluded from analyses. Likewise, tokens were excluded from analyses when they did not match the target referring expression verbatim. For example, in Narration 2, Speaker 16 (see Example 2) used *health issues* instead of *health problems*. Because the target referring expression was *health problems*, these tokens were not used.

- (1) “*this guy powers*” as opposed to *john powers*
- (2) “... was due to mainly health issues” instead of health problems

Tokens were also excluded from analyses when they did not appear in the same syntactic context. For example, at one point Speaker 21 said “*met with the president...*” and later in the same narration said “*him and president Johnson*”. In the latter use, *president* is used as a modifier for *Johnson*, whereas in the initial use *president* is the head of the noun phrase. Cases such as this, where the target referring expression is used a head in the first utterance, and a modifier in the second utterance, were not used in analyses. However, we retained cases in which a target noun was used as a subject in one sentence and an object in another sentence. For example, *president* had been used as an object of a preposition in one utterance (*he had a meeting with the president*), and a subject of the sentence in another (*he and the president met...*), these were both valid instances of *president*.

Duration coding for stimuli. As can be seen from the word list in Table 1, the referring expressions used are quite varied in a number of ways. The particular segments that comprise each referring expression and the environment in which each occurred, specifically the surrounding segments, differed for each speaker and each narration. This variation made coding the durations of each referring expression a challenge. To obtain a consistent measure of duration, despite the difference in environments for the referring expressions, we developed a set of heuristics to determine exact length. In this section we describe the coding protocols that we established for coding the durations of these words.

Words that begin or end in stop consonants, such as the /p/ in *prime minister*, and the /d/ and /t/ in *disagreement* and *president*, can be difficult to code in that there is often no evidence on a waveform of where the closure of the stop begins, or ends if there is no release of the stop. In general, we excluded the stop closure and measured duration beginning at the point of release, i.e., where there was visible aspiration on the waveform, as demarcated by the line at the beginning of *president* in Figure 1. In this utterance, *president* is preceded by the word *the*. From the waveform in Figure 1, one can see that that the vowel in *the* is finished before the aspiration of the /p/ in *president* begins. We chose the beginning of the release, as opposed to the beginning of the stop closure because not all of the environments preceding a voiceless stop were as clear with regard to the point of closure. These heuristics allowed for consistent coding of stops.

[insert Figure 1 about here]

For tokens that begin in a segment identical to the preceding segment, as in *become* *mayor*, we chose the midpoint in the closure, as demonstrated in Figure 2. In tokens such as these, there is no other information to determine the end of the first bilabial nasal with the beginning of the second bilabial nasal, thus we had to pick a somewhat arbitrary point to

demarcate the beginning of *mayor*. For consistency, we chose to begin *mayor* at the midpoint of the bilabial nasal consonant.

[Insert Figure 2 here]

For consistency, one coder coded all of the tokens. To insure that we did not bias the coding in favor of our hypothesis, all of the coding was done “blind” in that the coder did not know which condition each narration was in while she was coding. After coding for the duration of referring expressions, durations were then normalized according to the procedure discussed in the Dependant Variable section. The duration data were used for all three experiments.

Experiment 1

Method

Design. Experiment 1 was designed to determine whether referring expressions are shortened upon subsequent use within one narration. In this experiment, we compared the normalized log duration of the first use of each referring expression to the normalized log duration of the second use of the same referring expression in the same narration. The purpose of this experiment was to ensure that the design and materials used in this study were sensitive to the durational shortening effects that were found by Fowler and Housum (1987) and Bard et al. (2000).

Stimuli. Of the 36 participant groups discussed in the General Method, one was lost due to experimenter error, and two were lost due to equipment failure. Data from one participant group was not used because the speaker repeated the story in its entirety to the hearer. In other words, the speaker told the story twice to the same hearer within one narration. Data from two participant groups were not used because the speaker confused at least two of the referring expressions. For example, in one narration the speaker used *John Fallon* instead of *John Powers*

from Fallon. A total of 30 narrations (1 per speaker) were used in analyses. Within the 30 narrations, stimuli were only included for analyses if (a) the whole referring expression was used, (b) the referring expression was used at least twice within the narration, and (c) the referring expression occurred in similar syntactic and semantic roles in each repetition (as discussed in the General Method). These exclusions left 21 word pairs for analysis. The word pairs consisted of the first use of a referring expression and the second use of the same referring expression in one narration. The prediction is that if repeated words are shortened upon subsequent mention, the second use of the referring expression should be shorter than the first use.

Results

A set of analyses was conducted on the means of referring expressions for the log duration in ms normalized by number of syllables. A paired t test was conducted comparing the mean of the duration of the first use of the referring expression to that of the second use, with subjects as the random variable. This analysis demonstrates that the mean of the first use of referring expressions is longer than that of the second use; $t(20) = 3.701$, $p < .001$ (one-tailed). A separate paired t test indicates that the effect holds with items as the random variable as well, $t(10) = 3.042$, $p < .01$ (one-tailed). Figure 3 represents the means of the first and second use of referring expressions.

[insert Figure 3 here]

Discussion

The results reported here demonstrate that referring expressions do indeed shorten upon subsequent use. These findings duplicate earlier findings that referents that are repeated in conversation are durationally shorter than the first use of the same referring expression (Bard et

al., 2000; Fowler et al. 1988; and Fowler & Housum 1987). These results do not address whether durational shortening of repeated expressions is caused by assumption about hearer knowledge or speaker internal processes because the discourse pragmatic environment was not manipulated. In Experiment 2, we manipulated the discourse pragmatic environment to specifically address whether durational shortening is affected by assumptions speakers make concerning what their hearers know.

Experiment 2

Experiment 1 and previous research has demonstrated that referring expressions get shortened upon subsequent mention. Is this a fact about hearer knowledge influencing the articulatory system, or a fact about speaker internal processes such as articulatory priming or lexical access? Previous research appears to be conflicting with regard to the source of durational shortening. On one hand, Fowler (1988) claims that durational shortening is a product of hearer knowledge, while on the other hand, Bard et al (2000) claim that it is a byproduct of speaker internal processes. We argue, however, that these two studies are not in conflict with one another. Instead, each of these only targets one of the possible sources of durational shortening and they simply do not address other possible sources.

The Fowler and Housum (1987) and Fowler (1988) studies discussed are inconclusive with regard to speaker internal processes because they do not provide evidence that the role of hearer knowledge is the *sole* factor durational shortening. Their claim that hearer modeling is the source of durational shortening is based on their findings that speakers shorten referring expressions in discourse contexts, but not when they are repeated in lists. However, there are many confounding differences between words in a list and words in a discourse that were not controlled in their study. Foremost is the fact that referring expressions in a discourse context,

such as a narrative, occur in running speech. Each referring expression in running speech is part of a larger intonational phrase, which has consequences for the length of individual words.

Words tend to be lengthened at the boundaries (initial and final positions) of intonational phrases (Beckman & Ayers, 1997). Referring expressions produced in list are spoken in isolation, and thus are produced as individual, isolated intonational phrases. Boundary lengthening effects might mask the effects of repetition on subsequent uses of items in a list. Thus, the referring expressions in a list format may not exhibit shortening due to phrase initial and final lengthening processes, rather than being “outside a discourse context.”

Additionally, Fowler, Levy, and Brown (1997) found that during narratives about a movie, there are environments that block durational shortening of repeated referring expressions. Specifically, the presence of an episode boundary blocked the durational shortening of referring expressions. While it is not clear how an isolated referring expression in a list relates to episode boundaries, both the episode boundary effects and the intonation phrasing effects demonstrate that are many factors that can influence the duration of words produced in isolation. Thus, the lack of durational shortening of repeated referring expressions in a list format might well be due to other factors rather than a lack of discourse context.

Studies supporting speaker-based models of production provide evidence that speaker internal processes are factors during production, however, they do not provide evidence that hearer knowledge does not also play an additional role in production. The study conducted by Bard et al. (2000), in which they compared the length of the first use of a referring expression in one set of directions to the length of the first use of the same referring expression in a second set of identical instructions to a new hearer, provides evidence for speaker internal processes influencing duration. Bard et al. found that speakers attenuate referring expressions in second

mentions even though the second mentions of the referring expression were in the presence of a new hearer. However, their study is inconclusive with regard to hearer-based attenuation: Would the presence of an “old” hearer in the second set of directions cause even more attenuation?

Method

Design. In this experiment, we investigated the role of hearer knowledge in the attenuation of second mentions. We chose to utilize a similar methodology to Bard et al. (2000) because their experimental design has been demonstrated to be sensitive to speaker-based effects. However, we added a condition that manipulates hearer knowledge. Similar to Bard et al., we compared the durations of the first use of a referring expression in the first narration of a story (Narration 1) to the durations of the first use of the same referring expression in a second narration (Narration 2). Our design differs from that of Bard et al. in that we have two conditions for the second narration: *new hearer* and *repeated hearer*. Table 2 represents the two conditions of Experiment 2.

[insert Table 2 here]

This experiment utilized the same corpus as Experiment 1. The data collection procedure, previously described in the General Method, is reviewed here for convenience. Groups of two or three participants were used. In each group one person was chosen at random to be the speaker. The speaker repeated the story twice from memory: either twice to the same person (*repeated hearer* condition), or to two different people (*new hearer* condition). We had two hypotheses in this experiment. First, if there are speaker-based processes that underlie durational shortening, then we should see that the first use of a referring expression in the second narration is shorter than the first use of the same referring expression in the first narration for both conditions. Second, if hearer knowledge also has an effect on word durations, then we should see a greater

degree of shortening in referring expressions in the second narration in the repeated hearer condition, compared those in the new hearer condition.

Stimuli. Stimuli were only included for analyses if (a) the whole referring expression was used, (b) the referring expression was used at least once in each narration, and (c) the referring expression was in the same syntactic and semantic environment in both narrations. A total of 120 referring expressions were used in all analyses.¹

Results

A set of analyses was conducted on the means for the log duration in ms normalized by the number of syllables. First, a mixed multifactorial analysis of variance was conducted on the durations including the between-subjects variable of hearer (*repeated* vs. *new*) and the within-subject variable of narration (*first* vs. *second*). There was a main effect of narration, $F(1,25)=7.598$, $MSE = .008$, $p<.01$. There was also an interaction between hearer-condition and narration, $F(1,25)=5.468$, $MSE = .006$, $p<.05$, reflecting the fact that there was a larger difference across narrations for the *repeated hearers* than for the *new hearers*. A mixed multifactorial analysis of variance with items as the random variable showed a main effect of narration, $F(1,22)=6.851$, $MSE = .010$, $p<.05$. There was no significant interaction between hearer and narration.

With subjects as the random variable, separate paired t tests comparing the duration of the first use of the referring expression in Narration 1 to that in Narration 2 yielded a significant effect for the repeated hearers, $t(11)=3.966$, $p<.01$ (one-tailed), but not for the new hearer condition, $t(14)=0.410$, $p>.10$ (one-tailed). These results indicate that the mean durations of the words in Narration 1 were significantly longer than those of Narration 2 when Narration 2 was in the repeated hearer condition, but not when it was in the new hearer condition. Separate t -tests

also indicate that this result was true by items as well, $t(11)=3.183$, $p<.01$ for repeated hearers and $t(11)=1.069$, $p>.10$ for new hearers. Figure 4 represents the mean durations for referring expressions in Experiment 2.

[insert Figure 4 here]

Discussion

These data confirm that hearer knowledge has an effect on the reduction of repeated referring expressions. The first use of a referring expression in the first narration was significantly longer than the first use of the referring expression in the second narration when the hearer was repeated. However, this is not the case when the speaker was addressing a new hearer in Narration 2. The first use of a referring expression in Narration 1 was no longer, nor shorter, than the first use of the same referring expression in Narration 2 when the hearer was new. These results provide evidence that speakers' productions, specifically low level processes such as articulation, are influenced by their assumptions concerning what is given to the hearer.

The fact that we found that repeated referring expressions in Narration 2 in the *new hearer* condition were not significantly shorter than the same referring expressions in Narration 1 indicates that there is no attenuation across narrations that is purely speaker based. Such a conclusion is seemingly inconsistent with Bard et al. (2000) who found attenuation across narrations in their *speaker-old/hearer-new* condition. However, these results are not inconsistent with their findings for two reasons. First, Bard et al.'s conclusions concerning the attenuation of referring expressions in second narrations were based on the intelligibility of the referring expression. When they compared the durations of the referring expressions, they, too, found that the use of the referring expression in the second narration was not significantly shorter or longer than that in the first narration. Second, in this study, as well as Bard et al.'s, there is a possible

confound between durational shortening across narrations. Fowler and Housum (1987) found that repeated referring expressions were shortened upon second use, unless the second use of the referring expression crossed an episode boundary. They concluded that episode boundaries block the durational shortening of repeated referring expressions. This fact could explain why we do not see any durational shortening in the new hearer condition. Perhaps the speaker-based attenuation was blocked by the episode boundaries and only the attenuation from the additional contribution of hearer-based processes remained.

If the presence of the episode boundary does block durational shortening across narrations, this should be evident in both the new hearer and repeated hearer conditions. However, in Experiment 2 we compared the length of the *first* use of a referring expression in Narration 1 to the first use of the same referring expression in Narration 2. We know that referring expressions get shorter with subsequent use (see Experiment 1). And there is evidence that durational shortening is a gradient rather than a binary distinction. In a study of word duration and predictability, Gregory et al. (1999) found that words get shorter when they are repeated, and yet even shorter if they are repeated again. In this experiment, we found that the first use of a referring expression in Narration 2 is shorter than the first use of the referring expression in Narration 1 in the repeated hearer condition. But perhaps the first use of a referring expression in narration 2 is *longer* than the *last* use of the same referring expression in Narration 1. Thus, to test if the presence of an episode boundary blocks durational shortening, the proper comparison to make is between the *last* use of a referring expression in Narration 1 to the *first* use of the same referring expression in Narration 2.

Experiment 3

In Experiment 2, we compared the first use of a referring expression in the first narration to the first use of the same referring expression in a second narration. We demonstrated that when the second narration was uttered in the presence of a repeated hearer, the referring expression was shorter, while there was no significant difference between the duration of the two referring expressions if the hearer in the second narration was new. However, does it continue to get shorter even across narrations in the presence of a repeated hearer, or does the change in episode boundary block durational shortening? Can the lack of significant shortening in the new hearer condition be attributed to blockage by the episode boundary rather than lengthening in the presence of a new hearer? In this experiment, we compare the *last* use of a referring expression in Narration 1 to the *first* use in Narration 2.

Method

Design. The corpus used for comparisons in Experiment 3 has been described in the General Method. The comparisons for this experiment were very similar to Experiment 2, except that the comparison were between the duration of last use of the referring expression in Narration 1 to that of the first use of the referring in Narration 2 for both the *repeated hearers* and *new hearers* conditions. The example in 3 illustrates the experimental design. Note that the comparisons that were made were between the underscored referring expressions for each narration. There were a total of 30 narrations analyzed for this experiment.

(3) **Condition: New Hearer**

Narration 1: New Speaker/ New Hearer

okay this is story number two it was about the prime minister his name was marty johnson he resigned on tuesday after his announcement that he had health problems¹ and it was also due to a disagreement with the president over a nuclear plant so after five months in his job he had a meeting with the president and he announced his resignation and he said it was due to health problems²

Narration 2: Old Speaker/ New Hearer

okay um this is story number two there was a prime minister his name was marty johnson and he decided that he was going to resign on tuesday um after he had a meeting with the president about health problems¹ and a disagreement over a nuclear plant so after his meeting with the president he decided that he was going to resign after five months at his job so he resigned yes so um it was over like the nuclear plant whatever that means and also his health problems² is another reason so he ended up resigning because of his health problems³

Condition: Repeated Hearer

Narration 1: New Speaker/ New Hearer

well you see what happened is there's this prime minister marty johnson who resigned on tuesday after working five months as prime minister and um there was some debate whether like he resigned because of a problem with the president or because of health problems¹ and he clearly stated that there was uh that it was only because of health problem² had nothing to do with the president and his personal relationship with him he was the prime minister

Narration 2: Old Speaker/ Old Hearer

okay this is what happened there's this prime minister um marty johnson who resigned tuesday after working five months didn't say where he worked what country or anything but there had been speculation about him resigning but he was unclear why he was going to resign but after he resigned on tuesday he stated that it was because his health problems¹ and the speculation had been whether it was his health problems² or over um an argument with the president over a nuclear power plant

The prediction is that if the length of the first use of a referring expression in Narration 2 is *longer* than the last use of the same referring expression in Narration 1 in the *repeated hearer* condition, we can conclude that the presence of an episode boundary, blocks durational shortening, and could explain the lack of shortening in the new hearer condition of Experiment 2. However, if we find the length of the first use of a referring expression in Narration 2 is *shorter* than the last use of the same referring expression in Narration 1 in the repeated hearer condition, then we can conclude that durational shortening is not blocked by the presence of an episode boundary. Similarly, if we find that the first use of a referring expression is *longer* than the last use of Narration 1 in the new hearer condition, but not in the repeated hearer condition, then we can conclude that the presence of a new hearer blocks attenuation.

Stimuli. As described in the General Method, stimuli were only included for analyses if (a) the whole referring expression was used, (b) the referring expression was used at least once in

each narration, and (c) the referring expression was in the same syntactic and semantic environment in both narrations. A total of 120 referring expressions were used in all analyses.²

Results

A set of analyses was conducted on the means for the log duration in ms normalized by number of syllables. First, a mixed multifactorial analysis of variance was conducted on the normalized durations including the between-subjects variable of hearer (*repeated* vs. *new*) and the within subject variable of narration. With subjects as the random variable, there was no main effect of narration, but there was an interaction between hearer and narration, $F(1,25)=6.367$, $MSE = .024$, $p<.05$, reflecting the fact that there was an increase in duration across narrations for new hearers but a decrease for repeated hearers. A mixed multifactorial analysis of variance with items as the random variable showed similar results: no main effect of narration, but a significant interaction between hearer and narration, $F(1,113)=11.408$, $MSE = .051$, $p<.01$. The graph in Figure 5 represents the means of each narration, split by condition.

[insert Figure 5 here]

The graph in Figure 5 demonstrates that in the new hearer condition, the first use of a referring expression in Narration 2 is longer than the last use of the same referring expression in Narration 1. In the repeated hearer condition, the first use of a referring expression in Narration 2 is shorter than the last use of the same referring expression in Narration 1. Separate paired t tests with subjects as the random variable comparing the durations of the last use of a referring expression in Narration 1 to those of the first use in Narration 2 yielded a significant effect for the *new hearers*, $t(14)=-1.887$; $p<.05$ (one-tailed), and a significant effect for the *repeated hearer* condition, $t(14)=2.898$, $p<.05$.³ These results indicate that the mean durations of the last use of the referring expressions in Narration 1 were significantly shorter than those of the first use of

the referring expression in Narration 2 when Narration 2 was in the new hearer condition. But when the second narration was in the repeated hearer condition, the mean durations of the last use of the referring expressions in Narration 1 were significantly longer than the first use of the referring expression in Narration 2.

Discussion

The data from Experiment 3 demonstrate that referring expressions are shortened even across episode boundaries when the speaker and the hearer are both old. This finding is not true when only the speaker is old, but the hearer is new: Referring expressions were longer in their first use in Narration 2 than in their last use in Narration 1 when the hearer was new. These data provide evidence that the presence of an episode boundary alone does not block durational shortening: Referring expressions were shortened across narrations in the repeated hearer condition. This fact, coupled with the fact that the first use of a referring expression in Narration 2 is longer than the last use of that referring expression in Narration 1 in the new hearer condition, suggests that the presence of a new hearer blocks durational shortening.

General Discussion

In Experiment 1, we showed that words are shortened upon second use within a narration. This finding is consistent with the findings of Bard et al. (2000), Fowler (1988), Fowler and Housum (1987), and Fowler et al. (1997). In Experiments 2 and 3, we demonstrated that referring expressions are shortened in the presence of repeated hearers, but not in the presence of new hearers, despite an intervening episode boundary in both conditions. Furthermore, Experiment 3 demonstrates that durational shortening continues across episode boundaries for repeated hearers, but is blocked by the presence of a new hearer. The fact that referring expressions are shortened with repeated hearers but not new hearers lends support to the role of

mutual knowledge in production (*c.f.* Clark & Wilkes-Gibbs, 1986). These results indicate that speakers' productions are sensitive to information that is common ground between interlocutors, and more specifically speakers' knowledge, or assumptions, concerning what hearers know: Speakers have a model of hearer knowledge that affects the durations of words in discourse.

It is important to note that these results illustrate the role of mutual knowledge in the execution of utterances. These results have little bearing on the role of shortening caused by speaker internal processes. We have not demonstrated that speaker internal processes are *not* a factor in reduction: We have merely shown that hearer knowledge is *also* a source of reduction. The fact that Bard et al. (2000) found a loss of intelligibility, as measured by excising the word and getting judgments from participants as to the words identity, of referring expressions in second narrations with a new hearer suggests that the some aspects of articulation are due to speaker internal processes.

It would be very difficult to completely tease apart the effects from speaker internal processes and effects from mutual knowledge. In this study we have merely contributed to the growing body of research by demonstrating that hearer knowledge has an effect on articulation. Bard et al. (2000) had two conditions in their study: speaker-new/hearer-new and speaker-old/hearer-new. We added one more condition, speaker-old/hearer-old. For a better understanding of how mutual information is used by the production system, it seems to best tease apart speaker-based effects from hearer-based, a fourth condition, speaker-new/hearer-old, should be added to the design, but we leave this to future research.

Some might argue that the design of this study is methodologically awkward. One can imagine that telling the same story to two people is a rough approximation of natural conversation: People tell their stories to new audiences all the time. However, repeating a story

to the same audience is a bit unnatural: People generally only repeat stories to the same person when they forget that they have already told someone a story. However, we argue that the unnaturalness of the task is biased against finding evidence for hearer-based shortening. If participants were confused or affected by the awkwardness of the task, then we might not expect to see any shortening in the repeated hearer condition. But, despite the awkwardness of the task, we saw a significant effect of shortening in this condition in all three experiments. Thus, we feel the awkwardness of the task did not affect the results.

While we have demonstrated that mutual knowledge is a factor in articulation, our experiments did not target the specific details of a production model that are affected by mutual information. Thus, we cannot make any claims concerning the proposal made by Brown and Dell (1991), Dell and Brown (1994) & Horton and Keysar (1996) that a hearer model may be consulted during the monitoring and adjustment phase of articulation (see Levelt, 1993). Monitor and adjustment can occur both after the execution of an articulatory plan, resulting in a repair, or before the execution of the articulatory plan. It is possible that mutual knowledge is not consulted until after the articulatory plan is in place, but before execution. It is also possible that mutual knowledge may be used to help construct the articulatory plan, prior to the monitoring and adjustment phase. However, much more needs to be known about the details of the lower level, fast automatic processes of the production system before the exact nature of the role of a hearer model can be determined.

Dell & Brown (1994), while acknowledging that hearer knowledge may be a factor during the monitoring and adjustment phase of articulation, proposed that in lieu of speakers having a model of hearer knowledge that interacts with the production system, speakers rely on their own knowledge as a good approximation of the hearers' knowledge. While our study did

not directly address this issue, our results are inconsistent with such a proposal. Speakers' internal knowledge was the same in both the repeated hearer condition and the new hearer condition (the speaker was "old" in the second Narration). If a speaker's internal knowledge (as a rough approximation of hearer knowledge) is the only source of durational shortening, then we should have seen shortening in both conditions, or in neither. Yet, we observed a difference in durational shortening only in the repeated hearer condition. One still might argue that speaker internal knowledge is updated to reflect the change in the discourse context (the presence of a new hearer) and the differing results between the repeated hearer condition and the new hearer condition are a fact about speaker internal knowledge only. Such an explanation requires that "a good approximation" of hearer knowledge must include very specific details about the discourse situation for the speakers productions to be so accurate, "the hearer is new, therefore there is no common ground, so produce an utterance as if it were new to me." If speakers' approximations of hearer knowledge are good enough to update the articulatory system in the appropriate manner for a given discourse context, then we argue that hearer knowledge still plays a role in articulation, even if that knowledge serves only as a pointer to appropriate speaker internal knowledge for the discourse context.

The focus of this study is the role of mutual knowledge in production. However, the results reported here have some interesting implications for the debate concerning the role of mutual knowledge in comprehension. Clark and Carlson (1982) and Clark and Marshall (1981) argued that comprehension is optimized when a hearer restricts the interpretation of an utterance to mutual information. However, Keysar, Barr, Balin, and Paek (1998) (see also Barr & Keysar, 2002) argued that hearers do not differentiate between mutual information and information that was available only to them and not the speaker. In a cross-modal eye tracking study, Keysar et al.

(1998) found that hearers' responses to instructions given by speakers who were physically present were affected by knowledge that they heard through a recording transmitted through an earpiece. The hearers knew the speakers did not have access to information they heard through the earpiece. They concluded that the comprehension system is not restricted to mutual knowledge; hearers are egocentric and make use of whatever information is available in the speech context.

While this debate has not been completely resolved, it does raise some interesting issues for language processing in general. Our results demonstrate that mutual knowledge is a factor in production. What is the functional role of shortening words that hearers know? There are two obvious candidates. The first concerns the speaker's economy. Jespersen (1922) argued that predictable are shortened because the hearer has other information by which to identify the word, so speakers are free to produce them in an attenuated manner (see also Liberman, 1963). This explanation is independent of hearers and the comprehension system. By this explanation, the role of mutual in production serves only as an aid to speaker economy. However, a second explanation for the durational shortening of repeated words is that the speaker uses productions to signal information to hearers. In order for such an explanation to be correct, one must find evidence that hearers are sensitive to mutual knowledge, or at least sensitive to the production processes that are affected by mutual knowledge.

If hearers are in fact egocentric, as Barr et al. claim, then we can assume that the role mutual information in production serves only as a guide to speaker internal processes, such as speaker economy. Conversely, if hearers are sensitive to mutual knowledge, and specifically, word durations, then we can conclude that the durational shortening of repeated words serves to signal information to the hearer. One must keep in mind that such a conclusion does not entail

that speakers shorten repeated words *in order* to signal information to the hearer, only that the comprehension system is sensitive to the shortening caused by mutual knowledge. Such findings would suggest a similarity between the comprehension and production systems with regard to mutual knowledge: Mutual knowledge serves to shape productions in a manner that can be useful to hearers. However, before any such conclusions can be drawn, we must first know if the comprehension system is sensitive to the aspects of production (durational shortening in this case) that are affected by mutual knowledge.

References

- Balota, D. A., Boland, J. E., & Shields, L. W. (1989). Priming in pronunciation: Beyond pattern recognition and onset latency. *Journal of Memory and Language*, 28, 14-36.
- Bard, E. G., Anderson, A. H., Sotillo, C., Newlands, A., Doherty-Sneddon, G., & Aylett, M. (2000). Controlling the intelligibility of referring expressions in dialogue. *Journal of Memory and Language*, 42, 1-22.
- Barr, D. J., & Keysar, B. (2002). Anchoring comprehension in linguistic precedents. *Journal of Memory and Language*, 46, 391-418.
- Beckman, M., & Ayers, G. (1997). *Guidelines for ToBI labeling, version 3*. Columbus Ohio: Linguistics Dept. Ohio State University.
- Bock, K., & Levelt, W. (1994). Language production: Grammatical encoding. In M. A. Gernsbacher (Ed.), *Handbook of Psycholinguistics* (pp. 945-984). San Diego: Academic Press, Inc.
- Boersma, P., & Weenink, D. (1992-2001). Praat, a system for doing phonetics by computer. Retrieved Nov. 6, 2000, from <http://www.praat.org>.
- 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. (1987). Adapting production to comprehension: The explicit mention of instruments. *Cognitive Psychology*, 19, 441-472.
- Chafe, W. (1987). Cognitive constraints on information flow. In R. Tomlin (Ed.), *Coherence and Grounding in Discourse* (pp. 21-51). Amsterdam: John Benjamins.

- Chafe, W. (1976). Givenness, contrastiveness, definiteness, subjects, topics, and points of view. In C. Li (Ed.), *Subject and Topic* (pp. 25-55). New York, NY: Academic Press.
- Clark, H. H. (1994). Discourse in production. In M. A. Gernsbacher (Ed.), *Handbook of Psycholinguistics* (pp. 985-1021). Academic Press: San Diego.
- Clark, H. H. & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. Levine, & D. Behrend (Eds.), *Perspectives on socially shared cognition* (pp. 127-149). Washington DC: American Psychological Association.
- Clark, H. H. & Carleson, T. B. (1982). Speech acts and hearers beliefs. In N. V. Smith (Ed.), *Mutual Knowledge* (pp. 1-37). London: Academic Press.
- Clark, H. H. & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshe, B. Webber, & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 10-63). Cambridge: Cambridge University Press.
- Clark, H. H. & Wilkes-Gibbs, D. (1986). Referring as a Collaborative Process. *Cognition*, 22, 1-39.
- Cutler, A. (Ed.). (1982). Slips of the tongue and language production. New York: Mouton.
- Dell, G. S. (1984). Representation of serial order in speech: Evidence from the repeated phoneme effect in speech errors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 222-233.
- Dell, G. S., & Brown, P. M. (1991). Mechanisms for listener adaptation in language production: Limiting the role of the "model of the listener". In D. J. Napoli & J. A. Kegl (Eds.), *Bridges Between Psychology and Linguistics: A Swarthmore Festschrift for Lila Gleitman* (pp. 111-129). Hillsdale, NJ: Erlbaum.

- Fowler, C. (1988). Differential shortening of repeated content words produced in various communicative contexts. *Language and Speech*, 31, 307-319.
- Fowler, C., & Housum, J. (1987). Talkers' signaling of "new" and "old" words in speech and listeners' perception and use of the distinction. *Journal of Memory and Language*, 26, 489-504.
- Fowler, C., Levy, E. T., & Brown, P. (1997). Reductions of spoken words in certain discourse contexts. *Journal of Memory and Language*, 37, 24-40.
- Garret, M. F. (1975). The analysis of sentence production. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 133-177). New York: Academic Press.
- Gerrig, R. H., Brennan, S. E., & Ohaeri, J. O. (2001). What characters know: Projected knowledge and projected co-presence. *Journal of Memory and Language*, 44, 81-95.
- Gregory, M. L., Raymond, W., Bell, A., & Jurafsky, D. (1999, May). *Effects of informativeness on word duration in conversation*. Paper presented at the meeting of the Society for Text and Discourse Vancouver, B.C.
- Gundel, J., Hedberg, N., & Zacharski, R. (1993). Referring expressions in discourse. *Language*, 69, 274-307.
- Horton, W. S., & Keysar, B. (1996). When do speakers take into account common ground? *Cognition*, 59, 91-117.
- Jespersen, O. (1922). *Language: Its nature, development, and origin*. New York: Henry Holt and Co.
- Keysar, B., Barr, D. J., Balin, J. A., & Paek, T. S. (2002). Definite reference and mutual knowledge: Processes models of common ground in comprehension. *Journal of Memory and Language*, 39, 1-20.

- Lieberman, P. (1963). Some effects of semantic and grammatical contexts on the production and perception of speech. *Language and Speech*, 6, 172-187.
- Levelt, W. (1983). Monitoring and self repair in speech. *Cognition*, 14, 41-104.
- Levelt, W. (1989). *Speaking: From intention to articulation*. Cambridge: MIT Press.
- Nadig, A. S. & Sedivy, J. C. (2001). Evidence of perspective-taking constraints in children's on-line reference resolution. *Psychological Science*, 13, 329-336
- Prince, E. (1981). Toward a taxonomy of given-new information. In P. Cole (Ed.), *Radical Pragmatics* (pp. 223-255). New York: Academic Press
- Raymond, W. D. (2000). *Toward a theory of grammatical encoding in speech production: Evidence from speech errors*. Unpublished doctoral dissertation, University of Colorado, Boulder.
- Sedivy, J. C. (2002). The incremental interpretation of adjectives in referential contexts: Evidence for the computation of Gricean inferences in on-line processing. Manuscript submitted for publication.
- Shattuck-Hufnagel, S. (1979). Speech errors as evidence for serial order mechanism in sentence production. In W. E. Cooper & E. C. T. Walker, E. C. T. (Eds.), *Sentence processing: psycholinguistic studies presented to Merrill Garrett* (pp.295-342). Hillsdale NJ: Erlbaum.
- Stemberger, J. P. (1983). Speech errors and theoretical phonology. *Indiana University Linguistics Club*. ms.

Appendix

Story 1.

Longtime resident John Powers became Fallon's first black mayor. He promises to unite black and white residents of a city that played a pivotal role in the civil rights movement. Powers, who ran against his former boss, takes office this month. He has made history in a city where demonstrations led to the opening of Southern polling booths to blacks.

Story 2.

Prime Minister Marty Johnson resigned Tuesday, ending weeks of speculation that he would step down because of health problems and a disagreement with the president over a nuclear plant. Johnson, who took the job five months ago, announced his resignation after a brief meeting with the President. He said the move was prompted by health problems.

Footnotes

¹ The fact that this number is so much larger than the total number of items for Experiment 1 reflects the fact that items were often not repeated during one narration, but were used at least once in both narrations.

² Even though we are comparing the durations of the last use of a referring expression in Narration 1, if a particular referring expression were used only once in the first narration, that use was considered to be the last use.

³ Given the results from Experiment 2, we expected that either the first use of the referring expression in Narration 2 would be longer than the last use of Narration 1 in the new hearer condition, or there would be no difference at all. We had no evidence suggesting the length might be shorter in Narration 2. Thus, we used a one-tailed t-test. However, for the repeated hearers, we had no expectations concerning the length of the first use of the referring expression in Narration 2 compared to the last use of the same referring expression in Narration 1. Thus, we used a two-tailed t-test for that comparison.

Table 1

The Twelve Referring Expressions.

Story 1	Story 2
John Powers	Marty Johnson
resident	prime minister
Fallon	health problems
Mayor	disagreement
former boss	Meeting
civil rights	President

Table 2.

The two conditions in Experiment 2.

Condition	Narration 1		Narration 2	
	speaker	Hearer	speaker	hearer
New hearer	new	New	old	new
Repeated hearer	new	New	old	old

Figure Captions

Figure 1. Examples of voiceless stops, word initial and final

Figure 2. Adjacent bilabial nasal consonants.

Figure 3. Means of duration in ms. for Experiment 1.

Figure 4. Mean durations in ms. for referring expressions across narrations in Experiment 2.

Figure 5. Mean duration of the *last* referring expressions compared to the *first* in Experiment 3.

Figure 1.

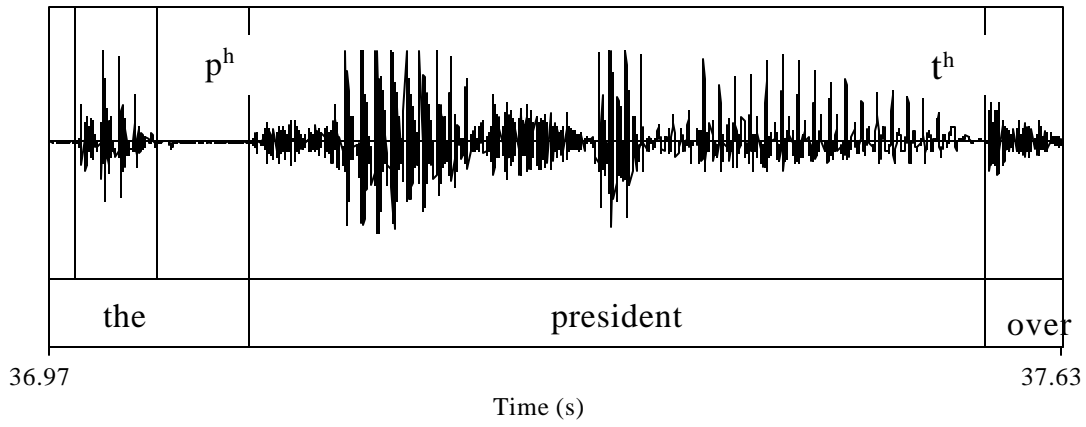


Figure 2.

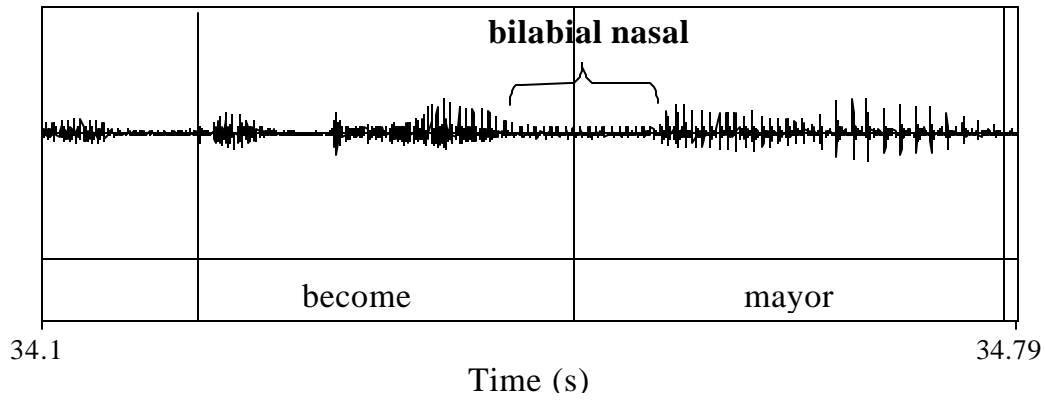


Figure 3.




Figure 4.




Figure 5.

