Published in final edited form as:

Lang Cogn Neurosci. 2015 January 1; 30(1-2): 149–166. doi:10.1080/01690965.2013.862285.

Interpreting prosodic cues in discourse context

Meredith Brown^{a,*}, Anne Pier Salverda^a, Christine Gunlogson^{a,b}, and Michael K. Tanenhaus^{a,b}

^aDepartment of Brain and Cognitive Sciences, University of Rochester, Rochester, NY, US

^bDepartment of Linguistics, University of Rochester, Rochester, NY, US

Abstract

Two visual-world experiments investigated whether and how quickly discourse-based expectations about the prosodic realization of spoken words modulate interpretation of acoustic-prosodic cues. Experiment 1 replicated effects of segmental lengthening on activation of onset-embedded words (e.g. *pumpkin*) using resynthetic manipulation of duration and fundamental frequency (F0). In Experiment 2, the same materials were preceded by instructions establishing information-structural differences between competing lexical alternatives (i.e. repeated vs. newly-assigned thematic roles) in critical instructions. Eye-movements generated upon hearing the critical target word revealed a significant interaction between information structure and target-word realization: Segmental lengthening and pitch excursion elicited more fixations to the onset-embedded competitor when the target word remained in the same thematic role, but not when its thematic role changed. These results suggest that information structure modulates the interpretation of acoustic-prosodic cues by influencing expectations about fine-grained acoustic-phonetic properties of the unfolding utterance.

*Corresponding author: mbrown@bcs.rochester.edu.

⁴Effects in visual-world experiments examining spoken-word recognition are often statistically more robust for competitors than for targets. A possible explanation comes from in-progress work on an event-based analysis framework that predicts both (1) the probability that a saccade will occur to a target, competitor and unrelated distractor and (2) the duration of the resulting fixation (Frank et al., in preparation).

In examining a data set (Salverda & Tanenhaus, 2010) we noticed the following pattern. Differences in fixation proportions to a strong and weak competitor were primarily due to fixations on the stronger competitor having longer durations than fixations on the weaker competitor. Thus the size of the competitor effect is due to participants being less likely to leave the stronger competitor within a specified time window. These “competitor-to-target” shifts are the fixations that are most likely to be sensitive to the relevant manipulations because participants nearly always shift from the competitor to the target by the end of the analysis window (which is typically chosen to coincide with the point in time at which competitor fixations merge with distractor fixations). The same pattern is observed in studies manipulating VOT (McMurray, p.c.). Dahan and Gaskell (2007) make a similar argument for frequency effects. Competitor-to target shifts also contribute to different patterns of looks to targets across conditions. However, target fixations within that same window also include fixations that are likely to be much less sensitive to experimental manipulations. These include initial fixations to the target, where participants never receive inconsistent segmental information. They also include distractor-to-target shifts, which may be less sensitive to experimental manipulations because distractor fixations are less likely to have been launched as a result of information from the target word.

If we assume that competitor-to-target fixations provide most of the signal with respect to the manipulation, competitor fixations will contain a higher signal-to-noise ratio than target fixations. If we further assume that noise-fixations are distributed similarly across conditions, we would expect the following patterns across the literature: (1) In most studies, competitor and target fixations will show the same patterns, with statistical analyses showing bigger effect sizes for competitor analyses compared to target analyses; (2) On some proportion of studies effects will be reliable in competitor analyses but not reliable (or even trending in the opposite direction) in target analyses. The reverse is possible, but much less likely; and (3) The likelihood of seeing reliable effects for both targets and competitors will be greater with manipulations with bigger effect sizes. This explanation is of course speculative and will require explicit tests with new classes of analyses and richer data sets. However, it provides a possible explanation for otherwise puzzling patterns in our data and in the larger set of visual-world studies of spoken-word recognition.

Keywords

information structure; prosodic structure; spoken-word recognition; eye movements

Introduction

Prosodic phrasing organizes spoken language into linear phrasal units based on prosodic properties such as pitch, rhythmic structure, and syllable duration. The acoustic signatures of prosodic phrasing are therefore useful cues during spoken language comprehension (e.g. Salverda, Dahan & McQueen, 2003; Frazier, Carlson & Clifton, 2006). Recent work has shown that the context provided by realization of prosodic phrasing in early stages of an utterance can shape the interpretation of downstream prosodic cues (e.g. Frazier et al., 2006; Dilley & McAuley, 2008; Dilley, Mattys & Vinke, 2010; Brown, Salverda, Dilley & Tanenhaus, 2011). Effects of context may also involve comparatively high-level representations known to influence prosodic realization, such as the status of a referent with respect to preceding discourse material (its *information status*). When a referent is mentioned multiple times in the same thematic role across utterances (e.g., when it remains the theme), repeated references tend to be acoustically reduced compared to when a new referent is introduced or a previously mentioned referent is assigned a different thematic role (Terken & Hirschberg, 1994). Here, we capitalize on acoustic overlap between cues to information status and cues to prosodic phrasing to examine whether and how quickly the information status of competing lexical alternatives modulates the interpretation of acoustic cues to word boundaries conditioned by prosodic phrasing. We hypothesize that relative activation of competing lexical candidates is influenced by expectations about their acoustic-phonetic realization that incorporate information from multiple aspects of linguistic structure, including prosodic phrasing and information structure.

Previous work has demonstrated that listeners use acoustic correlates of word boundaries, such as segmental lengthening, to assist in spoken-word recognition. In a visual-world study, Salverda et al. (2003) used utterances containing words like *hamster*, which contains the onset-embedded word *ham*. Tokens of *hamster* containing a longer *ham-* resulted in more fixations to a visually-displayed ham than tokens of *hamster* containing a shorter *ham-*. Salverda et al. concluded that this effect was related to prosodic phrasing: Because syllables are typically lengthened immediately before prosodic boundaries, and because prosodic boundaries are more likely to occur at word boundaries than word-medially, a relatively long *ham-* is more consistent with an immediately following prosodic boundary (i.e. the word *ham*) than a shorter *ham-*. The relative activation of monosyllabic and polysyllabic competitors is also dependent on position within an utterance (Salverda, Dahan, Tanenhaus, Crosswhite, Masharov, & McDonough, 2007). When *cap* occurs in utterance-medial position (e.g., *Put the **cap** next to the square*), *captain* competes more strongly for recognition than *cat*. However, when *cap* is utterance-final and therefore followed by a stronger prosodic boundary (e.g., *Now click on the **cap***), *cat* is the stronger competitor. Taken together, these results suggest that segmental duration influences perceived prosodic phrasing and that aspects of this phrasing in turn modulate the relative activation of competing lexical alternatives.

Segmental duration, together with pitch accenting or lack thereof, also signals information structure: Repeated mentions of an expression in the same thematic role within a discourse are typically deaccented, and therefore relatively short in duration. In contrast, discourse-new expressions and previously mentioned expressions in newly-assigned thematic roles are likely to be acoustically more prominent: longer in duration, containing a pitch accent, and associated with greater pitch excursion. Dahan, Tanenhaus and Chambers (2002) showed that listeners interpret the prosodic prominence of a spoken word as a cue to its information status. When participants encountered two successive instructions to move a candle within a visual display (1a), an acoustically prominent second mention of *candle* (which is infelicitous in this context) elicited stronger activation of a discourse-new cohort competitor (*candy*) than a reduced second mention (suggesting given information). Conversely, when the theme of the first instruction was *candy* rather than *candle* (1b), a prominent token of *candle* in the second instruction elicited weaker activation of the previously mentioned word *candy* than a reduced token. In another condition (1c), which is most relevant to the current study, *candle* was mentioned in both instructions but had different thematic roles in each instruction (goal in the first instruction; theme in the second). Contrastive pitch-accenting of the second mention of *candle* resulted in faster fixations to the correct target picture, compared to its repetition in the same thematic role in (1a). These findings suggest that listeners associate prosodic prominence with new information, and, crucially, that both discourse-new referents and previously mentioned referents in newly-assigned thematic roles contribute to new information.

- (1a) Put the candle below the triangle ... Now put the candle above the square.
- (1b) Put the candy below the triangle ... Now put the candle above the square.
- (1c) Put the necklace below the candle ... Now put the candle above the square.

Segmental duration is thus related to both prosodic word boundaries and information status. We assume that whereas acoustic cues to prosodic boundaries and information structure may not be identical, they do intersect with respect to lengthening. In particular, listeners might attribute segmental lengthening to the occurrence of a monosyllabic word in phrase-final position, or to it coinciding with prominence-lending pitch accenting typical of “new” information status, or both. Similarly, the relative brevity of a syllable may reflect its initial position within a polysyllabic word, or may be associated with the “given” information status of its containing expression, or both. It follows that situations may occur in which segmental lengthening signals partially conflicting information about an upcoming prosodic boundary vs. the information status of a referent. Such situations could shed light on how listeners interpret cues that could originate in diverse ways from diverse sources.

Previous work suggests that listeners simultaneously consider different levels of linguistic representation to constrain the activation of competing lexical alternatives (e.g. Dahan & Tanenhaus, 2004; Kukona, Fang, Aicher, Chen & Magnuson, 2011; Pirog Revill, Tanenhaus & Aslin., 2008; Magnuson, Tanenhaus & Aslin, 2008). We hypothesize that listeners likewise consider the acoustic characteristics of a word simultaneously as cues to prosodic phrasing and information structure, in order to rapidly establish the word’s likely identity. This hypothesis, which we will refer to as the *immediate-interaction hypothesis*¹, entails two

claims about spoken-word recognition: (1) prosodic cues are evaluated in parallel with respect to multiple levels of linguistic representation (in this case, prosodic phrasing and information structure); and (2) this process has an immediate impact on lexical activation. That is, listeners should simultaneously interpret durational lengthening with respect to prosodic phrasing (i.e., as pre-boundary lengthening) and with respect to information status (i.e., as prominence-induced lengthening). Listeners' initial interpretation of a spoken word is thus predicted to incorporate relevant information about potential lexical alternatives with respect to both prosodic phrasing and information structure. The interpretation of duration as a prosodic phrasing cue should therefore interact with the interpretation of overlapping acoustic cues to information status (pitch-accenting and lengthening).

One possible mechanism for the interaction of these different levels of representation involves predictive processes that forecast aspects of likely acoustic realization of different lexical candidates in context (Farmer, Brown & Tanenhaus, 2013). According to this perspective, listeners develop expectations about the acoustic realizations of words by generating internal models taking into account the words' prior realization(s) and current context. These expectations incorporate information from different levels of linguistic representation, including information structure and prosodic phrasing. These sources of information predict aspects of the phonetic realization of upcoming words, which can be used in the earliest moments of lexical processing. This expectation-based account has precedents in theories of other aspects of language processing (e.g. Levy, 2008; Altmann & Kamide, 1999) as well as in research in other cognitive and perceptual domains (Clark, 2013). Interaction in this framework occurs when a cue can be explained by hypotheses based on different subsystems and the interpretation of that cue simultaneously takes into account these hypotheses.

The phenomena we are investigating have not been explicitly addressed by existing models of spoken-word recognition. Therefore, we cannot directly map competing predictions onto specific models. Nonetheless, either or both of the assumptions underlying the immediate-interaction hypothesis might be incorrect, giving rise to at least two classes of alternative possibilities.

First, the interpretation of acoustic detail in spoken words might not depend simultaneously on both prosodic phrasing and information structure (the *non-interaction hypothesis*). Instead, listeners may interpret acoustic cues predominantly with respect to one level of linguistic representation, for instance prosodic phrasing. This scenario might arise, for example, if listeners prefer to attribute proximal cues like segmental lengthening to proximal causes (i.e. an upcoming prosodic boundary) as opposed to distal causes (e.g. discourse history). In this case, listeners might retrieve information about whether and in what role a referent was previously mentioned only when acoustic variation cannot be attributed to more proximal causes.

Second, the cues associated with different types of representation (e.g., prosodic phrasing cues vs. prosodic cues to information status) might operate at different time-scales, or

¹We use the terms "short-target" and "long-target" for ease of exposition. Each condition also involved F0 manipulations.

initially be considered by non-interacting subsystems. According to the *delayed-interaction hypothesis*, the interpretation of acoustic cues to prosodic boundaries is initially unaffected by the information status of competing lexical candidates. Conversely, the initial interpretation of prosodic cues to information status might not be modulated by prosodic phrasing. Instead, prosodic cues to different types of structures may be interpreted independently and with distinct time courses. For example, under certain assumptions of traditional feedforward models of language processing, processing proceeds in a hierarchical bottom-up order: The perception and interpretation of speech sounds precedes lexical access, which precedes situating a word within a syntactic representation, which in turn precedes retrieving semantic information and constructing information-structural representations (cf. Swinney, 1979; Tanenhaus, Leiman & Seidenberg, 1979; Norris, 1994). Extending these assumptions to the construction of prosodic-phrasing representations and information-structural representations would predict that cues to relatively low-level and/or local aspects of prosodic phrasing would have earlier or faster effects compared to cues to information status. Segmental lengthening should thus initially be interpreted as a cue to an upcoming prosodic boundary, favoring a monosyllabic interpretation, whereas inferences relating perceived prominence to previous mention in discourse context should have slower effects on lexical activation, resulting in delayed effects and interactions.

We used the visual-world paradigm (Cooper, 1974; Tanenhaus, Spivey-Knowlton, Eberhard & Sedivy, 1995) to assess effects of information structure on the interpretation of acoustic cues to prosodic phrasing. Experiment 1 evaluated the effect of segmental lengthening and pitch excursion on spoken-word recognition in isolated utterances (e.g., *Now put the pumpkin below the triangle*). In Experiment 2, an additional instruction at the beginning of each trial established a discourse context for each of the critical instructions used in Experiment 1, thereby manipulating the information status of the target word. Of interest was whether and how quickly the information status of the target word influenced the interpretation of segmental lengthening and pitch excursion.

Experiment 1

The goal of Experiment 1 was to evaluate effects of resynthetic manipulation of duration and F0 on spoken-word recognition.

Methods

Participants—Forty native English speakers from the University of Rochester and surrounding community participated in the experiment. All had normal hearing and normal or corrected-to-normal vision.

Materials and design—The 20 experimental items were instructions to manipulate pictures within a visual display (Table 1). In addition to four shapes in the corners of the display, which remained the same across trials, five pictures were displayed on each trial. These pictures included a target picture associated with a polysyllabic word (e.g. pumpkin), a competitor picture associated with the corresponding onset-embedded word (e.g. pump), and three phonologically-unrelated distractors (Figure 1). Distractor pictures were positioned such that all pictures were approximately evenly dispersed throughout the

display, with minimal displacement of the centroid of picture locations relative to the center of the grid.

All instructions were recorded by CG using a Marantz PMD660 digital recorder sampling at 44.1 kHz. Sentences were produced with consistent speech rate and intonational contours. The pitch synchronous overlap-and-add algorithm (PSOLA; Moulines & Charpentier, 1990) was used in the speech-editing software Praat (Boersma & Weenink, 2010) to create two partially-resynthesized versions of each critical instruction (e.g., *Now put the pumpkin below the triangle*). Both duration and F0 of the first syllable of the target word (*pumpkin*) were manipulated. We localized pitch excursion to the first syllable of the target word by flattening the F0 of the syllables immediately preceding and following it, setting their F0 level to match the initial F0 of the preposition following the target word (e.g. *below*). Crucially, the acoustic properties of these proximal context syllables, and of the rest of each critical sentence, were the same across the two resynthesized sentences. Only the acoustic properties of the first syllable of the target word differed across conditions.

In *short-target* items¹, the first syllable of the target word was shortened to 90% of its original duration, and its F0 was flattened to match that of the preceding and following syllables (Figure 2a). In *long-target* items, the first syllable of the target word was lengthened by 125%, and its F0 was altered to rise and fall approximately 45 Hz, beginning at vowel onset and peaking late in the rhyme (Figure 2b). These F0 manipulations were integral for the design of Experiment 2 because the prominence variations associated with information status affect both F0 and duration (e.g. Ladd, 2008). The goal of Experiment 1 was therefore to demonstrate that manipulating the duration plus F0 of the first syllable of the target word would still result in the effects of duration alone documented in previous studies, where it was shown to affect the interpretation of the syllable as corresponding to a monosyllabic vs. polysyllabic word (Davis, Marslen-Wilson & Gaskell, 2002; Salverda et al., 2003). As will become clear later, the logic of Experiment 2 required that these stimuli show a stronger monosyllabic competitor effect in the long compared to the short condition. The mean duration of the embedded word was 201 ms for short and 279 for long stimuli. The mean duration of the entire target word was 410 ms for short and 488 ms for long stimuli².

Forty-five filler sentences were constructed to mitigate prosodic and phonological contingencies in the stimuli. In all filler trials, the display contained pictures associated with two phonologically-related words and three unrelated distractor words. To discourage expectations that the target word was the longer of two phonologically-related words associated with displayed pictures, the target word was the shorter of two phonologically-related words for 20 fillers, and an unrelated word for the remaining 25. Target words in 10 of these 25 fillers underwent the same acoustic manipulations as target words in critical items, to further discourage expectations that acoustically-manipulated target words corresponded to the longer member of a phonologically-related pair.

²We elicited spoken instructions from five naive speakers in which the theme was either the polysyllabic target word or the monosyllabic competitor, intermixed with 96 syntactically varied sentences. The mean difference in duration between the monosyllabic competitor word and the initial syllable of the polysyllabic target word in these recordings was 78 ms (319 vs. 241 ms)--the same magnitude as the durational difference in our resynthetically-manipulated stimuli.

Five stimulus-lists were constructed by pseudo-randomizing picture positions within each trial (within the set of positions licensed by the spoken instruction and the even-dispersion constraint) and pseudo-randomizing trial order such that experimental trials were non-adjacent. Condition (short-target vs. long-target) was counterbalanced in a Latin square design, resulting in 10 lists completed by an equal number of participants. Each list began with five practice trials to familiarize participants with the task.

Procedure—Participants' eye-movements were recorded using a head-mounted SR Research Eyelink II eye-tracker sampling at 250 Hz. Drift-correction procedures were performed every five trials. Auditory stimuli were presented through headphones.

Each trial began with the visual display appearing in the center of the screen. After one second, the spoken instruction was presented. One second after the participant successfully executed the instruction, the next trial was initiated. The experiment lasted approximately 10 minutes.

Results and discussion

Trials on which participants required multiple attempts to select the correct object and drag it to the correct location were excluded from analysis (3.2%). Proportions of errors did not differ significantly across conditions.

Figure 3 shows mean proportions of fixations to target, competitor, and distractor pictures over time by experimental condition. At the onset of the target word, proportions of fixations to target and competitor pictures did not differ significantly between conditions. This was confirmed statistically by first transforming the proportions of fixations to each picture at target-word onset using the empirical logit function (Cox, 1970; Barr, 2008) and then conducting one-tailed paired t tests on proportions of competitor fixations (mean untransformed percentage of fixations=14% for short and 17% for long items; $t_1(39)=1.23$, $p>.1$; $t_2(19)=1.12$, $p>.1$) and target fixations (mean=16% for both short and long items; $t_1(39)=0.12$, $p>.1$; $t_2(19)=0.13$, $p>.1$).

Following Salverda et al. (2003), we hypothesized that the proportion of fixations to the picture depicting the embedded word (e.g. *pump*) as participants processed the target word would be higher for long than for short target words. The data were consistent with this prediction. Around 200 ms after target-word onset, the proportion of fixations to the target and competitor increased, whereas the proportion of fixations to the distractors decreased. Competitor fixations started to rise at the same rate in both conditions, but ultimately reached a higher peak value and decreased more slowly in the long-target condition than in the short-target condition. Around 1000 ms after target-word onset, competitor fixations merged with distractor fixations. One-tailed paired t tests on logit-transformed averaged proportions of competitor fixations between 200–1000 ms after the onset of the target word confirmed that there were more fixations to the competitor in the long-target condition than in the short-target condition ($t_1(39)=3.14$, $p<.005$; $t_2(19)=3.02$, $p<.005$). Complementary effects were observed in proportions of target fixations, which exhibited a faster increase in the short-target condition than in the long-target condition starting at approximately 550 ms after target-word onset ($t_1(39)=-1.88$, $p<.05$; $t_2(19)=-1.70$, $p=.053$).

The results demonstrate that resynthetic manipulation of duration and F0 across the initial syllable of the target word elicits effects similar to those obtained by Salverda et al. (2003) using materials with natural prosodically-conditioned acoustic variation. Target words with lengthened duration and F0 excursion were associated with stronger activation of embedded competitors, compared to target words that were compressed and had flat F0. This result sets the stage for Experiment 2, in which listeners heard the same sentences while we manipulated the information status of target and competitor words to adjudicate between the immediate-, delayed-, and noninteraction hypotheses. (In the general discussion and Appendix 1, we address concerns raised by one of the editors, Michael Wagner, that the competitor bias observed in response to long stimuli can be explained without appeal to prosodic phrasing because the longer initial syllable of the target word is consistent with both target and competitor interpretations for a longer period of time).

Experiment 2

The goal of Experiment 2 was to characterize the effects of information structure on the interpretation of acoustic cues to prosodic phrasing. An additional instruction preceding each of the critical sentences from Experiment 1 established a discourse context for each critical instruction. Thematic roles of the target and competitor words were varied within the context instruction, thereby manipulating the information status of competing lexical alternatives associated with the critical second instruction.

Methods

Participants—Eighty-four native English speakers from the University of Rochester and surrounding community took part in the experiment. They reported normal hearing and normal or corrected-to-normal vision.

Materials and design—The critical instructions in Experiment 2 were the same manipulated utterances used in Experiment 1. Context instructions were recorded during the same session as critical instructions. The display characteristics at the start of each critical instruction were identical to those of Experiment 1.

In each trial, participants heard one of two versions of the first instruction, as in (2), after one second of display preview. In same-role (SR) context instructions, the target word (e.g. *pumpkin*) was mentioned in the theme position, with the monosyllabic embedded word (e.g. *pump*) mentioned as part of the composite goal. In different-role (DR) context instructions, the monosyllabic embedded word was instead mentioned as the theme, and the target word was mentioned as part of the composite goal. Because both words were mentioned in both types of context sentences, the critical information-structural manipulation was whether the thematic role of the target word in the critical instruction remained the same as in the context instruction. In SR-context conditions, the target word had the same thematic role in both sentences, thus corresponding to *given* information in the critical instruction. In DR-context conditions, the critical target word had a newly-assigned thematic role in the second instruction, thus contributing to *new* information (compared to SR-context conditions)³. Further, in both types of context sentences, coordinate structures were used for the goal (e.g. *the pump and the circle*) to avoid differential effects of utterance-final lengthening on the

target word and embedded competitor word across the two context-types (cf. Salverda et al., 2007). A continuation rise (i.e. L-H% boundary tone; Pierrehumbert & Hirschberg, 1990) was used at the end of the first instruction to indicate prosodically that the second instruction should be interpreted with respect to the context established by the first instruction.

(2)

SR-context instruction:	Put the pumpkin between the pump and the circle...
	H* !H* !H* L-H%
DR-context instruction:	Put the pump between the pumpkin and the circle...
	H* !H !H* L-H%
Short critical instruction:	Now put the pumpkin below the triangle.
	H* !H* !H* L-L%
Long critical instruction:	Now put the pumpkin below the triangle.
	H* L+H* !H* L-L%

The 45 filler items used in Experiment 1 were also paired with context instructions. The resulting set included eight fillers in which the context instruction contained both members of a phonologically-related pair (like the context instructions for the experimental items) followed by an instruction whose theme was either the shorter member of the pair or an unrelated word. In addition, because the theme of each critical instruction was always mentioned in the prior instruction, the second theme of 17 fillers was discourse-new.

Experiment 2 used the same five pseudo-randomized trial lists as Experiment 1. Context-type (SR vs. DR) was crossed with target-word realization (short vs. long) and counterbalanced in a Latin square design, resulting in 20 lists.

Procedure—The procedure was similar to that of Experiment 1. Participants completed the initial instruction by moving the appropriate picture (e.g. pumpkin in SR-context conditions, pump in DR-context conditions) to its appropriate location. The successful completion of this action triggered the appearance of a central dot, which played the second instruction when clicked. This aspect of the procedure brought participants' gaze back to the center of the display prior to the critical instruction, which minimized baseline differences in fixations to target and competitor pictures at the onset of the critical instruction (e.g., continuing to fixate the most recently manipulated picture or the current location of the cursor). The successful completion of the critical instruction was required to proceed to the next trial. The experiment lasted approximately 15 minutes.

Predictions

The immediate-interaction hypothesis predicts that listeners simultaneously evaluate the acoustic realization of a word with respect to multiple dimensions of linguistic structure,

³We obtained recordings of the stimulus pairs and additional utterance pairs in which the critical themes were monosyllabic embedded words (e.g. *pump*), from the five speakers who provided supplementary recordings for Experiment 1. Following a context sentence in which a polysyllabic target word was mentioned as the theme and the embedded competitor word was part of the composite goal (i.e. a SR-context sentence), the mean duration of the first syllable of the polysyllabic target word in the critical instruction was 217 ms, whereas the duration of the monosyllabic word was 311 ms, a difference of 94 ms. When the roles of the polysyllabic and monosyllabic words were switched in the context sentence (i.e. a DR-context sentence), difference in durations between the syllables in the critical instruction was substantially reduced (248 ms for polysyllabic vs. 291 ms for monosyllabic words, a difference of 43 ms).

including prosodic phrasing and information structure, resulting in an interaction between context-type and target-word realization in the earliest moments of lexical activation.

In SR-contexts, the polysyllabic target word appears in the role of theme in the initial context-setting instruction, with the monosyllabic competitor as part of the goal phrase (e.g., *Put the pumpkin between the pump and the triangle*). If the theme of the critical instruction (e.g., *Now put the pumpkin below the triangle*) is acoustically reduced, listeners should be biased to interpret it as referring to *given* information (i.e. coreferential with the theme of the context instruction), favoring a polysyllabic interpretation. Conversely, if the target word is acoustically prominent, listeners should be biased to interpret it as referring to comparatively *new* information, favoring the monosyllabic competitor. These expectations are congruent with the durational cues to prosodic phrasing: Longer duration is a cue to an upcoming prosodic boundary, which should favor the competitor over the target. Both prosodic-boundary expectations and information-structural expectations should thus bias listeners toward an initial polysyllabic interpretation of short target words and an initial monosyllabic interpretation of long target words, as in Experiment 1.

In DR-contexts, the immediate-interaction hypothesis predicts a different pattern of results. Because DR-contexts introduce the polysyllabic target word as part of the goal and the monosyllabic competitor word as the theme (e.g., *Put the pump between the pumpkin and the triangle*), expectations based on information structure should favor the initial interpretation of an acoustically-reduced word in the theme-position of the second instruction as a less prominent instance of the monosyllabic word (consistent with its previous mention in the same thematic role), and of a lengthened target word as a contrastively accented reference to the polysyllabic target (in its new role as theme). The information-structural expectations conflict with prosodic-phrasing expectations regarding duration: Longer duration is a cue to an upcoming prosodic boundary, as in the SR condition, thus favoring the monosyllabic competitor over the target word. The immediate-interaction hypothesis therefore predicts that effects of duration on fixations to competitor and target pictures should be reduced in DR-context conditions, compared to SR-context conditions. In fact, proportions of fixations to the competitor picture might initially be higher for short target words than for long ones.

In contrast, the main prediction of the non-interaction hypothesis is that listeners consider acoustic cues primarily with respect to proximal aspects of linguistic structure (i.e. an immediately upcoming prosodic boundary), such that the effect of target-word realization on lexical activation is not affected by information structure. That is, listeners may attribute acoustic cues like segmental lengthening primarily to proximal prosodic phrasing, considering information from prior discourse history only when there is no obvious proximal cause for segmental lengthening. According to this hypothesis, effects of lengthening and pitch excursion on early fixations to target and competitor pictures should be similar to those observed in Experiment 1. When the initial syllable of the target word is short, listeners should not expect an upcoming prosodic boundary. When the initial syllable of the target word is long, listeners should expect an upcoming prosodic boundary, favoring the monosyllabic competitor. Thus, there should be more fixations to the monosyllabic

competitor (and slower convergence on the target) for long target words relative to short target words.

Finally, the delayed-interaction hypothesis predicts that listeners evaluate the acoustic realization of a word with respect to different aspects of linguistic structure, but with distinct time courses. According to this hypothesis, listeners evaluate the acoustic realization of a word first with respect to local prosodic phrasing and then with respect to information status. Thus, target words with long initial syllables should initially be associated with a higher proportion of competitor fixations regardless of context-type, followed by delayed effects of information status on fixation patterns (i.e., a delayed attenuation or reversal of duration effects in DR-context conditions, but not SR-context conditions).

Results and discussion

Four participants were excluded from analysis because they did not complete the experiment as instructed. Errors occurred in response to either or both instructions on 5.1% of experimental trials, which were excluded from analysis. Proportions of errors did not differ significantly across conditions.

Baseline analyses of fixation proportions—Figure 4 shows proportions of fixations to each picture in SR-context conditions, relative to the onset of the target word. Figure 5 shows proportions of fixations in DR-context conditions. In both conditions, the proportion of fixations to phonologically-unrelated distractor pictures began to drop between 100–200 ms following target-word onset. Assuming that it takes approximately 200 ms to program an eye-movement in multi-object displays in response to spoken language (Salverda et al., in press), this early effect suggests that participants anticipated repeated mention of pictures referred to in the context sentence. Further analyses of logit-transformed fixation proportions at the onset of the target word revealed that participants were more likely to be looking at the competitor picture in the SR-context conditions (mean untransformed percentage=10%) than in the DR-context conditions (mean=6%; $F_1(1,79)=8.445$, $p<0.005$; $F_2(1,19)=6.425$, $p<0.05$), and more likely to be fixating the target picture in the DR-context conditions (mean=15%) than in the SR-context conditions (mean=7%; $F_1(1,79)=17.174$, $p<0.0001$; $F_2(1,19)=14.466$, $p<0.0005$). Neither target-word realization nor the interaction between context-type and target-word realization significantly influenced competitor or target fixations at word onset.

The baseline differences in fixation proportions between context-types suggest that some aspect of our materials or task structure biased listeners to have a modest expectation for the theme of the second instruction to correspond to the goal of the first instruction rather than the theme of the first instruction. Similar baseline tendencies to fixate pictures corresponding to new information were observed by Dahan et al. (2002). Although listeners generally expect repeated mention across syntactically-parallel sentence positions (Chambers & Smyth, 1998), several factors could have reversed this expectation. For example, it is possible that the intonation of the first words of the critical instruction (*Now put the...*) led listeners to expect contrastive accenting on the theme (cf. Braun & Chen, 2010), resulting in an expectation for the upcoming theme to have a new thematic role.

However, the intonation of the carrier phrase cannot entirely account for baseline differences in fixations to the theme and goal of the context instruction, because they were present as early as 200 ms following onset of the carrier phrase (2.09% of fixations for themes vs. 4.86% for goals; $t_1(79)=2.890$, $p<.005$; $t_2(19)=3.781$, $p<.005$). A more likely explanation is that participants encoded more detailed visual and spatial information about the referent of the theme of the context instruction while moving it, resulting in a decrease in exploratory fixations to the previously mentioned theme relative to the goal during initial portions of the critical instruction (cf. Salverda, Brown & Tanenhaus, 2011). Listeners might also have been biased to fixate the most recently mentioned item. Crucially, however, baseline differences in fixation proportions did not prevent examination of effects of target-word realization on fixations to competitor and target pictures, which are described in the following sections.

Analysis of competitor fixations between 200–1000 ms—For statistical analysis, average proportions of competitor fixations were computed for the interval ranging from 200–1000 ms after target-word onset (as in Experiment 1), and transformed using the empirical logit function. Of primary interest was the interaction between context-type and target-word realization. The immediate-interaction hypothesis predicts that effects of acoustic manipulation on interpretation of the target word should be attenuated in DR-contexts, compared to SR-contexts. In contrast, the non-interaction hypothesis predicts that these two factors should not interact in the earliest moments of spoken-word recognition, whereas the delayed-interaction hypothesis predicts slower effects of information structure relative to prosodic boundaries. Comparison of competitor fixations in Figures 4 and 5 supports the predictions of the immediate-interaction hypothesis. In SR-contexts, long target words elicited a higher proportion of competitor fixations than short ones. In DR-contexts, however, proportions of competitor fixations followed a similar trajectory in short-target and long-target conditions.

A two-way ANOVA on logit-transformed proportions of competitor fixations confirmed a significant interaction between context-type and target-word realization ($F_1(1,79)=6.610$, $p<.05$; $F_2(1,19)=4.051$, $p<.05$). This interaction indicates that effects of acoustic manipulation on competitor activation differed as a function of context. In SR-contexts, a planned two-tailed paired t test revealed that competitor fixations were significantly higher for long-target items (mean=18%) than for short-target items (mean=13%; $t_1(79)=3.291$, $p<.005$; $t_2(19)=2.952$, $p<.01$). This result is consistent with all three hypotheses. In DR-contexts, however, competitor fixations were not significantly higher for long-target items (mean=12%) than for short-target items (mean=13%; $t_1(79)=-.686$, $p>.1$; $t_2(19)=-.947$, $p>.1$). In fact, competitor fixations were numerically slightly lower for long-target items than for short-target items.

Analysis of target fixations between 200–1000 ms—Although our hypotheses primarily focused on activation of the competitor word as a function of context-type and target-word realization, we expected proportions of target fixations to show a complementary pattern to proportions of competitor fixations (though see Footnote 4). We expected processing of short target words to be facilitated in SR-contexts, but not in DR-

contexts. However, participants fixated the target less upon hearing long-target items, regardless of context-type. Mean untransformed percentages of fixations in SR-context conditions were 41% for short-target items and 37% for long-target items, whereas mean percentages in DR-context conditions were 47% for short-target items and 43% for long-target items. A two-way ANOVA on logit-transformed proportions of target fixations revealed significant main effects of context-type ($F_1(1,79)=8.000, p<.01$; $F_2(1,19)=11.765, p<.005$) and target-word realization ($F_1(1,79)=4.394, p<.05$; $F_2(1,19)=4.203, p<.05$). However, the interaction between context-type and target-word realization was not significant ($F_1(1,79)=0.038, p>.1$; $F_2(1,19)=.167, p>.1$).

Adaptation analyses—One explanation for this unexpected result is suggested by an expectation-based account of the immediate-interaction hypothesis: Listeners may adapt to the prosodic characteristics of speakers, including their speech rate, pitch range, and consistency with which they signal information structure intonationally (see Kurumada, Brown & Tanenhaus, 2012, for supporting evidence). Because Experiment 2 used a counterbalanced crossed-factors design, the prosodic characteristics of half of the critical items were infelicitous or unexpected (e.g. SR-long items, in which a target word with a repeated thematic role was acoustically prominent). Therefore, interactions between context-type and target-word realization may have attenuated over the course of the experiment, due to increasing exposure to inconsistent prosodic usage.

This possibility was examined by analyzing logit-transformed competitor and target fixation proportions using multi-level linear regression models, with random intercepts and slopes for participants and items, using the *lmer* function within the *lme4* package (Bates, Maechler & Bolker, 2011) in R (version 2.15.0, R Development Core Team, 2012). Fixed effects and random slopes in the regression models included context-type, target-word realization, trial number (the position of the item in the sequence of trials encountered by the participant), and interactions between these factors. Context-type and target-word realization were sum-coded. Trial number was standardized by subtracting the mean value and dividing by the standard deviation. Final models were chosen by removing fixed effects stepwise and comparing each reduced model to the more complex model using a likelihood ratio test (Baayen, Davidson & Bates, 2008). Models included maximal random effect structure except as noted (Barr et al., 2013). We used the *kappa* function in R (R Development Core Team, 2010) to verify that the condition number was below 7 for all models. Condition numbers below 7 indicate a low degree of collinearity between predictors, which would not be expected to jeopardize the reliability of model estimates.

The final model predicting logit-transformed proportions of competitor fixations revealed a significant interaction between context-type and target-word realization (Table 2), confirming the results obtained by the two-way ANOVA. In addition, a main effect of trial number indicated an overall decrease in proportions of competitor fixations as the experiment progressed. The three-way interaction between context-type, target-word realization, and trial number was not included in the final model because it did not contribute significantly to explained variance.

For targets, the final model confirmed that the interaction between context-type and target-word realization was not significant (Table 3). However, this model revealed a significant three-way interaction between context-type, target-word realization, and trial number, indicating that the interaction between context-type and target-word realization changed as the experiment progressed. Effects of trial number on the interaction between context-type and target-word realization are illustrated in Figure 6, which shows fixation proportions averaged across the first and last third of experimental trials in SR-context and DR-context conditions, respectively. Early in the experiment, mean proportions of target fixations in SR-contexts were higher for short target words than for long ones (and, importantly, this effect originated early in the processing of the target word), whereas differences in proportions of target fixations by condition were not apparent in DR-contexts. Later in the experiment, however, mean proportions of target fixations were numerically higher in the DR-short condition than in the DR-long condition.

These patterns of target fixations were further evaluated statistically by comparing the regression lines for the two-way interaction between context-type and target-word realization when the trial number was set to different values (cf. Aiken & West, 1991). For this analysis, we fitted two additional models to the data, with trial number centered one standard deviation either below or above the mean, and examined the interaction between context-type and target-word realization within each model. The critical interaction was significant at one standard deviation below the trial-number mean (i.e., early in the experiment; Table 4), but not at one standard deviation above the mean (i.e., late in the experiment; Table 5). The significant interaction in the early trial-number model was driven by a significant difference in the interpretation of short vs. long items between SR and DR-contexts. Although fixations to the target picture rose significantly faster in the SR-short condition than in the SR-long condition early in the experiment ($\beta = -0.152$, $SE = 0.066$, $t = -2.282$, $p < 0.05$), target-word realization did not elicit significant differences in proportions of target fixations in DR-contexts early in the experiment ($\beta = 0.041$, $SE = 0.063$, $t = 0.658$, $p > 0.1$). That is, proportions of target fixations showed patterns that mirrored those observed for competitor fixations early in the experiment, and the interaction of context-type and target-word realization attenuated significantly across the experiment for target fixations but not for competitor fixations. Note, however, that the numerical size of the interaction for competitor fixations was smaller in the second half of the experiment⁵, suggesting similar adaptation effects for fixations to competitors.

To summarize, effects of segmental lengthening and F0 excursion on fixations to potential referents during the initial moments of processing the target word were modulated by context-type. In SR-contexts, monosyllabic embedded words competed for recognition more strongly when the initial syllables of target words were lengthened and associated with pitch excursion. However, this pattern was attenuated due to perceived prominence in DR-

⁵Although the three-way interaction was not significant for competitor fixations, visual inspection of competitor effects suggested a trend in the same direction: The interaction between context-type and target-word realization was numerically larger early in the experiment (Figure 6). The interaction between context-type and target-word realization was significant in a multilevel regression model with trial number centered one standard deviation below the mean (i.e. early trials; $\beta = -0.156$, $SE = 0.065$, $t = -2.397$, $p < 0.05$), but was not significant in a model with trial number centered one standard deviation above the mean (i.e. late trials; $\beta = -0.105$, $SE = 0.067$, $t = -1.556$, $p > 0.1$). Thus, the attenuation of target effects was qualitatively mirrored in competitor effects.

contexts, consistent with predictions of the immediate-interaction hypothesis. Similar patterns of results were observed for target fixations, but only early in the experiment.

General discussion

Our results suggest that the information status of competing referents affects interpretation of prosodically-conditioned detail in the speech signal. Segmental lengthening that is interpreted as signaling a following prosodic boundary when no discourse context is available (Experiment 1) can be interpreted as a consequence of pitch-accenting signaling the information status of a referent when an appropriate discourse context is available (Experiment 2). Further, listeners appear to evaluate the acoustic realization of a target word with respect to both prosodic phrasing and information structure upon processing its initial sounds (Experiment 2).

These results contribute to a growing body of work demonstrating that effects of prosodically-conditioned detail on lexical activation depend on the context in which a word is processed (e.g. Dilley & McAuley, 2008; Brown et al., 2011). For example, manipulations of prosodic information distal to the locus of processing (e.g. F0 and duration of preceding prosodic constituents or metrically prominent vs. non-prominent syllables) influence the extent to which phonemically overlapping words compete for recognition during spoken-word recognition (Brown et al., 2011, 2012a, 2012b). They are also consistent with work in speech perception showing that listeners evaluate phonetic cues with respect to speech context (e.g. Ladefoged & Broadbent, 1957; Cole, Linebaugh, Munson & McMurray, 2010; McMurray & Jongman, 2011). Taken together with the results of the present study, these context effects suggest that proximal prosodically-conditioned detail is interpreted dynamically with respect to a variety of broader contextual factors.

One potential concern about interpreting the present data involves the relative timing with which different aspects of linguistic structure are available to the language processing system. For example, if information status is available at an earlier point in time than prosodic-boundary information, then delayed effects of information status in conjunction with immediate (but later-occurring) effects of prosodic boundary information could result in what might appear to be an immediate interaction. Indeed, the presence of baseline differences between fixations to referents previously mentioned in a prominent (theme) vs. non-prominent (goal) position in Experiment 2 suggests the possibility that the discourse history influenced fixations prior to the onset of the target word (although, as previously discussed, baseline differences may also have resulted from task-based factors). Regardless, the information structure of previous utterances does not unambiguously specify the information structure of upcoming material. Listeners may have expectations about the relative likelihood of different referents being mentioned as the theme of the critical instruction—just as they may have expectations about the prosodic constituency of upcoming material (Brown et al., 2011)—but they cannot be certain of the identity of the theme during the processing of the first syllable of the target word. The crucial factor in Experiment 2 is the acoustic manipulation of the initial syllable of the target word. This manipulation results in patterns of fixations that are distinct from prior baseline effects, suggesting that listeners interpret the acoustic properties of this syllable as evidence for one referent or the other with

respect to prosodic phrasing and information structure. The onset of the target word is the earliest point at which acoustic cues to prosodic boundaries and information status are available to the listener, and the effects that we observed originated shortly after target-word onset.

As mentioned previously, another potential concern involves whether the patterns of competitor fixations observed in Experiment 1 reflect the interpretation of segmental lengthening as a cue to prosodic phrasing, or simply a delayed point-of-disambiguation effect. We address this concern in more detail in Appendix 1. However, even if one assumes that the competitor bias observed in Experiment 1 does not reflect the interpretation of lengthening as signaling an upcoming prosodic boundary, the general interpretation and theoretical conclusions from the data in Experiment 2 would be essentially the same. The data would still demonstrate that a higher-level factor (information structure established by preceding discourse context) has an immediate constraining effect on the interpretation of different acoustic realizations of the target word.

The finding that various and disparate aspects of preceding context influence the initial interpretation of segmentally identical material suggests that these effects arise from perceptual expectations. That is, listeners may evaluate incoming prosodic information with respect to expectations about aspects of the acoustic-prosodic realization of spoken words (based on, e.g., prior linguistic knowledge, prior experience, and the acoustic realization of words in preceding context). For example, listeners may expect a word to have a relatively short duration and small F0 excursion when it is mentioned twice in the theme position of sequential utterances (as in SR-context conditions in Experiment 2). The acoustic realization of the initial syllable of the target word in the SR-long condition is incongruent with this expectation. Likewise, following a DR-context sentence, listeners should expect a second mention of the target word in a newly-assigned thematic role to be associated with lengthening and F0 excursion, making the acoustic realization of the target word in the DR-short condition infelicitous within this context.

As in other cognitive domains, expectations about certain perceptual attributes of upcoming speech (e.g., prosodic characteristics) may provide a compelling explanation for the rapidity with which listeners are capable of evaluating acoustic information with respect to both prosodic phrasing and discourse-based information structure. According to this perspective, listeners develop provisional hypotheses (internal models) about the expected realization of different lexical alternatives (e.g. recently mentioned words). These hypotheses compete to explain the incoming linguistic input. The strength of support for competing lexical alternatives depends on the degree of congruence between the signal and the expected realization of each competing alternative in context. The primary determinant of whether different aspects of language processing interact, then, is the degree to which they contribute to alternative hypotheses about the realization of the same acoustic-phonetic cues, rather than their relative positions within a hierarchy of representations in the language-processing architecture. Because prosody is influenced by a wide range of linguistic factors, listeners' prosodic expectations may incorporate various sources of information from the surrounding context, including preceding prosodic patterning, information structure, and prosodic phrasing. The interaction between context-type and target-word realization observed in

Experiment 2 may have reflected expectations incorporating projected effects of both information status and upcoming prosodic boundaries on the acoustic realization of competing lexical alternatives.

The present results are incompatible with models of spoken language processing that assume that cues to relatively low-level or local aspects of linguistic structure are interpreted prior to cues to higher-level or more global structure. More specifically, such models would predict that the spoken-word recognition system does not rapidly take into account how the interpretation of a prosodic phrasing cue that is relevant to word recognition would be modulated by information structure. Instead, the findings of the present study (and other results demonstrating effects of different types of contextual constraint on spoken-word recognition) suggest that spoken-word recognition involves rapid integration of different types of linguistic information.

More generally, existing models of spoken-word recognition cannot straightforwardly account for our results. Simple effects of acoustic cues such as lengthening and F0 excursion on recognition of words in null or neutral contexts can be explained within the existing framework of several models of spoken-word recognition, such as localist connectionist models like TRACE and Shortlist (McClelland & Elman, 1986; Norris, 1994), by incorporating effects of such cues on the dynamics of lexical activation or on lexical selection criteria (see Salverda et al., 2003, for more discussion). However, distal-context effects observed in this and previous studies suggest that the interpretation of segmental duration as a proximal cue to spoken-word recognition cannot be attributed solely to dynamics of lexical competition elicited by the currently-processed word. Rather, segmental duration is interpreted dynamically with respect to a variety of broader contextual factors.

Effects of prosodic adaptation

The expectation-based perspective is further supported by the adaptation effects suggested by the significant three-way interaction between trial order, context-type, and target-word realization in the regression model predicting target fixation proportions. Numerous studies of phonetic processing have demonstrated that listeners' phonemic categories are continuously updated and adjusted in accordance with distributional properties of recent input (e.g. Norris et al., 2003; Clayards, Tanenhaus, Aslin & Jacobs, 2008; Kraljic & Samuel, 2007; Kraljic, Samuel & Brennan, 2008; Kraljic & Samuel, 2011; Kleinschmidt & Jaeger, 2011). The question of whether and how listeners adapt to prosodic properties of speech has received comparatively little attention. Nonetheless, observations that speakers vary considerably in their realization of intonational categories and in their use of intonational cues to signal linguistic or paralinguistic contrasts suggest a role for prosodic adaptation in language comprehension. Indeed, pragmatic interpretation of sentence-level prosodic contours can be modulated by exposure to a relatively small set of utterances sampled from novel distributions of prosodic contours (Kurumada et al., 2012).

Within-experiment adaptation is a methodological issue that has not received much attention within psycholinguistic domains in which effect sizes tend to be large and stable, such as syntactic processing (though see Farmer et al., 2011; Fine & Jaeger, 2011). For studies of prosody, however, adaptation may present methodological challenges, particularly for

standard experimental designs involving counterbalanced crossed factors (see also Bibyk, Heeren, Gunlogson & Tanenhaus, 2011; Carbary, Brown, Gunlogson, McDonough, Fazlipour, & Tanenhaus, under review). When participants' expectations are incongruent with the prosody of some stimuli, they may quickly learn that prosody is an unreliable cue within the experimental stimuli, given their expectations. As a result, participants may subsequently discount prosodic factors or interpret them differently when evaluating the acoustic realization of incoming words or phrases (Kurumada et al., 2012). Such learning effects are predicted by the expectation-based account: The perceived reliability of the association between information status and acoustic prominence, for example, should be related to the degree to which information structure contributes to listeners' expectations about the acoustic realization of competing lexical alternatives. Generally speaking, when perceptual expectations are incongruent with the actual realization of a word, listeners should update their beliefs about the cues that condition their perceptual expectations, resulting in adaptation that more closely aligns listeners' expectations with the characteristics of the signal in context.

Our results suggest that participants adapted to the reliability of perceived prominence as a referential cue over the course of Experiment 2, in which acoustic characteristics of our stimuli sometimes violated information-structural expectations. This apparently encouraged listeners to discount perceived prominence as a cue to information status over the course of the experiment. Acoustic cues to prosodic boundaries were also occasionally at odds with prosodic expectations in our materials, but the overall reliability of prosodic phrasing cues likely remained relatively high due to the comparatively large number of naturally-produced prosodic boundaries across the experimental materials. Prosodic adaptation would therefore be predicted to affect information-structural effects to a larger degree than prosodic phrasing effects in Experiment 2.

The effects of the information structure manipulation in Experiment 2 are reflected more strongly in the patterns of looks to competitors than to targets. This in itself is unsurprising: Although it is rarely commented upon in the literature, most studies using cohort competitors find stronger effects of experimental manipulations on looks to cohorts compared to looks to targets (e.g. McMurray, Tanenhaus & Aslin, 2002, 2009). Although a detailed explanation of the comparative sensitivity of competitor fixations is beyond the scope of this paper, we present a possible explanation in Footnote 3.

Conclusions

Our results provide evidence that listeners evaluate the acoustic realization of spoken words with respect to both prosodic phrasing and discourse-derived information structure, and that both sources of information have immediate effects on spoken-word recognition. This finding adds to a growing body of work on spoken-word recognition suggesting that information from multiple levels of linguistic representation simultaneously constrains the activation of competing lexical alternatives. The current study extends this line of inquiry to the interpretation of prosodic information at different levels of linguistic representation. Most importantly, these results suggest a crucial role for expectations about aspects of the acoustic-phonetic properties of unfolding utterances in language comprehension.

Acknowledgments

This research was supported by a National Science Foundation graduate research fellowship to MB and National Institutes of Health grants [HD27206] and [DC0005071/HD073890] to MKT. We thank Dana Subik for assistance with participant recruitment and testing; Sarah Bibyk, T. Florian Jaeger, and the audiences at Experimental and Theoretical Advances in Prosody II and CUNY 2011 for constructive discussions; and Delphine Dahan, Joe Toscano and the editors, Duane Watson and Michael Wagner, for cogent reviews which resulted in a much-improved paper.

References

- Aiken, LS.; West, SG. Multiple regression: Testing and interpreting interactions. Newbury Park: Sage; 1991.
- Altmann G, Kamide Y. Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition*. 1999; 73:247–264. [PubMed: 10585516]
- Baayen RH, Davidson DJ, Bates DM. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*. 2008; 59:390–412.
- Barr D. Analyzing ‘visual world’ eyetracking data using multilevel logistic regression. *Journal of Memory and Language*. 2008; 59:457–474.
- Barr DJ, Levy R, Scheepers C, Tily HJ. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*. 2013; 68:255–278.
- Bates, D.; Maechler, M.; Bolker, B. R package, version 0.999375–42. 2011. lme4: Linear mixed-effects models using S4 classes [Computer software].
- Bibyk, S.; Heeren, W.; Gunlogson, C.; Tanenhaus, M. Gotta conundrum, gotta solution. Poster presented at Experimental and Theoretical Advances in Prosody II; Montréal, QC. 2011.
- Boersma, P.; Weenink, D. Praat: doing phonetics by computer [Computer program]. 2010. Version 5.1.25, retrieved 20 January 2010 from <http://www.praat.org>
- Braun B, Chen A. Intonation of ‘now’ in resolving scope ambiguity in English and Dutch. *Journal of Phonetics*. 2010; 38:431–444.
- Brown, M.; Dilley, LC.; Tanenhaus, MK. In: Miyake, N.; Peebles, D.; Cooper, RP., editors. Real-time expectations based on context speech rate can cause words to appear or disappear; Proceedings of the 34th Annual Conference of the Cognitive Science Society; Austin, TX: Cognitive Science Society; 2012a. p. 1374-1379.
- Brown M, Salverda AP, Dilley LC, Tanenhaus MK. Expectations from preceding prosody influence segmentation in online sentence processing. *Psychonomic Bulletin & Review*. 2011; 18:1189–1196. [PubMed: 21968925]
- Brown, M.; Salverda, AP.; Dilley, LC.; Tanenhaus, MK. In: Miyake, N.; Peebles, D.; Cooper, RP., editors. Metrical expectations from preceding prosody influence spoken-word recognition; Proceedings of the 34th Annual Conference of the Cognitive Science Society; Austin, TX: Cognitive Science Society; 2012b. p. 1380-1385.
- Carbary, K.; Brown, M.; Gunlogson, C.; McDonough, JM.; Fazlipour, A.; Tanenhaus, MK. Anticipatory deaccenting in online language comprehension: A phonemic restoration study. (under review)
- Chambers GC, Smyth R. Structural parallelism and discourse coherence: A test of Centering Theory. *Journal of Memory and Language*. 1998; 39:593–608.
- Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*. 2013; 36:181–204. [PubMed: 23663408]
- Clayards M, Tanenhaus MK, Aslin RN, Jacobs RA. Perception of speech reflects optimal use of probabilistic speech cues. *Cognition*. 2008; 108:804–809. [PubMed: 18582855]
- Cole J, Linebaugh G, Munson CM, McMurray B. Unmasking the acoustic effects of vowel-to-vowel coarticulation: A statistical modeling approach. *Journal of Phonetics*. 2010; 38:167–184. [PubMed: 21173864]

- Cooper R. The control of eye fixation by the meaning of spoken language: A new methodology for the real-time investigation of speech perception, memory, and language processing. *Cognitive Psychology*. 1974; 6:84–107.
- Cox, DR. The analysis of binary data. London: Chapman and Hall; 1970.
- Dahan D, Gaskell MG. Temporal dynamics of ambiguity resolution: Evidence from spoken-word recognition. *Journal of Memory and Language*. 2007; 57:483–501. [PubMed: 18071581]
- Dahan D, Tanenhaus MK, Chambers CG. Accent and reference resolution in spoken language comprehension. *Journal of Memory and Language*. 2002; 47:292–314.
- Dahan D, Tanenhaus MK. Continuous mapping from sound to meaning in spoken-language comprehension: Immediate effects of verb-based thematic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2004; 30:498–513.
- Davis MH, Marslen-Wilson WD, Gaskell MG. Leading up the lexical garden path: Segmentation and ambiguity in spoken-word recognition. *Journal of Experimental Psychology: Human Perception and Performance*. 2002; 28:218–244.
- Dilley L, Mattys S, Vinke L. Potent prosody: Comparing the effects of distalprosody, proximal prosody, and semantic context on word segmentation. *Journal of Memory and Language*. 2010; 63:274–294.
- Dilley L, McAuley J. Distal prosodic context affects word segmentation and lexical processing. *Journal of Memory and Language*. 2008; 59:291–311.
- Farmer TA, Brown M, Tanenhaus MK. Prediction, explanation and the role of generative models in language processing [Commentary]. *Behavioral and Brain Sciences*. 2013; 36:211–212. [PubMed: 23663410]
- Farmer TA, Monaghan P, Misyak JB, Christiansen MH. Phonological typicality influences sentence processing in predictive contexts: A reply to Staub et al. (2009). *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2011; 37:1318–1325.
- Fine, AB.; Jaeger, TF. In: Carlson, L.; Hölscher, C.; Shipley, T., editors. Language comprehension is sensitive to changes in the reliability of lexical cues; Proceedings of the 33rd Annual Conference of the Cognitive Science Society; Austin, TX: Cognitive Science Society; 2011. p. 925-930.
- Frank, A.; Salverda, AP.; Pontillo, D.; Jaeger, TF.; Tanenhaus, MK. An event-based framework for analyzing visual world data. (in preparation)
- Frazier L, Carlson K, Clifton C Jr. Prosodic phrasing is central to language comprehension. *Trends in Cognitive Sciences*. 2006; 10(6):244–249. [PubMed: 16651019]
- Kleinschmidt, D.; Jaeger, TF. A Bayesian belief updating model of phonetic recalibration and selective adaptation. *ACL Workshop on Cognitive Modeling and Computational Linguistics*; Portland, OR. 2011.
- Kraljic T, Samuel AG. Perceptual adjustments to multiple speakers. *Journal of Memory and Language*. 2007; 56:1–15.
- Kraljic T, Samuel AG, Brennan SE. First impressions and last resorts: How listeners adjust to speaker variability. *Psychological Science*. 2008; 19:332–338. [PubMed: 18399885]
- Kraljic T, Samuel AG. Perceptual learning evidence for contextually-specific representations. *Cognition*. 2011; 121:459–465. [PubMed: 21939965]
- Kukona A, Fang S, Aicher KA, Chen H, Magnuson JS. The time course of anticipatory constraint integration. *Cognition*. 2011; 119:23–42. [PubMed: 21237450]
- Kurumada, C.; Brown, M.; Tanenhaus, MK. In: Miyake, N.; Peebles, D.; Cooper, RP., editors. Prosody and pragmatic inference: It looks like adaptation; Proceedings of the 34th Annual Conference of the Cognitive Science Society; Austin, TX: Cognitive Science Society; 2012. p. 647-652.
- Ladd, DR. Intonational phonology. 2. 2008. Cambridge Studies in Linguistics
- Ladefoged P, Broadbent DE. Information conveyed by vowels. *Journal of the Acoustical Society of America*. 1957; 29:98–104.
- Levy R. Expectation-based syntactic comprehension. *Cognition*. 2008; 106:1126–1177. [PubMed: 17662975]

- Magnuson JS, Dixon J, Tanenhaus MK, Aslin RN. The dynamics of lexical competition during spoken word recognition. *Cognitive Science*. 2007; 31:133–156. [PubMed: 21635290]
- Magnuson JS, Tanenhaus MK, Aslin RN. Immediate effects of form-class constraints on spoken-word recognition. *Cognition*. 2008; 108:866–873. [PubMed: 18675408]
- McClelland JL, Elman JL. The TRACE model of speech perception. *Cognitive Psychology*. 1986; 18:1–86. [PubMed: 3753912]
- McMurray B, Jongman A. What information is necessary for speech categorization? Harnessing variability in the speech signal by integrating cues computed relative to expectations. *Psychological Review*. 2011; 118:219–246. [PubMed: 21417542]
- McMurray B, Tanenhaus MK, Aslin RN. Gradient effects of within-category phonetic variation on lexical access. *Cognition*. 2002; 86:B33–42. [PubMed: 12435537]
- McMurray B, Tanenhaus MK, Aslin RN. Within-category VOT affects recovery from “lexical” garden paths: Evidence against phoneme-level inhibition. *Journal of Memory and Language*. 2009; 60(1): 65–91. [PubMed: 20046217]
- Moulines E, Charpentier F. Pitch-synchronous waveform processing techniques for text-to-speech synthesis using diphones. *Speech Communication*. 1990; 9:453–467.
- Norris DG, McQueen JM, Cutler A. Perceptual learning in speech. *Cognitive Psychology*. 2003; 47:204–238. [PubMed: 12948518]
- Norris D. Shortlist: A connectionist model of continuous speech recognition. *Cognition*. 1994; 52:89–234.
- Pierrehumbert, J.; Hirschberg, J. The meaning of intonational contours in the interpretation of discourse. In: Cohen, P.; Morgan, J.; Pollack, M., editors. *Intentions in Communication*. MIT Press; Cambridge MA: 1990. p. 271–311.
- Pirog Revill K, Tanenhaus MK, Aslin RN. Context and spoken word recognition in a novel lexicon. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 2008; 34:1207–1223.
- R Development Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing; 2012. <http://www.R-project.org>
- Salverda AP, Brown M, Tanenhaus MK. A goal-based perspective on eye-movements in visual-world studies. *Acta Psychologica*. 2011; 137(2):172–180. [PubMed: 21067708]
- Salverda AP, Dahan D, McQueen J. The role of prosodic boundaries in the resolution of lexical embedding in speech comprehension. *Cognition*. 2003; 90:51–89. [PubMed: 14597270]
- Salverda AP, Dahan D, Tanenhaus MK, Crosswhite K, Masharov M, McDonough J. Effects of prosodically modulated sub-phonetic variation on lexical competition. *Cognition*. 2007; 105:466–476. [PubMed: 17141751]
- Salverda AP, Kleinschmidt D, Tanenhaus MK. Immediate effects of anticipatory coarticulation in spoken-word recognition. *Journal of Memory and Language*. (in press).
- Salverda AP, Tanenhaus MK. Tracking the time course of orthographic information in spoken-word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2010; 36:1108–1117.
- Swinney D. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*. 1979; 18:645–660.
- Tanenhaus MK, Leiman JM, Seidenberg MS. Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*. 1979; 18:427–441.
- Tanenhaus M, Spivey-Knowlton M, Eberhard K, Sedivy J. Integration of visual and linguistic information in spoken language comprehension. *Science*. 1995; 268:1632–1634. [PubMed: 7777863]
- Terken J, Hirschberg J. Deaccentuation of words representing ‘given’ information: Effects of persistence of grammatical function and surface position. *Language and Speech*. 1994; 37:125–145.

Appendix 1. Analysis of distractor-to-target and distractor-to-competitor saccades as a function of target-word realization

Here, in response to concerns raised by Michael Wagner, we consider an alternative explanation for the effect of lengthening on competitor fixations that we observed in Experiment 1 that does not appeal to effects of prosodic phrasing on lexical activation. Instead, this account considers how lengthening of the initial syllable of the target word affects the temporal distribution of segmental information, and how this may affect the lexical interpretation of the speech signal. We will refer to this account as the point-of-disambiguation account.

When the first syllable of the target word is lengthened, there is a longer period of time during which the speech signal is phonemically consistent with both target and competitor. Listeners therefore have more time to process the initial syllable of a lengthened target word, and more time to launch signal-driven saccades to both target and competitor pictures (resulting in increasing proportions of looks to both pictures) as the proportion of looks to the distractors decreases. This could result in a higher overall proportion of competitor (and target) fixations, with a later peak, when the initial syllable is lengthened. Prior to the availability of phonetic information from the second syllable, the proportion of fixations to the target and competitor picture should increase at approximately the same rate. Once listeners begin processing the second syllable of the target word, they start to favor the target over the competitor. This results in more fixations to the target. Importantly, the increase in target fixations (relative to competitor fixations) due to processing the second syllable of the target word would start earlier when the initial syllable of the target word is short compared to when that syllable is long. Therefore, in Experiment 1, when proportions of fixations to target and competitor are analyzed across a relatively broad temporal window (i.e., a window extending beyond the period of time during which listeners are presented with the initial syllable of the target word), the point-of-disambiguation account predicts (a) a higher overall proportion of competitor fixations and delayed peak in the long condition, and (b) an earlier rise in and/or higher overall proportion of fixations to the target in the short condition—from the point in time where the second syllable of the target word is processed. The statistical tests that we report in the main body of the text focus on a broad time window, and therefore do not necessarily distinguish between the predictions of the point-of-disambiguation account and the prosodic-boundary account.

However, the two accounts differ in the predictions that they make regarding listeners' eye movements prior to the point of disambiguation. According to the point-of-disambiguation account, as listeners process the initial syllable of the target word, they should launch signal-driven saccades from distractor pictures to competitor and target pictures (which are both phonemically consistent with the input during this window), resulting in decreases in proportions of distractor fixations and increases in proportions of both competitor and target fixations. Therefore, the point-of-disambiguation account predicts that lengthening will increase fixations to competitor and target pictures to a similar extent. In contrast, the prosodic-boundary account predicts that lengthening will favor a short-word interpretation of the syllable (e.g. *pump*) over a long-word interpretation (e.g. *pumpkin*), because listeners

anticipate an upcoming prosodic boundary on the basis of cues to duration contained within the initial syllable (Salverda et al., 2003), and such a boundary is consistent with a short-word interpretation but not with a long-word interpretation. Thus, only the prosodic-boundary account predicts that lengthening will result in more fixations to the competitor relative to the target during the processing of the initial syllable of the target word.

To distinguish between predictions of the point-of-disambiguation and prosodic-boundary accounts, we examined saccades from signal-irrelevant screen positions (i.e., any location on the screen other than the target or competitor picture) to the target or competitor during the processing of the initial syllable of the target word (see Table A1). These data were assessed for each individual stimulus, taking into account a 200 ms delay for programming and executing an eye movement. Because the number of such events was relatively low, we combined the data from both experiments, collapsing across discourse context in Experiment 2. There were substantially more trials with a saccade from a distractor to the competitor when the initial syllable of the target word was long than when it was short. However, the opposite pattern, although numerically small, was found for saccades from a distractor to the target. A Chi-square test showed that lengthening of the initial syllable of the target word had a different effect on the likelihood of competitor vs. target saccades ($\chi^2=6.6, p<.05$).

To address the possibility that effects of lengthening on target vs. competitor fixations were merely skewed in some way by the discourse manipulation in Experiment 2, we verified that differences in distractor-to-competitor vs. distractor-to-target saccades as a function of target-word realization were observed when the analysis included only events from SR-contexts ($\chi^2=6.0, p<.05$) and only events from DR-contexts ($\chi^2=3.2, p<.10$). In addition, data from Experiment 1 alone show the same numerical pattern as the combined dataset. Statistical analysis of these data failed to reach significance, but this was likely due to insufficient power.

Taken together, these exploratory analyses suggest that the increase in competitor fixations observed in response to lengthened stimuli in the current study is not simply due to an effect of increased processing time on lexical activation, but instead results at least in part from the interpretation of segmental lengthening as a cue to prosodic phrasing.

	Saccade endpoint	
	target	competitor
Experiment 1+Experiment 2		
short	164	112
long	161	170
Experiment 1+Experiment 2 (SR-context)		
short	107	67
long	108	114
Experiment 1+Experiment 2 (DR-context)		
short	96	80

	Saccade endpoint	
	target	competitor
long	88	109
Experiment 1 alone		
short	41	38
long	41	58

Figure 1.
Example of visual display at the start of a trial in Experiment 1.

Figure 2. Waveforms and F0 of a resynthesized critical instruction containing (a) short and (b) long target words. The F0 of the original instruction is depicted in gray across the acoustically altered syllables.

Figure 3. Proportions of fixations to the target, competitor, and averaged distractors for short-target and long-target items in Experiment 1, relative to the onset of the target word.

Figure 4. Proportions of fixations to the target, competitor, and averaged distractors for short-target and long-target items in same-role contexts in Experiment 2, relative to the onset of the target word.

Figure 5.

Proportions of fixations to the target, competitor, and averaged distractors for short-target and long-target items in different-role contexts in Experiment 2, relative to the onset of the target word.

Figure 6.

Proportions of fixations to target, competitor and distractor pictures for short-target and long-target items in Experiment 2 for: (a) the first third of the trials in same-role contexts; (b) the first third of the trials in different-role contexts; (c) the last third of the trials in same-role contexts; (d) the last third of the trials in different-role contexts.

Table 1

Target and competitor pictures displayed in experimental trials.

target	competitor
antlers	ant
beaker	bee
boulder	bowl
candy	can
captain	cap
carpet	car
dolphin	doll
hamster	ham
leaflet	leaf
nectarine	neck
panda	pan
pirate	pie
pumpkin	pump
reindeer	rain
soldier	sole
spider	spy
taxi	tack
toaster	toe
tractor	track
welder	well

Table 2

Parameters and formula of final linear regression model of logit-transformed competitor fixations, averaged between 200–1000 ms following target-word onset.

	β	SE	t	p
intercept	−0.781	0.023	−33.53	<0.0001
context=DR	−0.062	0.040	−1.53	n.s.
target=long	0.043	0.022	1.90	<0.10
trial-number	−0.035	0.011	−3.11	<0.005
context:target	−0.129	0.044	−2.93	<0.005

Formula:

```
lmer(CFIX~context+target+context:target+trialnumber+(1+context+target+context:target|subj)+(
1+context+target+context:target|item), data=data)
```


Table 3

Parameters and formula of final linear regression model of logit-transformed target fixations averaged between 200–1000 ms after target-word onset.

	β	SE	t	p
intercept	−0.173	0.040	−4.30	< 0.0001
context=DR	0.135	0.051	2.64	< 0.01
target=long	−0.081	0.033	−2.46	< 0.05
trial-number	0.019	0.020	0.93	n.s.
context:target	0.027	0.068	0.40	n.s.
context:trial-number	−0.023	0.042	−0.55	n.s.
target:trial-number	−0.037	0.032	−1.15	n.s.
context:target:trial-number	−0.169	0.062	−2.72	< 0.01

Formula:

```
lmer(TFIX~context+target+context:target+trialnumber+context:trialnumber+target:trialnumber
+target:trialnumber:context+(1+context+target+trialnumber+context:target+context:trialnumber
+target:trialnumber|subj)+(1+context+target+trialnumber|item), data=data)
```

Table 4

Parameters and formula of linear regression model of averaged logit-transformed target fixations with trial number centered one standard deviation below the mean.

	β	SE	t	p
intercept	-0.192	0.042	-4.58	<0.0001
context=DR	0.162	0.054	2.98	<0.005
target=long	-0.044	0.045	-0.99	n.s.
trial-number	0.020	0.020	1.00	n.s.
context:target	0.197	0.091	2.16	<0.05
context:trial-number	-0.028	0.041	-0.68	n.s.
target:trial-number	-0.038	0.032	1.18	n.s.
context:target:trial-number	-0.175	0.065	-2.69	<0.01

Formula:

```
lmer(TFIX~context*target*lowtrialnumber+(1+context+target+lowtrialnumber+context:target+
context:lowtrialnumber+target:lowtrialnumber+context:target:lowtrialnumber|subj)+(1+context
+target+lowtrialnumber|item), data=data)
```

Table 5

Parameters and formula of linear regression model of averaged logit-transformed target fixations with trial number centered one standard deviation above the mean. Because the model failed to converge with random effects by items in addition to by participants, we ran an additional model containing only random effects by items (intercepts and slopes for context-type, target-word realization, trial number, and the interaction between context-type and target-word realization). Effects were qualitatively similar between the two models, though the interaction between context-type and target-word realization failed to reach significance, marginal or otherwise ($\beta = -0.114$, $t = -1.163$).

	β	SE	t	p
intercept	-0.170	0.040	-4.27	<0.0001
context=DR	0.112	0.053	2.12	<0.05
target=long	-0.112	0.046	-2.44	<0.5
trial-number	0.006	0.019	0.299	n.s.
context:target	-0.167	0.095	-1.75	<0.10
context:trial-number	-0.025	0.038	-0.66	n.s.
target:trial-number	-0.028	0.032	-0.86	n.s.
context:target:trial-number	-0.186	0.067	-2.76	<0.01

Formula:

```
lmer(TFIX~context*target*hightrialnumber+(1+context+target+hightrialnumber+context:target
+context:hightrialnumber+target:hightrialnumber+context:target:hightrialnumber|subj), data=data)
```