Developmental Science 8:1 (2005), pp 1–7

REPORT

Changing the tune: the structure of the input affects infants' use of absolute and relative pitch

Jenny R. Saffran, Karelyn Reeck, Aimee Niebuhr and Diana Wilson

Waisman Center and Department of Psychology, University of Wisconsin–Madison, USA

Abstract

Sequences of notes contain several different types of pitch cues, including both absolute and relative pitch information. What factors determine which of these cues are used when learning about tone sequences? Previous research suggests that infants tend to preferentially process absolute pitch patterns in continuous tone sequences, while other types of input elicit relative pitch use by infants. In order to ask whether the structure of the input influences infants' choice of pitch cues, we presented learners with continuous tone streams in which absolute pitch cues were rendered uninformative by transposing the tone sequences. Under these circumstances, both infants and adults successfully tracked relative pitches in a statistical learning task. Implications for the role played by the structure of the input in the learning process are considered.

The infant's environment is filled with an extraordinary array of information. Even a small sampling of the environment renders myriad cues for young learners to encode and use in the learning process. Like all learners, infants are presumably biased to preferentially use some cues more readily than others given the structure of their perceptual system; for example, given the same aversive conditioning stimulus, rats focus on its taste while quail focus on its visual structure in learning to subsequently avoid it (Garcia & Koelling, 1966). At the same time, features of the input itself likely drive the types of cues detected and used by learners (e.g. Jusczyk, Bertoncini, Bijeljac-Babic, Kennedy & Mehler, 1990; Maye, Werker & Gerken, 2002). In the current study, we sought to determine whether the consistency of a cue influences the degree to which it is exploited by learners. To the extent that this is the case, we may find explanations for how learners avoid the 'richness of the stimulus' problem, that is, how learners are able to focus on relevant information without being misled by the countless irrelevant cues available in the input.

Pitch perception is an ideal domain in which to assess the relative contributions of pre-existing perceptual biases as opposed to the structure of the input. Any given sequence of notes contains two different types of pitch information: absolute pitches (AP), the encoding of a pitch independent of its relation to other sounds, and relative pitches (RP), the intervals between pitches. For example, given a particular rendition of 'Here comes the bride', the listener could represent the first two notes either in terms of absolute pitches (e.g. DG), or relative pitches (P4↑), or both types of cues. While infants show evidence of processing relative pitches in some types of tasks, particularly those in which infants are familiarized with repetitions of brief melodies played in transposition (e.g. Lynch & Eilers, 1992; Trainor & Trehub, 1992, 1993), other tasks elicit absolute pitch perception (Saffran, 2003; Saffran & Griepentrog, 2001).

Why might some sets of input lead infants to show one type of perceptual processing while other tasks elicit the use of a different set of perceptual cues? One possibility concerns an interaction between learners' existing perceptual sensitivities and the structure of the input presented during the learning tasks themselves. This hypothesis has been previously explored in the domain of avian pitch processing (MacDougall-Shackleton & Hulse, 1996). When European starlings were presented with a pitch discrimination task that could be performed using either AP or RP cues, the birds initially solved the task using AP cues, suggesting that birds first use AP cues to categorize novel auditory stimuli. However, when the task was changed to require transfer of the pitch sequences, the birds instead used RP cues. The authors suggest that the birds had access to both types of pitch

Address for correspondence: Jenny R. Saffran, Department of Psychology, University of Wisconsin–Madison, Madison, WI 53706, USA; e-mail: jsaffran@facstaff.wisc.edu

cues, but that the demands of the task, particularly with respect to the structure of the input, determined which dimensions of the auditory stimuli were used during the test discriminations.

It is possible that a similar process occurs for pitch learning in human infants, i.e. infants may have access to both types of pitch cues, and the structure of the input to be learned and/or the test discrimination affects which aspects of pitch are tracked. To test this hypothesis, we adapted a task previously demonstrated to elicit AP perception but not RP perception by infants (Saffran, 2003; Saffran & Griepentrog, 2001). Infants in these experiments listened to a 3-minute sequence of continuous tones (e.g. E C# F G# C G A B C# . . .), and were then tested to determine whether they had represented the tones with respect to their AP or RP patterns. When tested on their knowledge of the AP patterns, infants successfully discriminated the test items; however, they failed to show any evidence of acquiring the RP patterns to which they had been exposed. Adult learners showed a different pattern of performance, consistently tracking RP patterns more readily than AP patterns.

One possible explanation for these results is that infants, like the avian listeners described above, are capable of tracking both AP and RP patterns. However, their perceptual predispositions, combined with the structure of the input presented during the learning task, led them to focus preferentially on the AP cues in the continuous streams of tones. If this hypothesis is correct, then stimuli in which AP cues are rendered unreliable should lead infants to track RP cues instead of AP cues. We thus created tone sequences that, while otherwise very similar to those used in our previous experiments, no longer contained consistent AP sequences. RP sequences, however, remained highly predictable. By essentially removing AP cues as a signal of structure in the input, we can test the hypothesis that infants will now reliably discriminate the test materials using RP cues, something they did not do when AP cues were available during learning.

Our materials consisted of a continuous sequence of tones, organized into three-tone sequences by virtue of the statistical properties of the tone patterns. As in our previous studies, the question of interest was which pitch cues entered into listeners' computations: did listeners track the probabilities with which absolute pitches followed one another, or the probabilities with which relative pitches followed one another? Unlike our prior studies, however, the three-tone sequences were presented in transposition, such that the intervals between tones were consistent across the input corpus but the absolute pitch sequences were highly variable. To test whether listeners had tracked RP patterns, the test discrimination included novel pitch sequences that either followed the same relative pitch patterns as the input materials or instead consisted of novel RP sequences.

Before testing infants with these materials, it was necessary to ascertain whether learners *known* to track relative pitches in tone sequence segmentation tasks could do so given these stimuli. We thus began by testing a group of adult non-musicians. When presented with continuous tone sequences, adults consistently use RP cues even when absolute pitches are available for learning (Saffran, 2003; Saffran & Griepentrog, 2001). If these materials tap relative pitch ability, then adults should show successful learning. Experiment 1 was thus designed to verify that these stimuli are learnable by a group known to exploit RP cues. Experiment 2 then addressed the question of infant access to RP cues for learning about tone sequences.

Experiment 1

Unlike our prior studies, in which both AP and RP cues provided strong statistical pointers to 'tone words' – consistent three-tone sequences – AP cues were far less informative than RP cues in the current experiment. When considered from the perspective of RP sequences, the tone stream contained only three words. However, these words were continuously transposed, so that from the perspective of AP sequences, the tone stream contained nine words; this is more than twice as many words as used in any prior tone sequence segmentation experiment. RP cues were thus a far better predictor of structure than AP cues in the current materials. To ensure that the weak AP cues could not be used for the test discrimination, all of the test items were novel transpositions of the sequences played during familiarization. We thus hypothesized that learners attuned to RP cues, like adults, would readily succeed at this task.

Method

Participants

Eight adults participated in this experiment in exchange for a small gift. All participants were identified as non-musicians (did not self-identify as a musician and had not played an instrument, sung in choruses, or studied music theory since the seventh grade). Each participant was randomly assigned to hear the stimuli from either Condition One or Condition Two.

Materials

The tone sequences included the chromatic pure tones of the octave starting at middle C, including the C above middle C. Each tone was 0.33 seconds in duration, generated using the sine-wave tone generator in CoolEdit on the PC. We constructed two counterbalanced tone streams to serve as the familiarization stimuli for Conditions One and Two. Each stream contained a sequence of tone 'words'; each word consisted of three sine wave tones. The tone words may be described in terms of both their AP patterns and their RP patterns. When considered in terms of absolute pitches, each tone stream consisted of nine tone words. However, when considered in terms of relative pitches, each tone stream consisted of only three tone words. As shown in Table 1, the nine AP sequences in each condition were transpositions of three RP sequences. Thus, from the perspective of RP patterns, each condition included only three words. For clarity of explanation, we will refer to the nine tone words that differ in absolute pitch patterns as AP words, and the three tone words that differ in relative pitch patterns as RP words. The tone words were not constructed in accordance with the rules of standard musical composition and did not resemble any paradigmatic melodic fragments.

To generate the tone stream for each condition, 15 tokens of each AP word (45 repetitions of each RP word) were concatenated together in random order, with the stipulation that the same AP and RP words never occurred twice in a row. Each tone stream was 2 minutes and 13 seconds in length. As in the linguistic and tone sequence materials used in prior segmentation studies (Saffran, 2003; Saffran & Griepentrog, 2001; Saffran, Johnson, Aslin & Newport, 1999; Saffran, Newport & Aslin, 1996), there were no acoustic markers signaling word boundaries. An orthographic representation of the tone stream is analogous to the following: A B F# A# D# G E C G B E G# . . . , etc. Only statistical cues were available to indicate the beginnings and ends of the tone words; AP and RP sequences that occurred within words were more probable than AP and RP sequences that crossed word boundaries.

We constructed a two-alternative forced-choice test designed to determine whether learners acquired the RP words during familiarization. Four new three-tone sequences were generated as test stimuli (see Table 1). All test items were transpositions of RP words from the tone sequences played during familiarization. Specifically, two of the test items contained interval patterns identical to RP words in Condition One, while the other two test items contained interval patterns identical to RP words in Condition Two. For example, one of the RP words in Condition One was P5↓ M3↑, while one of the RP words in Condition Two was M3↑ M2↓. The test items included A D F#, a novel exemplar of the Condition One RP word P5↓ M3↑, and F A G, a novel exemplar of the Condition Two RP word M3↑ M2↓. For a learner in Condition One, A D F# contains a familiar relative pitch sequence while F A G contains a novel relative pitch sequence, which we will call a non-word; the opposite would be true for a learner in Condition Two. Each test trial consisted of a novel exemplar of a Condition One RP word paired with a novel exemplar of a Condition Two RP word. Thus, each trial consisted of a familiar RP word and a non-word; which item was familiar and which was novel depended on the learner's familiarization condition. This between-subjects counterbalanced design ensured that successful performance on the forced-choice task across the two conditions was due to learning, and not to any arbitrary preferences for certain tone sequences.

Table 1 *Tone words and test items for Experiments 1 and 2*

	Absolute pitches	Relative pitches
Condition One:	A# D# G, B E G#, G C E	P5↓ M3↑
	F# D A, E C G, G D# A#	M3↓ P5↑
	G# A# F, F# G# D#, A B F#	M2↑ P4↓
Condition Two:	D# G F, D F# E, F# A# G#	M3↑ M2↓
	C# F# D, D# G# E, F# B G	P4↑ M3↓
	G# A# D#, F G C, A# C* F	M2↑ P5↓
Test items:	A D F#	P5↓ M3↑
	G# E B	M3↓ P5↑
	F A G	M3↑ M2↓
	D G D#	P4↑ M3↓

Note: Items notated with respect to absolute pitch and relative pitch patterns. All pitches are taken from the octave starting at middle C except for C*, which is an octave above middle C.

Crucially, all of the test items contained novel AP sequences, because they were transpositions of the familiarization sequences. This manipulation ensured that learners could not use AP cues to determine which word was more familiar. In addition, we balanced the frequency of individual tones in the test items relative to both familiarization corpora: the tones used in the test tone words occurred an average of 6.5 times in the two exposure corpora, while the tones used in the test non-words occurred an average of 6 times in the two exposure corpora. Information about absolute pitch frequencies of occurrence was thus unavailable to serve as discrimination cues.

The forced-choice test consisted of eight test trials. Each test trial consisted of a novel RP word from Condition One and a novel RP word from Condition Two; each pair was repeated twice, once with the Condition One item first and once with the Condition Two item first. For a learner in Condition One, the novel RP words from Condition One served as test words while the novel RP words from Condition Two served as non-words; the opposite was the case for learners in Condition Two. The two items presented on each trial were separated by a

0.5-second pause with a 5-second interval between trials. Learners in both conditions received the same test.

Apparatus

The study was conducted in a small quiet room. The tone stream and the test were presented directly from .wav files synthesized on the PC and were heard via Sony headphones.

Procedure

Participants were instructed that they would hear a recording of continuous tones with a subsequent test, but were not told that the tone stream contained units of any sort. Participants received the test immediately after listening to the 2-minute, 13-second familiarization tone stream. Participants were instructed to indicate which of the two tone sequences played on each trial was most similar to the familiarization tone stream by circling either 1 or 2 on their answer sheet, corresponding to whether the familiar sequence was played first or second on that trial. Because of the constant interval between test trials (5 seconds), there was no additional warning prior to the onset of each test item. Each participant was tested individually.

Results and discussion

We first compared the performance of listeners assigned to the two counterbalanced familiarization conditions (all tests two-tailed). As there was no significant difference in performance [$t(6) = 1.26$, n.s.], the two groups were combined in the subsequent analysis. To determine whether learners successfully discriminated the tone words from non-words, we compared their performance on the forced-choice test to chance (4 out of 8 possible). Participants chose the words an average of 5.88 times (73%), a rate significantly better than would be expected by chance: $t(7) = 3.64$, $p < .01$. These results suggest that our stimuli are learnable by adults, despite a lack of consistent AP cues. The next experiment asked whether infant learners, who have not previously been shown to use RP cues when presented with continuous streams of tones, will do so when RP cues are better indicators of structure than AP cues.

Experiment 2

The results of Experiment 1 demonstrate that these materials are learnable by adults, who are known to use relative pitch in tracking tone sequences. Our central question is whether we can induce infants, who have heretofore shown no evidence of the use of relative pitch in tone sequence segmentation, to use RP cues when AP cues become less predictive. These materials, in which the tone words are continuously transposed, provide an opportunity to determine whether the structure of the input can induce infants to shift from the type of primitives that enter into their computations in segmentation tasks with continuous tone streams – absolute pitches – to instead prioritize the more informative relative pitch cues.

Method

Participants

Two groups of 12 full-term 8-month-old infants were tested (mean age = 7 months, 4 weeks; range = 7:2 to 8:1). Seven additional infants who were older than 8 months and 1 week were excluded from the analysis. Twenty-five additional infants were tested but not included in the analysis for the following reasons: fussiness (8), not looking at the side lights (2), looking times averaging less than 3 seconds to one or both sides (9), parental interference (4), and experimenter error (2). Infants were free of ear infections at the time of testing, and no hearing deficits were reported. Infants were solicited from local birth announcements; parental consent was obtained prior to testing in accordance with the guidelines of the local human subjects review committee and the principles of ethical treatment established by the American Psychological Association.

Materials

The familiarization materials were identical to Experiment 1. The test materials consisted of the four novel RP words from Conditions One and Two used in Experiment 1. We used the Head-turn Preference Procedure to test learning. Infants were tested by repeatedly presenting a single test item on each test trial. We then compared the infants' responses to the two different types of items (words versus non-words) over a series of test trials.

Procedure

Infants were randomly assigned to Conditions One or Two. Each infant was tested individually while seated in a caregiver's lap in a sound-attenuated booth. An observer outside the booth monitored the infant's looking behavior on a closed-circuit television system, and coded the infant's behavior using a button-box connected to the computer. This button-box was used to initiate trials and to enter the direction of the infant's head-turns, which controlled the duration of each test trial. Both the caregiver and the observer listened to masking music over headphones to eliminate bias. At the beginning of the 2-minute, 13-second familiarization phase, the infant's gaze was first directed to a blinking light on the front wall in the testing booth. Then the sound sequence for one of the two tone streams was presented without interruption from two loudspeakers (one located on each of the two side walls in the booth). During this familiarization period, to keep the infants' interest, a blinking light above one of the two loudspeakers (randomly selected) was lit and extinguished dependent on the infant's looking behavior. When this blinking side light was extinguished, the central blinking light was illuminated until the infant's gaze returned to center; another blinking side light was then presented to elicit the infant's gaze. During this entire familiarization phase there was no contingency between lights and sound, which played continuously.

Immediately after familiarization, 12 test trials were presented (three trials for each of the four test items, presented in random order). Six of these trials were novel RP words from Condition One and six were novel RP words from Condition Two; the novel RP words from Condition One served as words for infants in Condition One and non-words for infants in Condition Two, with the opposite pattern for the novel RP words from Condition Two. Each test trial began with the blinking light on the front wall. When the observer signaled to the computer with a button press that the infant was fixating this central light, one of the lights on the two side walls began to blink and the central light was extinguished. When the observer judged that the infant had made a head-turn of at least 30 degrees in the direction of the blinking side light, another button press signaled to the computer that one of the test items should be presented from the loudspeaker adjacent to the blinking light. This test item was repeated with a 500-millisecond interstimulus interval until the observer coded the infant's head-turn as deviating from the blinking light for 2 consecutive seconds, with a maximum of 15 repetitions per trial. When this look-away criterion was met, the computer extinguished the blinking side light, turned off the test stimulus, and turned on the central blinking light to begin another test trial. The computer accumulated total looking time to each of the four test items.

Results and discussion

We first compared the performance of infants assigned to the two counterbalanced familiarization conditions. As there was no significant difference in performance [$t(22) = .70$, n.s.], the two groups were combined in the subsequent analysis. To determine whether the infants successfully discriminated words from non-words based on relative pitch cues, we compared their listening times to the two types of test items (words versus non-words). This difference was significant: $t(23) = 2.01$, $p < .05$. Listening times to words ($M = 6.57$ seconds, SE = .61) exceeded listening times to non-words ($M = 5.41$, SE = .43). Because all of the test items contained novel absolute pitch sequences, AP cues could not have been used for the test discrimination. We hypothesize that under these conditions, infants were able to capitalize on relative pitch cues, listening longer to test items that contained familiar RP sequences.

On this view, infants were able to use relative pitch information in this statistical learning task. How detailed these representations are remains unknown. For example, infants may be responding to the similarity of pairs of relative pitches in a three-tone sequence when discriminating novel from familiar sequences, or they might perform a simpler comparison, noting particular relative pitches that had or had not occurred within words in the familiarization materials. In either case, the test discrimination requires the use of relative pitches. The fact that infants failed to perform similar test discriminations in previous studies (Saffran, 2003; Saffran & Griepentrog, 2001) strongly suggests that the information content of these familiarization materials affected infants' choice of perceptual primitives for learning.

Why did the infants in this experiment prefer to listen to the familiar tone words, rather than the non-words? Studies using this methodology with both linguistic and non-linguistic stimuli have typically elicited novelty preferences (Saffran *et al.*, 1996, 1999; Saffran & Griepentrog, 2001; although see Saffran, 2003). However, more complicated tasks – those including less familiarization with the exposure stimuli, or where there is a mismatch between the familiarization and test stimuli – typically elicit a familiarity preference (e.g. Aslin, 2000; Jusczyk & Aslin, 1995; Thiessen & Saffran, 2003). The current task is likely to be more complex than prior tasks using the tone sequence segmentation methodology. In our prior studies using tone sequences (Saffran & Griepentrog, 2001; Saffran *et al.*, 1999), only four words were included, and those words shared neither their absolute pitch patterns nor their relative pitch patterns with one another. In the current study, absolute and relative pitch information diverge from one another; there were three words played during familiarization from the perspective of relative pitch, but nine words from the perspective of absolute pitch, making the exposure array quite complex. Moreover, there was a mismatch between the familiarization and test items; the test items were physically quite dissimilar from the exposure items, as they consisted of transpositions of the exposure items. Thus, it is likely that the increased complexity of the stimulus sequences and the more distant match with the test items resulted in the observed familiarity preference.

General discussion

In several prior studies using tasks in which infants hear continuous sequences of tones, we found that while infants readily track sequences of absolute pitches, evidence did not emerge to support the claim that infants can also track sequences of relative pitches (Saffran, 2003; Saffran & Griepentrog, 2001). The current results suggest that our previous inability to demonstrate the use of relative pitch information may have been due to the structure of the learning tasks used in those experiments. The only difference between the current study and the prior studies is that we removed absolute pitch as a useful cue to word boundaries by continually transposing the tone words. In Experiment 1 we found that adult learners could learn these materials using RP cues, which is consistent with prior results with adult participants (Saffran, 2003; Saffran & Griepentrog, 2001).

With the results of Experiment 2, we now have evidence that 8-month-old infants can track relative pitch patterns given unsegmented input. These results dovetail nicely with other recent findings suggesting that infant learners can represent musical materials in terms of their relative pitch patterns. For example, infants appear to represent familiar pieces of music in long-term memory using relative pitch cues (Trainor, Wu, Tsang & Plantinga, 2002). Generalization to represent auditory information in terms of intervals – a more abstract mode of representation than absolute pitches – may be influenced by the types of information to be learned. The same point emerges from studies of non-human primates (Wright, Rivera, Hulse, Shyan & Neiworth, 2000): rhesus monkeys performed octave generalization only given a restricted set of musical passages: tonal melodies (as opposed to random or atonal melodies).

If relative pitch information is so useful for representing the auditory world, why did infants in our previous studies fail to show the use of RP cues? One possibility is that infants, like the starlings studied by MacDougall-Shackelton and Hulse (1996), use AP as a primitive, basic strategy for encoding auditory information (see also Trehub, Schellenberg & Hill, 1997). Relative pitch is also available, however, and is tracked when AP cues are unreliable, as in the current materials. Alternatively, infants have equal access to both types of cues, and the parameters of the listening experience determine which cues are prioritized; that is, there is no default cue. Similar considerations pertain to adults, who also have access to both types of pitch cues (e.g. Halpern, 1989; Levitin, 1994; Schellenberg & Trehub, 2003); the interesting question is which factors drive the use of each type of information.

The results of the current study suggest that aspects of the learning task itself shape the outcome of learning. That is to say, the choice of which perceptual primitives are to be the target of learning (here, absolute versus relative pitches) is driven partly by the structure of the input, rather than solely by the internal structure of the learner. To the extent that this is the case, we can consider the manner in which the structure of the world shapes the way we perceive the world. Infants may begin to process the auditory environment by tracking the two pitch cues that are easiest for the inexperienced brain to process: absolute pitches, which are represented via tonotopic maps in the auditory cortex and elsewhere in the auditory system, and pitch contour, a domain-general coding of up–down pitch changes. At this point, relative pitch information is available to infants as well, but it may be more difficult to compute: unlike absolute pitches, relative pitches require the detection of ratios between individual pitches rather than encoding single pitches, and unlike pitch contour, relative pitches require the detection of exact distances between pitches. By adulthood, the brain is readily able to detect such information, even in the absence of formal musical training (Trainor, McDonald & Alain, 2002), but this may require more extensive experience with auditory information than infants have yet obtained. Moreover, the structure of the auditory environment may help infants to learn that relative pitch is, in general, a far more effective cue to structure than either absolute pitch or pitch contour. The latter may provide categories that are either too specific (absolute pitch) or too general (pitch contour) to adequately capture the types of information that matter to learners across domains including both speech and music, whereas auditory learning based on relative pitch may serve to account for the most variance in the environment (for further discussion, see Saffran & Griepentrog, 2001). The process of learning about auditory events and objects may tune the perceptual system such that relative pitches become primary, without sacrificing the ability to process either absolute or contour cues. Understanding the structure of the environment is likely to be critical for understanding the processes underlying perceptual learning.

Acknowledgements

This research was funded by grants to Jenny R. Saffran from the International Foundation for Music Research, NICHD (RO1HD37466), and NSF (BCS-9983630). Thanks to Susan Bunton, Brianne Karabetsos, Joanna Miller, Becky Seibel, Erik Thiessen and Kim Zinski for their assistance in conducting this research, to two anonymous reviewers for helpful comments, and to Sandra Trehub for ongoing discussion.

References

Aslin, R.N. (2000). Interpretation of infant listening times using the headturn preference technique. Talk presented at the International Conference on Infancy Studies, Brighton, England.

Garcia, J., & Koelling, R.A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, **4**, 123–124.

Halpern, A.R. (1989). Memory for the absolute pitch of familiar songs. *Memory and Cognition*, **17**, 572–581.

Jusczyk, P.W., & Aslin, R.N. (1995). Infants' detection of the sound patterns of words in fluent speech. *Cognitive Psychology*, **29**, 1–23.

Jusczyk, P.W., Bertoncini, J., Bijeljac-Babic, R., Kennedy, L.J., & Mehler, J. (1990). *Cognitive Development*, **5**, 265–286.

Levitin, D. (1994). Absolute memory for musical pitch: evidence from the production of learned melodies. *Perception and Psychophysics*, **56**, 414–423.

Lynch, M.P., & Eilers, R.E. (1992). A study of perceptual development for musical tuning. *Perception and Psychophysics*, **52**, 599–608.

MacDougall-Shackleton, S.A., & Hulse, S.H. (1996). Concurrent absolute and relative pitch processing by European starlings. *Journal of Comparative Psychology*, **110**, 139–146.

Maye, J., Werker, J.F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, **82**, B101–111.

Saffran, J.R. (2003). Absolute pitch in infancy and adulthood: the role of tonal structure. *Developmental Science*, **6**, 37–45.

Saffran, J.R., & Griepentrog, G.J. (2001). Absolute pitch in infant auditory learning: evidence for developmental reorganization. *Developmental Psychology*, **37**, 74–85.

Saffran, J.R., Aslin, R.N., & Newport, E.L. (1996). Statistical learning by 8-month-old infants. *Science*, **274**, 1926–1928.

Saffran, J.R., Newport, E.L., & Aslin, R.N. (1996). Word segmentation: the role of distributional cues. *Journal of Memory and Language*, **35**, 606–621.

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**, 27–52.

Schellenberg, E.G., & Trehub, S.E. (2003). Good pitch memory is widespread. *Psychological Science*, **14**, 262–266.

Thiessen, E.D., & Saffran, J.R. (2003). When cues collide: statistical and stress cues in infant word segmentation. *Developmental Psychology*, **39**, 706–716.

Trainor, L.J., McDonald, K.L., & Alain, C. (2002). Automatic and controlled processing of melodic contour and interval information measured by electrical brain activity. *Journal of Cognitive Neuroscience*, **143**, 430–442.

Trainor, L.J., & Trehub, S.E. (1992). A comparison of infants' and adults' sensitivity to Western musical structure. *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 394–402.

Trainor, L.J., & Trehub, S.E. (1993). Musical context effects in infants and adults: key distance. *Journal of Experimental Psychology: Human Perception and Performance*, **19**, 615–626.

Trainor, L.J., Wu, L., Tsang, C.D., & Plantinga, J. (2002). Long-term memory for music in infancy. Poster presented at the International Conference on Infant Studies, Toronto, Canada.

Trehub, S.E., Schellenberg, G., & Hill, D. (1997). The origins of music perception and cognition: a developmental perspective. In I. Deliege & J. Sloboda (Eds.), *Perception and cognition of music* (pp. 103–128). Hove: Psychology Press.

Wright, A.A., Rivera, J.J., Hulse, S.H., Shyan, M., & Neiworth, J.J. (2000). Music perception and octave generalization in rhesus monkeys. *Journal of Experimental Psychology: General*, **129**, 291–307.

Received: 24 January 2003
Accepted: 17 February 2004