Published in final edited form as:

Infancy. 2019 ; 24(3): 300–317. doi:10.1111/infa.12280.

A cross-linguistic examination of toddlers' interpretation of vowel duration

Daniel Swingley,

Department of Psychology, University of Pennsylvania, 425 S. University Ave., Philadelphia, PA 19104 USA, swingley@psych.upenn.edu, 215 898 0334

Suzanne Van der Feest

Department of Linguistics, University of Texas at Austin, 305 E. 23rd Street, Austin TX 78712 USA, suzanne@austin.utexas.edu

Abstract

Languages differ in their phonological use of vowel duration. For the child, learning how duration contributes to lexical contrast is complicated because segmental duration is implicated in many different linguistic distinctions. Using a language-guided looking task, we measured English and Dutch 21-month-olds' recognition of familiar words with normal or manipulated vowel durations. Dutch but not English learners were affected by duration changes, even though distributions of short and long vowels in both languages are similar, and English uses vowel duration as a cue to (for example) consonant coda voicing. Additionally, we found that word recognition in Dutch toddlers was affected by shortening but not lengthening of vowels, matching an asymmetry also found in Dutch adults. Considering the subtlety of the crosslinguistic difference in the input, and the complexity of duration as a phonetic feature, our results suggest a strong capacity for phonetic analysis in children before their second birthday.

Keywords

language acquisition; word recognition; speech perception; phonetics; phonology

All spoken languages have conventions governing how sounds relate to meaning. Because the conventions vary from language to language, children must learn them to acquire their language. Most of what is known about this early phonological learning concerns the development of a finely tuned sensitivity to language-specific distinctions between similar-sounding consonants and vowels, a development manifested in its precocious effects on speech-sound discrimination between the ages of 6 and 12 months. The classic outcome of dozens of speech discrimination experiments is that infants learn to classify together speech sounds that their language treats as phonologically identical, while keeping separate the sounds that their language uses to distinguish words (e.g., Werker, Yeung, & Yoshida, 2012). This adaptation might be viewed loosely as providing a language-specific analog-to-digital

Correspondence should be directed to Daniel Swingley at 425 S. University Ave., Department of Psychology, University of Pennsylvania, Philadelphia PA 19104; phone, 215-898-0334; swingley@psych.upenn.edu.

The authors declare the absence of any conflict of interest.

conversion: sound goes in, and categories come out, ready to use in distinguishing words in the growing lexicon.

A problem with this caricature is that it neglects one of the most interesting and complex features of phonological systems, namely that phonological structure is found at multiple levels, not just the phoneme level. These structures use some of the same phonetic raw materials to multiple ends. Learning a language's consonants and vowels cannot be done in isolation, because some of the features that distinguish them also vary for indicating linguistic events like prosodic boundaries or lexical stress, and marking aspects of speaker attitude and identity. Thus, the speech-sound classification problem is embedded in a larger problem of *attribution*: working backward from the signal toward the various linguistic features that led to that signal.

One way to study the development of this language-specific attribution process is to compare how children treat a given phonetic property when learning languages that use that property differently. For example, vowel duration is lexically contrastive in Japanese, but is not lexically contrastive in Spanish (e.g., Hirata, 2004; Mendoza et al., 2003). As a result, the Spanish-learning baby should learn to attribute vowel-duration variation to variables like speaking rate, lexical stress, and the termination of intonational phrases, but should not attribute it to vowel identity to the same degree that a Japanese-learning baby should.

At present, the development of this attribution process can only be roughly characterized. In the case of linguistic pitch features, English-learning children transition from a willingness to treat tone as lexically distinctive at 14 months, to resisting this interpretation at 19 months (Hay, Graf Estes, Wang, & Saffran, 2015), and appearing to encode pitch-independent representations of words by 30 months (Quam & Swingley, 2010). But children learning Mandarin, a tone language, maintain tone as an important lexical feature throughout this developmental period (Singh, Tam, Chan, & Golinkoff, 2014; Singh, Goh, & Wewalaarachchi, 2015). Although one might argue that the mechanism behind the English learners' apparent disregard for tone variation is an acquired difficulty in discriminating tones at all (e.g., Mattock, Molnar, Polka, & Burnham, 2008), this was not the case in Hay et al., who showed pure discrimination of tone contrasts at 19 months, or in Quam & Swingley (2010) who tested huge prosodic pitch sweeps (also see e.g. Chen & Kager, 2016, who found some maintenance of tone discrimination; and Liu & Kager, 2014). By some procedure that is not yet understood, children recognize that pitch variation, which is ubiquitous in all languages, in their language serves a particular set of roles that may or may not include lexical contrast.

The current study asks similar questions of vowel duration variation. As mentioned above, vowel duration has many different linguistic functions. The languages that concern us here, English and Dutch, are phonologically similar in several respects: both allow consonant clusters as onsets and as codas; both mark lexical stress with elevated pitch and elongated duration; both commonly reduce vowels to schwa in unstressed syllables (though English more so than Dutch; Warner & Cutler, 2017); and in both languages the child-directed vocabulary is dominated by monosyllables and trochees (e.g. Finegan, 1987; Kooij, 1987). However, Dutch listeners seem to weight vowel duration more heavily than English listeners

do when judging vowel identity, as described below. Phonological feature theories typically describe Dutch but not English as having phonologically contrastive vowel duration, even if duration is not the only acoustic cue signaling vowel identity (e.g. Booij, 1999; Gussenhoven, 1999). The similarities between the two languages, coupled with sometimes divergent way listeners treat vowel duration, make this comparison interesting for evaluating the nature of children's learning abilities. Here, we measured Dutch and English 21-month-olds' encoding of vowel duration in familiar words, using a word-recognition task.

Both English vowels and Dutch vowels vary in their typical or "intrinsic" duration. In English, tense vowels in focal positions are for the most part longer than lax vowels. For example, the tense vowel [i] is typically 25-30% longer than the lax vowel [ɪ] in citation form (e.g., Hillenbrand, Getty, Clark, & Wheeler, 1995). English speakers without special linguistic training may informally refer to "long vowels" and "short vowels," though this distinction is not clearly represented in the orthography. When listeners are asked to identify spectrally similar vowels under controlled conditions, vowel duration can have a strong effect on vowel identification (Hillenbrand, Clark, & Houde, 2000). In addition, vowel duration in English is one of the cues signaling coda consonant voicing, with longer vowel durations before voiced consonants than before voiceless consonants (e.g., Crowther & Mann, 1992; House & Fairbanks, 1953; Raphael, 2008). Thus, English listeners hearing a lengthened *bit* might plausibly interpret the longer duration as signaling the consonant /d/ rather than /t/, or the vowel /i/ rather than /ɪ/. (This interpretation is not available in Dutch, where voicing contrasts in codas are neutralized.)

Dutch vowels are reported to have more divergent intrinsic durations. A set of short vowels including [ɪ, e, a, u] may be only half as long as another set including [ɑ, o, ɛ], again in citation form (e.g., Adank, van Hout, & Smits, 2004; Booij, 1999). These differences are substantially larger than those recorded under similar conditions in the Hillenbrand et al. work on English, and as mentioned above, in Dutch the "long" vowels have been argued to have a different featural representation for duration than "short" vowels. Also, Dutch orthography clearly marks durationally opposed tense and lax vowels (e.g., *beek* [be:k], 'stream', vs. *bek*, [bek], 'mouth/beak'), using rules with which Dutch adults are conversant. Like English, Dutch is not generally considered to be a language that has durational oppositions in its vowel system, unlike languages such as Japanese, Thai, or Hungarian, which have phonologically distinct long and short versions of several vowels that are similar in quality. Still, as with English, Dutch adults' judgments of vowel identity can be affected by manipulations of vowel duration (e.g., Chládková, Escudero, & Lipski, 2015; Tillman, Benders, Brown, & van Ravenzwaaij, 2017). In particular, Nootboom and Doodeman (1980) found that for the word pair /tak/-/ta:k/, Dutch listeners perceived artificially shortened [a:] as /ɑ/, but not lengthened [ɑ] as /a:/. This asymmetry was also revealed in an electrophysiological study of Dutch adults by Chládková et al. (2015).

Two prior studies have directly compared Dutch and English listeners' interpretation of vowel duration, one with adults and one with children. Van der Feest and Swingley (2011) asked adult native speakers of each language to transcribe a set of spoken words whose vowel had been lengthened or shortened (or unaltered). In the Dutch materials, long vowels were shortened and short vowels lengthened; in the English materials, vowels were

lengthened if they preceded voiceless codas and shortened if they preceded voiced codas, so that listeners could also reveal a sensitivity to vowel length were they to interpret it as a cue to coda voicing. In this study, Dutch listeners were more strongly affected by the durational manipulation than English listeners were, substituting one vowel for another that more closely matched the manipulated token's duration. Replicating Nooteboom and Doodeman (1980), vowel shortening changed Dutch listeners' interpretation more often than lengthening, in particular for the pairs [a:--α] and [i:--i]. Listeners also sometimes exchanged [ɔ] and [o:] when durations were altered. English listeners made errors too, though significantly less often. Errors included several instances of shortened [æ] interpreted as [ɛ], and sometimes the reverse. In some cases vowel duration manipulation led English listeners to re-interpret the voicing value of the coda consonant, e.g. hearing *leave* as *leaf* when the vowel was shortened. In summary, listeners' specific reactions to vowel duration changes were mostly in line with prior literature, but Dutch listeners showed a stronger influence than English listeners did.

The other study to have compared Dutch and English listeners' interpretation of vowel duration differences tested 18-month-olds (Dietrich, Werker, & Swingley, 2007). Children were taught two novel words, using a standardized habituation procedure, and tested to see if they kept the words distinct (Werker, Cohen, Lloyd, Casasola, & Stager, 1998). The words differed only in vowel duration, with one word having been naturally recorded and the other a digital manipulation. In one experiment, the source materials were Dutch (the nonce words [tam] and [tɑ:m]); in another, the materials were Canadian English ([tæm] and [tæ:m]). Dutch-learning toddlers, but not English-learning toddlers, were successful in this task, whether the stimuli were originally Dutch or English. Both groups performed equally well on native vowel contrasts (English, [æ,ɛ]; Dutch, [α,ɛ]). Thus, by 18 months, Dutch and English learners seem to have drawn different conclusions about vowel duration, at least for low vowels (see also Mugitani et al., 2009, for related work comparing English and Japanese; Ko, Soderstrom, & Morgan, 2009; and Heeren, 2006).

The fact that English and Dutch toddlers interpreted vowel duration differently is striking given how similar the surface distributions of English and Dutch vowel durations seem to be in infant-directed speech. One might expect that Dutch speech would exhibit bimodal durational distributions, i.e. that measurement of a corpus of Dutch vowels would reveal a mass of short instances (the short vowels) and a separate mass of long instances (the long vowels). Instead, not only do the short and long Dutch vowels fail to form two modes, they also do not appear to be markedly more bimodal than English vowels are, even when considering infant-directed speech (Swingley, in press). A similar observation has been made for Japanese infant directed speech, where there is no straightforward bimodal distribution of vowel duration, despite the fact that Japanese uses vowel duration phonemically (Bion, Miyazawa, Kikuchi & Mazuka, 2013). Thus, if children learn their language's phonological system by attending to the statistical properties of the speech they hear (which seems incontrovertible), the data reviewed so far suggest that infants must do more than simply compute modes in the distribution of vowel durations they experience.

The present study was designed to explore this phenomenon experimentally, with slightly older English- and Dutch-learning children. Our goals were (a) to see if the cross-linguistic

difference in Dutch and English vowel duration interpretation would extend to real words young children know from their ordinary experience with language, as opposed to novel words known only from a short-term training experience; (b) to test a somewhat broader range of vowels and phonetic contexts; and (c) to evaluate whether Dutch toddlers would show a similar asymmetry to that shown by Dutch adults, being more strongly affected by vowel shortening than vowel lengthening (cf. Nooteboom & Doodeman, 1980).

To do this, we used a “mispronunciation” design in the context of a language-guided looking procedure (Swingley & Aslin, 2000; Swingley, 2012). This method is based on the premise that toddlers fixate named pictures less when the name is pronounced in a non-canonical way, provided that they detect the deviation and consider it as such; thus, picture fixation behavior can be used to measure children’s interpretation of phonetic variation. In this study, Dutch-learning and English-learning children viewed pairs of pictures on a screen, and one of the pictures was named in a sentence. That name was produced with either a normal-duration vowel, or an unusually long or short vowel (the “mispronunciation”). In Experiment 1 we tested Dutch-learning toddlers on Dutch language materials, and in Experiment 2 we tested American English-learning toddlers on American English materials.

Experiment 1: Dutch children

Methods

Participants.—Participants were 24 monolingual Dutch children (13 girls), with a mean age of 21;19 (months;days; range, 21;03 to 22;08). An additional 3 children were tested but not included in the final sample because the child refused to watch the display at all, or did not complete at least 16 of the 24 trials. A trial was considered as “completed” if the child looked at the pictures for at least 16 video frames of the 42-frame test window of each trial. The study was conducted at the Baby Research Center, Radboud University, Nijmegen. Children were recruited using contact information provided by the city of Nijmegen, and excluded if reported by a parent to be hearing less than 75% Dutch speech in a typical week. Socioeconomic data from the parents were not obtained. No parents reported any birth complications or abnormalities of hearing or vision when queried at the intake interview. This study (and Expt. 2) were conducted following guidelines set by the Declaration of Helsinki, with informed consent obtained from a parent or guardian before data collection. All procedures were approved by the IRBs at Radboud Universiteit (Expt. 1) and the University of Pennsylvania (Expt. 2).

Visual stimuli.—The visual stimuli were photographs of objects on a light background, presented side by side (about 20 cm apart) on a 192-cm diagonal LCD video projection screen. Paired pictures were of similar sizes, averaging about 23 cm wide. In an effort to maintain children’s attention throughout the procedure, every 5 trials a “filler” animation was added, consisting of few-second-long clips showing events like a fish swimming across the screen or a duck moving around. Images in these animations did not overlap with any of the test items.

Auditory stimuli.—The words used as targets were all words we expected the 21-month-olds to know, based on vocabulary-inventory data collected in prior studies of somewhat

younger children drawn from the same population (e.g., Swingley, 2007). This prior data included 259 Dutch 16–19-month-olds. The average proportion of parents indicating that their child knew or said our six words was 68%, with a minimum of 41% for *boot* ('boat') and all the others above 62%, with a maximum of 98% for *bal* ('ball'). In the lab, parents were queried informally about each test word, and most parents said their children knew all six (one parent reported that she was unsure if her child knew the word *boot*).

The speech stimuli were digitally recorded in a sound-attenuated room by a female native speaker of Dutch (SvdF). Her speaking rate was slow and in an "infant-directed" register. The sentences were of the form *Kijk naar de/het [target word]*. ('Look at the... '). After a pause of 750 msec, a brief and encouraging second phrase was presented, like *Leuk, hè?* ('Nice, right?') or *Kun je hem vinden?* ('Can you find it?').

Target words presented in the lengthened or shortened conditions were derived from the same recorded tokens as those presented in the canonical-duration condition. To determine how much to stretch or compress the vowels, in the interest of operating only within a normal range of duration variation, we first recorded a set of nouns in the Dutch "Look at the... " sentence frame, including words of a variety of phonologically long and short vowels. The durations of these natural tokens were measured, and these were used to guide our target durations for the manipulated vowels. We found that long vowels were about 1.8 times longer than short vowels in this context, and so we avoided exceeding this ratio in manipulating the stimulus tokens. Although vowel duration is affected by many different factors and is hence quite variable, a 1.8x difference in short and long vowels is in line with the durational difference found in phonetic studies of Dutch (e.g., Adank, van Hout, & Smits, 2004; Nootboom & Doodeman, 1980; Rietveld, Kerkhoff, & Gussenhoven, 2004). In our stimuli, durations were manipulated at a ratio of about 1:1.6 (shorter:longer); for some tokens we found that larger manipulations made the words sound artificial. This is similar to the duration ratio that was used with adult listeners in Van der Feest & Swingley (2011). In order to make sure pitch patterns of altered stimuli matched the words with original vowel duration, pitches were manipulated in Praat (Boersma & Weenink, 2005) using its implementation of the PSOLA algorithm, where needed. Additional editing was done to remove minor vocal sounds like pops. The full set of stimuli is described in Table 1.

Apparatus and procedure.—The experiment consisted of 24 trials. On each trial, the two images were presented in silence for 2.5 seconds, and then the auditory stimulus was initiated. Trials ended 4s after the onset of the auditory stimulus. Between trials, an animation of a blinking white star shape centered on a black background was displayed for 500 ms.

Children saw each picture pair eight times. Each picture of each pair was named four times, twice with the canonical pronunciation and twice with the manipulated-duration vowel. Each target word was presented once in each condition in each half of the experiment, and each picture was the target on canonical-pronunciation and deviant-pronunciation trials equally often on the right and left halves of the screen. The ball was paired with the boat, the bed with the horse, and the bear with the doll. Four stimulus orders were created, by creating one base order and exchanging the left and right sides to produce the second, and then inverting

the trial orders of these to produce the third and fourth. Due to an error, eleven children were assigned to the first order at the expense of the fourth.

The experiment was conducted in a sound-insulated room housing a three-sided booth 2 m tall, 1.3 m wide, and 1.2 m deep, with the screen forming the back end of the booth. The parent sat on a desk chair with the child on her lap, facing the screen at a distance of about 1 m. The speech stimuli were produced at a level of about 70 dB from the speakers of the video screen. Children were videotaped onto digital videocassettes using a low-light videocamera placed about 15 cm beneath the screen. Room lighting was dim. Before the procedure began, parents were instructed to refrain from speaking and to close their eyes and orient their faces downward during the procedure; in this way, parents were blind to target side (compliance was monitored and, rarely, enforced, by the experimenter). Parents also wore Sennheiser Noisegard headphones that played a masking stimulus made of music mixed with repetitions of the spoken stimulus materials.

Parents were given a vocabulary questionnaire to complete before their session (the Netherlands-Communicative Developmental Inventory; Zink & Lejaegere, 2002). Because 7 of the parents failed to complete to return any vocabulary data (almost 30%), the CDI scores were not considered further in Experiment 1.

Coding.—Videorecordings of the child's face were digitized into mpeg format. Several trained coders used custom software (Hollich, 2008) to step frame by frame through the video, noting for each 40 ms frame whether the child was looking at the right picture, the left picture, or in any other location. The beginning and end of each test trial was identifiable on the video by a change in background light. The coder was blind to target side and test order. This timestream of eye movements was sorted into a separate timestream of stimulus events (trial onset and acoustic target-word onset) so that their relative timing could be computed. Coders were drawn from a pool trained to an assessed reliability of 96.7% agreement (mean Cohen's kappa .93).

Trials on which children's fixation to the target and distracter summed to fewer than 16 of the 41 video frames between 360 and 2000 ms following target-word onset were discarded (Swingley, 2016). This removed 12.7% of trials. On about 2/3 of these trials children did not fixate either picture at all.

Results, Experiment 1

In evaluating whether changes to vowel duration affect word recognition, it must first be demonstrated that the children recognized the target words. For this purpose, each child's mean level of performance on each word, collapsing over conditions, was compared against chance levels. In similar experiments children this age often fixated named pictures above chance levels whether the words were slightly mispronounced or not (e.g., Swingley, 2009). Because children appeared to prefer looking at some pictures over others, for this analysis we used fixation to target arithmetically corrected for picture preferences.

Picture preferences were computed for each child by taking the set of trials involving a given pair of pictures, and over that set, computing the proportion of time the child fixated each of

the pictures on each trial, during the period of the trial before the target word was spoken. These proportions can be considered a baseline from which target looking is expected to increase when children recognize the word. This method is imperfect (Swingley, 2012), but subtractive correction procedures of this sort are widely applied because they help correct for potential imbalances in individual children's interest in the pictures. Target looking was assessed over the window extending from 360 ms after the acoustic onset of the target word, to 2000 ms after this onset. This window is standard (Swingley, 2012). Before about 360 ms children's eye movements are unlikely to be responses to the spoken word, and after about 2000 ms children often drift away from the target picture.

By this measure (and also by raw percent-to-target), the item *boot* (boat) was not reliably recognized by children; mean salience-corrected target looking for *boot* was -8.3% with a standard deviation of 20.8 (and raw percent-to-target 50.1% , $sd\ 23.4$). Salience-corrected target looking for this item when correctly pronounced was -9.9 , and when mispronounced -7.8 . This was a surprise to us, and we suspect the fault lies with the particular boat picture we used. Given that there was no indication that children understood this word, this item was eliminated from the set. The other five words were all numerically above (salience-corrected) zero and also above 50.0 on the raw percent-to-target measures. As a whole, the remaining words were recognized by children: by items, mean 12.2% ($sd\ 7.6$), $t(4) = 3.57$, $p(\text{one-tailed}) = 0.012$; by subjects, mean 11.9% ($sd\ 9.7$), $t(23) = 5.99$, $p(\text{one-tailed}) < 0.0001$.

Turning to the question of whether children were affected by the vowel duration manipulation, we considered a slightly later window of analysis compared to mispronunciation detection studies where changes in the onsets of words were made. For children to determine that the vowel had been lengthened or shortened, they would need to have heard at least some portion of the vowel and detected that it was anomalous. The timing of children's eye movement responses reflects the timing of the phonetic information that permits evaluation of the fixated picture (Mahr, McMillan, Saffran, Ellis Weismer, & Edwards, 2015; Swingley, 2009; Swingley, Pinto, & Fernald, 1999). For example, in Swingley (2009), children's rejection of the target picture when it was named with a mispronounced word came about 750 ms after the mispronunciation occurred in the speech signal, whether that was at word onset or word offset. We should expect, then, that mispronunciation effects would be delayed relative to word onset when the mispronunciations themselves are delayed too. Here, it is not clear when children might in principle come to consider a vowel too long or too short, but it is certainly not at word onset. We conservatively began our analysis window at 560 ms, i.e. assuming a 200 ms delay beyond the usual 360 ms time to program an eye movement. (In fact differences started to emerge rather later, perhaps because detecting an anomalous vowel duration is different from detecting an incorrect consonant, for example; we return to this in the discussion. The outcomes of the analyses are essentially unchanged under the usual 360–2000 window.)

Figure 1 summarizes children's performance on the originally long words (*paard*, *beer*; left panel) and the originally short words (*bal*, *bed*, *pop*; right panel). The taller rectangles show mean target looking on correct-pronunciation trials (left bar) and mispronunciation trials (right bar). Individual children's means on these measures are indicated by the connected

pairs of dots. The shorter, darker bars below 0.20 indicate target fixation with picture salience subtracted out (this measure was not used in the analyses, but shown here for comparison with other studies). As the plot shows, children overall looked longer to the pictures of the originally long words than the originally short words; and for the originally long words, target fixation was substantially greater when the words were correctly pronounced with their canonical vowel durations. This mispronunciation effect was present for *paard* (mean CP, 80.9% target fixation; MP, 67.2%) and *beer* (mean CP, 62.6%, mean MP, 54.1%).

The condition differences were analyzed using a binomial regression model predicting target fixation trial by trial. Predicting fixation trial by trial permits entering predictor variables that vary at the trial or item level, such as the proportion of time children fixated a given image before it was named on a given trial. A binomial model was used because looking time proportions are strongly non-normal, ruling out ordinary linear regression. Because standard logistic models may yield overly conservative standard error estimates when used to model non-binary outcomes, standard errors were estimated using a bootstrapping technique following the procedure described by Humphrey & Swingley (2018), based on 600,000 runs.

The outcome measure for a given trial was the proportion of target fixation in the 560 to 2000 ms window. Predictors were Subject (a random effect), Salience (over the trials on which a given target picture appeared, how much time a given child fixated that picture in the time before the spoken target word began), Condition (CP, MP), Length (whether the vowel is canonically short or long), and the interaction between Condition and Length. (Target word was not included as a random effect because there were only three or two words per cell of the condition X length interaction.) The model specification in R was

$$\text{target.looking} \sim \text{condition} * \text{original.length} + \text{salience} + (1 | \text{subject}).$$

The results of the model are given in Table 2. The results are all as would be expected given Figure 1 and as described above. The exponentiated coefficients give an estimate of the multiplicative change in the odds of fixating the target relative to the distracter given (for example) a change from the reference CP condition to the MP condition. The 0.563 odds ratio for the Condition effect means that in the long-vowel condition, if none of the other effects was present, we would expect that a child looking at the target 75% of the time given a correct pronunciation (odds of 3, or $0.75 / (1 - 0.75)$) would look at the target $(.563 * 3 = \text{odds of } 1.689)$ or $(1.689 / (1 + 1.689)) = 62.8\%$ of the time given a shortened-vowel mispronunciation.

Children's longer target-looking times on trials with words that normally have long vowels, *paard* ('horse') and *beer* ('bear'), was unexpected. Children often appear to do better on animates than inanimates, but this is probably a picture preference effect much of the time, and here we attempted to account for such preferences statistically. In any event, the important comparisons here are both within-item and within-subject.

To summarize, Dutch 21-month-olds' performance in recognizing words was consistent with prior research on adults, in which shortening long vowels changes interpretation of the vowel, but lengthening short vowels does not, or to a lesser degree. The present results extend to real words the findings of Dietrich et al. (2007), in which Dutch children were capable of linking two objects to two novel syllables varying only in vowel duration. In the next experiment we test whether English learning children respond similarly to vowel duration changes.

Experiment 2

Methods

Participants.—Participants were 24 American 21-month-olds (12 girls) with a mean age of 21;12 (range 20;29 to 21;24). An additional eight children visited the lab but either refused to begin ($n=5$) or did not complete at least 16 of the 24 trials (cf. Experiment 1; $n=3$). The study was conducted in the city of Philadelphia at the University of Pennsylvania. Children were recruited primarily via mailed invitations to parents, to addresses purchased from a data provider. Children were not recruited if their parent reported the child hearing less than 75% English during a typical week. About 29% of participants' parents identified their child as Black / African-American, 58% White, and the remainder Asian or Hispanic. No birth complications or abnormalities of hearing or vision were reported when queried during the recruiting process.

Stimuli and procedure.—Speech stimuli were recorded by a native English-speaking American woman, in the same sound-insulated room where the materials for Experiment 1 were recorded. Sentences were realized in an “infant-directed” register and were of the form “Look at the... [target]” followed by a 750 msec pause and an encouraging second sentence like “Do you like it?” or “Can you find it?”

To ensure that our duration manipulation of the English materials was at least as large as that of the Dutch materials, target words' vowels were lengthened or shortened at a ratio of 1:1.8 (shorter:longer), a slightly larger manipulation than tested in the Dutch case. This duration ratio is the same as the ratio used to test adults in Van der Feest and Swingley (2011). The vowels in *cup*, *juice*, and *sock* were lengthened (allowing the possibility that children would take the longer duration as signaling a mispronounced voiced coda consonant), and the vowels in *bed*, *dog*, and *keys* were shortened (which in principle could have been interpreted as signaling an unvoiced coda). The stimulus set is described in Table 3.¹

Trial orders were constructed in the same way as in Experiment 1, with the same counterbalancing constraints. As in the first experiment, test trials were interspersed with brief “filler” animations every 5 trials to help keep children interested in the task. The bed was paired with the cup, the dog with juice, and the keys with the sock.

¹The duration of the vowel in “juice” tends to fall between the American English long and short vowels phonetically. Here we treated it as short, and altered its pronunciation by lengthening it. This was motivated by Van der Feest & Swingley (2011) finding that adults sometimes treated shortened /o/ as /u/, and sometimes incorrectly treated lengthened-/u/ syllables as having voiced codas. Mispronouncing /u/ by lengthening it allowed us to probe for similar effects.

Parents were given a vocabulary questionnaire to complete before their session (the Macarthur-Bates Communicative Development Inventory, Words & Sentences or CDI; 1994); all 24 were completed before the session or within a few days afterward.

Results, Experiment 2

As in Experiment 1, we began by confirming that children recognized the test words, by testing whether salience-subtracted fixation proportions significantly exceeded zero in the standard window from 367 to 2000 ms following target onset.² All six words were recognized at significantly above-chance levels (mean, 12.8%, mean sd, 20.7; min, 8.4; min(t) = 2.1). As a group, children recognized the words: by items, $t(5) = 6.0$, $p(\text{one-tailed}) < 0.001$; by subjects, $t(23) = 7.20$, $p(\text{one-tailed}) < 0.0001$.

Figure 2 summarizes children's performance on the originally long words (*bed*, *dog*, *keys*; left panel) and the originally short words (*cup*, *juice*, *sock*; right panel), following the same conventions as Figure 1. The plot shows children's apparent indifference to the duration manipulation. To model these (non)effects we performed an analysis parallel to that of Experiment 1.

The outcome measure was the proportion of target fixation in the 567–2000 ms window. Predictors were Subject, spoken vocabulary score (from the CDI), Salience, Condition, Length, and the interaction between Condition and Length. Word counts divided by 100 were used in place of the raw CDI totals to balance the magnitude of the predictors. The results of the model are given in Table 4.

English-learning children showed no sign of being affected by changes in vowel duration, whether lengthening or shortening. Their overall performance was somewhat better for the long words (*bed*, *dog*, and *keys*) than the short words (*cup*, *juice*, and *sock*), with *dog* leading the set; as for the Dutch children, we attribute this to the words as individual items and not to their vowels or codas per se. Children with higher vocabularies performed better overall, but tests in which interactions of CDI scores with any other variables did not come close to improving the model, and no such interactions were significant (all $p > 0.5$).

To compare the Dutch and English results directly, the effects of mispronunciation in each group, for lengthened and shortened words, are presented in Figure 3. Each child's mean target fixation in each cell of the 2×2 condition X length matrix was computed, and the difference in performance on correct and mispronunciation trials was calculated. These difference scores are shown in the Figure, along with the cell means and standard errors. The difference in effects for the originally-long words was significant by two-sample t-test (mean 12.4%, $t(45.2) = 2.30$, $p = 0.026$); the difference for originally-short words was not (mean 1.0%, ns).

²The frame rate of the PAL (European) video recording standard is 25 per second (40 ms intervals), whereas the NTSC (US) standard is about 30 (33 ms intervals). This accounts for the slight difference in where the Dutch and American test windows began.

General discussion

We found that word recognition in English-learning toddlers was not hindered when vowel duration was altered. In contrast, we found that for Dutch-learning toddlers words were more difficult to recognize when (originally long) vowels were shortened, but not when (originally short) vowels were lengthened. This result is in line with and extends the cross-linguistic difference first documented in young children by Dietrich et al. (2007). The Dietrich et al. finding did not necessarily show that Dutch children would spontaneously treat the vowel duration distinction as significant, because they might have learned it during the experiment (Yeung, Chen, & Werker, 2014). The present results show that shortening a long vowel in the tested prosodic context is unacceptable to Dutch 21-month-olds, even in the absence of any laboratory training, and extending beyond the /ɑ/ - /a:/ contrast.

Why did changing English vowel durations have no apparent effect on children's word recognition? This result might be predicted by phonological feature theories that describe Dutch as having phonologically contrastive vowel duration, but English as not having an independent vowel duration feature represented in the lexicon (e.g., Booij, 1999; Gussenhoven, 1999). Of course, the task of the child is to discover this property of the language. One possibility is that English learners detect systematic duration variation (for example, in the intrinsic differences between [i] and [ɪ]) but take some time to gain confidence in their model of duration because they also have to work out the systematic duration variation attributable to coda voicing, a regularity that is absent in Dutch because Dutch codas are unvoiced. It may be that English learners take longer to work out the quantity features of vowels because the system they have to learn is more complex. On this account, Dutch children make stronger attributions based on duration because their mental model of duration is more precise. This seems unlikely to be the full story, though, given Dutch adults' apparently greater reliance on duration in making lexical judgments (Van der Feest & Swingley, 2011).

Dutch children matched Dutch adults in being less strongly affected by vowel lengthening than vowel shortening (Chládková et al., 2015; Nooteboom & Doodeman, 1980; Van der Feest & Swingley, 2011). Why might this be? One possibility is that fundamentally, the vowel's quality features (e.g., formants) weigh more strongly than quantity, and lengthening renders vowel quality more readily perceptible whereas shortening renders it less perceptible. Thus, for example, a child hearing *ba/* with a lengthened vowel would have ample evidence of its quality, placing the vowel in the /ɑ/ category despite its length, while a child hearing *paard* with a short vowel could have a less secure grasp on the vowel's quality, perhaps leading to greater reliance on the short duration to identify the category. It is common in speech perception to reweight attention to different cues as a function of their apparent reliability in a given situation; to take a familiar example, listeners rely on visual cues more when speech is partially masked by noise (Sumby & Pollack, 1954; for a different example, Holt & Lotto, 2006). A difficulty for this account in the present case is that the shortened vowels were not particularly short in absolute terms, and in this task young children make distinctions among vowels that are comparably short phonetically (e.g., Ramon-Casas, Swingley, Sebastián-Gallés, & Bosch, 2009).

A second possibility that could in principle account for the asymmetry is that in Dutch children's experience, vowels that are typically short may sometimes be lengthened, perhaps more often than long vowels are shortened—for example, if a parent wants to emphasize a word strongly when speaking to a child. If the distribution of phonologically-long vowels' durations is relatively peaked (having low variance), a vowel that falls outside of the norm could be interpreted as a member of a neighboring phonologically shorter category, especially given that the stimuli in our experiment were produced in a child-directed register. Conversely if short vowels' duration distributions are more spread out, duration values just as far from the average might nonetheless have relatively little impact on interpretation. But the sparse evidence available suggests that in fact the phonologically long vowels can be realized with short durations in prosodically weak positions, arguing against this interpretation (Gussenhoven, 1999; Rietveld et al., 2004). Furthermore, Swingley (in press) found that in a sample of Dutch infant-directed speech the long vowels of Dutch were no less variable than the short vowels, and overall quite similar in their distribution to the durations of *English* long and short vowels. If, indeed, children learning English and Dutch hear similar vowel duration distributions in their daily lives, the explanation for children's language-specific treatment of duration manipulations must lie elsewhere. Swingley's analyses also showed that the long/short pairs of Dutch infant-directed speech are not, in general, more difficult to distinguish from one another using formant measurements than the long/short pairs of English are. This argues against the notion that Dutch children attend more to duration in order to compensate for a greater degree of overlap in vowel quality between the long and short vowels.

As described in the Introduction, Dutch long/short vowel pairs show greater duration differentiation in citation form than English pairs do. But it seems that once vowels are taken out of these relatively pristine contexts, the differences between the languages become harder to identify statistically by observing surface distributions. Yet toddlers seem to have succeeded in doing so. Perhaps they discover the durational features of their vowels by learning a prosodic model of their language in which multiple sources of variation control vowel duration, and in Dutch a statistical difference then emerges that children can correctly attribute to intrinsic differences in vowel duration. Or perhaps specific words in children's experience exhibit distinctive duration features, and it is the lexicon that guides Dutch children to weighing duration more heavily (the case made by Swingley, in press).

It is these possibilities that make the rather esoteric question of Dutch and English vowel duration a valuable test case for understanding children's capacity for phonological learning. Although Dutch and English do not present a perfectly controlled, minimally different pair, as might be constructed in a laboratory intervention experiment, the languages are phonetically similar, and our experimental samples were drawn from similar populations. Thus, our participants might be viewed as if they were subjects in a 21-month-long intervention experiment with imperfectly controlled familiarization materials. From their ordinary language experience they have learned to draw somewhat different conclusions about how duration relates to lexical identity, in a way that aligns (grossly, at least) with adults' performance. Infant speech perception research often focuses either on speech sounds or on suprasegmental features, but languages use duration both segmentally and suprasegmentally, and in different ways. Children cannot simply uncover segmental

phonology and prosody as separate problems. Our results suggest that by 21 months, children can make substantial progress in untangling the phonetic attribution problem that vowel duration presents.

Acknowledgements:

This work was supported by NIH grant R01-HD049681 to D. Swingley, and by an NWO Rubicon grant to S.V.H. van der Feest. The authors thank Angela Khadar and the Baby Research Center in Nijmegen; members of the Penn Infant Language Center, particularly Jane Park; and Sarah Johnstone. We are especially grateful to the many children and parents who participated.

References

- Adank P, van Hout R, & Smits R (2004). An acoustic description of the vowels of Northern and Southern Standard Dutch. *Journal of the Acoustical Society of America*, 116, 1729–1738. [PubMed: 15478440]
- Bion RA, Miyazawa K, Kikuchi H, & Mazuka R (2013). Learning phonemic vowel length from naturalistic recordings of Japanese infant-directed speech. *PLoS one*, 8(2), e51594. doi: 10.1371/journal.pone.0051594 [PubMed: 23437036]
- Boersma P, & Weenink D (2005). *Praat: doing phonetics by computer [computer program]* (Tech. Rep.). retrieved 2005 from www.praat.org: Version 4.1.
- Booij G (1999). *The phonology of Dutch*. Oxford University Press.
- Chen A, & Kager R (2016). Discrimination of lexical tones in the first year of life. *Infant and Child Development*, 25(5), 426–439. doi: 10.1002/icd.1944
- Chládková K, Escudero P, & Lipski SC (2015). When “aa” is long but “a” is not short: speakers who distinguish short and long vowels in production do not necessarily encode a short–long contrast in their phonological lexicon. *Frontiers in Psychology*, 6, 438. doi: 10.3389/fpsyg.2015.00438 [PubMed: 25964766]
- Crowther CS, & Mann V (1992). Native language factors affecting use of vocalic cues to final consonant voicing in English. *The Journal of the Acoustical Society of America*, 92(2), 711–722. [PubMed: 1506526]
- Dietrich C, Swingley D, & Werker JF (2007). Native language governs interpretation of salient speech sound differences at 18 months. *Proceedings of the National Academy of Sciences of the USA*, 104, 16027–16031. doi: /10.1073/pnas.0705270104 [PubMed: 17911262]
- Fenson L, Dale PS, Reznick JS, Bates E, Thal DJ, & Pethick SJ (1994). Variability in early communicative development. *Monographs of the Society for Research in Child Development*, 59. doi: /10.2307/1166093
- Finegan E (1987). English In Comrie B (Ed.), *The world’s major languages* (p. 77–109). Oxford.
- Gussenhoven C (1999). Illustrations of the IPA: Dutch In *Handbook of the international phonetic association* (p. 74–77). Cambridge: Cambridge University Press.
- Hay JF, Graf Estes K, Wang T, & Saffran JR (2015). From flexibility to constraint: The contrastive use of lexical tone in early word learning. *Child Development*, 86, 10–22. doi: 10.1111/cdev.12269 [PubMed: 25041105]
- Heeren WFL (2006). *Perceptual development of phoneme contrasts in adults and children* Unpublished doctoral dissertation. University of Utrecht.
- Hillenbrand J, Getty LA, Clark MJ, & Wheeler K (1995). Acoustic characteristics of American English vowels. *Journal of the Acoustical Society of America*, 97, 3099–3111. [PubMed: 7759650]
- Hillenbrand JM, Clark MJ, & Houde RA (2000). Some effects of duration on vowel recognition. *Journal of the Acoustical Society of America*, 108, 3013–3022. [PubMed: 11144593]
- Hirata Y (2004). Effects of speaking rate on the vowel length distinction in Japanese. *Journal of Phonetics*, 32(4), 565–589. doi: 10.1016/j.wocn.2004.02.004

- Hollich G (2008). Supercoder: A program for coding preferential looking (v. 1.7.1). [Computer Software]. (West Lafayette: Purdue University)
- Holt LL, & Lotto AJ (2006). Cue weighting in auditory categorization: implications for first and second language acquisition. *Journal of the Acoustical Society of America*, 119, 3059–3071. [PubMed: 16708961]
- House AS, & Fairbanks G (1953). The influence of consonant environment upon the secondary acoustical characteristics of vowels. *The Journal of the Acoustical Society of America*, 25(1), 105–113.
- Humphrey C, & Swingley D (2018). Regression analysis of proportion outcomes with random effects. [arXiv.org](https://arxiv.org/abs/1805.08670). doi: arXiv:1805.08670
- Ko E-S, Soderstrom M, & Morgan J (2009). Development of perceptual sensitivity to extrinsic vowel duration in infants learning american english. *The Journal of the Acoustical Society of America*, 126(5), EL134–EL139. doi: 10.1121/1.3239465 [PubMed: 19894788]
- Kooij JG (1987). Dutch In Comrie B (Ed.), *The worlds major languages* (p. 77–109). Oxford.
- Liu L, & Kager R (2014). Perception of tones by infants learning a non-tone language. *Cognition*, 133(2), 385–394. doi: 10.1016/j.cognition.2014.06.004 [PubMed: 25128796]
- Mahr T, McMillan BT, Saffran JR, Weismer SE, & Edwards J (2015). Anticipatory coarticulation facilitates word recognition in toddlers. *Cognition*, 142, 345–350. doi: 10.1016/j.cognition.2015.05.009 [PubMed: 26072992]
- Mattock K, Molnar M, Polka L, & Burnham D (2008). The developmental course of lexical tone perception in the first year of life. *Cognition*, 106(3), 1367–1381. doi: 10.1016/j.cognition.2007.07.002 [PubMed: 17707789]
- Mendoza E, Carballo G, Cruz A, Fresneda MD, Muñoz J, & Marrero V (2003). Temporal variability in speech segments of spanish: Context and speaker related differences. *Speech Communication*, 40(4), 431–447. doi: 10.1016/S0167-6393(02)00086-9
- Mugitani R, Pons F, Dietrich C, Werker JF, & Amano S (2009). Perception of vowel length by Japanese- and English-learning infants. *Developmental Psychology*, 45, 236–247. doi: 10.1037/a0014043 [PubMed: 19210005]
- Nooteboom SG, & Doodeman GJN (1980). Production and perception of vowel length in spoken sentences. *Journal of the Acoustical Society of America*, 67, 276–287. [PubMed: 7354196]
- Quam C, & Swingley D (2010). Phonological knowledge guides two-year-olds' and adults' interpretation of salient pitch contours in word learning. *Journal of Memory and Language*, 52, 135–150. doi: /10.1016/j.jml.2009.09.003
- Ramon-Casas M, Swingley D, Bosch L, & Sebastián-Gallés N (2009). Vowel categorization during word recognition in bilingual toddlers. *Cognitive Psychology*, 59, 96–121. doi: /10.1016/j.cogpsych.2009.02.002 [PubMed: 19338984]
- Raphael LJ (2008). Acoustic cues to the perception of segmental phonemes In Pisoni DB & Remez RE (Eds.), *The handbook of speech perception* (p. 182–206). Oxford: Wiley Blackwell. doi: 10.1002/9780470757024.ch8
- Rietveld T, Kerkhoff J, & Gussenhoven C (2004). Word prosodic structure and vowel duration in Dutch. *Journal of Phonetics*, 32, 349–371.
- Singh L, Goh HH, & Wewalaarachchi TD (2015). Spoken word recognition in early childhood: Comparative effects of vowel, consonant and lexical tone variation. *Cognition*, 142, 1–11. doi: 10.1016/j.cognition.2015.05.010 [PubMed: 26010558]
- Singh L, Hui TJ, Chan C, & Golinkoff RM (2014). Influences of vowel and tone variation on emergent word knowledge: a cross-linguistic investigation. *Developmental Science*, 17(1), 94–109. doi: 10.1111/desc.12097 [PubMed: 24118787]
- Sumby WH, & Pollack I (1954). Visual contribution to speech intelligibility in noise. *The Journal of the Acoustical Society of America*, 26(2), 212–215.
- Swingley D (2007). Lexical exposure and word-form encoding in 1.5-year-olds. *Developmental Psychology*, 43, 454–464. [PubMed: 17352552]
- Swingley D (2009). Onsets and codas in 1.5-year-olds' word recognition. *Journal of Memory and Language*, 60, 252–269. doi: /10.1016/j.jml.2008.11.003 [PubMed: 20126290]

- Swingley D (2012). The looking-while-listening procedure In Hoff E (Ed.), Research methods in child language: A practical guide (pp. 29–42). Oxford: Wiley-Blackwell.
- Swingley D (2016). Two-year-olds interpret novel phonological neighbors as familiar words. *Developmental Psychology*, 52, 1011–1023. doi: 10.1037/dev0000114 [PubMed: 27337510]
- Swingley D (in press). Learning phonology from surface distributions, considering Dutch and English vowel duration as a test case. *Language Learning and Development*.
- Swingley D, & Aslin RN (2000). Spoken word recognition and lexical representation in very young children. *Cognition*, 76, 147–166. [PubMed: 10856741]
- Swingley D, Pinto JP, & Fernald A (1999). Continuous processing in word recognition at 24 months. *Cognition*, 71, 73–108. doi: 10.1016/s0010-0277(99)00021-9 [PubMed: 10444905]
- Tillman G, Benders T, Brown SD, & van Ravenzwaaij D (2017). An evidence accumulation model of acoustic cue weighting in vowel perception. *Journal of Phonetics*, 61, 1–12. doi: 10.1016/j.wocn.2016.12.001
- Van der Feest SVH, & Swingley D (2011). Dutch and English listeners' interpretation of vowel duration. *The Journal of the Acoustical Society of America*, 129(3), EL57–EL63. doi: 10.1121/1.3532050 [PubMed: 21428468]
- Warner N, & Cutler A (2017). Stress effects in vowel perception as a function of language-specific vocabulary patterns. *Phonetica*, 74, 81–106. [PubMed: 27710964]
- Werker JF, Cohen LB, Lloyd VL, Casasola M, & Stager C (1998). Acquisition of word-object associations by 14-month-old infants. *Developmental Psychology*, 34, 1289–1309. [PubMed: 9823513]
- Werker JF, Yeung HH, & Yoshida KA (2012). How do infants become experts at native-speech perception? *Current Directions in Psychological Science*, 21, 221–226. doi: 10.1177/0963721412449459
- Yeung HH, Chen LM, & Werker JF (2014). Referential labeling can facilitate phonetic learning in infancy. *Child development*, 85(3), 1036–1049. [PubMed: 24936610]

Figure 1.

Dutch children's target fixation proportions when hearing words that normally have long vowels (left panel) or short vowels (right panel). Responses to correct pronunciations (CP) are darker bars, mispronunciations (MP) lighter ones. Subject means are shown for percentage to target (where chance is about 50%) and for this percentage less each child's "salience" score for the target picture (so chance is about 0%; see text). Individual child means for each condition are given as dark linked dots. Error bars are standard errors of the mean.

Figure 2.

American children's target fixation proportions when hearing words that normally have long vowels (left panel) or short vowels (right panel). Responses to correct pronunciations (CP) are darker bars, mispronunciations (MP) lighter ones. Subject means are shown for percentage to target (broad bars) and for this percentage minus each child's "salience" score for the target picture (narrower, darker bars). Individual child means for each condition are given as dark linked dots. Error bars are standard errors of the mean.

Figure 3.

For each child, the difference between his or her mean target fixation on correct-pronunciation trials and target fixation on mispronunciation trials. The left panel shows performance for words with canonically long vowels; the right panel words with canonically short vowels. Means and standard errors are shown with open circles and error bars.

Table 1

Characteristics of Dutch target words.

word	IPA	English	orig. ms	manip. ms	ratio
bal	bɑl	ball	167	240	1.43
bed	bɛt	bed	95	165	1.74
pop	pɔp	doll	117	169	1.45
boot	bo:t	boat	230	132	1.74
paard	pa:lt	horse	276	176	1.57
beer	bɛ:ɾ	bear	150	97	1.54

Note. The table shows, for each Dutch target word, its International Phonetic Alphabet transcription, an English gloss, the duration of the vowel as originally produced, the duration of the manipulated vowel, and the ratio formed by dividing the longer duration by the shorter for each stimulus.

Table 2

Regression coefficients and other statistics in analysis of Dutch children's gaze patterns.

predictor	coef	exp(coef)	std.error	p(boot)
(intercept)	0.276	1.318	0.514	.2554 (ns)
condition (MP)	−0.575	0.563	0.338	<.001
orig. length (short)	−0.594	0.552	0.318	<.001
salience	1.431	4.182	0.83	.004
condition * length	0.661	1.936	0.424	.002

Note. *Coef* refers to the estimated beta coefficient from the binomial logistic regression. *Exp(coef)* provides the number by which the odds of fixating the target would be multiplied, given an increase of one in the predictor's value. Intervention levels are indicated in parentheses (i.e., the coefficient for condition shows the change expected in going from CP to the intervention MP). *Std.error* indicates the standard error of the predictor in the base (non-bootstrapped) regression model. *p(boot)* gives a significance level, computed from the 2-tailed confidence intervals of the effects estimate from the bootstrap analysis (600,000 runs). *ns* indicates a result not significant at the 0.10 level (2-tailed).

Table 3

Characteristics of English target words.

word	IPA	orig. ms	manip. ms	ratio
bed	bed	297	167	1.78
keys	kiz	414	227	1.82
dog	dɔg	344	187	1.84
cup	cʊp	144	274	1.91
juice	ʒus	186	302	1.63
sock	sak	170	276	1.62

Note. The table shows, for each English target word, its International Phonetic Alphabet transcription, the duration of the vowel as originally produced, the duration of the manipulated vowel, and the ratio formed by dividing the longer duration by the shorter for each stimulus.

Table 4

Regression coefficients and other statistics in analysis of American children's gaze patterns.

predictor	coef	exp(coef)	std.err	p value
(intercept)	−0.248	0.781	0.446	.415 (ns)
condition (MP)	−0.120	0.988	0.277	.939 (ns)
orig. length (short)	−0.435	0.647	0.273	.016
salience	1.456	4.289	0.710	.001
spoken vocab. (/100)	0.207	1.230	0.097	.0133
condition X length	0.098	1.103	0.380	.645 (ns)

Note. *Coef* refers to the estimated beta coefficient. *exp(coef)* is the coefficient exponentiated; a value of 1 is no effect at all. Intervention levels are indicated in parentheses (i.e., the coefficient for condition shows the change expected in going from reference CP to the intervention MP). *Std.err* indicates the standard error of the predictor in the base (non-bootstrapped) regression model. *p value* gives a significance level, computed from the 2-tailed confidence intervals of the effects estimate from the bootstrap analysis (600,000 runs).