This manuscript is currently in press at Developmental Science.

The final article is available upon publication, at DOI: [10.1111/desc.12896](https://doi.org/10.1111/desc.12896)

Statistical learning of multiple speech streams: A challenge for monolingual infants

Viridiana L. Benitez^{a1}, Federica Bulgarelli^{b,c}, Krista Byers-Heinlein^c, Jenny R. Saffran^d, &

Daniel J. Weiss^e

^aArizona State University

^bDuke University

^cConcordia University

^dUniversity of Wisconsin-Madison

^ePennsylvania State University

Corresponding author

Viridiana Benitez, 950 S. S. McAllister Ave., Tempe, AZ 85287, viridiana.benitez@asu.edu, 480-965-5425

Conflict of interest statement

The authors do not have any conflicts of interest to report.

Acknowledgements: This work was funded by grants from the National Institute of Child Health and Human Development to JRS (R37HD037466), DJW (R01HD067250) and the Waisman Center (U54 HD090256), grants from the National Science Foundation to VLB (NSF SPRF 1513834) and FB (NSF GRFP), a grant from the Natural Sciences and Engineering Council of Canada under award number 402470-2011 to KBH, and support from the Concordia University Research Chairs program to KBH. We would like to thank the research assistants and staff at the Infant Learning Lab for their help in participant recruitment and testing.

Data Availability Statement: The data that support the findings of this study are openly available in the Open Science Framework at <https://osf.io/6xjqc/>.

Data citation:

[dataset] Bulgarelli, F., Benitez, V.L., Byers-Heinlein, K., Saffran, J.R., & Weiss, D.J.; 2019; Statistical learning of multiple speech streams: A challenge for monolingual infants; <https://osf.io/6xjqc/>

¹ Viridiana L. Benitez and Federica Bulgarelli should be considered joint first authors

Research highlights

- We tested monolingual 8-month-old infants' ability to segment two artificial speech streams that shared syllables
- Infants exhibited learning when presented with a single stream, but did not exhibit learning of the first or second stream when presented in succession
- The addition of contextual cues (a pitch shift or accent + speaker cues) did not lead infants to learn the second presented stream
- Slightly older infants exhibited the same difficulties, confirming that monolingual infants find it difficult to segment two speech streams presented in succession

Abstract

Language acquisition depends on the ability to detect and track the distributional properties of speech. Successful acquisition also necessitates detecting changes in those properties, which can occur when the learner encounters different speakers, topics, dialects, or languages. When encountering multiple speech streams with different underlying statistics but overlapping features, how do infants keep track of the properties of each speech stream separately? In four experiments, we tested whether 8-month-old monolingual infants ($N = 144$) can track the underlying statistics of two artificial speech streams that share a portion of their syllables. We first presented each stream individually. We then presented the two speech streams in sequence, without contextual cues signaling the different speech streams, and subsequently added pitch and accent cues to help learners track each stream separately. The results reveal that monolingual infants experience difficulty tracking the statistical regularities in two speech streams presented sequentially, even when provided with contextual cues intended to facilitate separation of the speech streams. We discuss the implications of our findings for understanding how infants learn and separate the input when confronted with multiple statistical structures.

Keywords: statistical learning, multi-stream segmentation, multiple representations, infants, word learning, language development, bilingualism

Introduction

Statistical learning, the process by which learners track the distributional properties of the input to determine the underlying structure, operates at many levels of language. Studies with infants suggest that they are sensitive to statistical regularities in linguistic domains including phonetic learning (e.g., Maye, Werker, & Gerken, 2002), word learning (e.g., Smith & Yu, 2008), grammar learning (e.g., Gomez & Gerken, 1999), and in non-linguistic domains such as visual sequences (Fiser & Aslin, 2001) and musical tones (Saffran, Johnson, Aslin, & Newport, 1999). Infants' statistical learning ability has been most extensively studied in the domain of speech segmentation: how infants are able to identify words in running speech and the boundaries between them. Typically, researchers create artificial languages where the only signal to word boundaries is syllable co-occurrences. Hidden in the stream are "words" where transitional probabilities across the syllables are high (i.e., these syllables often follow one another) and boundaries between words where the transitional probabilities are low (i.e., these syllables rarely follow one another, as they are not part of the same "word"). After a few minutes of exposure to a continuous stream of speech, infants can discriminate these statistically-defined "words" from novel combinations of the same syllables (Aslin, Saffran, & Newport, 1998; Saffran, Aslin, & Newport, 1996). Research over the past two decades has revealed infants' sensitivity to these statistical patterns in both artificial and natural speech (for a recent review, see Saffran & Kirkham, 2018).

Notably, the vast majority of research on infant statistical learning to date has assumed, at least implicitly, that the task of the learner is to acquire a single underlying structure. Equally important, however, is the ability to detect a change in the underlying regularities and to separate the regularities stemming from different contexts. This problem is relevant for infants raised in monolingual environments, who may encounter changes in the regularities of speech across

instances, speakers, accents, or dialects (Gonzales, Gerken, & Gómez, 2015). This problem is also relevant for infants growing up bilingual, who are exposed to two different languages, each with its own set of regularities, and who must learn their languages as distinct systems (Byers-Heinlein, 2014). In any language-learning environment, any single exposure to speech (ranging from individual words or sentences to full days of discourse) may not be fully representative of the underlying statistics of the language or languages that infants must acquire. Instead, infants must detect the statistics within a single exposure, remain flexible to changes to those statistics across different experiences, and determine which speech streams to aggregate over and which to keep separate (Qian, Jaeger, & Aslin, 2016).

Detecting and tracking the statistics of two speech streams can be a challenge, especially when the two speech streams share features. This issue is most prominent when considering multiple languages. For example, the syllable “zeh” means “this” in Hebrew, often signaling a word boundary. In English, however, the syllable “zeh” is more likely to occur within a word, such as “zealous” or “zealot”. Thus, the challenge for a listener that encounters both English and Hebrew is that similar syllables can pattern differently across the two languages (Weiss, Gerfen & Mitchel, 2009). If a learner aggregates the statistics across languages, syllable cues to word boundaries become noisy and unreliable. However, if a learner aggregates the statistics separately, they can notice the reliable word boundaries within each language. This issue is not limited to learners in bilingual environments, as similar challenges for making statistical inferences confront learners even in first language learning. For example, monolingual learners can encounter changes in dialect or accent (among other features) and thus must also contend with variability and make implicit decisions about when to aggregate statistics separately. Consequently, there have been several proposals suggesting a unified framework for statistical inference that spans monolingual and multilingual acquisition (e.g., see Qian, Jaeger, and Aslin,

2012; Pajak, Fine, Kleinschmidt, & Jaeger, 2016). Understanding how young infants contend with variability in the statistics of the input, such as encountering multiple speech streams, can lend deep insights into how the process of statistical learning unfolds during development.

Statistical learning from multiple speech streams has largely been studied in monolingual adults by manipulating the syllable overlap and congruence between the speech streams. Two speech streams are considered congruent when the transitional probabilities within each language, as well as the combined transitional probabilities across languages, are higher within words than between words. Congruent streams of speech can be created by using completely separate syllable inventories across speech streams (meaning no syllables are shared), or by reusing syllables in the same contexts across speech streams, such that syllable-to-syllable transitional probabilities are maintained across streams. When two streams are congruent, statistical regularities consistently cue word boundaries, regardless of whether learners collapse statistics across speech streams or track each speech stream separately. Indeed, when monolingual adults are presented with two congruent artificial speech streams, they learn both (Weiss, et al., 2009).

Two speech streams are considered incongruent when the transitional probabilities within each stream are higher within words than between words, but result in uninformative transitional probabilities when combined across speech streams. Incongruence in the statistics of two speech streams occurs when there is partial or full overlap in syllable inventory across the two streams, and additionally, when the overlapping syllables are used in different contexts. This produces different syllable-to-syllable transitional probabilities across streams. Therefore, aggregating statistical information across two incongruent speech streams results in noisier information about word boundaries. For example, the syllables /pa/ and /bi/ may occur together within a word in one stream with a 1.0 transitional probability, but span word boundaries in another stream (if /pa/

is the end of a word and /bi/ is the beginning of a word) and thus have a lower transitional probability in the second stream (because any other word from the stream could occur after a word boundary). In contrast to congruent speech streams, successfully learning the statistics of two incongruent speech streams requires learners to aggregate transitional probabilities separately for each speech stream.

Studies to date suggest that adult monolingual learners have difficulty separating and segmenting incongruent speech streams. When presented with two interleaved, incongruent speech streams without any cue to distinguish them, monolingual adults fail to learn either stream (Mitchel & Weiss, 2010, Weiss et al., 2009), or show a primacy effect, acquiring the first of two consecutive streams but not the second (Gebhart, Newport, & Aslin, 2009). Adults can, however, successfully learn two interleaved, incongruent speech streams when there is a correlated contextual cue that differentiates the streams, such as hearing each stream in a different voice (Weiss et al., 2009), or associated with a different talking face (Mitchel & Weiss, 2010). Adults can also successfully learn two incongruent speech streams when exposure is increased substantially (i.e. the second stream is tripled in length), or when provided with explicit instructions that there are two languages in the input in combination with a pause between the streams (Gebhart et al., 2009). Together, these studies suggest that for monolingual adult learners, segmenting two incongruent speech streams is difficult. Successful learning in monolingual adults appears to require an effective contextual cue that corresponds to the change in statistical structure, or extending the exposure to the second stream.

Adults are sophisticated listeners, who, for better or worse, bring years of prior language experience to statistical learning. How do infants solve this type of task? In particular, can they separate and segment speech streams that contain incongruent statistics? Some evidence suggests that learning multiple sets of linguistic rules poses a challenge for infant learners,

particularly those growing up in monolingual environments (Kovács & Mehler 2009a, 2009b; Potter & Lew-Williams, 2019). Most relevant to the current study, Antovich and Graf-Estes (2018) compared monolingual and bilingual 14-month-old infants who were exposed to two congruent speech streams made up of non-overlapping syllable inventories. The streams were interleaved during familiarization, and additionally, were presented with a contextual cue to separate the speech streams (one stream was spoken by a female and the other by a male). The study found that only the bilingual infants exhibited learning; monolingual infants did not acquire either speech stream. Given that the two speech streams were congruent and non-overlapping in syllable inventory, it is possible that the bilingual infants acquired the input as a single, larger language as opposed to distinct sets of words drawn from two smaller languages (Antovich & Graf-Estes, 2018). Therefore, it remains unknown whether infants can successfully track the statistical properties of two speech streams when they are incongruent. This requires infants to track each speech stream separately and is representative of the type of conflicting regularities that young listeners may encounter within multilingual, multidialectal, or multi-accent environments.

In the current set of experiments, we asked whether monolingual infants can separately track the statistics of two incongruent speech streams. We focused on 8- to 10-month-old infants, given that at this age, infants can use statistics to discover language-specific cues for segmentation (Johnson & Jusczyk, 2001; Saffran & Thiessen, 2003), and this is the most commonly tested age in the statistical word learning literature (Black & Bergmann, 2017; Saffran & Kirkham, 2018). Further, we tested monolingual infants given that even in monolingual environments, learners encounter variation in the statistics across instances and speakers, requiring them to figure out when the statistics have changed and over what inputs the statistics should be aggregated (Gonzales et al., 2015; Qian, et al., 2012; Qian et al., 2016).

Understanding whether and how monolingual infants detect changes across incongruent speech streams is a necessary and informative step for beginning to understand the larger issue of how infants deal with changes in the statistical input, a factor that is relevant for understanding learning in different language environments. Finally, to better understand the contexts in which monolingual infants may be able to segment two incongruent speech streams separately, we also manipulated the presence of contextual cues that signaled two separate streams.

In a series of four experiments, we exposed monolingual infants to two incongruent artificial speech streams. The streams shared a subset of the syllables, but each contained their own transitional probability cues to word boundaries. The syllables' distributions yielded speech streams that were incongruent: when aggregated together, transitional probabilities would not provide consistent evidence for word boundaries. Thus, in order to successfully detect the word boundaries, infants would need to detect and learn the statistical patterns of each speech stream separately.

In Experiment 1 we presented infants with just one of the two speech streams, to ensure that infants were successful at learning when presented with a single language. In Experiment 2, we presented infants with the two languages in succession, with no contextual cues corresponding to the change in structure (both were spoken by the same speaker and there was no pause between the first and second language). Infants were subsequently tested on their learning of either the first language (Initial Language condition) or the second language (Final Language condition). In Experiment 3, we added two types of contextual cues, a pitch shift (Pitch condition) and a change in speaker and accent (Accent condition) to highlight the presence of two languages; following exposure, we tested infants on the second presented language (Final Language condition). In Experiment 4, based on results from exploratory analyses in Experiment 3, we conducted a pre-registered study with a group of older infants (<https://osf.io/7bru4>). Infants

in all experiments were tested on their preference for statistically-conforming sequences of syllables in the tested language (words) compared to statistically-nonconforming novel sequences of syllables from the tested language (nonwords), using the same set of test materials. The key question was whether infants exhibit differences in looking times to words and nonwords when listening to the two languages in succession.

Experiment 1

The purpose of Experiment 1 was to verify that monolingual infants could learn when presented with a single language during exposure, using each of the novel languages used in subsequent experiments (Languages A and B; see Table 1). Half of the infants were randomly assigned to listen to Language A, while the other half listened to Language B. We then tested infants on their ability to learn the familiarized language by recording their looking times to words and nonwords using the Headturn Preference Procedure. As these languages were created following guidelines from previous experiments, we expected that infants should exhibit differences in looking time between words and nonwords.

Participants

A total of 24 infants (11 females) participated in Experiment 1. This sample size was selected as it is commonly used in studies of infant word segmentation. Infants ranged in age from 7.3 to 8.6 months (mean age = 8.1, $SD = 0.39$) and were randomly assigned to be familiarized with either Language A (13) or Language B (11). Infants who were born full-term, were from monolingual, English speaking homes, and who did not report a recent history of ear infections were included in the study. Additionally, only infants who completed at least 8 out of the 12 test trials were included in the final sample. An additional 9 infants were tested but excluded due to failing to complete at least 8 test trials (6) or experimenter error (3). The

research was approved by the University of Wisconsin IRB and parental informed consent was obtained for all infants in all experiments.

Stimuli

We created two artificial languages, each consisting of three trisyllabic words (see Table 1). Four of the 9 syllables from each language were the same across the two languages, a feature relevant for Experiments 2-4. We recorded and manipulated the syllables using the pseudosynthesis method previously used by Graf Estes, Evans, Alibali, & Saffran (2007; see also Antovich & Graf Estes, 2016; Graf-Estes & Lew-Williams, 2015; Lew-Williams & Saffran, 2012). To create the speech streams, each syllable of each word was recorded by a single female native English speaker in a monotone voice in every trisyllabic context in which it would appear in the artificial language. Doing so allowed us to preserve coarticulation while ensuring that acoustic cues to word boundaries were not provided. Each syllable was then manually extracted from the recording using Praat (Boersma, 2002), normalized for pitch, duration, and intensity, and concatenated together into a continuous stream. All speech streams used in this study are available at: <https://osf.io/6xjqc/>.

Speech streams were created by concatenating the words from the language in a random order (with the constraint that the same word could not be presented twice in a row) without any acoustic cues to word boundaries. Each speech stream included 30 repetitions of each word. The only cue to word boundaries was the transitional probabilities (TPs) between syllables, which were 1.0 within words. TPs for syllables at word boundaries ranged from 0.43 to 0.53 (mean of 0.49 for each language). The transitional probabilities in each language reliably marked the word boundaries (i.e., the TPs at word boundaries were always the lowest in the stream; see Table 1). The beginning and end of the recording was faded in and out to avoid any additional cues to word boundaries. Both speech streams were approximately 1.5 minutes in length.

Test items (see Table 2) consisted of words and positionally-sensitive nonwords presented in isolation. Positionally-sensitive nonwords were created from strings of syllables that never occurred together during familiarization, but maintained the correct syllable position of the words from which they were drawn. For example, the nonword ‘pabutu’ is a novel string that maintains the syllable position of each syllable from the original words (i.e. ‘pa’ only occurred at the onset, ‘bu’ occurred medially, and ‘tu’ occurred in word final position). Thus, although the syllable positions were maintained across words and nonwords, the transitional probabilities between the syllables in nonwords were always zero. Discriminating between the two types of items therefore required infants to track syllable-to-syllable probabilities and not within-word syllable positions.

Different words and nonwords were used in the test of each language. The test items were recorded in isolation and were normalized for pitch, intensity, and duration. Two words and two nonwords were presented during the test phase. Each word and nonword occurred once within each block of four test trials. There were three test blocks, for a total of 12 test items. All test stimuli are available at <https://osf.io/6xjqc/>.

Procedure

Infants sat on their caregiver’s lap facing a center monitor in a sound attenuated testing booth, with two additional monitors to the left and right of the infant. Caregivers wore headphones that played music to ensure they could not hear the experimental stimuli. The experiment consisted of two phases, a 90-second familiarization phase in which the speech stream was played over speakers in the booth, and a 12-item test phase that measured looking behavior using the Headturn Preference Procedure. The experimenter sat outside the booth and used custom software (WISP) to code the infants’ head turns using corresponding button presses. The experimenter was unaware of which stimulus was played on each trial.

During familiarization, while participants listened to the continuous speech stream, a light flashed on one of the 3 monitors in response to the infant's looking behavior to expose them to the contingency between their looking behavior and the activation of each monitor. The flashing light would always begin in the center monitor. When the infant looked to the light flashing on the center monitor, the light would deactivate and start flashing on one of the two side locations. If the infant looked to the side monitor, the light on that monitor remained active until the infant disengaged attention for longer than 2 seconds, at which point the light activated at the center location again. The test phase was similar; once participants looked to the light flashing on the center monitor, the light would activate on one of the two side locations. In contrast to familiarization, the test item played only when the infant looked to the side monitor, and the length of time the infant looked was recorded. Similar to familiarization, when infants looked away for longer than 2 seconds, the presentation of the test item ended. If an infant attended to a test item for less than 2 seconds, the item was repeated at the end of the test phase. Side of stimulus presentation was counterbalanced.

Results and Discussion

Summary spreadsheets of all data as well as an R script of the analyses are available at <https://osf.io/6xjqc/>. Figure 1 presents infants' mean looking times to both types of test items. To assess learning, we compared infants' mean looking preferences to words ($M = 6.85$, $SD = 2.03$) vs. nonwords ($M = 7.91$, $SD = 3.32$) using a paired samples t-test. Results showed infants looked longer to nonwords than to words [$t(23) = -2.47$, $p = 0.02$, $d = -0.50$]. These results provide evidence that infants were sensitive to the statistical regularities in the speech stream when presented with only a single language during exposure.

Exploratory analyses. We conducted several additional analyses to examine if other factors beyond the main hypothesis were related to infants' looking behavior. As these were unplanned, these are considered exploratory and should be interpreted with caution.

We examined whether looking preferences were affected by familiarized language, age, or outliers. To do so, we calculated a mean difference score for each infant by subtracting looking times to words from looking times to nonwords. To examine if there were differences in preference between the two languages tested, we compared difference scores for Language A [Words: $M = 6.87$, $SD = 2.70$; Nonwords: $M = 8.24$, $SD = 4.33$; Diff. Score: $M = 1.38$, $SD = 2.43$] to difference scores for Language B [Words: $M = 6.83$, $SD = 0.87$; Nonwords: $M = 7.51$, $SD = 1.62$; Diff. Score: $M = 0.68$, $SD = 1.66$] using an independent samples t-test. There were no significant differences in preferences between the two languages [$t(22) = 0.8$, $p = 0.43$, $d = 0.33$]. We also assessed if difference scores varied by age. We found no significant correlation between age and looking time difference scores [$r(22) = 0.005$, $p = 0.98$]. Finally, we identified the presence of outliers by assessing if difference scores were 2 SDs above or below the mean difference score of the sample. One infant was identified as an outlier. When we removed this infant, the results of the analyses were unchanged.

Consistent with previous infant statistical word learning research (e.g., Saffran, et al., 1996), the results from Experiment 1 demonstrate that infants can successfully track the statistical patterns in a single language, as infants showed a novelty preference for the positionally-sensitive nonwords relative to words at test. In Experiment 2, we used these same materials to test infants' learning when the two languages were presented in succession.

Experiment 2

Having established that infants could learn a single language after 1.5 minutes of familiarization, Experiment 2 tested learning when two languages were presented in succession.

Importantly, the two languages were incongruent, and were structured such that when treated as a single input source, their statistical patterns were noisier and less reliable with respect to word boundaries. Successful learning therefore required infants to separately track the statistics of each language.

All infants were presented with both languages from Experiment 1 back-to-back, with no facilitative manipulations: the languages were spoken by the same speaker and there was no pause between the languages. The first presented language will hereafter be referred to as the *initial language*, and the second presented language will hereafter be referred to as the *final language*. There were two between-subject conditions that varied only in which language was tested. In the Initial Language condition, infants were tested on their learning of the initial language presented during familiarization, while in the Final Language condition, infants were tested on their learning of the final language presented during familiarization.

If infants can track each set of regularities separately, then participants in both conditions should exhibit learning. However, it was possible that participants would exhibit a primacy effect, only exhibiting learning in the Initial Language condition (as observed in similarly-structured adult studies; Gebhart et al., 2009; Bulgarelli & Weiss, 2016). Alternatively, infants might exhibit a recency effect, only exhibiting learning in the Final Language condition. Finally, if infants collapse the transitional probabilities over the input from both languages, amalgamating the statistics across the familiarization phase, then we would not expect to see any evidence of learning.

Participants

Two new groups of infants were tested in Experiment 2. A total of 24 infants (11 females) ranging in age from 7.5 to 9.3 months (Mean Age = 8.11, SD = 0.49) were included in the Initial Language condition, and a separate group of 24 infants (12 females) were tested in the

Final Language condition, ranging in age from 7.43 to 8.7 months (Mean Age = 8.08, SD = 0.4). All inclusion criteria were set the same as in Experiment 1. Across both conditions, an additional 25 infants were tested but excluded for fussing out before reaching the test phase (7), not meeting our 8 test trial minimum criterion (16), or computer error (2).

Stimuli

Stimuli were the same as those used in Experiment 1, except that the two familiarization languages were concatenated. There was no break between the two languages, and the entire familiarization stream was 3 minutes in length (1.5 min for each language). The beginning and end of the recording was faded in and out to avoid any additional cues to word boundaries. In contrast to Experiment 1, since the languages shared four of the nine syllables, the transitional probabilities collapsed across both languages were unreliable, as a subset of within-word probabilities matched those found at word boundaries. As can be seen in Table 2, while ‘pa’ could only be followed by a single syllable within each language (‘bi’ in language A, ‘di’ in language B), if the statistics were collapsed across the two languages, then the transitional probabilities declined (as ‘pa’ could now be followed by one of two syllables). Thus, this was a stringent test of whether infants could track syllables separately across each of the two languages, as infants could not succeed if they collapse statistics across the two languages. Test items were the same as those used in Experiment 1.

Procedure

The procedure was the same as Experiment 1, except that infants were exposed to each language, one after the other. Half of the infants heard Language A followed by Language B, and the other half heard Language B followed by Language A. As infants were only tested on one of the two languages (depending on condition), participants heard the same number of test items as

in Experiment 1. As language order was counterbalanced across participants, half of the infants in each condition were thus tested using test trials from Language A and half from Language B.

Results and Discussion

As in Experiment 1, we compared infants' looking times to words to looking times to nonwords, for each condition separately. Figure 1 presents the looking time means. In the Initial Language condition, infants listened to words from the initial language for 6.18 sec ($SD = 2.17$) and nonwords for 6.79 sec ($SD = 1.95$), a difference which was not statistically significant [$t(23) = -1.37, p = 0.18, d = -0.28$]. In the Final Language condition, infants listened to words in the final language for 6.59 sec ($SD = 1.87$), and nonwords for 7.12 sec ($SD = 2.35$), a difference which was also not statistically significant [$t(23) = 1.41, p = 0.17, d = -0.29$]. Thus, we did not find evidence that infants were able to successfully track the syllable sequences from either the initial or the final language.

Exploratory analyses. Next, we explored if the results were affected by familiarization language, age, or outliers. For the Initial Language condition, looking time difference scores were similar across the two language orders [Language A first vs. Language B first: $t(22) = 0.14, p = 0.89, d = 0.06$] and there was no significant correlation between age and difference scores [$r(22) = 0.04, p = 0.84$]. There were no outliers in the sample of infants tested in the Initial Language condition. For the Final Language condition, we found no differences in looking time difference scores between the two language orders [$t(22) = -0.22, p = 0.83, d = -0.09$] and no significant correlation between age and difference scores [$r(22) = -0.08, p = 0.69$]. We identified one outlier whose difference score was more than two standard deviations from the mean; the results remained unchanged when this outlier was removed from the analysis.

The results from Experiment 2 suggest that infants were unable to track the statistics of the initial and the final language separately when the two languages were presented consecutively. There are multiple factors that could explain a failure to learn in our two conditions. In the Initial Language condition, one possibility is that the delay between the end of the initial language and testing, which was due to the presentation of the final language, may have led to forgetting. However, infants also failed to demonstrate learning in the Final Language condition, where there was no delay between the presentation of the final language and the test (similar to Experiment 1). Another possibility is that infants remembered what they encountered in both languages, but aggregated the statistics across the two speech streams. Given the incongruent structure of the two languages, aggregating the transitional probabilities across the speech streams would have resulted in noisier statistics that would pose a challenge for statistical learning, which would explain infants' failure at test. Although it is not clear what led to infants' failure at test, the results suggest that infants did not separate and track the statistics of each language independently.

Given the role of contextual cues in supporting adults' learning of two artificial languages (Gebhart et al., 2009; Weiss et al., 2009), it may not be surprising that we failed to find evidence for segmentation in a condition where no contextual cues were provided. Thus, in Experiment 3, we added contextual cues to aid infants in separating the two languages during learning. In addition, as it may be difficult for 8-month-old infants to retain information learned during the initial language through exposure to the final language, we only tested infants on learning of the final language (Final Language condition). If contextual cues support infants' segmentation of two speech streams, as previously found in studies with adults, then infants in Experiment 3 should be successful at learning the final language.

Experiment 3

In Experiment 3 we exposed infants to the same two incongruent languages, but provided infants with contextual cues that may help learners track multiple sets of statistical regularities (Gebhart et al., 2009). As previous studies with adults suggest that not all correlated cues allow learners to separately track multiple streams of speech, we provided infants with different types of cues in two between-subjects conditions. In the Pitch condition, one of the languages was pitch shifted to sound as if it was spoken by a different speaker. In adults, pitch shifting one of the languages increased learning of the second presented language, but was not sufficient to overcome the primacy effect (Gebhart et al., 2009). As 7.5-month-old infants exhibit difficulty generalizing across speakers of different genders (Houston & Jusczyk, 2000), this type of cue may be sufficient to signal a difference in the speech stream for 8-month-old infants. In addition to the pitch shift, the two speech streams were also separated by a brief pause. For adults, a pause inserted between two speech streams, together with explicit instructions cueing adults that they would be learning two languages, led to learning of the second presented speech stream (Gebhart et al., 2009). As it is not possible to explicitly inform 8-month-olds that they will be learning two languages, we chose to combine the pitch shift and pause cues.

In the Accent condition, we used a more ecologically valid set of cues to aid infants in tracking each speech stream separately. While previous research with adults (for speech segmentation) and infants (for rule learning) has shown that speaker cues are useful for learning multiple regularities, particularly when these regularities overlap in features (Weiss et al. 2009; Kovacs & Mehler, 2009b), talkers on their own are not typically a good cue to language identity, as individuals hear a multitude of talkers produce the same language. Changes in phonological inventory, however, may be more indicative of the presence of different languages. In this condition, while one of the languages continued to be presented in an English-accented voice, we recorded the same syllables using a Spanish-accented voice produced by a different female

native speaker of Spanish. This manipulation provided a combination of speaker and accent cues, which may be more indicative of the presence of multiple speech streams, and therefore a strong signal to track the syllable statistics of each stream separately. In addition to the speaker and accent cues, we maintained the short pause between the two languages as in the Pitch condition. The key question was whether these cues, which are relevant to the presence of different structures in real-world environments, can support infants' ability to find the words in the second language.

In both conditions, we only manipulated the initial language so as to keep the final language the same as in Experiment 2. We also only tested learning of the final presented language (as in Experiment 2, Final Language condition), which allowed us to use the same test materials as in Experiment 1 and the Experiment 2-Final Language condition. If the presence of contextual cues aids infants' ability to separately track the statistics of the two speech streams, as they do in adults, then we should find successful learning of the second language.

Participants

Two new groups of infants were included in Experiment 3. Twenty-four infants (12 females, ranging from 7.7 – 8.7 months, mean = 8.08, $SD = 0.33$) were tested in the Pitch condition and 24 infants (17 females, ranging from 7.47-8.97 months, mean = 8.18, $SD = 0.4$) were tested in the Accent condition. The same exclusion criteria as Experiments 1 and 2 were used. An additional 19 infants were tested but excluded due to fussing out before the test phase (3), failing to meet our 8 test trial minimum criterion (12), experimenter error (3), or inattentiveness (1).

Stimuli

Pitch condition. The stimuli were the same as those used in Experiment 2, except that the first presented language was pitch shifted down by 60 Hz (using Audacity) to sound as if it was spoken by a different speaker with a lower voice. A 5 second pause was also inserted between the initial and the final language. The beginnings and ends of both languages were faded in and out to avoid perceptual cues to word boundaries.

Accent condition. The stimuli were the same as those used in Experiment 2, except for the first presented language which was recorded by a new native Spanish speaking female. The speaker was instructed to produce the target phonemes using their native language phonology of Spanish. The recordings and speech stream were created in the same manner as the ones used in previous experiments. This resulted in a speech stream comprised of Spanish-accented phonemes that was otherwise identical to the streams used in Experiments 1 and 2.

In both conditions, participants heard both languages in succession, with the manipulated language (pitch shifted or accented) presented first. This ensured that across all experiments, participants were tested on the same version of the language (those recorded by the original native speaker of English) using the same test items as Experiments 1 and 2. After listening to both languages, participants were again tested using words and nonwords from the final language.

Procedure

The procedure was the same as that used in the Final Language condition in Experiment 2. Infants were tested on their listening preferences using the Headturn Preference Procedure. Half of the participants heard Language A first, while the other half heard Language B first, and in all cases only the initial language was either pitch shifted or accented.

Results and Discussion

First, we assessed performance in the Pitch condition (words: $M = 7.74$, $SD = 2.12$; nonwords: $M = 7.52$, $SD = 2.21$; see Figure 1). Results showed no difference in looking between words and nonwords [$t(23) = 0.63$, $p = 0.53$, $d = 0.13$]. We next assessed performance in the Accent condition (words: $M = 7.18$, $SD = 2.17$; nonwords: $M = 7.65$, $SD = 2.38$). As in the Pitch condition, results in the Accent condition showed no difference in looking between words and nonwords [$t(23) = -1.24$, $p = 0.23$, $d = -0.25$].

Exploratory analyses. For the Pitch condition, the difference scores revealed no differences between language order [Language A first vs. Language B first: $t(22) = 0.79$, $p = 0.44$, $d = 0.32$] and no significant correlation with age and difference scores [$r(22) = 0.21$, $p = 0.32$]. However, when we removed 2 outliers whose difference scores were more than 2SD from the mean, the results revealed statistically significantly longer looking times to nonwords vs. words [$t(21) = 2.10$, $p = 0.05$, $d = 0.5$] and additionally, a positive and statistically significant correlation between age and difference scores [$r(20) = 0.46$, $p = 0.03$]. Older infants exhibited a larger preference for nonwords. All other results remained the same.

In the Accent condition, exploratory analyses revealed similar difference scores between language orders [$t(22) = -0.65$, $p = 0.52$, $d = -0.27$] and no correlation with age [$r(22) = 0.27$, $p = 0.2$]. However, when we removed 2 outliers, we found a positive and significant correlation between age and difference scores [$r(20) = 0.43$, $p = 0.04$], such that older infants exhibited a larger preference for nonwords (all other results remained the same).

Given that both groups of infants exhibited a positive correlation between age and difference scores when outliers were removed, we combined the data for the two contextual cue conditions, without outliers ($N = 44$; see Figure 2). We then explored whether there was a significant difference between words and nonwords separately for older infants and younger

infants. We grouped infants based on age using a median split (median age of the two conditions without outliers = 8.12). Paired sample t-tests for the younger infants revealed a significant preference for words over nonwords [$t(21) = 2.37, p = 0.03, d = 0.51$]. In contrast, results for the older infants revealed a significant preference of nonwords over words [$t(21) = -2.14, p = 0.04, d = -0.46$].

In sum, Experiment 3 provided mixed results with respect to infants' abilities to segment two speech streams with the addition of contextual cues. With all infants included, there was no evidence of tracking the statistics of the second presented language. However, excluding outliers revealed a slightly different story. In the Pitch condition, infants as a group demonstrated successful learning of the statistics in the second language, exhibiting a preference for nonwords over words. Additionally, in both the Pitch and Accent conditions, there was a significant correlation between age and preference, with older infants showing a preference for nonwords over words and younger children showing a preference for words over nonwords.

These results suggest that contextual cues may help infants track each language separately, but that the direction of preference may be related to age. In particular, the differences in preference between the younger and older infants may have led to the null results found when the entire group of infants was assessed. However, given that we arrived at this conclusion through unplanned analyses examining correlations and subsets of the infants after outliers were excluded, our goal in the last experiment was to confirm if indeed age may be a factor in infants' ability to learn the second of two speech streams with the presence of contextual cues. In Experiment 4, we conducted a pre-registered study to test older infants' ability to track the statistics of the second language with accent cue added.

Experiment 4

The exploratory analyses of Experiment 3 yielded the prediction that older infants would be successful at segmenting the incongruent speech streams, providing the input contained salient contextual cues corresponding to the change in structure. We therefore designed and pre-registered Experiment 4 to test older infants using the same materials as the Accent condition in Experiment 3 [<https://osf.io/7bru4/>]. We chose 8.5 to 10.5-month-old infants, as this age group encompasses the older infants from the previous three experiments. If older infants are successful at using the contextual cues to track the statistics of the two speech streams separately, then we should find differences in looking times between words and nonwords.

Participants

A total of 24 infants (16 females) participated in Experiment 4. Participants ranged in age from 8.6 to 10.6 months (Mean Age = 9.4, $SD = 0.49$). An additional 12 infants were tested but excluded due to not completing at least 8 test trials (10), experimenter error (1), or computer error (1).

Stimuli and procedure

The stimuli and procedure were the same as that used in the Accent condition in Experiment 3. Half of the infants heard Language A first, and half heard Language B first. The first language was always the Spanish accented speech stream. As in the prior studies, only the second language was tested using the same test materials as in Experiments 1-3.

Results and Discussion

Comparison of words ($M = 7.44$, $SD = 2.1$) to nonwords ($M = 7.33$, $SD = 2.34$) revealed no difference in looking time [$t(23) = 0.25$, $p = 0.81$, $d = 0.05$]. In pre-registered confirmatory analyses examining the effects of language, age, and outliers, results revealed that difference scores did not significantly differ between language orders [Language A first vs. Language B

first: $t(22)=1.3$, $p = 0.21$, $d = 0.53$], and that there was no significant correlation with age [$r(22) = 0.05$, $p = 0.81$]. All results remained the same when 1 outlier was removed from the sample.

Contrary to our predictions, the results suggest that older infants did not succeed at learning the words of the second language when a pause, speaker, and accent cue was added to the first presented language. We return to these results in the General Discussion.

General Discussion

In this study, we sought to investigate whether monolingual infants can separately track two incongruent speech streams when the speech streams were presented in succession. Overall, our results failed to demonstrate strong evidence of learning of either of the speech streams presented to infants. When presented with a single language (Experiment 1) we found that 8-month-old infants were sensitive to the statistical properties of that language after 1.5 minutes of exposure. However, when infants were presented with both languages back-to-back without any cues to signal the presence of two different speech streams (Experiment 2), infants did not exhibit learning of either the first or the second presented language. Further, the addition of contextual cues (Pitch condition or Accent condition) distinguishing the two languages did not appear to facilitate infants' learning of the Final Language. Finally, although exploratory analyses in Experiment 3 suggested that age might play a role in boosting infants' ability to track the transitional probabilities in the second language with contextual cues, our final study with 8.5-10.5-month-old infants did not confirm this hypothesis. Across all experiments we used the same test materials, changing only what infants listened to at training. Thus, infants' failure to show significant test discrimination in Experiments 2-4 is unlikely to be due to the test stimuli themselves, since infants successfully discriminated these stimuli in Experiment 1. Overall, our findings indicate that infants raised in monolingual homes have difficulty acquiring the statistics of multiple incongruent speech streams.

A key feature of the artificial languages that infants heard in our studies was that they were incongruent, meaning that the combined statistics of the two languages reduced the transitional probability cues that marked word boundaries. One explanation for infants' failure in our task is that they aggregated the statistics of these two speech streams, rather than tracking them separately. Monolingual adults find it difficult to track the statistics of two incongruent speech streams in the absence of contextual cues, exhibiting either a primacy effect (Gehbart et al., 2009) or no learning (Weiss et al., 2009). The results of Experiment 2 demonstrate that infants also struggle in this task, learning neither the initially presented nor the final language. In this experiment, infants were not given any salient cues to signal the presence of two separate speech streams and aggregating them might be a reasonable strategy. In Experiments 3 and 4, however, infants were given different cues – including pitch and accent changes, as well as a pause between speech streams – and they nonetheless failed to provide robust evidence of learning. For adults, such contextual cues appear sufficient to trigger them to track the statistics of two incongruent speech streams separately, but such cues do not appear to be sufficient for monolingual infants of this age.

Research using other paradigms has reported convergent evidence that monolingual infants find it difficult to track two sets of linguistic regularities. For example, in a rule learning paradigm, 12-month-old infants were taught two linguistic rules: that syllable sequences with an AAB structure would predict a reward on the left side of the screen, and syllable sequences with an ABA structure would predict a reward on the right side of the screen. Monolinguals found this task challenging, only learning one rule, while bilinguals were able to learn both rules (Kovacs & Mehler, 2009b). However, when provided with a contextual cue of a change in speaker gender, monolingual infants were also able to learn both rules. In a statistical word segmentation task akin to the current study, 14-month-old monolinguals failed to segment two speech streams,

while bilinguals succeeded (Antovich & Graf Estes, 2017). This study was different from ours in several ways: infants were somewhat older, the speech streams were congruent and presented in alternating blocks during training, speaker gender was used as a contextual cue, and infants were tested on words vs. partwords. Our study makes a unique contribution by showing that monolinguals' difficulty generalizes beyond this specific testing context. We found that this failure is consistent even for younger infants, when the speech streams are incongruent, when they are presented in a blocked fashion, and when infants are exposed to the speech streams across a variety of conditions, with and without contextual cues (including pitch-shifting and accents). Our findings, together with this previous research, provide growing evidence that monolingual infants between 8-14 months old find it difficult to track the statistics of two speech streams, at least in lab tasks with brief exposure durations and disembodied voices.

While our main interest concerned testing infants' ability to learn two incongruent speech streams, we cannot rule out the possibilities that some of the incidental features of our design might have contributed to infants' learning difficulties, and thus our null results across several experiments. First, while Experiment 1 demonstrated that infants could learn each of our two languages separately, we do not know whether the pitch-shifted and accented versions of these languages used in subsequent experiments might have been more difficult to learn, despite having the same statistical regularities. If these manipulated languages were harder to learn, presenting these initially could have hindered learning of the second language (similar to Jungé, Scholl, & Chun, 2007). Second, we chose to block exposure to each language rather than interleave the two languages. Recent work with both infants and adults suggests that multiple transitions between languages can alert learners to the presence of multiple structures (Orena & Polka, 2019; Polka, Orena, Sundara, & Worrall, 2016; Zinszer & Weiss, 2013), although related work found that infants' learning was better with long blocks relative to interleaved stimuli

(Gonzalez et al., 2015). Third, our exposure phase was 1.5 minutes for each language. The results from Experiment 1 show that this is sufficient for statistical learning when infants hear a single speech stream, but we cannot rule out the possibility that this exposure length is insufficient when there is increased variance in the input due to the presence of two incongruent speech streams with twice as many regularities to track.

Aspects of our test phase design could also have inadvertently impacted our findings. Because of infants' limited attention span, it was only possible to test each infant on one of the two languages they encountered. In Experiment 2, we tested infants on either the initial or the final language, and found no evidence of learning in either case. In Experiments 3 and 4 we chose to focus on testing the final language, as it was encountered immediately prior to the test. Thus, it is possible that infants were able to learn the initial, untested language in Experiments 3 and 4, when contextual cues were added. Moreover, our test items included words and positionally-sensitive nonwords. Although we found that infants were successful at discriminating positionally sensitive nonwords from words in Experiment 1, it is possible that these types of items are particularly difficult to discriminate from words in more complex statistical learning tasks, such as when two speech streams are present. In sum, this discussion highlights the many decisions that researchers confront when designing statistical learning studies involving multiple speech streams, with very little research to guide methodological choices. Examining the effects of exposure length, presentation format, and choice of test items during statistical learning of multiple speech streams are avenues for future research and may provide insights into the capacities and limits of infants' statistical learning.

Our investigation explored two of many possible contextual cues that might be available to infants: pitch shifting and speaker accent, which are both examples of talker-specific information. The literature on whether talker-specific information is used to track multiple

statistical regularities is mixed, with some studies finding that talker-specific information is helpful (Kovacs & Mehler, 2009b; Weiss et al., 2009) while others, including our study, finding that it does not benefit learning (Potter & Lew-Williams, 2019; Tsui, Erickson, Thiessen, & Fennell, 2017). In particular, speaker gender is often a perceptually salient cue in speech, but it typically does not signal a change in statistical information. Gender appears to affect infants' speech processing, albeit sometimes in non-adaptive ways. For example, Houston and Jusczyk (2000) showed that 7.5-month-old infants had no difficulty recognizing speech produced by a novel speaker of the same gender, but struggled with speakers of different genders. Thus, whether speakers are the same or different genders may be particularly important for infants to determine whether to segregate or combine the input. In our Pitch Shift condition, it is unclear whether infants perceived the streams as being spoken by different speakers, which may have led to the marginal effect in our exploratory analyses in the Pitch Shift condition. Even if participants perceived these as different speakers, aggregating across multiple speakers may be a more adaptive strategy, in particular for monolinguals who don't have prior experience with a change in speaker resulting in a change in statistics. An avenue for future research is understanding the conditions under which infants use talker-specific information to track the regularities of multiple speech streams.

In contrast, speaker accent is a cue in speech that is more likely to signal a change in statistical information, as speakers who use different phonological inventories are often speaking different dialects or languages. It is therefore surprising that the novel phonological inventory provided in the Accent condition did not lead to infants exhibiting learning. However, it is possible that the monolingual infants we tested lacked everyday experience with varying phonological inventories, and thus did not know how to use the accent cue.

While our study focused on speaker cues including pitch and accent, we note that there are many other potential contextual cues that infants might use to track multiple speech streams. For example, Mitchel & Weiss (2010) found that using different synchronized talking faces helped adults track the statistics of two speech structures. It is possible that the addition of visually correlated cues, such as faces, may also benefit infants' tracking of two presented streams of speech. At the same time, not all cues are equally informative for adults: in the same study, encountering still images of different faces and different background colors on a computer monitor did not support the segmentation of multiple speech streams (Mitchel & Weiss, 2010). Given that the complexity of the learning environment and talker-specific information may interact, future research should take both factors into consideration to better understand how infants and adults track the statistics of multiple inputs.

Our study tested monolingual infants who may have had limited experience tracking multiple linguistic structures in their everyday speech input. In contrast, bilinguals must separately track the statistics in each of their languages in order to successfully acquire them. Some studies have found experience learning multiple languages may enable both infants (e.g., Antovich and Graf-Estes, 2018; Kovacs & Mehler, 2009b) and adults (e.g. Wang & Saffran, 2014) to track multiple sets of regularities. Bilingualism may benefit the tracking of transitional probabilities in a complex word segmentation task (Wang & Saffran, 2014), as well as the acquisition of multiple mappings in statistical word-referent paradigms (Benitez, Yurovsky, & Smith, 2016; Poepsel & Weiss, 2016). However, several studies examining multi-language segmentation in adult bilinguals have failed to find an advantage for bilinguals relative to monolinguals (e.g., Bogulski, 2013; Bulgarelli & Weiss, 2016). Whether bilingual experience benefits infants' statistical learning of multiple incongruent speech streams is an open question for future research.

One potentially interesting type of language experience that has received less empirical attention is exposure to multiple dialects. Dialects are close varieties of the same language, with statistical structures that likely contain both congruent and incongruent elements. For example, in African American English, it is common to say “dis” for the Standard American English “this”. The syllable “dis” may therefore signal a word boundary within African American English but not in Standard American English. Further, across both dialects, the syllable “dis” can occur as the first syllable of a multisyllable word (e.g., “discover”, “distinct”, “discuss”). Language learners exposed to both types of dialects must therefore be able to both aggregate and separate the statistics across encounters with speech in both dialects (Qian, et al., 2016). Indeed, one recent study that included bidialectal adults indicated that experience with multiple dialects affected whether listeners treat regularities from each of their dialects separately (Lundquist & Vangsnes, 2018). While our study did not include bilingual or bidialectal infants, future work would benefit from continuing to examine how infants with different language experiences (multiple languages, dialects, or accents) adapt their statistical learning to accommodate multiple sets of regularities. Moreover, this research would benefit from testing infant second language or second dialect learners (for example infants with exposure to a new language or dialect upon entering daycare) who are confronted with the real-life task of detecting and learning a new set of regularities.

In conclusion, we set out to ask if monolingual English learning infants were capable of segmenting two incongruent speech streams presented in succession with and without contextual cues. Our findings suggest that monolingual infants found this task to be difficult, as we did not find reliable evidence of successful learning even when contextual cues were added to signal the presence of two speech streams. These findings highlight the need to better understand how

infants with different linguistic experiences track the statistical information of multiple speech streams separately.

References

- Antovich, D. M., & Graf Estes, K. (2018). Learning across languages: Bilingual experience supports dual language statistical word segmentation. *Developmental Science*, 21(2), e12548.
- Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, 9(4), 321-324.
- Benitez, V. L., Yurovsky, D., & Smith, L. B. (2016). Competition between multiple words for a referent in cross-situational word learning. *Journal of Memory and Language*, 90, 31-48.
- Black, A. & Bergmann, C. (2017). Quantifying infants' statistical word segmentation: A meta-analysis. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, pp. 124-129, Austin, TX: Cognitive Science Society.
- Boersma, P. (2002). Praat, a system for doing phonetics by computer. *Glott international*, 5.
- Bogulski, C. A. (2013). Are bilinguals better learners. *A neurocognitive investigation of the bilingual advantage (Doctoral dissertation)*. The Pennsylvania State University, PA. Retrieved from ProQuest Dissertations and Theses Database.(Accession Order No. AAT 3585561).
- Bulgarelli, F., & Weiss, D. J. (2016). Anchors aweigh: The impact of overlearning on entrenchment effects in statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(10), 1621.
- Byers-Heinlein, K. (2014). Languages as categories: Reframing the “one language or two” question in early bilingual development. *Language Learning*, 64(s2), 184-201.
- Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, 12(6), 499-504.

- Gebhart, A. L., Newport, E. L., & Aslin, R. N. (2009). Statistical learning of adjacent and nonadjacent dependencies among nonlinguistic sounds. *Psychonomic bulletin & review*, 16(3), 486-490.
- Gomez, R. L., & Gerken, L. (1999). Artificial grammar learning by 1-year-olds leads to specific and abstract knowledge. *Cognition*, 70(2), 109-135.
- Gonzales, K., Gerken, L., & Gómez, R. L. (2015). Does hearing two dialects at different times help infants learn dialect-specific rules?. *Cognition*, 140, 60-71.
- Graf Estes, K., & Lew-Williams, C. (2015). Listening through voices: Infant statistical word segmentation across multiple speakers. *Developmental psychology*, 51(11), 1517.
- Graf Estes, K., Evans, J. L., Alibali, M. W., & Saffran, J. R. (2007). Can infants map meaning to newly segmented words? Statistical segmentation and word learning. *Psychological Science*, 18(3), 254-260.
- Houston, D. M., & Jusczyk, P. W. (2000). The role of talker-specific information in word segmentation by infants. *Journal of Experimental Psychology: Human Perception and Performance*, 26(5), 1570.
- Johnson, E. K., & Jusczyk, P. W. (2001). Word segmentation by 8-month-olds: When speech cues count more than statistics. *Journal of Memory and Language*, 44(4), 548-567.
- Jungé, J. A., Scholl, B. J., & Chun, M. M. (2007). How is spatial context learning integrated over signal versus noise? A primacy effect in contextual cueing. *Visual Cognition*, 15(1), 1-11.
- Kovács, Á. M., & Mehler, J. (2009a). Cognitive gains in 7-month-old bilingual infants. *Proceedings of the National Academy of Sciences*, 106(16), 6556-6560.
- Kovács, Á. M., & Mehler, J. (2009b). Flexible learning of multiple speech structures in bilingual infants. *Science*, 325(5940), 611-612.

- Lew-Williams, C., & Saffran, J. R. (2012). All words are not created equal: Expectations about word length guide infant statistical learning. *Cognition*, 122(2), 241-246.
- Lundquist, B., & Vangsnes, Ø. A. (2018). Language Separation in Bidialectal Speakers: Evidence From Eye Tracking. *Frontiers in Psychology*, 9, 1394.
doi:10.3389/fpsyg.2018.01394
- Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, 82(3), B101-B111.
- Mitchel, A. D., & Weiss, D. J. (2010). What's in a face? Visual contributions to speech segmentation. *Language and Cognitive Processes*, 25(4), 456-482.
- Orena, A. J., & Polka, L. (2019). Monolingual and bilingual infants' word segmentation abilities in an inter-mixed dual-language task. *Infancy*, 24(5), 718-737.
- Pajak, B., Fine, A. B., Kleinschmidt, D. F., & Jaeger, T. F. (2016). Learning additional languages as hierarchical probabilistic inference: insights from first language processing. *Language Learning*, 66(4), 900-944.
- Poepsel, T. J., & Weiss, D. J. (2016). The influence of bilingualism on statistical word learning. *Cognition*, 152, 9-19.
- Polka, L., Orena, A. J., Sundara, M., & Worrall, J. (2017). Segmenting words from fluent speech during infancy—challenges and opportunities in a bilingual context. *Developmental Science*, 20(1), e12419.
- Potter, C. E., & Lew-Williams, C. (2019). Infants' selective use of reliable cues in multidimensional language input. *Developmental Psychology*, 55(1), 1.
- Qian, T., Jaeger, T. F., & Aslin, R. N. (2012). Learning to represent a multi-context environment: more than detecting changes. *Frontiers in Psychology*, 3, 228.

- Qian, T., Jaeger, T. F., & Aslin, R. N. (2016). Incremental implicit learning of bundles of statistical patterns. *Cognition*, 157, 156-173.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27-52.
- Saffran, J. R., & Kirkham, N. Z. (2018). Infant statistical learning. *Annual Review of Psychology*, 69.
- Saffran, J. R., & Thiessen, E. D. (2003). Pattern induction by infant language learners. *Developmental Psychology*, 39(3), 484.
- Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106(3), 1558-1568.
- Surrain, S., & Luk, G. (2017). Describing bilinguals: A systematic review of labels and descriptions used in the literature between 2005–2015. *Bilingualism: Language and Cognition*, 401-415.
- Tsui, A. S. M., Erickson, L. C., Thiessen, E. D., & Fennell, C. T. (2017). Statistical Learning from Accented Speech: A Bilingual Advantage. *Proceedings of the 41st annual Boston University Conference on Language Development*. (Ed. Maria LaMendola and Jennifer Scott). Somerville, MA: Cascadilla Press, 679-690.
- van Heugten, M., Krieger, D. R., & Johnson, E. K. (2015). The developmental trajectory of toddlers' comprehension of unfamiliar regional accents. *Language Learning and Development*, 11(1), 41-65.
- Wang, T., & Saffran, J. R. (2014). Statistical learning of a tonal language: The influence of bilingualism and previous linguistic experience. *Frontiers in psychology*, 5, 953.

- Weiss, D. J., Gerfen, C., & Mitchel, A. D. (2009). Speech segmentation in a simulated bilingual environment: A challenge for statistical learning?. *Language Learning and Development*, 5(1), 30-49.
- Zinszer, B., & Weiss, D. (2013). When to hold and when to fold: detecting structural changes in statistical learning. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 35, No. 35).

Table 1: Language structure and test items

Language	Words	Test words	Test nonwords
A	P abiku	Pabiku	Pabutu
	T ibudo	Tibudo	Gobido
	G olatu		
B	B ugofay	Bugofay	Bupaso
	T ifaso	Mupadi	Tigodi
	M upadi		

*Bolded syllables overlap across languages

Table 2: Transitional probabilities if each language is tracked separately (within language) or if the statistics are collapsed across the two languages (across languages). The final number in each column represents the transitional probabilities across word boundaries.

Language	Words	Transitional probabilities within language	Transitional probabilities across languages
A	Pabiku	1.0, 1.0, 0.5	0.5 1.0 0.5
	Tibudo	1.0, 1.0, 0.5	0.5 0.5 0.5
	Golatu	1.0, 1.0, 0.5	0.5 1.0 0.5
B	Bugofay	1.0, 1.0, 0.5	0.5 0.5 0.5
	Tifaso	1.0, 1.0, 0.5	0.5 1.0 0.5
	Mupadi	1.0, 1.0, 0.5	1.0 0.5 0.5

Figure 1. Mean looking time (and standard error of the mean) to Words vs. Nonwords for each condition in each experiment. * $p < 0.05$

Figure 2. Mean looking time difference scores (nonwords – words) by Age for each infant in the Pitch and Accent conditions in Experiment 3. Outliers are not plotted.