Published in final edited form as:

J Exp Psychol Hum Percept Perform. 2003 April ; 29(2): 333–342. doi:10.1037/0096-1523.29.2.333.

Change Deafness: The Inability to Detect Changes Between Two Voices

Michael S. Vitevitch

Michael S. Vitevitch. Department of Psychology. Indiana University Bloomington, and Spoken Language Laboratory, Department of Psychology. University of Kansas.

Abstract

A shadowing task was used to demonstrate an auditory analogue of change blindness (the failure to detect a change in a visual scene), namely *change deafness*. Participants repeated words varying in lexical difficulty. Halfway through the word list, either the same or a different talker presented the words to participants. At least 40% of the participants failed to detect the change in talker. More interesting is that differences in shadowing times were found as a function of change detection. Alternative possibilities to the change detection phenomenon were ruled out. The results of these experiments suggest that the allocation of attention may influence the detection of changes as well as the processing of spoken words in complex ways.

Change blindness is a counterintuitive phenomenon (Levin, Momen, Drivdahl, & Simons, 2000; Simons & Levin, 1998) in which observers in a variety of paradigms (e.g., Henderson & Hollingworth, 1999; Levin & Simons, 1997; Rensink, O'Regan, & Clark, 1997) fail to detect what may be described as obvious changes in a visual scene. For example, Grimes (1996) found that participants noticed only 30% of the changes in photographs that occurred during an eye movement—even changes as obvious as two heads switching bodies. Simons and Levin (1998) dramatically demonstrated that only 33% of the participants in a real-life interaction noticed that the person asking them for directions was exchanged when a door being carried by confederates momentarily interrupted the discussion. Rensink et al. (1997) hypothesized that the inability to detect changes in a visual scene was related to the allocation of attention during visual processing; attention must be focused on an object for a change in that object to be detected.

Rensink et al. (1997; see also Klein, Kingstone, & Pontefract, 1992; Posner, 1980) further suggested that attention may be “pushed” by high-level influences (i.e., volitional control) or “pulled” by low-level changes toward relevant objects or features in a visual scene, thereby facilitating the detection of a change in that scene. Indeed, Werner and Thies (2000) found that the high-level knowledge of experts in American football facilitated their detection of changes in football images compared with participants with less expertise in American football, suggesting that domain-specific expertise may push attention toward relevant objects or features in a visual scene (see also Hollingworth & Henderson, 2000). Alternatively, low-level changes such as those that produce transient motion may pull attention to that area in the scene, unless a mask of some sort is used to eradicate this cue (e.g., the flicker technique of Rensink et al., 1997).

The work of Cherry (1953) provides some evidence from the auditory domain to support the hypothesis of Rensink et al. (1997). Among the studies described in Cherry (1953) is an experiment in which participants had to shadow, or repeat out loud, passages of continuous speech presented to them over one channel of a set of headphones; the concurrent message in the other channel of the headphones was to be ignored. Cherry found that listeners were quite accurate in their repetition of the attended passages. Changes to the message presented in the unattended ear—such as switching from English to German or switching from one male speaker to another—went relatively undetected, as predicted by the hypothesis of Rensink et al. (1997). That is, participants appeared to be deaf to changes in the environment.

Although the hypothesis regarding the role of attention in detecting changes in the environment is supported by (the rather vague descriptions provided in) the work of Cherry (1953), the hypothesis appears to suffer from circular reasoning. A change in a scene is not detected because attention is not directed at that object; the proof that attention is not directed at the object is the failure to detect a change. To break the circularity, one must use alternative means to assess the allocation of attention and the detection of a change in the environment. The use of an integral stimulus, such as spoken language, affords researchers the unique opportunity to examine the hypothesis that attention must be allocated to the relevant dimension to detect a change in a complex stimulus (Rensink et al., 1997) without circularity encroaching into the argument.

An integral stimulus comprises two stimulus dimensions. In the case of spoken language, there is a linguistic and an indexical dimension. The *linguistic* dimension of the speech signal conveys propositional information about objects in the world. For example, the linguistic message of a spoken utterance may be a request to close a window. *Indexical* information refers to acoustic correlates in the speech signal that provide information on various characteristics of the talker, including identity, emotional state, age, dialect, and gender (Pisoni, 1997). If one hears a request to close the window, the indexical information conveyed concurrently with that linguistic message would allow him or her to identify whether the speaker was a shivering old woman that was unknown to him or her, or a more familiar interlocutor, like “Uncle Joe from western Pennsylvania.”

Although different aspects of the acoustic signal are correlated with linguistic (Zue & Schwartz, 1980) and indexical attributes (Bricker & Pruzansky, 1976; Hecker, 1971), evidence from several studies using speeded classification tasks (Garner, 1974) suggests that spoken language is an integral stimulus with these two dimensions (Mullennix & Pisoni, 1990). For example, Jerger and colleagues (Jerger, Martin, Pearson, & Dinh, 1995; Jerger, Pirozzolo, et al., 1993; Jerger, Stout, et al., 1993) had participants selectively attend to indexical information, in this case the gender of the talker, while ignoring the linguistic dimension, in this case the word being spoken, and vice versa. In both cases, the classification performance of listeners (with normal hearing) for the relevant dimension was affected by variation in the irrelevant dimension, suggesting that spoken language is an integral stimulus.

If one now considers the interference across dimensions that occurs in an integral stimulus (such as spoken language) and the attention allocation hypothesis of Rensink et al. (1997), a unique opportunity to more precisely evaluate the role of attention in the detection of changes can be seen. That is, if attention is allocated to dimension *X* in an integral stimulus, then a change in dimension *X* should be detected. More important, if attention is allocated to dimension *X*, then a processing cost should be observed in dimension *Y*, thereby providing a means of evaluating the allocation of attention to dimension *X* other than the subjective detection of the change.

The results from several studies investigating talker variability in spoken word recognition suggest that this method of independently evaluating the allocation of attention is valid (e.g., Goldinger, 1998; Mullennix, Pisoni, & Martin, 1989; Nygaard, Sommers, & Pisoni, 1994). In such experiments, a change on each trial was made in the indexical dimension (by using a different voice) while processing of the linguistic information (response times to words differing in various lexical characteristics) was assessed. Changing the voice from trial to trial resulted in less accurate identification of words presented in noise (Mullennix et al., 1989; Nygaard et al., 1994; for effects on recall and recognition, see also Goldinger, Pisoni, & Logan, 1991; Martin, Mullennix, Pisoni, & Summers, 1989; Palmeri, Goldinger, & Pisoni, 1993). Thus, changes in the indexical dimension adversely affected processing of the linguistic dimension. Although the participants in the talker-variability studies were never explicitly asked whether they detected the change in the speakers, it is most likely that the continuous low-level changes in the stimuli (such as the fundamental frequency of the different voices presented on each trial) pulled attention toward the indexical dimension of the signal to the detriment of the linguistic dimension.

To better examine the attention-directed detection hypothesis of Rensink et al. (1997), I used the integral stimulus of spoken language in a situation that minimized the pull of attention to the indexical dimensions that occurred in the talker-variability studies. This was accomplished by changing the voice of the talker only once—halfway through the list of words—rather than on every trial, as in the talker-variability studies. Furthermore, in contrast to the continuous, meaningful passages used by Cherry (1953), isolated words were presented to participants to shadow as quickly and as accurately as possible. The connected speech passages that Cherry used as stimuli may have engaged higher level processes that pushed attention toward the meaning of the passages to the detriment of other characteristics of the auditory input, such as the gender of the talker or the language of the message in the other channel.

If attention is directed toward the indexical dimension of the spoken stimulus, then the change in voices should be detected. However, processing of the linguistic dimension should suffer, as demonstrated in the studies of talker variability (e.g., Mullennix et al., 1989; Nygaard et al., 1994). In contrast, if attention is directed toward the linguistic dimension of the spoken stimulus, then a detriment should be observed in the processing of the indexical dimension. Namely, the change in voices should not be detected. Also, the detriment to processing in the linguistic dimension should be attenuated in this case. Listeners in the present experiment, like observers in change blindness studies, were directly questioned at the end of the experiment as to whether they detected the change (in voice) in the experiment. In this way, I could assess subjectively as well as objectively how attention was allocated when listeners were deaf to the change in voices.

To assess how attention to the indexical dimension affected processing of the linguistic dimension, I designed the experiments such that the words to be repeated varied in lexical difficulty (with regard to ease in recognizing the words). Lexically *easy* words have a high word frequency, and few words sound similar to them. In addition, the similar-sounding words have a low frequency of occurrence. In contrast, lexically *hard* words have a low word frequency, many words sound similar to them, and the similar-sounding words have a high frequency of occurrence (Torretta, 1995). Lexically easy words are generally recognized more quickly and more accurately than lexically hard words (e.g., Kirk, Pisoni, & Miyamoto, 1997; Luce & Pisoni, 1998; Sommers, 1996). By examining changes to the overall and differential response times to easy and hard words, I could better determine how attention to the indexical dimension affected processing of the linguistic dimension. Measuring latencies to the easy and hard words is also an important extension over the pioneering work of Cherry (1953), in which only accuracy rate (not latency) was measured in the shadowing task performed by participants. (Note that there were no actual data reported in Cherry, 1953,

making it difficult to determine exactly what was found.) By measuring response latencies to the linguistic dimension of an integral stimulus, subtle processing difficulties that participants have (which may have escaped detection in Cherry, 1953) may be observed.

Experiment 1

Method

Participants—A group of 24 native speakers of English who reported no history of hearing or speech disorders participated in the experiment for partial fulfillment of an introductory psychology research requirement.

Stimuli—One hundred words with a familiarity rating of 6 or higher on a 7-point scale (Nusbaum, Fisoni, & Davis, 1984) were selected from the Indiana Multi-Talker Speech Database (Totretta, 1995). The same 100 words were selected from two different male talkers in the database (Talkers M0 and M9). Fifty words were lexically easy, and 50 words were lexically hard. The three variables of word frequency, number of similar sounding words (neighbors), and frequency of the neighbors were statistically different in this subset of stimuli: for mean word frequency (based on word counts from Kuèera & Francis, 1967), $F(1, 96) = 41.75, p < .01$; for the number of neighbors, $F(1, 96) = 238.71, p < .01$; for the mean frequency of the neighbors, $F(1, 96) = 58.42, p < .01$. Although the easy and hard words were significantly different from each other on the relevant lexical characteristics overall, the easy words in the first half of the experiment did not differ on the relevant lexical characteristics from the easy words in the second half of the experiment ($F_s < 1$). The same was true for the lexically hard words ($F_s < 1$). The means for each of these variables overall and in the first and second halves of the experiment can be found in Table 1.

Procedure—Participants were tested one at a time on a Macintosh Quadra 950 running PsyScope 1.2.2 (Cohen, MacWhinney, Flatt, & Provost, 1993), which controlled stimulus randomization and presentation and collection of response latencies. A headphone-mounted microphone (Beyerdynamic DT109) was interfaced with a PsyScope button box that acted as a voice key. A typical trial proceeded as follows: A stimulus word was presented over the headphones at 70 dB SPL to a participant who had been instructed to repeat the word as quickly and as accurately as possible. Response latency, measured from the beginning of the stimulus, was triggered by the onset of the participant's verbal response. Another trial began 1 s after a response was made. Responses were recorded on audiotape for later accuracy analyses.

Each participant received a total of 100 trials. In the first half of the experiment, 25 easy words and 25 hard words were presented. In the second half of the experiment, the remaining easy and hard words were presented. Each participant was presented with the same words in each half of the experiment but in a different random order. Halfway through the experiment, the phrase *TAKE A ONE MINUTE REST BREAK* appeared on the computer screen and remained until 1 min had elapsed. Participants remained in front of the computer during the rest break. When the experiment resumed, one half of the participants heard the same talker present the rest of the stimuli, whereas the other half of the participants heard the other talker present the stimuli. The order of presentation for the talkers was counterbalanced in the cases in which listeners heard two different voices.

When each participant finished the auditory shadowing task, they were asked three questions in the following order: (a) *Did you notice anything unusual about the experiment?* (b) *Was the first half of the experiment the same as the second half of the experiment?* (c) *Was the voice in the first half of the experiment the same voice that said the words in the second half of the experiment?* These questions were adapted from the naturalistic change blindness experiment by Simons and Levin (1998). I recorded the responses to each question.

Results

Detection of change in the indexical dimension—Of the 12 participants who heard the same voice in both halves of the experiment, all responded “yes” to Question 3, indicating that they had indeed heard the same voice in both halves of the experiment. Of the 12 participants who heard different voices in both halves of the experiment, 7 noticed the change in the talker either by stating that the voice was different in response to Question 1 or 2 or by answering “no” to Question 3. The remaining 5 participants (42%) did not state that the talker changed when asked Questions 1 and 2 and answered “yes” in response to Question 3, indicating that they failed to detect the change in the talker.

Concomitant influence on the linguistic dimension—Mixed 2 (lexical difficulty) $\times$ 3 (talker condition) analyses of variance (ANOVAs) were used to examine the response latencies of the correctly repeated words in the experiment. Lexical difficulty, a within-participants factor, refers to the easy-hard manipulation among the words. Talker condition, a between-participants factor, was determined by whether a change in talker was presented and whether that change was detected. *SAME* refers to the group of participants that heard the same talker throughout the entire list. *DIFFERENT–YES* refers to the 7 listeners who heard two different talkers and who detected that they received two different talkers. *DIFFERENT–NO* refers to the 5 listeners who heard two different talkers but failed to detect the change in talkers. Separate analyses were conducted for the words in the first block (i.e., before the break) and for the words in the second block (i.e., after the break). The means for each condition before and after the break are displayed in Figure 1.

Before the break—A main effect of lexical difficulty was found, $F(2, 21) = 5.59, p < .05$, such that easy words ($M = 928$ ms) were repeated more quickly than hard words ($M = 941$ ms). This result replicates those in previous studies examining the effects of lexical difficulty in spoken word recognition (e.g., Kirk et al., 1997; Luce & Pisoni, 1998; Sommers, 1996). No other main effects or interactions were significant ($F_s < 1$), suggesting that there was no difference among the three groups of listeners in how quickly they responded. Furthermore, no differences in accuracy rates were found ($F_s < 1$), suggesting that participants did not sacrifice accuracy for speed in their responses to the words.

After the break—A main effect of lexical difficulty was found, $F(2, 21) = 24.76, p < .01$, such that easy words ($M = 908$ ms) were repeated more quickly than hard words ($M = 934$ ms). This result again replicates those in previous studies examining the effects of lexical difficulty in spoken word recognition. The main effect of talker condition was again not statistically significant ($F < 1$), suggesting that there was no difference among the three groups of listeners in how quickly they responded. Finally, there was an interaction between lexical difficulty and talker condition, $F(2, 21) = 3.37, p < .05$. In particular, the participants who failed to detect the change in the talker (*DIFFERENT–NO*) had a larger difference between easy and hard words (52 ms) than the participants who detected the change in the talker (*DIFFERENT–YES*; 13 ms) and the participants that received the same talker throughout the experiment (*SAME*; 22 ms). No differences in accuracy rates were found ($F_s < 1$), suggesting that participants did not sacrifice accuracy for speed in their responses to the words.

Additional comparisons—In studies investigating the effects of talker variability on spoken word recognition, it is often observed that hearing the same voice throughout the experiment results in faster and more accurate responses than hearing different voices throughout the experiment. In the present experiment, the more efficient processing of hearing the same voice throughout the experiment was assessed in two ways. First, for the listeners that heard the same voice throughout the experiment (*SAME*), the easy and hard words heard before the break were compared with the easy and hard words heard after the break. The easy

words heard after the break were repeated more quickly (883 ms) than the easy words heard before the break (924 ms), $F(1, 11) = 41.30, p < .01$. The hard words heard after the break were repeated more quickly (905 ms) than the hard words heard before the break (936 ms), $F(1, 11) = 23.61, p < .01$. These results suggest that listeners who heard the same voice throughout the experiment (i.e., those for which the indexical dimension did not change) became more efficient in processing the linguistic information.

The second way in which the more efficient processing of receiving the same voice throughout the experiment was assessed was by comparing the SAME condition after the break with the condition in which listeners heard and detected two different voices (DIFFERENT–YES; after the break). This comparison is usually the one made in studies of talker variability. Although the observed difference in the present experiment was in the expected direction (words in the SAME condition were repeated more quickly than words in the DIFFERENT–YES condition), this difference was not statistically significant ($F < 1$). Note, however, that in the present experiment, the voice was changed once rather than on every trial, as in many studies of talker variability. Furthermore, the present experiment used only two voices, in contrast to the 20 different voices (including male and female talkers) used in some studies of talker variability (e.g., Palmeri et al., 1993). These differences may account for the lack of statistical significance observed in the comparison between the SAME and DIFFERENT–YES conditions even though the difference between the two conditions was in the expected direction.

The easy-hard comparison in the SAME condition demonstrated that processing the same voice throughout the experiment facilitated processing of the linguistic dimension. The (nonsignificant) tendency to respond to words more slowly when a change in the indexical dimension was detected suggested that detecting a change in the indexical dimension of the stimulus may be detrimental to processing information in the linguistic dimension of the integral speech stimulus.

It is interesting that Goldinger (1998, Experiment 3A) showed that participants hearing isolated nonwords in a training and a test session shadowed nonwords in the test session more quickly when those items were produced by the same voice used in the training session than when different talkers were used in the training and test sessions. In the case of the present experiment, if participants failed to detect the change in the voice, they may have treated the new voice as if it were the same as the old voice. If listeners treated the different voice as the same voice they had been hearing, some facilitation in processing of the linguistic dimension should be observed, as in the same voice conditions of studies of talker variability (Goldinger, 1998). Although real words rather than nonwords were used in the present experiment, the results from comparisons of the easy and hard words before and after the break for the DIFFERENT–NO condition partially support this hypothesis. Listeners that received a different voice but did not detect it repeated easy words after the break more quickly (905 ms) than they repeated easy words before the break (946 ms), $F(1, 4) = 8.89, p < .05$. The same listeners, however, did not differ in the speed with which they repeated hard words ($M_s = 958$ ms; $F < 1$). Recall that the easy and hard words used in each part of the experiment were comparable in their lexical characteristics.

Discussion

Several important results were observed in Experiment 1. First, the inability to detect obvious changes in the environment can be found in the auditory as well as the visual modality. Of the 12 participants who were switched to a different voice halfway through the experiment, 5 of them (42%) did not detect the change in the indexical dimension of the spoken stimulus. Thus, participants may be deaf as well as blind to changes in the environment.

Second, the results of Experiment 1 support the hypothesis that attention must be focused on an object for a change in that object to be detected (Rensink et al., 1997). The integral nature of spoken language allowed the observation of a change in the processing of linguistic information as a function of focused attention on the indexical dimension of the stimulus, thereby avoiding the potential circularity inherent in the relationship between attention and detection. It was predicted that if attention focused on the indexical dimension of the stimulus, the change in the voice would be detected. Furthermore, because of the integral nature of spoken language, if attention was focused on the indexical dimension, some cost to processing would be observed in the linguistic dimension. Conversely, if attention was focused on the linguistic dimension of the stimulus, the change in the voice (i.e., the indexical dimension) would go undetected. Because of the integral nature of spoken language, if attention was not focused on the indexical dimension, some benefit might be observed in the linguistic dimension.

The significant interaction between lexical difficulty and talker condition and the additional comparisons between the first and second halves of the experiment support the hypothesis that attention must be focused on an object for a change to be detected. Focusing attention on the indexical dimension resulted in detection of the change in voices (the DIFFERENT–YES group) but came at a cost to processing in the linguistic dimension. Recall that the words in the DIFFERENT–YES condition were repeated more slowly than the words in the SAME condition (although the difference was not statistically significant), as is typically reported in studies of the effects of talker variability on spoken word recognition (e.g., Mullennix et al., 1989).

Not focusing attention on the indexical dimension resulted in a failure to detect the change in voices (the DIFFERENT–NO group), but efficient processing of the linguistic dimension resulted. The DIFFERENT–NO listeners, similar to the listeners in the SAME condition, showed some evidence of facilitated processing of the linguistic dimension (i.e., for easy but not for hard words) as a function of not attending to the indexical dimension. Recall that listeners who did not have a change in voice to attend to (the SAME group) were more efficient at processing both easy and hard words in the second half of the experiment than in the first half of the experiment. This pattern of results was similar to the advantage observed by Goldinger (1998) when the same compared with a different talker was used in training and test conditions.

Thus, the results of Experiment 1 suggest that change deafness (as well as change blindness; see Rensink et al., 1997) may occur as a function of (not) attending to the relevant stimulus dimension. Before elaborating on the implications of this finding for studies of change detection, attention, and spoken word recognition, several other possibilities that may explain why observers or listeners fail to detect changes must be ruled out. In many change blindness experiments as well as in the present experiment, detection of the change is assessed via subjective responses on the part of the participants. Could the results of change blindness (and change deafness) studies simply be due to some form of experimenter bias? That is, experimenters in change detection studies may unknowingly frame questions about the experiment in a way that influences participants' responses. An extensive amount of research has examined the formation of false memories (e.g., Roediger & McDermott, 1995) and the malleability of representations in memory as a function of suggestion (e.g., Loftus, 2000a, 2000b). Furthermore, work by Porter, Birt, Yuille, and Lehman (2001) found that interviewers in false memory experiments who were more extroverted were more successful at planting false memories than interviewers who were more introverted. Perhaps mechanisms that are related to (false) memory rather than to attention are producing failures to detect changes in the environment.

The inclusion of the SAME condition in the present experiment, however, rules out such an implanted memory explanation. In the present experiment, 12 participants did not receive a change in the talker producing the stimulus words. All of the participants in this group (SAME)—just like the participants in the group that did receive a change in the talker producing the stimulus words (DIFFERENT–NO and DIFFERENT–YES)—were asked whether they detected a change in the talker. If the effects of change detection were due to influences of suggestibility on memory, then one would predict that participants who did not hear a change would falsely report that there was a change. None of the listeners in the SAME group falsely reported that a change had occurred. The inclusion of this condition is an important extension of previous studies of change detection, which often fail to include catch trials or other types of control conditions to assess the influence of the experimenter and other related biases on the detection of changes.

Experiment 2

The results of Experiment 1 provide some support for an attention-based account of detecting changes in the environment (Rensink et al., 1997). Furthermore, Experiment 1 used a control condition to rule out the possibility that the reported failure to detect changes was due to influences of the experimenter (e.g., Porter et al., 2001) or to influences of suggestion on memory (e.g., Loftus, 2000a, 2000b). However, Experiment 1 did not rule out another memory-based explanation for the failure to detect an auditory change. Namely, the 1-min silence-filled rest break that occurred halfway through the auditory shadowing task in Experiment 1 was much longer in duration than the millisecond delays often used in visual change detection tasks. For example, Rensink et al. (1997) developed a *flicker paradigm* that used an 80-ms delay between Picture A and Picture A' (which introduced a change in the visual scene depicted in Picture A). To rule out the possibility that the 1-min delay that appeared between Talker A and Talker B in Experiment 1 outstripped a short-term memory representation of voice-related information, I ran the same experiment again, except that there was no delay in switching from Talker A to Talker B.

Method

Participants—A group of 14 native speakers of English who reported no history of hearing or speech disorders participated in the experiment for partial fulfillment of an introductory psychology research requirement. None of the participants in the present experiment took part in Experiment 1.

Stimuli—The same stimuli that were used in Experiment 1 were also used in the present experiment.

Procedure—The same equipment and procedure used in Experiment 1 were also used in the present experiment, with two exceptions. First, there was no delay in switching from Talker A to Talker B. That is, there was no 1-min rest break halfway through the experiment. Instead, participants repeated all 100 words with no obvious break in timing. Second, half the participants heard Voice A followed by Voice B, whereas the other half heard Voice B followed by Voice A. No participants heard only Voice A or only Voice B in both halves of the experiment. That is, there was no same condition in the present experiment.

The same three questions used in Experiment 1 to assess explicit detection of the change in talker were also used in the present experiment: (a) *Did you notice anything unusual about the experiment?* (b) *Was the first half of the experiment the same as the second half of the experiment?* (c) *Was the voice in the first half of the experiment the same voice that said the words in the second half of the experiment?*

Results

Detection of change in the indexical dimension—Of the 14 participants who heard different voices in both halves of the experiment, 6 noticed the change in the talker either by stating that the voice was different in response to Question 1 or 2 or by answering “no” to Question 3. The remaining 8 participants (57%) did not state that the talker changed when asked Questions 1 and 2 and answered “yes” in response to Question 3, indicating that they failed to detect the change in the talker.

Concomitant influence on the linguistic dimension—Mixed 2 (lexical difficulty) $\times$ 2 (talker condition) ANOVAs were used to examine the response latencies of the correctly repeated words in the experiment. Lexical difficulty, a within-participants factor, refers to the easy-hard manipulation among the words. Talker condition, a between-participants factor, was determined by whether participants explicitly detected the change in talkers. DIFFERENT-YES refers to the 6 listeners who explicitly stated that they detected two different talkers. DIFFERENT-NO refers to the 8 listeners who failed to detect the change in talkers. Separate analyses were done for the words in the first block (i.e., before the change) and for the words in the second block (i.e., after the change). The means for each condition before and after the change are displayed in Figure 2.

Before the change—A main effect of lexical difficulty was found, $F(1, 12) = 6.39, p < .05$, such that easy words ($M = 845$ ms) were repeated more quickly than hard words ($M = 858$ ms). No other main effects or interactions were significant ($F_s < 1$), suggesting that there was no difference between the two groups of listeners in how quickly they responded. Furthermore, no differences in accuracy rates were found ($F_s < 1$), suggesting that participants did not sacrifice accuracy for speed in their responses.

After the change—A main effect of lexical difficulty was found, $F(1, 12) = 28.29, p < .01$, such that easy words ($M = 819$ ms) were repeated more quickly than hard words ($M = 845$ ms). The main effect of talker condition was again not statistically significant ($F < 1$), suggesting that there was no difference between the two groups of listeners in how quickly they responded.

Of greatest interest was the interaction between lexical difficulty and talker condition, $F(1, 12) = 5.77, p < .05$. In particular, the participants who failed to detect the change in the talker (DIFFERENT-NO) had a larger difference between easy and hard words (36 ms) than the participants who detected the change in the talker (DIFFERENT-YES; 14 ms). Finally, no differences in accuracy rates were found ($F_s < 1$), suggesting that participants did not sacrifice accuracy for speed in their responses.

Additional comparisons—As in Experiment 1, additional comparisons were performed to examine the influence on processing linguistic information as a function of not attending to the indexical dimension of the stimulus. That is, in the DIFFERENT-NO condition, easy words before the change were compared with easy words after the change, and hard words before the change were compared with hard words after the change. Listeners that received a different voice but did not detect it repeated easy words after the break more quickly (808 ms) than they repeated easy words before the break (836 ms), $F(1, 8) = 38.54, p < .01$. The same listeners, however, did not differ significantly in the speed to repeat hard words, although the difference was in the predicted direction (hard words before the change = 850 ms, hard words after the change = 844 ms; $F < 1$).

Discussion

The results of Experiment 2 showed that 8 out of 14 participants (57%) were deaf to the change in talkers that occurred halfway through the experiment, providing an important replication

and extension of the results of Experiment 1. In Experiment 1, a 1-min rest break separated the presentation of Voice A and Voice B. Typically, in experiments examining change detection in the visual domain (e.g., Rensink et al., 1997; Simons & Levin, 1998), a shorter duration separates Image A and Image A' (the image containing a difference). It is possible that the long duration separating Voice A and Voice B played a role in listeners not being able to detect the change in the talker's voice observed in Experiment 1. That is, representations in memory of the voice characteristics may have decayed during the long duration between Voice A and Voice B in Experiment 1. However, the fact that the results in Experiment 2, in which the switch was immediate between Voice A on the 50th trial and Voice B on the 51st trial, were similar to those in Experiment 1 shows that the 1-min silence-filled rest break (in Experiment 1) was not the cause of the reported change deafness effects.

Rather, the results of Experiment 2 provide additional evidence that attention must be focused on an object for a change to be detected (Rensink et al., 1997). As in Experiment 1, when attention was focused on indexical information the change in the voice was detected at the expense of efficient processing in the linguistic dimension (the DIFFERENT-YES condition). In contrast, when attention was not focused on the indexical dimension of the integral spoken stimulus, the change in the voice was not detected. Furthermore, when attention was not focused on the indexical information, processing of the linguistic dimension increased in efficiency (as observed in the significantly faster responses to easy words in the DIFFERENT-NO condition). Together the results of Experiments 1 and 2 show that participants may fail to detect changes in the auditory as well as in the visual domain if they fail to focus attention on the relevant stimulus dimensions.

However, even if attention is focused on an object, a change in that object may still go undetected. Recall the pioneering psychophysical work of Weber (1846/1912) on the difference threshold. Weber noticed that the amount a stimulus must increase (or decrease) in magnitude remained relatively constant for a given sensory domain. For weight, the Weber fraction is 1/20, which means that if someone was to lift a 100-g weight as well as weights that were 101 g, 102 g, 103 g, and 104 g, he or she would perceive them all as being equally heavy. Only a weight of 105 g would be perceived as being heavier than the 100-g weight. Even though attention is focused on the weights, the change in the weights would not be detected until a certain physical threshold is crossed.

In the context of the present study, perhaps there is an analogous difference threshold for indexical information. Two voices may be perceived as the same voice until a certain physical threshold in the two voices is crossed, allowing one to perceive a difference between the two voices. Thus, the observed effects of change deafness in Experiments 1 and 2 may have been due to the two male voices not exceeding a difference threshold for voice characteristics. Although this possibility is unlikely given the detection of the change in voice by some listeners and the concomitant influences on lexical processing observed in both experiments, I used an AX matching task in Experiment 3 to empirically demonstrate that listeners can discriminate between the two male voices used in Experiments 1 and 2.

Experiment 3

A number of participants in Experiments 1 and 2 were unable to detect that a change in the talker occurred halfway through the experiment. The present experiment was performed to better identify the cause of the observed change deafness effect. If the two male voices used in the experiments were too similar perceptually (i.e., if the voices did not exceed a difference threshold of some sort), perhaps listeners were merely guessing whether the voices were the same or different throughout the experiment. To assess the perceptual similarity of the two voices used in Experiments 1 and 2, I used an AX matching task in the present experiment. In

this experiment, listeners were presented with a word (e.g., *cat*) over a set of headphones. This word was followed by the same word (i.e., same linguistic message) spoken in either the same voice or by the other voice. For one group of participants, the time between the two words was 50 ms, whereas for another group of participants, the time between the two words was 1,500 ms. Two different interstimulus intervals (ISIs) were used to demonstrate that relevant voice characteristics can remain in memory for a relatively long time (i.e., this information is not represented in *echoic* memory). Participants then indicated whether the voice that spoke the words was the same in each instance or whether two different voices said the same word.

Method

Participants—A group of 20 native speakers of English who reported no history of hearing or speech disorders participated in the experiment for partial fulfillment of an introductory psychology research requirement. None of the participants took part in the previous experiments. Ten participants heard the two words with a 50-ms ISI, and 10 participants heard the two words with a 1,500-ms ISI.

Stimuli—The same stimuli used in Experiments 1 and 2 were used in the present experiment.

Procedure—Participants were tested one at a time on a Macintosh Quadra 950 running PsyScope 1.2.2 (Cohen et al., 1993), which controlled stimulus randomization and presentation and collection of response latencies via a PsyScope button box. A typical trial proceeded as follows: A stimulus word was presented over the headphones at 70 dB SPL to a participant. An ISI of 50 ms (or 1,500 ms) elapsed, and a second word was presented over the headphones. Note that in all cases, participants received the same word (i.e., the same linguistic message; e.g., *cat*). What varied in each trial was whether the voice saying the word was the same or different.

For each ISI, two lists were created, with half the words on each list being spoken by the same talker and the other half being spoken by different talkers. The words on each list were counterbalanced such that the words that were spoken by the same talker on List 1 were spoken by different talkers on List 2 and the words that were spoken by different talkers on List 1 were spoken by the same talker on List 2. In the different-voice word pairs on each list, half the pairs were presented in the order of Talker A followed by Talker B, whereas the other half were presented in the order of Talker B followed by Talker A. Another trial began 1 s after a response was made.

Results and Discussion

At the 50-ms ISI, participants were 92.2% accurate in responding that the voices were different when a word was presented in two different voices. When the word was presented in the same voice separated by 50-ms, participants were 98.6% accurate in responding that the voices were the same. Comparable results were obtained when the words were separated by a 1,500-ms ISI: When a word was presented in two different voices, participants were 90.2% accurate in responding that the voices were different; when a word was presented in the same voice, participants were 96.2% accurate in responding that the voices were the same. These results suggest that the two male voices used in this set of experiments were highly discriminable perceptually.

The ability of listeners to discriminate between two unrelated voices is perhaps not surprising given that Gedda, Bianchi, and Bianchi-Neroni (1955) found that monozygotic and dizygotic twins could discriminate between recordings of their own voice and that of their twin. The ability of monozygotic and dizygotic twins to discriminate between their own voice and that of their twin is particularly interesting given that Luchsinger (1940) found that the voice range

of monozygotic twins is more similar than that of dizygotic twins. One might also expect that the voice range of dizygotic twins is more similar than the voice range of two randomly selected speakers of the same gender, such as the male voices used in this set of experiments. Thus, it is perhaps somewhat unremarkable that listeners can easily discriminate between the two voices used in the present experiments.

The ability to discriminate between voices at a shorter (50-ms) and a longer (1,500-ms) ISI should also not be surprising given the work of Craik and Kirsner (1974), who found that voice details may persist in memory for 2–3 min (see Palmeri et al., 1993, for evidence that voice details persist for even longer amounts of time). In short, the results of Experiment 3 suggest that the failures to detect the change in voice in Experiments 1 and 2 were not due to the perceptual similarity of the two voices or to limits in memory for voice-related information.

General Discussion

The results of the present set of experiments suggest that participants are deaf as well as blind to changes in the environment, further generalizing the inability to detect changes in the environment. More important, the results of Experiments 1–3 rule out a number of possibilities for the failure to detect changes. As demonstrated in Experiment 1, the inability to detect changes in the auditory (and presumably the visual) environment does not seem to be due to the influence of the experimenter on participants, as in studies of implanted or false memories (e.g., Loftus, 2000a, 2000b; Roediger & McDermott, 1995; see also Porter et al., 2001). Experiments 2 and (to a certain extent) 3 demonstrated that the observed change deafness effect was not due to the 1-min delay used in Experiment 1 outstripping some sort of short-term memory store for voice-related information. Finally, Experiment 3 showed that the observed change deafness effect was not due to the two male voices that were used in the experiments being perceptually indiscriminable from each other.

The evidence from these experiments using an integral stimulus supports the hypothesis of Rensink et al. (1997) that suggests attention must be directed toward an object for a change in that object to be detected. Rensink et al. (1997; see also Nosofsky, 1987; Rensink, 2000; Shapiro, 2000) hypothesized that attention aids in the detection of changes in the environment by modulating the specificity of representations. If attention is allocated to a particular characteristic, a more detailed representation is made and stored in memory. If attention is not allocated to a particular characteristic, a less complete, less detailed representation is made and stored. Detection of a change in a scene is more likely when a new representation of that scene is compared with a detailed representation (i.e., one that was attended) than when it is compared with a less detailed representation of a given scene in memory.

Variability in representational specificity provides an interesting account of the present experiments investigating change deafness. Consider the results of Remez, Fellows, and Rubin (1997), which suggest that a single representation of a spoken utterance—on the basis of the phonetic properties of the utterance—is used to discriminate lexical as well as indexical information. An increase in attention to the phonetic properties of an utterance will increase the specificity of the representation of the utterance. A more specific representation will not only allow one to detect that a change has been made in the voice producing the utterance (i.e., indexical information), but it may also aid in the recognition of the word (i.e., lexical or linguistic information). Consider the number of words described by the general representation *plosive-vowel-plosive* and the number of words that are confusable with the more detailed representation /kæt/. There are fewer words in the latter case, making it more likely that a word will be perceived quickly and accurately when it is represented in a more specified manner.

Note, however, that an increase in representational specificity may not affect all words equally. Those words with few similar-sounding words (i.e., easy words) may benefit more from an increase in specificity than words with many similar-sounding words (i.e., hard words). In the case of easy words, even a small increase in specificity may be enough to enable one to differentiate a word from a potential competitor and to respond more quickly or accurately. However, in the case of hard words, a great deal of overlap among representations may still exist despite an increase in specificity, making any benefit to spoken word recognition difficult to measure. For example, the slightly more specified representation /k/-vowel-plosive still leaves several neighbors among which to discriminate: *cap*, *cad*, *cab*, *cub*, *cod*, *code*, and *keep*. This may explain why, in the present set of experiments, only easy words seemed to receive significant improvements.

The processing advantage observed for easy words in the present experiment contrasts with the processing advantage observed for hard words by Sommers and Danielson (1999) in the context of investigating the influence of semantic information on spoken word recognition. The difference in processing advantages may be related to differences in the kind of specificity added to a representation by top-down information, as in Sommers and Danielson, versus the kind of specificity added to a representation by stimulus-based, bottom-up information, as in the present experiments. Continuing with the example from above, semantic information constraining the conversation to animals would reduce the size of the still dense /k/-vowel-plosive neighborhood. Words such as *cap*, *cab*, *keep*, and *code* would be eliminated, leaving *cub* and *cod* as phonologically and semantically relevant. Such semantic information may do little to accelerate the already rapid processing of a word in a sparse phonological neighborhood.

The present results support the hypothesis that attention must be directed toward a particular dimension of a stimulus for a change in that dimension to be detected. They do not, however, provide extensive insights into what directs attention in a given scene. When factors that pull attention to certain features of a scene are absent, characteristics of the listener or observer or the task at hand may still result in attention being pushed toward other features of a scene by higher level cognitive processes. Nonsense words may need to be used in future studies of change deafness to further attenuate top-down biases to process the linguistic content found in a spoken utterance (cf. Williams & Simons, 2000). Alternatively, comparisons could be made between expert and novice language users to observe the effects of attention being pushed to different degrees in future change deafness experiments. Such experiments would parallel the work of Werner and Thies (2000), who found that experts in American football were more likely to detect changes in football images than participants with less expertise in American football. Future studies in the auditory domain producing results concordant with Werner and Thies would provide additional support for the hypothesis that domain-specific expertise may influence the allocation of attention.

Even though attention may be pushed or pulled to certain dimensions of a stimulus, not all dimensions of a stimulus may engage attention. Recall that indexical and linguistic information are integral dimensions of the speech signal (Jeger et al., 1995; Jeger, Pirozzolo, et al., 1993; Jeger, Stout, et al., 1993; Mullennix & Pisoni, 1990). Also recall that Remez et al. (1997) suggested that phonetic representations underlie both indexical and linguistic information. Modifications to a speech signal that do not result in a modification of phonetic information may not engage attention and therefore may not influence the processing of (or the detection of changes in) linguistic or indexical dimensions.

As evidence of the relevance of only certain stimulus dimensions, consider the results of Sommers, Nygaard, and Pisoni (1994). In Experiment 1, Sommers et al. (1994) asked a set of talkers to produce the same words at three different rates (slow, medium, and fast), resulting

in three tokens for each word with slightly different phonetic patterns (Miller, 1981). When these recordings were played to another group of listeners, differences in the speaking rate as well as the talker affected processing. However, when Sommers et al. (1994) digitally modified the amplitude of the recordings (in Experiment 2), only differences in talkers influenced processing. Note that digitally increasing the amplitude of the signal does not modify the phonetic information in the signal, in contrast to naturally producing tokens at various amplitudes (e.g., whisper, normal conversation level, and scream). Thus, only certain dimensions of a stimulus may be cognitively relevant—perhaps only those dimensions that are naturally correlated (McBeath & Neuhoff, 2002)—and only certain dimensions may be engaging to attention and may affect processing. Future research of change detection in the auditory domain may lead to insights about the general mechanisms of attention and cognition as well as to specific insights into the processes and representations involved in spoken word recognition.

Acknowledgements

This research was supported in part by Training Grant T32 DC 00012 and Research Grant R03 DC 04259 from the National Institute on Deafness and Other Communication Disorders, National Institutes of Health. Some of this work was conducted while I was a postdoctoral research fellow at Indiana University Bloomington. Portions of these data were presented at the 141st meeting of the Acoustical Society of America, Chicago, Illinois (June 2001).

I thank Paul Luce, Robert Nosofsky, David Pisoni, Mitch Sommers, and Holly Storkel for their helpful comments, support, and suggestions. I also thank Stacy Greenbaum for her assistance in collecting additional data at the University of Kansas.

References

- Bricker, PD.; Pruzansky, S. Speaker recognition. In: Lass, NJ., editor. Contemporary issues in experimental phonetics. New York: Academic Press; 1976. p. 295-326.
- Cherry EC. Some experiments on the recognition of speech, with one and two ears. *Journal of the Acoustical Society of America* 1953;25:975–979.
- Cohen JD, MacWhinney B, Flatt M, Provost J. PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments, and Computers* 1993;25:257–271.
- Craik FIM, Kirsner K. The effect of speaker's voice on word recognition. *Quarterly Journal of Experimental Psychology* 1974;26:274–284.
- Garner, W. The processing of information and structure. Potomac, MD: Erlbaum; 1974.
- Gedda L, Bianchi A, Bianchi–Neroni L. Voce dei gemelli: I. Prova di identificazione intrageminale della voce in 104 coppie [The voices of twins: I. Proof of intrageminal identification of voice in 104 twin pairs]. *Acta Geneticae Medicae et Gemmologiae* 1955;4:121–130.
- Goldinger SD. Echoes of echoes? An episodic theory of lexical access. *Psychological Review* 1998;105:251–279. [PubMed: 9577239]
- Goldinger SD, Pisoni DB, Logan JS. On the nature of talker variability effects on recall of spoken word lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1991;17:152–162.
- Grimes, J. On the failure to detect changes in scenes across saccades. In: Akins, K., editor. *Perception: Vancouver studies in cognitive science. 2*. Oxford, England: Oxford University Press; 1996. p. 89-110.
- Hecker, M. Speaker recognition: An interpretive survey of the literature: ASHA Monographs, No. 16. Washington, DC: American Speech and Hearing Association; 1971.
- Henderson JM, Hollingworth A. The role of fixation position in detecting scene changes across saccades. *Psychological Science* 1999;10:438–443.
- Hollingworth A, Henderson JM. Semantic informativeness mediates the detection of changes in natural scenes. *Visual Cognition* 2000;7:213–235.
- Jerger S, Martin R, Pearson DA, Dinh T. Childhood hearing impairment: Auditory and linguistic interactions during multi-dimensional speech processing. *Journal of Speech and Hearing Research* 1995;38:930–948. [PubMed: 7474984]

- Jerger S, Pirozzolo F, Jerger J, Elizondo R, Desai S, Wright E, Reynosa R. Developmental trends in the interaction between auditory and linguistic processing. *Perception & Psychophysics* 1993;54:310–320. [PubMed: 8414890]
- Jerger S, Stout G, Kent M, Albritton E, Loiselle L, Blondeau R, Jorgenson S. Auditory Stroop effects in children with hearing impairment. *Journal of Speech and Hearing Research* 1993;36:1083–1096. [PubMed: 8246474]
- Kirk KI, Pisoni DB, Miyamoto RC. Effects of stimulus variability on speech perception in listeners with hearing impairment. *Journal of Speech and Hearing Research* 1997;40:1395–1405.
- Klein, R.; Kingstone, A.; Pontefract, A. Orienting of visual attention. In: Rayner, K., editor. *Eye movements and visual cognition: Scene perception and reading*. New York: Springer; 1992. p. 45-65.
- Kuèera, H.; Francis, WN. *Computational analysis of present-day American English*. Providence, RI: Brown University Press; 1967.
- Levin DT, Momen N, Drivdahl SB, Simons DJ. Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition* 2000;7:397–412.
- Levin DT, Simons DJ. Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin & Review* 1997;4:501–506.
- Loftus, EF. Remembering what never happened. In: Tulving, E., editor. *Memory, consciousness, and the brain: The Tallinn conference*. Philadelphia: Psychology Press; 2000a. p. 106-118.
- Loftus, EF. Suggestion, imagination, and the transformation of reality. In: Stone, A.; Turkkan, JS., editors. *The science of self-report: Implications for research and practice*. Mahwah, NJ: Erlbaum; 2000b. p. 201-210.
- Luce PA, Pisoni DB. Recognizing spoken words: The neighborhood activation model. *Ear and Hearing* 1998;19:1–36. [PubMed: 9504270]
- Luchsinger R. Die Sprache und Stimme von ein- und zweieiigen in Beziehung zur Motorik und zum Erbcharakter [The speech and voice of one- and two-egg twins in reference to motor mechanism and hereditary character]. *Archiv der Julius Klaus Stiflung fuer Vererbungsforschung, Sozialanthropologie, und Rassenhygiene* 1940;15:66–76.
- Martin CS, Mullennix JW, Pisoni DB, Summers WV. Effects of talker variability on recall of spoken word lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1989;15:676–684.
- McBeath MK, Neuhoff JG. The Doppler effect is not what you think it is: Dramatic pitch change due to dynamic intensity change. *Psychonomic Bulletin & Review* 2002;9:306–313. [PubMed: 12120793]
- Miller, JL. Effects of speaking rate on segmental distinctions. In: Eimas, PD.; Miller, JL., editors. *Perspectives on the study of speech*. Hillsdale, NJ: Erlbaum; 1981. p. 39-74.
- Mullennix JW, Pisoni DB. Stimulus variability and processing dependencies in speech perception. *Perception & Psychophysics* 1990;47:379–390. [PubMed: 2345691]
- Mullennix JW, Pisoni DB, Martin CS. Some effects of talker variability on spoken word recognition. *Journal of the Acoustical Society of America* 1989;85:365–378. [PubMed: 2921419]
- Nosofsky RM. Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1987;13:87–108.
- Nusbaum, HC.; Pisoni, DB.; Davis, CK. Sizing up the Hoosier mental lexicon: Measuring the familiarity of 20,000 words (Research on Speech Perception, Progress Report No. 10). Bloomington: Speech Research Laboratory, Department of Psychology, Indiana University; 1984.
- Nygaard LC, Sommers M, Pisoni DB. Speech perception as a talker-contingent process. *Psychological Science* 1994;5:42–46.
- Palmeri TJ, Goldinger SD, Pisoni DB. Episodic encoding of voice attributes and recognition memory for spoken words. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1993;19:309–328.
- Pisoni, DB. Some thoughts on “normalization” in speech perception. In: Johnson, K.; Mullennix, JW., editors. *Talker variability in speech processing*. San Diego, CA: Academic Press; 1997. p. 9-32.
- Porter S, Bin AR, Yuille JC, Lehman DR. Negotiating false memories: Interviewer and rememberer characteristics relate to memory distortion. *Psychological Science* 2001;6:507–510.

- Posner MI. Orienting of attention. *Quarterly Journal of Experimental Psychology* 1980;32:3–25. [PubMed: 7367577]
- Remez RE, Fellows JM, Rubin PE. Talker identification based on phonetic information. *Journal of Experimental Psychology: Human Perception and Performance* 1997;23:651–666. [PubMed: 9180039]
- Rensink RA. Visual search for change: A probe into the nature of attentional processing. *Visual Cognition* 2000;7:345–376.
- Rensink RA, O'Regan JK, Clark JJ. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science* 1997;8:368–373.
- Roediger HL, McDermott KB. Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1995;21:803–814.
- Shapiro KL. Change blindness: Theory or paradigm? *Visual Cognition* 2000;7:83–91.
- Simons DJ, Levin DT. Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review* 1998;5:644–649.
- Sommers MS. The structural organization of the mental lexicon and its contribution to age-related declines in spoken-word recognition. *Psychology and Aging* 1996;11:333–341. [PubMed: 8795062]
- Sommers MS, Danielson SM. Inhibitory processes and spoken word recognition in young and older adults: The interaction of lexical competition and semantic context. *Psychology and Aging* 1999;14:458–472. [PubMed: 10509700]
- Sommers MS, Nygaard LC, Pisoni DB. Stimulus variability and spoken word recognition: I. Effects of variability in speaking rate and overall amplitude. *Journal of the Acoustical Society of America* 1994;96:1314–1324. [PubMed: 7962998]
- Torretta, GM. The “easy-hard” word multi-talker speech database: An initial report (Research on Spoken Language Processing, Progress Report No. 20). Bloomington: Speech Research Laboratory, Department of Psychology, Indiana University; 1995.
- Weber, EH. The sense of touch and common feeling. In: Rand, B., translator and editor. *The classical psychologists*. Boston: Houghton Mifflin; 1912. p. 557-561. (Original work published 1846)
- Werner S, Thies B. Is “change blindness” attenuated by domain-specific expertise? An expert-novices comparison of change detection in football images. *Visual Cognition* 2000;7:163–173.
- Williams P, Simons DJ. Detecting changes in novel, complex three-dimensional objects. *Visual Cognition* 2000;7:297–322.
- Zue, VW.; Schwartz, RM. Acoustic processing and phonetic analysis. In: Lea, WA., editor. *Trends in speech recognition*. Englewood Cliffs, NJ: Prentice-Hall; 1980. p. 101-124.

Figure 1. Reaction times from the first (top) and second (bottom) halves of Experiment 1 to easy and hard words from participants who received the same talker in both halves of the task (SAME), from participants who detected the change in the talker (DIFFERENT-YES), and from participants who failed to detect a change in the talker (DIFFERENT-NO).

Figure 2. Reaction times from the first (top) and second (bottom) halves of Experiment 2 to easy and hard words from participants who detected the change in the talker (DIFFERENT-YES) and from participants who failed to detect a change in the talker (DIFFERENT-NO).

Table 1
Mean Values for the Easy and Hard Words Used in Experiments 1–3

Block	Word frequency		Neighborhood density		Neighborhood frequency	
	Easy	Hard	Easy	Hard	Easy	Hard
1	192.32	5.76	14.20	26.88	39.03	255.45
2	153.72	11.24	12.52	27.60	30.34	348.95
Overall	173.02	8.50	13.36	27.24	34.68	302.19