

RUNNING HEAD: Comprehension After Oh

To cite, please refer to: Fox Tree, J. E., & Schrock, J. C. (1999). Discourse markers in spontaneous speech: Oh what a difference an oh makes. *Journal of Memory and Language*, 40, 280-295.

Discourse Markers in Spontaneous Speech:

Oh What a Difference an Oh Makes

Jean E. Fox Tree and Josef C. Schrock

University of California, Santa Cruz

Abstract

Discourse markers are usually studied from the vantage point of corpora analyses. By looking at where they fall in spontaneous talk, hypotheses can be made about their possible functions. But direct tests of listeners' uses of these expressions are rare. In five experiments, we looked at the on-line spontaneous speech comprehension effects of one discourse marker, oh. We found that recognition of words was faster after oh than when the oh was either excised and replaced by a pause or excised entirely. We also found that semantic verification of words heard earlier in the discourse was faster after oh than when the oh was either excised and replaced by a pause or excised entirely, but only when the test point was downstream from the oh. Results demonstrate that oh is not only a potential signal to addressees, as has been suggested by corpora analyses, but that it is in fact used by addressees to help them integrate information in spontaneous talk.

Oh What a Difference an Oh Makes

One of the ways spontaneous talk differs from planned talk is the presence of discourse markers such as well, I mean, like, and oh. Discourse markers are rarely found in prepared or rehearsed speeches, but rarely absent in conversations. Many novelists make use of this distinction; their descriptions of the setting or plot have no you knows and so anyways, but the characters' dialogues do. The presence of these words creates a naturalistic conversational effect. But they may do even more, including highlighting important elements in a narrative, helping listeners follow a speaker's train of thought, helping listeners recover from a repair, and explicitly showing the relationship between two utterances, such as whether the upcoming utterance is in contrast to the prior utterance, is an elaboration of the prior utterance, or is unrelated to the prior utterance (Erman, 1986, 1987; Flowerdew & Tauroza, 1995; Jucker, 1993; Levelt, 1989; Redeker, 1990; Schiffrin, 1987; Schourup, 1985). Discourse markers may be infrequent in prepared manuscripts and speeches because these forms of communication allow advance planning and extensive revision time. In contrast, they may be frequent in spontaneous talk because that is where people need help organizing and expressing ideas.

In this paper, we take a closer look at the discourse marker oh. First we will review what analyses of spontaneous talk suggest about how oh is used, and then we will present results of five experiments where we tested how oh affects listeners' comprehension.

The Function of Oh from a Discourse Analysis Perspective

One way to find out how oh is used is to look at where it occurs and infer from that its likely function. Several researchers have taken this approach in investigating the conversational use of oh. Aijmer (1987) has done perhaps the most extensive study of oh. She looked at every oh in the London-Lund corpus of speech and identified a myriad of ways that ohs could be used, including the following: (1) as a sudden reaction to new or surprising information, such as a sudden epiphany, a sudden recollection, a surprise

question, or a surprise offer; (2) as a disjunct marker informing the addressee that something said earlier is necessary to understand the upcoming utterance; (3) as a signal of an upcoming ambiguous, nonserious, or elliptical thought; (4) as a signal of an upcoming repair; (5) as an enquoting device to mark a shift between the speaker and the character in the speaker's narrative; (6) to add emphasis; (7) to indicate an upcoming emotional or evaluative utterance, such as one showing frustration, annoyance, repentance, agreement, or politeness; (8) as an interactional device to indicate that an utterance has been accepted to common ground, that an utterance has not been accepted to common ground, or that a speaker should keep talking to complete an idea; or (9) as part of conventionalized phrases without any meaning whatsoever (as in "oh I beg your pardon" p. 82). Schiffrin (1987), in another corpus, identified yet more uses including: (10) as an attention-getting device, (11) as a floor keeping device, (12) to demonstrate the speaker's engagement in the conversation, or (13) to demonstrate that an interlocutor's emotions are either less intense or more intense than expected. And other researchers have identified other uses: (14) as a device to elicit information from a speaker (Redeker, 1991; James, 1972), (15) as a mitigator (Redeker, 1991), (16) to show that the speaker is choosing what to say next or hedging (James, 1972), and (17) to show that the speaker was present at the time that a quote prefaced by oh was made (Schourup, 1985).

This desultory array of functions can be reduced to some regularities. In perhaps the earliest in-depth investigation of oh, Heritage (1984) proposed that speakers say oh to let their addressees know that the speaker's model of the communicative exchange is undergoing a change of state. This change of state often comes about from new information, such as new facts or new perspectives on prior information. In some circumstances, the use of oh may be even more precise, and may mark unanticipated new information, as opposed to anticipated new information (Schiffrin, 1987, p. 89). Of course, every utterance will add something new to the developing discourse (Clark & Haviland, 1974; Haviland & Clark, 1974). But a change of state is different from implicit

given-new structure because a change of state need not refer to adjacent utterances.

Given-new structure describes relationships where following utterances are interpreted with respect to preceding utterances. Oh can mark that the following utterance contains an updating of information that was presented much earlier in the speech stream, not necessarily only an updating of the immediately prior utterance. In contrast to given-new structure, oh explicitly marks a disjuncture between adjacent utterances, halting the normal integrative processes.

Positing that oh has an underlying change-of-state function can help make sense out of some of the seemingly disparate functions of oh researchers have identified. By indicating a change of state, oh implies that the preceding information has been accepted into the developing discourse record. This implication can both bluntly end an argument and show politeness. When said prematurely, oh can have the effect of cutting short a narrative, as in the following example (adapted from Heritage, 1984, p. 332; asterixes indicate overlapping speech):

(1) N: Does he have his own apart*ment*?

H: *hhh* Yeah

N: Oh (1 sec) How did you get his number?

The questioner's oh stops the answerer from elaborating the yeah or providing further information. But at other times, replying with oh can have the effect of identifying an interlocutor's utterance as relevant, unanticipated, and newsworthy. This can make a reply with oh more polite than a reply with a mere acknowledgment. Saying yes only recognizes that the interlocutor has provided understandable information. Similar to its politeness function, oh can be used deceptively to indicate that information has brought about a change of state even when it hasn't. Likewise, speakers can deliberately withhold an oh to indicate that information is not informative when it is (Heritage, 1984). Oh doesn't have to indicate a genuine change of state, or to precede genuinely new or unanticipated information. It just has to be used to imply a change of state by the oh-producer.

The change-of-state function can also help explain why oh is not acceptable in some environments. Oh cannot occur in an idiom (*John kicked...oh...the bucket; James, 1972, p. 164), a curse (*Fuck...oh...John Wayne; James, 1972, p. 165), or in some sentence constructions (*Wisely...oh...Sue arranged to meet Carol; James, 1972, p. 165). But note that in all of these cases, speakers are likely to be able to complete the sentences without a change of state, so they shouldn't need an oh. Using oh would be pragmatically implausible. Consider the following minimal pair: With a hammer...oh...Bill hit Fred (James, 1972, p. 165) is unacceptable, but With a hammer...oh...you can build a stepladder (James, 1972, p. 166) is fine. In the first, speakers must already have an idea about the relationship between the parts of the utterance, so they shouldn't need a change of state marker. But in the second, oh does represent a change of state, from not knowing what to do with a hammer to knowing that you can build a stepladder with it.

Finally, three common environments that oh occurs in can be easily viewed as environments where a change of state is likely. Oh commonly occurs with repairs, after question answers, and with quotations. A repair entails a change of state from the original version to the corrected version. When a repair is made, some information already in the discourse record has to be updated to reflect additional or corrected information. In the following example taken from a corpus of spontaneous speech collected by Herbert Clark at Stanford University, a speaker is relating an interview between a journalist and Gary Hart's campaign manager. The speaker initially fails to recall the name of the fourth president mentioned by the campaign manager. A moment later the speaker does recall the name, and updates the discourse to reflect this change in available information:

(2) and then he cited four presidents, um Jefferson Eisenhower Kennedy, and uh someone else for um . . . for being adulterers, but- but being very great presidents, oh and Roosevelt.

When the speaker recalls the name Roosevelt, she begins the utterance with oh. By introducing the information with oh, the speaker foreshadows an upcoming change of state, in this case, an update to the previous discourse.

Questions are often asked in order to solicit new information. When the question askers say oh after hearing answers, they are providing overt evidence of their updated knowledge states. Without ohs, answerers might not know if they had provided sufficient or appropriate information. In the following example adapted from Heritage (1984, p. 314), a caller to a radio show tries to guess “the second group to enter the British ‘Top 20’ at No.1” and doesn’t realize she has it wrong:

(3) Caller: The Jam. “Going Underground.” Nineteen eighty.

DJ: Uh no Carla. That’s why I asked you if you thought Slade were the first.

Caller: Yes.

DJ: ‘Cos the Beatles were first. Slade were second and the Jam were third.

Caller: Yes.

DJ: No. The Jam were the third group to go straight in at number one. Yeah?

Caller: O:h.

The caller answers yes to information she couldn’t have known based on her wrong answer to the question; the disc jockey expects an oh and continues providing information until the caller produces it. Whereas yes could mean simply “I’m listening; go ahead,” oh shows that the caller recognized the DJ’s information as new. Of course, oh is not mandatory. The caller could also have demonstrated understanding by continuing the conversation in a way that incorporated the new information, or by asking another question using the information. There are many ways to show an updated knowledge state; oh is just one useful device.

As an enquoting device, oh can indicate a change of state in the person being quoted (Schourup, 1985). In the following example taken from the Clark corpus, the speaker reports an exchange between Alfie, who is British, and Jane, an American:

(4) they begin talking about the high taxes, and- um and- uh Jane comments on how high the taxes are, and Alfie says well what do you get with your high taxes? and she says oh we get wonderful things like moon probes

The oh is Jane’s, not the speaker’s, and it shows Jane moving from a state of not having thought

about the benefits of high taxes to a state of having thought about it. Direct quotations are rarely verbatim accounts of what the quoted speaker actually said (Schourup, 1985; Wade & Clark, 1993). Instead they demonstrate something about the quoted information, such as the language used or the tone of voice (Wade & Clark, 1993). We think oh adds particular information to the demonstration; that is, that the enquoted person is undergoing a change of state.

So the function of oh can either be described as an exhaustive list of each environment that oh occurs in, with each use treated independently, or as a single underlying construct that in combination with different environments manifests itself somewhat differently. We choose to take the underlying construct path, partly because the exhaustive list yields some contradictory claims about oh's use, and partly because the exhaustive list makes it difficult to make predictions about which function applies to a particular example. But whether the underlying construct approach will ultimately be satisfactory for all occasions of use we leave to future analyses.

Although we have hypotheses about what oh means for addressees, its actual effect on listeners' comprehension has never been directly tested. Speakers' productions of oh may correspond with changes of state, but this doesn't mean that listeners use this information. Corpora analyses help generate hypotheses for what function oh might serve for addressees, but only direct testing can demonstrate the actual use of such a potential signal. In this paper we will focus on the use of oh in the environment of repair and investigate how these ohs affect listeners' understanding of the discourse.

When speaking spontaneously, people sometimes need to correct a previous statement, insert new information, or redirect their utterances. Oh may serve an important function for listeners in these self-initiated repairs (see also discussion in Schiffrin, 1987). In the following example from the Clark corpus, the speaker initially omitted information that she later realized was important to the story, which is about buying pet ants:

(5) and for a really champion one, you can- it's gonna be twenty cents, she says well okay how about ten cent range [sic]? so he tips out- oh no then he says um okay what

size were you thinking of? and she said oh I guess medium size. he said okay, and he tips out a bunch of medium si- medium-sized ants on the counter.

Before discussing tipping out the ants, the appropriate box, ranged by size, had to be selected. Repairs usually consist of stopping the current flow of speech, inserting a pause or editing expression, and then providing new or modified information (Clark, 1996; Fox Tree & Clark, 1997; Levelt, 1983). By hypothesis, the speaker uses the editing expression to forewarn the upcoming repair. But furthermore, the selection of oh indicates that the repair will be a change of state, or a disjuncture from the immediately preceding utterance. In our examples, the repair may be updating information in the prior utterance, but it may also be updating information from much further back in the discourse. In the current experiments, we tested whether hearing oh is informative to listeners.

Testing How Listeners Use Oh

There are at least three ways oh could influence listeners' comprehension. It could have no effect, a negative effect, or a beneficial effect.

Our argument about the function of oh and other discourse markers in spontaneous speech coupled with the discourse analysis research just reviewed supports the idea that oh benefits comprehension. There are two possible mechanisms by which oh might benefit comprehension. We will call them the disjuncture mechanism and the expectation mechanism. The disjuncture mechanism is that oh informs listeners to halt integrative processes across adjacent utterances. By indicating a disjuncture between the immediately preceding and immediately following utterances, oh informs listeners that they should process the upcoming information independently from the immediately preceding information. Without the oh listeners would attempt ordinary integration across adjacent utterances that may not be integratable, leading to confusion. The expectation mechanism is that oh might cause listeners to expect a change of state in the upcoming speech. This expectation might benefit listeners' ability to integrate upcoming information. Both mechanisms lead to the prediction that comprehension will be better when oh is present than when oh is absent.

Although we predict that oh will benefit comprehension, listeners' actual use of oh has not been tested, and it is possible to argue that oh has no effect or a detrimental effect.

Comprehension models usually tacitly treat discourse markers as ignorable particles of speech (for discussion, see Brennan & Williams, 1995; Clark, 1997; Fox Tree, 1995). Most syntactic parsing models, lexical access models, and speech parsing models do not even discuss them. And when they are discussed they are often treated as semantically empty (O'Donnell & Todd, 1991; see also discussion in Segal, Duchan, & Scott, 1991). Perhaps this treatment is warranted. It is possible that oh has no meaning for a listener, no matter what it might mean to a speaker. If oh is ignored by listeners, it should have no effect on comprehension.

It is also possible that oh has a negative effect on comprehension. Because it isn't an ordinary word with a conventional grammatical role, it might create havoc in the listener's attempt to figure out the syntactic relationship among words in an utterance. What node of a traditional syntactic tree would oh attach to? As an unconventional word that might not have a meaning stored in the lexicon, oh might also create problems with lexical access. Furthermore, the updating function of oh, marking the change from some prior state to a different state, could be viewed as a kind of a repair itself, as opposed to just an editing expression occurring between a reparandum and a repair. Some repairs are disruptive and can slow comprehension (Fox Tree, 1995). Oh might be one of them.

To test what effect oh actually has on comprehension we compared comprehension of discourse containing an oh to comprehension of the same discourse with the oh digitally spliced out. We tested comprehension in two ways. In one way we measured the availability of words from the discourse following oh. If subsequent words are more available when oh is heard than when it is not, this could indicate that oh helps listeners integrate speech. But it could also indicate that oh merely serves as a general attention-getting device. In our second type of test, we measured the availability of words from the discourse preceding oh. If prior information is more available when oh is heard than when it is not, then listeners must be using oh to help them integrate the discourse as a whole. With the disjuncture mechanism, if oh signals that listeners

should process the upcoming information independently, then hearing oh puts listeners in a better position to reconnect the repair information after the oh with the reparandum preceding the oh. Without the oh, listeners may attempt to integrate two nonintegratable utterances, leading them to confusion which could prevent their being able to connect the repair and reparandum. With the expectation mechanism, if oh focusses attention on an upcoming change of state, then hearing it will lead listeners to think about how upcoming information fits with prior information, activating earlier concepts.

To test the availability of information after oh, we used an identical word monitoring task (Marslen-Wilson & Tyler, 1980). In this task, people listen to speech and press a button if they hear a particular word that they've been told beforehand to monitor for. Their reaction times are measured from the onset of the word to the time they hit the button. This task has been used in other studies of spontaneously spoken language understanding (Fox Tree, 1995). The more interpretable the information prior to the target word is, the faster people can monitor for the target word. To investigate the effect of oh on upcoming speech, we compared peoples' response latency to a word after oh to their response latency to the same word with the oh either replaced by a pause of equal length (Experiment 1) or excised entirely (Experiment 2).

To test the availability of information before oh, we used a semantic verification task. In this task, people see a visual target word while listening to speech and then press a button corresponding to whether or not that word had occurred in the narrative. Their reaction times are measured from the onset of the visually presented word to the time they hit the button. This task has been used in other studies of spoken talk to evaluate the activation of discourse concepts (Gernsbacher & Jescheniak, 1995). The more active the concept is, the faster people can verify that the word was present. If oh is only highlighting upcoming information, then visually presented words from the prior discourse will be verified equally quickly whether or not they are presented after an oh. But if oh helps listeners integrate the discourse as a whole, then the words should be verified faster after an oh. To investigate the effect of oh on discourse integration, we compared peoples' response latencies to visually presented target words selected from the

discourse preceding the ohs with their response latencies to the same word with the oh either replaced by a pause of equal length (Experiment 3) or excised entirely (Experiment 4). The targets were presented at the same point in the repair phrases following either the oh or the splice point.

Experiment 1: Word Monitoring With Pause

Method

Participants. Thirty-four UCSC students participated in this experiment in partial fulfillment of a course requirement. All participants in this and the following experiments were native English speakers. Six participants were removed because they failed to respond to at least 25% of the trials.

Materials. Materials were taken from a corpus of spontaneous narratives collected by Herbert Clark at Stanford University. This corpus consists of pairs of students telling stories to each other in a face-to-face situation. Stories were based on previously heard audiotapes, so some of the students were recalling the same story. Materials consisted of a stretch of speech from one speaker that contained a spontaneously occurring oh. The stretch of speech always included enough information so that the information after the oh could be fit in to the prior discourse where it belonged. Stimuli varied from 20 seconds to 1 minute 46 seconds in length, and from 41 words to 247 words. Three examples follow, with the critical oh in italics:

(6) And um she said that- that to see um Gary Hart cheating on his wife, how could one- one could be led to believe that um that Gary Hart might be inclined to cheat upon the public too. And he said that- and uh the press secretary responded by saying that there have been um many people that have been very faithful to their wives, but um were just absolutely atrocious presidents. And then he cited four presidents um Jefferson Eisenhower Kennedy and uh someone else for um for being adulterers but- but being very great presidents. Oh and Roosevelt.

(7) Okay this story um was supposed to be an interview. Johnny Carson was

supposed to be interviewing Mary Lou Hehner from Taxi. And they started talking about- oh he asks her about um her having bought a house.

(8) So he pulls out a couple ants, some ten cent ones, and he names them off. Uh one of them's an English bitch or something, different kinds of names, and she decides on an Afghan. and then they start talking. She asked him what um you feed an ant, and he says you don't, and- and she says won't they die? And she goes yeah- and he goes um yeah they will, if you don't feed them. And he discussed you know um that it's much cheaper not to feed, just buy- keep buying new ones. And then he goes you need an ant house. And so she goes okay. Uh and he pulls out an ant house um with all this ant furniture in it and stuff, a little ladder for the ant to com- climb up to and ring a bell, and little just- I don't know, kinda like what you'd see in a gerbil cage I guess just for ants. And so um she goes uh you know it- oh it's resembling a lion's cage or something. Anyway she goes um well won't the ant just crawl right out ?

Each stimulus began at the beginning of an idea and finished with a completed thought.

There were 20 critical stimuli, 20 filler stimuli, and 4 practice stimuli. In the critical stimuli, the first content word after the oh that had not occurred earlier in the passage was selected as the target word. On average, target words occurred 4.5 words after the oh, with a range of 1 to 11 words, and 3.03 sec after the oh, with a range of 0 to 8.45, as measured from the offset of the oh to the onset of the target word. Target words were from a variety of form classes and were from one to four syllables long. Because the same tokens of target words were heard across conditions, any effect found must result from the presence or absence of oh and not spurious variations in word pronunciation, word frequency, or grammatical constructions surrounding the target words. The filler stimuli were selected from the same corpus as the critical stimuli and were similar in length and content. Filler targets were created by identifying words that occurred in the filler passages about where the critical targets occurred in the critical passages and selecting closely related words that did not occur in the passages.

As can be expected with spontaneous speech, oh and other discourse markers did occur elsewhere in the target and filler stimuli. In fact, there were 11 other ohs in the materials people heard. There were also 22 occurrences of well, 25 occurrences of so, 9 occurrences of you know, and 3 occurrences of I mean. The presence of many other ohs in the stimuli prevents participants from adopting a strategy of responding merely upon hearing an oh. Further evidence that participants were not using this sort of strategy is that in debriefing, no participant detected the oh manipulation, even when they were explicitly asked “Did anything sound strange or unnatural in any of the speech?”.

The materials were digitized onto a Power Macintosh computer using SoundEdit 16 software. For each of the critical stimuli, a second version was created where the critical oh was excised and replaced by a pause of the same length. An example of the relevant section of an item in the two conditions follows:

(9) Unedited: but- but being very great presidents, oh and Roosevelt

Edited: but- but being very great presidents, [pause equivalent to length of _____] o

The pause was created by copying another pause from part of the speaker’s speech stream. This helped ensure that the pause was seamless with the surrounding speech. A pure silence, white noise, or indeed pauses from other speaker’s background noise would create a noticeable difference in the speech flow. Original pauses preceding or following the oh were retained in the edited stimuli. Ohs were on average 255 ms long, with a range from 147 ms to 418 ms. Pauses before oh were on average 928 ms long, with a range from 0 ms to 2994 ms. Pauses after oh were on average 94 ms long, with a range from 0 ms to 615. Although all but two stimuli had a pause preceding oh, 13 of the 20 stimuli had no pause following oh. In a follow up study we tested whether or not listeners could detect editing in our stimuli. From each of the critical trials, we recorded a 10 second segment that included either the original oh or the inserted pause. Eight participants who did not participate in any of the current experiments listened to one version of each of the 20 trials. Participants were told specifically that a number of the trials were spliced, and their job was to judge each trial as being spliced or not spliced. Participants performed no

better than chance at identifying the edited speech.

Design

Two lists were created. List 1 contained the practice stimuli, 10 unedited critical stimuli, 10 edited critical stimuli, and the 20 fillers. List 2 contained the same practice and filler stimuli, but had the edited versions of the List 1 unedited critical stimuli and the unedited versions of the List 1 edited critical stimuli. The order of presentation was the same across lists and was pseudo-randomized with the constraint that retellings of the same story were separated by at least four trials (even though the stories were retold by different speakers), that no more than three filler or critical trials occurred in succession, and that trials with the same speaker were separated by at least five trials.

Each list of digitized stimuli was recorded onto the right channels of separate digital audio tapes. Between each trial, there was a 400 ms tone, followed by 3500 ms of silence during which time the target word for the upcoming trial appeared on the screen. Participants heard only the right channels of the tapes. On the left channels of the tapes, tones were recorded to control the presentation of the target words. A program written in Superlab software received the left-channel signals from the tapes, presented the words on the screen, and recorded reaction times with millisecond accuracy.

Procedure

Each participant was randomly assigned to listen to one tape, and so heard only one version of each stimulus. Participants were tested individually in a session lasting about 40 minutes. They were given written instructions and then seated in front of a computer monitor in a sound-attenuated booth. After a 400 ms warning tone, participants had 1500 ms to focus their attention on the computer screen to read the target word. The target word was presented on the screen for 1000 ms, followed by 1000 ms of silence, after which the speech began. Participants monitored for the occurrence of the word in the speech and were instructed to press a response button as quickly as possible if they heard the target word. If the word did not occur in the trial, participants did not respond. Responses were timed out after 1500 ms. All stimuli played to the

finish regardless of whether or not a button was pressed. Following each trial, there was 1500 ms of silence, and then the warning tone occurred indicating that the target for the next trial would be presented.

Results

Here and in the experiments to follow, responses between 0 ms and 150 ms were excluded from the analyses because they were considered to have undershot the minimum time necessary to process a word and initiate a physical response; that is, responses faster than 150 ms reflected something other than target recognition. Four items were eliminated from the analysis because more than 35% of the responses were errors, with participants responding either outside the time limits or not at all. On one of these trials the target word was not clearly spoken, and it appears participants had trouble detecting it. The other three eliminated trials were among the longest in the study. In one case, the target did not occur until 90 seconds after being visually presented. When questioned, participants reported that they often forgot the target words in the long trials. After removing the outlying fast responses, responses more than two standard deviations from the mean reaction time were treated as errors. Error rates did not vary across conditions (unedited: 19.6%, 44/224; edited: 22.3%, 50/224; $t_1(1,27) = -1.15$, $p = n.s.$; $t_2(1,15) = -1.07$, $p = n.s.$). We believe the high error rate is due to forgetting of the target word. The length of time participants mentally rehearsed the visually presented target words before they heard it was on average 44 seconds over the original 20 items, and 38 seconds over the 16 retained items.

Participants were about 62 ms faster at responding to the target word when it followed an oh than when it didn't, an effect that is reliable over both participants and items ($t_1(1,27) = 2.73$, $p < .025$; $t_2(1,15) = 2.17$, $p < .05$). A summary of the means for all experiments is presented in Table 1.

There was no correlation between the differences between conditions and either how many words after the oh the target word was heard ($r = .15$, $p = ns$) or how many ms after oh the target word was heard ($r = .13$, $p = ns$).

Discussion

People were faster at detecting a word in a speech stream when that word followed an oh than when the oh had been digitally excised and replaced by a pause. So people are able to make use of at least this one discourse marker while listening to speech. These data are consistent with the predictions of the disjuncture mechanism where hearing oh reduces listeners' confusion about the following discourse. They are also consistent with the expectation mechanism where listeners recognize oh as a change-of-state marker and pay more attention to upcoming speech in anticipation of hearing some new or unexpected information.

Closer analyses of the phenomenon help rule out one alternative explanation, the trigger hypothesis. The trigger hypothesis is that hearing oh provides no specific information, but is instead a general indicator of trouble. Listeners are triggered to pay more attention when they hear an oh (as they would be with any kind of triggering sound, such as a warning tone) and therefore respond more quickly. But if listeners are triggered to pay more attention, then the effect of the trigger should decay over time as the target moves away from the trigger. That is, listeners' heightened attention after oh should show some decrement the further away from oh the target word occurs. But it turns out that the benefit of hearing oh was the same on words that were closer and further away from the oh, as tested from 1 word after oh up to 11 words after oh. Attention did not diminish as the trigger hypothesis would predict. So, oh is not merely triggering listeners to pay more attention to what immediately follows the oh. Instead, oh provides more specific information about upcoming speech. Both the disjuncture mechanism and the expectation mechanism predict that the benefit of oh should be equally strong for targets immediately after the oh and several words downstream because all targets used were either part of the phrase that could cause confusion or the phrase containing change of state information.

Related to the trigger hypothesis is the hypothesis that listeners adopted a strategy to respond after hearing an oh. This alternative is unlikely given that there were many noncritical ohs in the materials. Strategically responding after hearing oh would lead to false alarms. Also, in only 25% (10/40) of the materials presented did the target actually follow oh. Perhaps most convincingly, debriefing confirmed that no listener made a conscious connection between the

presence of oh and the likelihood of a positive response.

Another alternative explanation is that the benefit of oh really resulted from the detrimental effect of the pause in the comparison condition: it was not that word recognition was faster with an oh, but that it was slower after a pause. Of course, the ohs were relatively short in comparison to the average preceding pauses, so it is unlikely that substituting them with a pause would make any difference to the results, either positively or negatively. But to be sure of this, we repeated Experiment 1 but did not insert a pause into the oh-excised stimuli. In addition, we selected new target words and shortened some stimuli in an effort to improve performance.

Experiment 2: Word Monitoring Without Pause

Method

Participants. Thirty-three UCSC students participated in this experiment in partial fulfillment of a course requirement. None of them had participated in Experiment 1. Three participants were removed because they failed to respond to at least 25% of the trials.

Materials. The same materials were used as in Experiment 1, except that we removed the oh and did not replace it with a pause. Any pause that occurred before or after the oh remained. Additionally, several long trials were shortened to reduce the likelihood of participants' forgetting the target words while listening to the trials and more prominent target words were selected. An example of the relevant section of an item in the two conditions follows:

(10) Unedited: but- but being very great presidents, oh and Roosevelt

Edited: but- but being very great presidents, and Roosevelt

Design and Procedure

The design and procedure is the same as in Experiment 1.

Results

Two items were eliminated from the analyses because more than 35% of the responses were errors. Both items were among the four removed in Experiment 1. The speaker had a tendency to slur words, so it's possible that the new target words selected for these items were still difficult to detect. In addition, one of these items was a long item that could not be

shortened, so participants may simply not have been able to keep any target word in memory. Participants did make fewer errors over the other items. Once again, error rates did not vary across conditions (unedited: 14.8%, 40/270; edited: 19.6%, 53/270; $t_1(1,29) = 1.63$, $p = n.s.$; $t_2(1,17) = 1.47$, $p = n.s.$).

Participants were about 79 ms faster at responding to the target word when it followed an oh than when it didn't, an effect that is reliable over both participants and items ($t_1(1,29) = 4.26$, $p < .01$; $t_2(1,17) = 4.51$, $p < .01$).

There was no correlation between the differences between conditions and either how many words after the oh the target word was heard ($r = .23$, $p = ns$) or how many ms after oh the target word was heard ($r = .20$, $p = ns$).

Discussion

As expected, people are still faster at detecting a word in a speech stream when the word followed an oh than when it didn't even when the excised oh is not replaced by a pause. The seeming benefit of oh is a real one, and not a chimeric result of a detrimental effect of a pause in the comparison condition. Oh does seem to help listeners comprehend the discourse, either by helping them integrate information or by signalling that they should expect an upcoming change of state. As in Experiment 1, there was no relationship between the speed of target recognition and the distance the target was from the splice, as would be expected by the trigger hypothesis. Oh is not merely priming listeners to pay more attention.

As a second test of oh's contribution to discourse comprehension, we tested how active prior concepts were at a point in the speech stream following oh, and at the same point with the oh excised and replaced by a pause of the same length. If oh helps listeners integrate upcoming information into the discourse, either by anticipating a change of state or by reducing the confusing effects of a disjuncture, listeners should be faster at identifying concepts that occurred prior to the information marked by oh.

Experiment 3: Semantic Verification With Pause

Method

Participants. Thirty-nine UCSC students participated in this experiment in partial fulfillment of a course requirement. None of them had participated in either Experiments 1 or 2. Five participants were removed because they failed to respond to at least 25% of the trials.

Materials. The same materials were used as in Experiment 1.

Design. The design is the same as in Experiments 1 and 2.

Procedure. All aspects of the procedure are the same as in Experiments 1 and 2 except that participants did not rehearse a visually presented word in memory while listening for that word in the speech stream, but watched the screen as they listened until a word appeared. Because the target word no longer needed to be presented in advance, there was only 1000 ms after the warning tone instead of 3500 ms. When they saw the word appear on the screen, participants immediately pressed a button if they had heard the word they saw in the prior discourse. If the word did not occur in the trial, participants did not respond. The word remained on the screen for 2500 ms. As before, all stimuli played to the finish regardless of whether or not the button was pressed.

Target words to be verified were selected by picking a content word from the section of the discourse that the information after the oh could refer back to. This judgement was made by one of the authors and two research assistants. Content words were chosen because they were thought to present the least problem with forgetting. Any one of a number of words could have been chosen for any particular stimulus. The content words that were chosen represented a variety of form classes including nouns, verbs, adjectives, and adverbs. Information related to the discourse following oh was thought to provide the strongest opportunity for an effect to be found.

Semantic verification was tested several words after the oh, at the point where the to-be-verified information should be active. In the following example, the asterix indicates when the italicized target word was presented on the computer screen for verification:

(11) and then he cited four presidents, um Jefferson Eisenhower Kennedy, and uh someone else for um . . . for being *adulterers*, but- but being very great presidents, oh and

Roosevelt *

In the critical trials, the word to be verified was always present in the prior discourse. Target words preceded the oh by an average of 12 words, with a range of 2 to 27. They were displayed on the screen an average of 10 words after the oh, with a range of 2 to 16. The filler trials were the same as those used in the previous experiments. Filler targets were created by identifying words that occurred in the filler passages about where the critical targets occurred in the critical passages and selecting closely related words that did not occur in the passages. Eight of the filler trials contained an oh followed by a negative probe. Three of the critical trials contained an oh in an early section of the passage. Targets were presented at the end of the passages, after the completion of the utterances marked by the critical ohs in the critical trials or the last phrases in the filler trials. The completion of an utterance was taken to be the first point where speech could be terminated without truncating an idea. Because semantic verification requires more thinking time than word monitoring, responses were timed out a second later, after 2500 ms.

Results

Eight items were eliminated from the analyses because more than 35% of the responses were errors. Error rates did not vary across conditions (unedited: 17.6%, 36/204; edited: 20.5%, 42/204; $t_1(1,33) = -.48$, $p = n.s.$; $t_2(1,11) = -.53$, $p = n.s.$).

Participants were about 97 ms faster at verifying the presence of the target word when an oh preceded the test point than when it didn't, an effect that is reliable over both participants and items ($t_1(1,33) = 2.45$, $p < .025$; $t_2(1, 11) = 2.63$, $p < .025$).

Discussion

When tested at the end of the repair, concepts from the prior discourse were more active when an oh was heard previously than when the oh was replaced by a pause. These data are consistent with the hypothesis that hearing oh prompts listeners to halt normal integration processes and to treat the utterance after the oh independently. By not trying to integrate the utterance after the oh with the immediately prior utterance, the utterance after the oh can be processed more quickly, and can consequently be connected with the reparandum more

effectively, causing information from the reparandum to be more active after an oh is heard. The more information after the oh is heard, the more information listeners have to connect repairs with their reparandums. From another angle, the more information after the splice is heard, the more confused listeners might become from trying to integrate two nonintegratable utterances. The data are also consistent with the hypothesis that oh signals a change of state, prompting listeners to pay attention to how upcoming information connects with prior information, causing a heightened awareness of prior information at the moment it becomes relevant.

Experiment 4 addresses an alternative explanation encountered earlier, that the benefit of oh really resulted from the detrimental effect of a pause in the comparison condition: it is not that hearing an oh is beneficial, but that hearing a pause in the comparison condition disrupts processing, causing confusion and slower semantic verification. In Experiment 4 we also tried to reduce the number of errors caused by participants' responding incorrectly or after the timeout. One cause of participant error might have been the length of the items. Recalling a word that occurred 20 or more words back might have been too difficult. Verification might also have been hindered by selection of target words that were not prominent in the discourse and easily forgotten. In addition, if listeners did not watch the screen while they were listening, they would not see the target word that they had to verify. Listeners' attention might be more likely to wander when the stimulus is long, causing the visual target word to be missed, and making any response impossible. We adjusted the items by making them shorter and by selecting more noteworthy target words.

Experiment 4: Semantic Verification Without Pause

Method

Participants. Thirty-six UCSC students participated in this experiment in partial fulfillment of a course requirement. None of them had participated in either Experiments 1, 2, or 3. Four participants were removed because they failed to respond to at least 25% of the trials.

Materials. The same materials were used as in Experiment 3, except that the pause was removed from the edited critical stimuli. Additionally, several of the trials were shortened and

more prominent target words were selected.

Design and Procedure. The design and procedure is the same as in Experiment 3.

Results

Six items were eliminated from the analyses because more than 35% of the responses were errors. These six were a subset of the eight eliminated in Experiment 3. Error rates did not vary across conditions (unedited: 14.7%, 33/224; edited: 16.5%, 37/224; $t_1(1,31) = -.77$, $p = n.s.$; $t_2(1,14) = -.73$, $p = n.s.$).

The six items that were eliminated were evenly distributed throughout the experiment. The target words varied in form class in the same way as the unmissed items. They also had similar syntactic environments: both missed and unmissed targets sometimes occurred in false starts, had restarts between the auditory target and oh, and had ohs that fell both between and within sentences. Furthermore, the eliminated items had low responses in both conditions. It was not the case that participants missed targets in the eliminated items only when the oh was present or absent. Although we cannot explain why these stimuli had such low responses, it is important to bear in mind that they do not show an opposite result to the one found with the remaining 14 stimuli: they show no result whatsoever.

Participants were about 58 ms faster at verifying the presence of the target word when an oh preceded the test point than when it didn't, an effect that is reliable over both participants and items ($t_1(1,31) = 2.52$, $p < .025$; $t_2(1,13) = 2.73$, $p < .025$).

Discussion

Once again, concepts from the prior discourse were more active when they were tested after the utterance marked by oh than when they were tested at the same point in the speech stream but with the preceding oh excised. So it is not the case that the pause in the comparison condition in Experiment 3 was slowing semantic verification. It doesn't matter if there is a pause or not, oh will have the same effect.

With repair ohs, information after the oh always refers back to some information in the prior discourse. By hypothesis, oh explicitly informs listeners to halt normal cross-sentential

integration and prepare for a disjuncture or anticipate a change of state. Both processes can improve integration of prior information either by preventing confusion that might otherwise arise from two adjacent disjunct utterances or by priming listeners that they may need to work harder to integrate the upcoming change of state information. But both processes' effects on the availability of prior information should only be detected downstream from the point where oh is heard. With the disjuncture mechanism, it is at a point downstream that the absence of oh creates enough confusion to interfere with activation of prior concepts. With the expectation mechanism, it is at a point downstream that the connection between the change of state and the prior discourse is made, and it is this connection that enhances the availability of prior information.

In contrast to the two mechanisms proposed, it is possible that oh does have an immediate effect on the availability of prior information. Oh might serve as a cue to listeners to search back in the discourse record for information that might need updating. If this is the case, then immediately after hearing an oh, all prior information should be more active than when the oh is not heard. Immediate cuing would predict the same results obtained in Experiments 3 and 4. In Experiment 5, we ran the semantic verification experiment with the test points shifted to immediately after the ohs. If immediate cuing can account for the results, then Experiment 5 should produce the same results as Experiments 3 and 4. But if immediate cuing cannot account for the results, then Experiment 5 should produce different results.

In Experiment 5 we also modified the procedure slightly to further investigate the trigger hypothesis: we had participants explicitly press a no key when the word to be verified did not occur in the prior discourse. In Experiments 3 and 4, we predicted that people would be faster to respond yes to a word from the prior discourse after it follows an oh. But this result could be due to listeners' heightened propensity to respond after hearing an oh trigger. However, we would predict no difference in peoples' speed at responding no to a word not in the prior discourse, whereas the trigger hypothesis would predict people would be faster to respond no after hearing an oh trigger.

Experiment 5: Semantic Verification Immediately after Oh

Method

Participants. Thirty-eight UCSC students participated in this experiment in partial fulfillment of a course requirement. None of them had participated in any of the previous experiments.

Materials. Materials consisted of the 40 critical stimuli from Experiment 4 (20 unedited stimuli and 20 corresponding edited versions) and 40 new critical stimuli from the same corpus (20 unedited stimuli and 20 corresponding edited versions). Targets were selected using the same criteria as in Experiments 3 and 4.

Design. Two counterbalanced lists were created, each containing 10 trials in which oh occurred and the target word occurred, 10 trials in which oh occurred and the target word did not occur, 10 trials in which oh was removed and the target word occurred, and 10 trials in which oh was removed and the target word did not occur. The old and new critical stimuli were distributed evenly across conditions.

Procedure. The procedure is the same as in Experiments 3 and 4 except that participants pressed a yes key if the word had occurred in the prior discourse and a no key if it did not, and that semantic verification was tested immediately after the offset of the oh.

Results

Participants were equally fast at verifying the presence of the target word when an oh preceded the test point as when it did not ($t_1(1,37) = .22, p = n.s.$; $t_2(1,19) = .60, p = n.s.$). They were also equally fast at verifying the absence of the target word when an oh preceded the test point as when it did not ($t_1(1,37) = .12, p = n.s.$; $t_2(1,18) = .57, p = n.s.$; one item was eliminated from the analyses because more than 35% of the responses were errors). A summary of the means is presented in Table 2. Error rates did not vary across conditions (target present unedited: 15.8%, 60/380; target present edited: 12.5%, 45/380; $t_1(1,37) = 1.39, p = n.s.$; $t_2(1,19) = 1.21, p = n.s.$; target absent unedited: 18.2%, 66/361; target absent edited: 14.1%, 51/361; $t_1(1,37) = 1.9, p = n.s.$; $t_2(1,18) = .57, p = n.s.$).

Discussion

Listeners do not have an advantage immediately after hearing oh. They have not heard enough information to either become confused by a following disjoint utterance, or said another way, to gain an integration advantage by using the oh as a signal to halt normal integration and treat the information after the oh independently, thereby processing that information faster and reconnecting it to the reparandum faster. Nor have they heard enough information for the expectation mechanism to demonstrate its advantage. With the expectation mechanism, it's only after figuring out how the change of state connects with prior information that the prior information gains heightened activity. The benefit of hearing oh cannot be detected if the test point is right after the oh.

The lack of an effect for both the no responses and the yes responses lends further evidence to the correlations discussed in Experiments 1 and 2 that the trigger hypothesis cannot be an alternative explanation for the effects found in the main experiments.

General Discussion

Thoughts are often unstructured. One idea can shift to another without any inherent order, and sometimes without even any obvious connection. Sometimes thoughts are organized hierarchically, such as when a person has one main goal that has several subgoals, or when a person is describing the rooms of a house and uses the main hall as a point of departure before entering each room and describing it. When people speak spontaneously, they face the problem of turning these hierarchically organized or even unorganized thoughts into a linear stream of speech. This linear speech has to somehow capture the relative organization of the thoughts. In addition, speaking spontaneously often precludes carefully planning speech in advance, so some ideas will be spoken out of turn, or will need to be corrected or elaborated. Investigations of when discourse markers are used in conversation suggest that they might provide information about the overall structure of the discourse, about how one utterance is related to a subsequent utterance, about how a repair is related to a reparandum, or other pragmatic or social processes. This information might help listeners create coherency out of linearly presented complex ideas, or even incoherently presented ideas.

Flowerdew & Tauroza (1995) provide some evidence for the general usefulness of discourse markers in the processing of spontaneous speech. They found that a spoken lecture with discourse markers present is comprehended better than the same lecture with the discourse markers edited out. Although the most frequent discourse markers in this study did not include oh (they were all right, and, because, but, now, OK, right, so, then, and well; Flowerdew & Tauroza, p. 438), one of the reasons suggested for discourse markers' general usefulness does apply to oh. That is that the explicit marking of discontinuity helps listeners avoid the confusion that would arise if they tried to connect two disjointed utterances (Flowerdew & Tauroza, 1995).

In a series of five experiments, we demonstrated that oh can help overcome some of the inevitable pitfalls of spontaneous speech. Unlike what some theories might predict, ohs are neither detrimental nor irrelevant to comprehension. When listeners hear an oh, they can recognize a word in the upcoming speech faster than when they do not hear an oh. They can also verify that a word had occurred in the prior discourse faster after hearing an oh, as long as the test point is far enough after the oh for the benefit of the oh to manifest itself. Oh's posited underlying change of state function can help explain these results. Upon hearing an oh, listeners know that they should focus attention on the upcoming utterance in anticipation of a change of state. They also know that the information immediately preceding the oh might be disjoint from the information following the oh, and so they should halt normal integrative processes and treat the upcoming information independently. This causes the upcoming utterance to be processed faster, and thereby integrated with the prior information faster. The oh helps avoid the confusion that would arise had the speaker merely stopped one train of thought and begun another without a marker.

Analogies can be drawn between the effects of oh in spontaneous speech and the effects of other discourse markers in text comprehension. Listeners notice and use oh in making sense of speech and readers notice and use other discourse markers in making sense of texts (Bestgen & Vonk, 1995; Millis & Just, 1994; Segal, Duchan, & Scott, 1991). Using off-line tasks, Segal, Duchan, and Scott (1991) found that readers interpret then, and, so, because, and but as

indicators of either continuity or discontinuity. They also found that readers interpret these markers as doing more than relating adjoining sentences; they interpret them as relating sentences to the discourse as a whole. The usefulness of written markers for expressing discontinuity and for expressing discourse relationships beyond the intersentential level is similar to our findings for oh.

Millis and Just (1994) investigated readers' on-line uses of because and although. They found that the use of one of these words between adjacent sentences helps readers integrate information better across sentences; people were faster at verifying that a word had been read in the prior text when a marker was present than when it wasn't. Similar to our results, Millis and Just did not find a beneficial effect when semantic verification was tested right after the marker (we tested immediately after the offset of the marker, they tested one word after the marker). They found an effect only when the test point was at the end of the sentence after the marker, the point at which information from the prior sentence is reactivated to allow integration, with the reactivation being prompted by the presence of marker. In addition to missing the reactivation prompt when markers are absent, listeners lose explicit information about how to relate adjacent utterances. They are then forced to rely on more subtle information, and this takes processing time. Part of the reason because and although may have the effects Millis and Just found is because they explicitly mark a causal relationship between adjacent utterances. This contrasts with the hypothesized function of oh, which is to explicitly mark a non-relationship between adjacent utterances. So although the results are similar for auditory oh and visual because and although, the reasons for the results may be different.

Bestgen and Vonk (1995) looked explicitly at markers that anticipate discontinuity in text. Using a semantic verification task, they found that then and explicit references to time like around 2 o'clock can inhibit the accessing of preceding information. Bestgen and Vonk argue that these markers focus attention on following information, thereby closing off processing of prior information and making prior information less accessible. These effects of the disjunct markers then and explicit references to time contrasts with the results we found here with the

disjunct marker oh, which does not make prior information less accessible. If it did make prior information less accessible, semantic verification should have been slower after an oh than at the same point with the oh excised, both when tested immediately after the oh, as in Experiment 5, and when tested several words downstream, as in Experiments 3 and 4. The difference between our results and Bestgen and Vonk's could arise because of a difference between the markers studied (then and explicit references to time versus oh), the presentation form (visual versus auditory), or the nature of the discourse (spontaneous versus prepared). Further study is needed to identify which of these factors contributed to the variation in effects.

Although analogies can be drawn between the effects of discourse markers in speech and the effects in text, it might be misleading to draw too close parallels. With spontaneous talk, speakers organize ideas on the fly, correct errors publicly, and monitor interlocutor's responses. These activities are absent in prepared written discourse. Discourse markers in speech and writing might be used to solve different problems. Not surprisingly, the markers found in written and oral communication differ (Flowerdew & Tauroza, 1995; Stenström, 1990) and some have observed that their use also differs (Flowerdew & Tauroza, 1995). The spoken ohs we investigated arise in a way that is unnecessary in text, where ample revision time avoids the problems that cause oh to be used in the first place. So were an oh to be used in prepared text, it is likely to be used for some purpose that differs substantially from its use in spontaneous speech.

In spontaneous speech, oh helps listeners integrate discourse. It could serve this function in two ways. One is that oh indicates a disjuncture between the immediately preceding and immediately following utterances. This informs listeners to halt integrative processes across adjacent utterances and to process upcoming information independently from immediately preceding information. A second way oh might help listeners integrate discourse is by priming listeners to expect a change of state in the upcoming speech. This expectation might benefit listeners' ability to integrate upcoming information. The average oh was only a quarter of a second

long, but it played an important role in aiding listeners' on-line processing of talk. A little oh can make quite a difference.

References

- Aijmer, K. (1987). Oh and ah in English conversation. In Meijs (Ed.), Corpus Linguistics and Beyond. Amsterdam: Rodopi.
- Bestgen, Y., & Vonk, W. (1995). The role of temporal segmentation markers in discourse processing. Discourse Processes, 19(3), 385-406.
- Brennan, S. E., & Williams, W. (1995). The feeling of another's knowing: Prosody and filled pauses as cues to listeners about the metacognitive states of speakers. Journal of Memory and Language, 34, 383-398.
- Clark, H. H. (1996) Using Language. Cambridge University Press; Cambridge, England.
- Clark, H. H. (1997). Dogmas of understanding. In press Discourse Processes.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), Perspectives on Socially Shared Cognition. Washington, D. C.: American Psychological Association.
- Clark, H. H., & Haviland, S. E. (1974). Psychological processes as linguistic explanation. In D. Cohen (Ed.), Explaining Linguistic Phenomena. Washington, D. C.: Hemisphere Publishing Corporation.
- Erman, B. (1986). Some discourse markers in English conversation. In G. Tottie & I. Bäcklund (Eds.), English in Speech and Writing.
- Erman, B. (1987). Pragmatic Expressions in English: A Study of You Know, You See and I Mean in Face-to-Face Conversation. Stockholm: Almqvist & Wiksell.
- Flowerdew, J., & Tauroza, S. (1995). The effect of discourse markers on second language lecture comprehension. Studies in Second Language Acquisition, 17, 455-458.
- Fox Tree, J. E. (1995). The effects of false starts and repetitions on the processing of subsequent words in spontaneous speech. Journal of Memory and Language, 34, 709-738.
- Fox Tree, J. E., & Clark, H. H. (1997). Pronouncing "The" as "Thee" to Signal Problems in Speaking. Cognition, 62, 151-167.
- Gernsbacher, M. A., & Jescheniak, J. D. (1995). Cataphoric devices in spoken discourse.

Cognitive Psychology, 29, 24-58.

Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process in comprehension. Journal of Verbal Learning and Verbal Behavior, 13, 512-521.

Heritage, J. (1984). A change-of-state token and aspects of its sequential placement. In J. M. Atkinson & J. Heritage (Eds.), Structures of Social Action: Studies in Conversation Analysis. Cambridge: Cambridge University Press.

James, D. (1972). Some aspects of the syntax and semantics of interjections. Papers from the Eighth Regional Meeting. Chicago Linguistics Society.

Jucker, A. H. (1993). The discourse marker 'well': A relevance theoretical account. Journal of Pragmatics, 19, 435-452.

Levelt, W. J. M. (1983). Monitoring and self-repair in speech. Cognition, 14, 41-104.

Levelt, W. J. M. (1989). Speaking: From Intention to Articulation. Cambridge, MA: The MIT Press.

Marslen-Wilson, W., & Tyler, L. K. (1980). The temporal structure of spoken language understanding. Cognition, 8, 1-71.

Millis, K. K., & Just, M. A. (1994). The influence of connectives on sentence comprehension. Journal of Memory and Language, 33(1), 128-147.

O'Donnell, W. R., & Todd, L. (1991). Variety in Contemporary English (2nd ed.). New York, NY: Harper Collins Academic.

Redeker, G. (1990). Ideational and pragmatic markers of discourse structure. Journal of Pragmatics, 14, 367-381.

Redeker, G. (1991). Linguistic markers of discourse structure. Linguistics, 29, 1139-1172.

Schiffrin, D. (1987). Discourse Markers. Cambridge: Cambridge University Press.

Schourup, L. C. (1985). Common discourse particles in English conversation. New York: Garland.

Segal, E. M., Duchan, J. F., & Scott, P. J. (1991). The role of interclausal connectives in narrative structuring: Evidence from adults' interpretations of simple stories. Discourse Processes, 14, 27-54.

Stenström, A.-B. (1990). Lexical items peculiar to spoken discourse In J. Svartvik (Ed.), The London-Lund Corpus of Spoken English: Description and Research. Lund, Sweden: Lund University Press.

Wade, E., & Clark, H. H. (1993). Reproduction and demonstration in quotations. Journal of Memory and Language, 32, 805-819.

Author Note

This research was supported by faculty research funds granted by the University of California, Santa Cruz, a Social Sciences Divisional Research Award granted by the University of California, Santa Cruz, and a Regents Affirmative Action Career Development Award granted by the University of California. We thank Herbert Clark for sharing his corpus and tapes with us and Paul Meijer for insightful discussions. We also thank the many students who helped with preparing and running the experiments, with a special thanks to Jennifer Black, Duane Cleghorn, Duane Cobb, Ben Nisenbaum, and Taylor Norrish. Finally, we thank three anonymous reviewers for their thoughtful comments. Transcripts of the materials can be obtained by writing to either of the authors. Correspondence concerning this paper should be addressed to Jean E. Fox Tree or Josef C. Schrock, Psychology Department, Social Sciences II, University of California, Santa Cruz, CA 95064.

Table 1

Summary Statistics for Experiments 1 through 4 with Mean (SD) in Milliseconds				
Experiment		With Oh	Without Oh	Difference
1	Word Monitoring With Pause	472 (239)	534 (243)	62
2	Word Monitoring Without Pause	486 (210)	565 (251)	79
3	Semantic Verification With Pause	1136 (370)	1233 (395)	97
4	Semantic Verification Without Pause	987 (305)	1045 (335)	58

Table 2

Summary Statistics for Experiment 5 with Mean (SD) in Milliseconds				
Condition		Reaction Time		
1	Oh Present: Target Occurred	1174 (208)		
2	Oh Removed: Target Occurred	1179 (212)		
3	Oh Present: Target Did Not Occur	1325 (216)		
4	Oh Removed: Target Did Not Occur	1328 (207)		