

Simultaneous grouping and auditory continuity

C. J. DARWIN

University of Sussex, Brighton, England

Are the conditions for illusory auditory continuity entirely local in frequency, or are judgments of continuity made on auditory objects? Listeners made continuous/pulsating judgments on a variety of complex tones that repeatedly alternated with a 100- to 500-Hz bandpass noise. A sufficiently quiet complex tone was heard as continuous when all its harmonics fell within the frequency range of the noise. Adding harmonics outside the noise's frequency range substantially reduced the impression of continuity, which was largely restored when these additional components were given a different fundamental frequency. Judgments of auditory continuity thus appear to be based on entire simultaneously grouped objects, rather than being determined solely by local criteria based on individual frequency channels.

When a pure tone alternates with a noise, it can be heard as continuous, provided that the noise is sufficiently intense and the tone's frequency is well within the noise's frequency band (Thurlow, 1957; Vicario, 1960). With monaural presentation, the tone will be heard as continuous, provided that the listener has no evidence that the tone has stopped during the noise. A convenient way to envisage this criterion is to think how the output of an auditory filter at the frequency of the tone changes as the tone gives way to the noise. When the noise is presented at a low level, the output of this filter to the noise will be less than its output to the tone, so listeners will have evidence that the tone has stopped and will hear the tone stop as the noise starts, giving an overall percept of a pulsing tone alternating with noise. But if the noise is made sufficiently intense, there will be no dip in energy when the tone stops and the noise starts; the percept is then one of the tone's continuing through the noise. For convenience, we will refer to this criterion as the *masking criterion*. It is local in frequency, since it involves comparing energy levels at different times within the auditory filter at the frequency of the tone.¹ But what happens with more complex sounds?

The related phenomenon of phonemic restoration (Warren, 1970) is influenced by the acoustic relationship between the perceptually restored sound and the noise that has replaced it (Layton, 1975; Samuel, 1981), indicating that the masking criterion is still relevant for speech sounds. There have also been experiments on the continuity and intelligibility of speech or complex sounds that have been masked or interrupted by a noise that covers only some of the sound's frequency components (Carlyon, Deeks, Norris, & Butterfield, 2002; Howard-Jones &

Rosen, 1993a, 1993b; Plomp, 1982). Plomp (1982; 2002, p. 38), for example, reported that a complex sound can be heard as continuous when the lower and higher harmonics are alternately replaced by low- and high-frequency noise, respectively. Carlyon et al. (2002) reported similar findings for vowels and also noted that vowel intelligibility was increased when the continuity illusion was present.

Although these experiments show that the masking criterion can apply to part of a complex sound, all of them have involved sounds where the part that was not being replaced by noise was continuous. But what happens when some parts of a complex sound that alternates with noise meet the masking criterion but others do not? Are those that do meet it heard as continuous? One possible principle is that decisions about auditory continuity are always made locally on the basis of the masking criterion. Here, any component of a complex that met the masking criterion *would* be heard as continuous, regardless of the fact that it was part of a complex. But perhaps a more likely possibility is that the continuity decision is not made about individual frequency components but about auditory objects—the results, at least in part, of auditory grouping principles. Now the question that the perceptual system is faced with is whether the masking criterion is fulfilled for *all* the frequencies that make up a particular auditory object. If this criterion is not met for the whole auditory object, the whole object would be heard as pulsating, rather than continuous, despite the fact that some of its component frequencies *did* meet the masking criterion. This article asks whether continuity is determined by local application of the masking criterion or more globally, across auditory objects.

Auditory grouping principles may conveniently be divided into sequential and simultaneous (Bregman, 1990). Simultaneous principles determine which of the frequency components that are present at a particular time should be grouped together as belonging to the same auditory object. Sequential principles determine whether a group of frequencies at one time belong to the same auditory object as a group of frequencies at a different time.

The experimental data were collected by Anne Heddy and were presented as a poster to the Short Papers Meeting of the British Society for Audiology, London, September 2004. Correspondence concerning this article should be addressed to C. J. Darwin, Department of Psychology, University of Sussex, Brighton BN1 9QG, England (e-mail: cjd@sussex.ac.uk).

There is already evidence that continuity judgments depend on sequential grouping principles. Bregman (1990, chap. 3; Bregman, Colantonio, & Ahad, 1999; Ciocca & Bregman, 1987) has argued that, for continuity to occur, the sounds before and after the alternating noise must be groupable into the same stream, so that one hears the same sound source before, during, and after the interrupting noise. As an example from the experiment of Ciocca and Bregman, when two consecutive frequency glides are separated by a burst of noise, listeners will hear the first glide continuing behind the noise, provided that the second glide is a plausible continuation of the trajectory of the first.

A close similarity between one of the factors affecting sequential grouping and continuity has been shown in Bregman et al. (1999). Consider two brief narrowband noises centered on different frequencies that alternate and have a brief silence separating one from the other. If the noises' center frequencies are sufficiently similar, listeners will hear them as a single stream similar to a musical trill. But if their center frequencies are sufficiently dissimilar, listeners will hear two separate streams, illustrating the sequential grouping principle of frequency proximity. If now the silences between the narrowband noises are filled with a sufficiently intense broadband noise, then, when the narrowband center frequencies are close, listeners will hear a continuous narrowband noise sweeping in frequency through the wideband noise. But when the narrowband center frequencies are sufficiently far apart, listeners will hear no continuous sound. Continuity is thus intimately linked to the temporal cohesion of sound sources and so to *sequential* grouping principles, such as frequency proximity.

The purpose of the experiment reported in the present article was to establish the role of *simultaneous* grouping principles in continuity. Do simultaneous grouping principles, such as harmonicity, influence continuity? This question was investigated by looking at continuity judgments for complex tones that alternate with a noise that only fulfills the masking criterion for a subset of the tone's components. This question will now be illustrated with a specific example.

The upper panel of Figure 1 shows a complex tone alternating with a band-limited noise. The complex tone consists of the first three harmonics of 140 Hz (140, 280, and 420 Hz), and the noise contains frequencies between 100 and 500 Hz. The harmonics are heard as a single (dull) complex tone that, if the noise is sufficiently loud, is continuous through the noise. A more interesting case is shown in the lower panel. Here, we have added three more harmonics (5–7), which are outside the frequency range of the bandpass noise. There are now two plausible alternative perceptual interpretations, depending on the influence that simultaneous grouping mechanisms have on perceived continuity. If continuity is determined solely by criteria that apply locally in frequency to individual frequency components, one should hear the lower harmonics as continuous just as before, with the upper harmonics as a separate pulsating sound. On the other hand, if decisions about auditory continuity are made about simultaneously




Figure 1. The upper panel shows the first 1.5 of the 8 cycles of the dull stimulus. The tonal part consists of the first three harmonics of 140 Hz, which alternates with a noise band-limited between 100 and 500 Hz. The lower panel similarly shows the rich stimulus, whose tonal part additionally has harmonics 5–7 of 140 Hz.

grouped auditory objects, one might hear all the harmonics as a single (rich) complex that is pulsating by virtue of the fact that part of it is *not* in the same frequency region as the bandpass noise and so does not meet the masking criterion. The second alternative would be a more sensible strategy for the auditory system to adopt and would also sit comfortably with Bregman's (1990; Bregman et al., 1999; Ciocca & Bregman, 1987) view on continuity's relationship to sequential grouping.

Pilot experiments did indeed show that the sound in the lower panel gave a single, rich, pulsating tone, supporting the idea that continuity judgments are based on objects that are the result of simultaneous grouping. Consequently, in the following experiment, it was also asked whether introducing a difference in fundamental frequency (f_0) between the upper and the lower harmonics that is sufficient for the lower harmonics to be heard as a separate sound source enables them to be heard as continuous.

The experiment was performed to investigate what perceptions of continuity would occur for sounds such as that illustrated in the lower panel of Figure 1, where all the frequency components are harmonics of the same fundamental, and for those in the related condition in which the upper three components of the complex have a fundamental frequency different from that for the lower three components. The basic paradigm varies the relative level of the target sound and the noise, and listeners are asked

to judge whether some aspect of the sound is continuous or pulsing. Provided that the noise fulfills the masking criterion for the appropriate frequency components, continuity judgments should increase as the level of the target decreases relative to the level of the noise. The experiment is performed to investigate what these *appropriate frequency components* are.

When listeners make continuity judgments, it is important that they can reliably identify what it is that they hear as continuous. In experiments with simple tones, the problem is trivial, but for complex sounds such as those used here, which can be parsed in a variety of ways, the issue is substantive. Consequently, the main experiments are preceded by training sessions which familiarize listeners both with the continuity judgment task itself and with the different sounds that they will be hearing. These training and experimental sessions will be described in detail below but will now be briefly summarized.

We first (part 1) made listeners familiar with the continuity judgment task, using the first three harmonics of 140 Hz, as illustrated in the upper panel of Figure 1. They produced a dull-sounding complex, all of whose components satisfied the masking criterion when the noise was sufficiently intense. So when the level of the target complex was low, relative to the noise, the masking criterion would predict that listeners would hear the complex as continuous. We then (part 2) introduced the listeners to a richer sounding complex, which had six harmonics of 140 Hz, as shown in the lower panel of Figure 1, and asked them to judge whether this richer sound was heard as continuous when alternating with the same noise as before. For this sound, only the three lower harmonics were covered by the masking criterion at low levels. The remaining three, higher harmonics were well outside the frequency range of the noise and so would not meet the masking criterion even at very low signal levels. If continuity judgments are based on whether the masking criterion holds for all the components of a complex, we would expect that this rich complex would not be heard as continuous, even at low signal levels. Our pilot data had suggested that this was indeed the case.

As preparation for carrying out the analogous experiment with sounds that had the upper frequency components on a different f_0 from the lower components, we then (parts 3 and 4) introduced the listeners to two different sounds, both of which had a thin timbre. These sounds consisted of three upper harmonics of either 140 or 160 Hz. We also introduced the listeners to a mixed sound that had the lower harmonics of 140 Hz, together with the upper harmonics of 160 Hz. The listeners were first trained to identify these thin and mixed sounds reliably in identification experiments that also contained the previously used dull and rich sounds.

When they could reliably identify these sounds, they then heard (part 5) the dull, rich, and mixed sounds at different levels, alternating with the noise used in Figure 1. They were asked to listen for either the rich tone or the dull-sounding tone and to judge whether it was continuous or pulsating. If the sequential grouping cue

of harmonicity was being used to determine the auditory objects over which continuity judgments were made, we would expect more reports of continuity for the mixture than for the rich complex.

In the final stage of the experiment, the listeners again made continuity judgments of the various sounds, but this time, presentation was blocked by stimulus type.

METHOD

Stimuli

The basic stimulus configuration used in the experiment was similar to that shown in Figure 1, but with a sequence of eight alternations of a target sound and a noise. Each sound segment, whether target or noise, had a duration of 120 msec, including 10-msec rise–fall ramps. The ramps between successive sounds completely overlapped, so that the entire sequence lasted 1.77 sec.

Five different target sounds were used. Three of the targets (dull, thin140, and thin160) consisted of a single group of three consecutive harmonics of a fundamental frequency (f_0) of either 140 or 160 Hz. The dull sound consisted of the first three harmonics of 140 Hz, whereas the other two consisted of three higher harmonics of either 140 or 160 Hz (chosen to have a similar frequency range):

1. Dull: $f_0 = 140$ Hz; Harmonics 1–3, 140, 280, 420 Hz
2. Thin140: $f_0 = 140$ Hz; Harmonics 5–7, 700, 840, 980 Hz
3. Thin160: $f_0 = 160$ Hz; Harmonics 4–6, 640, 800, 960 Hz

The other two targets were formed by adding to the lower harmonics that formed the dull sound either the higher group with the same f_0 (thin140), to give the rich sound, or the higher group with the different f_0 (thin160), to give the mixture sound.

4. Rich: Dull + Thin140
5. Mixture: Dull + Thin160

Note that the rich sound's frequency components (140, 280, 420, 700, 840, and 980 Hz) are all harmonics of 140 Hz, whereas although the mixture's three lowest components are harmonics of 140 Hz, its three higher components are harmonics of 160 Hz. We would expect the rich sound to be heard as a single rich complex, whereas, because of the different f_0 between its lower and higher harmonics, we would expect the mixture to be heard as two separate sounds, a dullish low-pitched sound together with a higher pitched thin sound.

The noise that alternated with all of these sounds was designed to be able, when sufficiently intense, to meet the masking criterion for the dull sound (the first three harmonics of 140 Hz). It was a frozen segment of Gaussian noise, bandpass filtered (in the frequency domain using 100-Hz-wide skirts) between 100 and 500 Hz. The relative amplitude of target and noise was set using the dull sound, so that a nominal 0-dB level of the target gave the same energy as did the noise through a gammatone filter centered at 140 Hz with a 40-Hz bandwidth. Within each group of three harmonics, the harmonic amplitudes decreased at -6 dB/oct, relative to their fundamental frequency. The reference levels of the 140- and 160-Hz fundamental frequencies were equal.

The sounds were delivered at 22.05 kHz from an Audiomeia III board through TDT attenuators and headphone amplifiers to Sennheiser HD-414 headphones in a double-skinned IAC booth. The nominal 0-dB level of the dull sound corresponded to 66 dB SPL (unweighted).

Subjects

Sixteen audiotically normal listeners took part in all phases of the experiment. They were predominantly female undergraduate students at the University of Sussex, mainly in their early 20s. Those who played a musical instrument tended to complete the training phases of the experiment more easily and provide cleaner data but did not give qualitatively different results in the main experiments.

Two of the 16 were dropped for inconsistent identification in the training part of the experiment.

Training

In order to establish which sounds listeners hear as being continuous, the subjects were trained, before taking part in the main experiment, to identify reliably the different sounds and mixtures used in the experiment. They were also given practice at making continuous versus pulsating judgments. The subjects controlled a sequence of five PowerPoint presentations, which allowed them to hear individual sounds at will during the demonstration phases of the presentation before taking part in the linked identification experiments.

1. The first presentation introduced the listeners to a pulsing and a continuous dull tone and to the pulsing noise. They then took part in a short 50-trial experiment in which they had to judge whether the dull tone, alternating with the noise, was continuous or pulsing (*dull-initial* condition). The tone varied in level in 2-dB steps between +8 and -10 dB, relative to the calibrated 0-dB level referred to above. The method of constant stimuli was used to generate a psychometric function from which the pulsation threshold (50% pulsating responses) could be calculated. After each sound, the listeners pressed "c" if the tone was continuous or "p" if it was pulsing. They were asked to ignore the pulsing of the noise and any pulsing of the tone that just occurred near the beginning of the sequence. If the listeners gave a reasonably smooth psychometric function, they progressed to the second presentation; otherwise, they repeated the first.

2. The second presentation introduced them to the rich tone and, again, was followed by a 50-item test similar to the earlier one but using sequences of the rich tone, instead of the dull (*rich-initial* condition). The listeners then progressed immediately to the next presentation.

3. The third presentation introduced the listeners to the two thin sounds and the mixture, which was described as mixture (dull + thin). They then took part in a 25-trial identification experiment in which they had to label 0-dB exemplars of sequences of the target sound (alternating with silence instead of noise) as dull (d), thin (t), rich (r), or mixture (m). Both types of thin sound were used (thin140 and thin160), but they were not labeled differentially either in the initial familiarization or in the test. If the listeners achieved 90% correct or better, they progressed to the next presentation; otherwise, they repeated this one.

4. The fourth presentation was similar to the third but varied the level of the target sounds, including the -10-dB targets, as well as the 0-dB ones.

5. In the final presentation, the listeners were asked again to classify sounds as pulsing or continuous. They heard three different types of tones: dull, rich, and mixed. They were instructed to listen for the dull tone. If they heard it either by itself or as part of the mixture, they were then asked to make a decision as to whether the dull tone was continuous (c) or pulsing (p). They were to ignore the thin tone in the mixture. If they heard the rich tone, they were asked to indicate whether it was continuous (c) or pulsing (p). These instructions bias listeners toward hearing the dull tone, either by itself or as part of the mixture or the rich sound (which, if perceptually decomposed into dull plus thin, would be perceptually very similar to the mixture). These instructions thus bias the experiment toward listeners' hearing the dull sound, with the potential for it to be heard as continuous. The listeners heard five repetitions of the three types of sound, with the target sound varying in level in 2-dB steps between +8 and -10 dB, giving 150 trials in all. If the listeners produced sufficiently smooth psychometric functions for the dull targets, they progressed to the main experiment.

Main Experiment

The main experiment measured the listeners' boundaries by the method of constant stimuli between continuous and pulsing targets for the three different types of sounds used in the final stage of training: dull, rich, and mixed. Presentation was blocked by sound

type, and the order of the blocks was counterbalanced within and between listeners. The gradations in level of the target tone were reduced to 1 dB, rather than the 2-dB steps used in training, and the range of levels was restricted to +5 to -10 dB. The listeners heard 10 repetitions of each of the 16 sounds within a block and took each block twice (the order of the stimulus types for the second blocks being the reverse of that for the first). The total experiment thus consisted of six blocks of 160 trials. In order to check that the listeners did not hear the thin sounds as continuous, a single short block was run after the main experiment, using just the two thin sounds, with 2-dB steps between 0 and -10 dB and 5 repetitions of each sound. For the blocks that contained either the rich or the mixed sounds, the listeners were instructed as follows. "If you hear the Rich tone, then indicate whether it is continuous (c) or pulsing (p). If you hear the Mixture (Dull + Thin), then make a decision on whether the *dull tone* is continuous (c) or pulsing (p). Ignore the thin tone in the mixture."

RESULTS

Training Experiments

Figure 2 shows all the data from continuous versus pulsating judgments during the training phase of the experiment. The "dull-initial" and "rich-initial" conditions show the data from separate blocks at the beginning of training, whereas the data for the remaining three conditions (dull, rich, and mixed) were from a single block of trials at the end of training. For each of the two dull conditions, a clear pulsation threshold was shown at between a 0- and a 3-dB tone level. The less sharp psychometric function for the dull-initial condition probably reflected the earlier stage of training. For the two rich conditions, a different pattern was shown. When the rich condition was played as a separate block early in training (when the listeners had heard only the dull sound previously and were not instructed to listen for a dull sound within the rich), it was almost always heard as pulsating. Later, when the rich sound was randomized with the dull and the mixed sounds, it was heard more frequently as continuous. Here, the listeners were instructed to respond "continuous" if they could hear the dull sound within any of the sounds as continuous. It may be that these instructions, together with the trial-by-trial context of the dull and the mixed, allowed some listeners to decompose the rich sound. There was, however, a clear difference between the number of continuous judgments made to the mixed sound and the number made to the rich, even though they were presented under similar conditions. So the difference in f_0 between the upper and the lower harmonics of the mixed sound allowed the listeners to hear the lower harmonics as a dull complex and to classify this complex as continuous at low levels. The sound in the mixed condition was heard slightly less often as continuous than was the dull, perhaps reflecting some listeners' difficulty in attending to the dull sound within the mixed.

Main Experiment

The number of "continuous" responses given by each subject in each condition was averaged across subjects to give the data shown in Figure 3. The sounds in the dull condition were heard as continuous at low levels, but as




Figure 2. Mean percentage of continuous judgments during the training sessions (± 1 SEM). The dull-initial and rich-initial data were collected as separate blocks of 50 trials in Training Presentations 1 and 2, respectively. The data for the remaining three conditions were collected in a single block of 150 trials in the fifth training presentation (see the Method section for details).

pulsing at high levels, with a boundary at about -1 dB. All the rich sounds, however, were predominantly heard as pulsing, with the proportion of continuous responses approaching 50% only at the lowest tone level (-10 dB). The mixed sounds, which were similar to the rich, except that their upper harmonics were on a different f_0 , gave a pattern of results that was very similar to that shown by the dull sounds. Overall, the sounds in the mixed condition were heard as slightly less continuous than were those in the dull condition [$F(1,13) = 4.9, p < .05$], but their

boundaries were very similar, and the two lots of data diverged only at low signal levels. However, the sounds in the rich condition were heard as less continuous than the dull or the mixed: A repeated measures ANOVA gave a highly significant main effect of stimulus condition [dull, rich, or mixed; $F(2,26) = 41.9, p < .0001$] and of its interaction with the 16 intensities of the tone [$F(30,390) = 18.4, p < .0001$]. The smaller amount of data collected for the thin condition, as a check that it was heard as pulsating, shows that the sounds in this condition were, indeed,

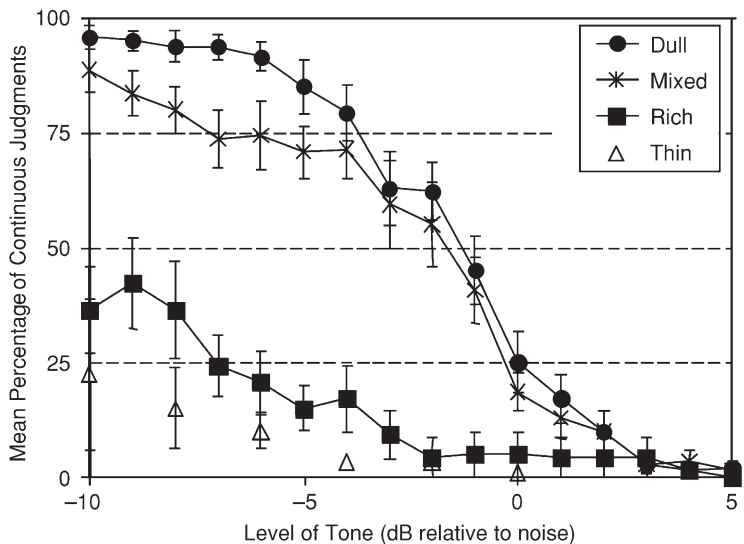


Figure 3. Mean percentage of continuous judgments for the main experiment (± 1 SEM). The data from the different conditions were collected in separate blocks.

rarely heard as continuous, with a maximum of fewer than 25% continuous responses occurring at the lowest tone level. These continuous responses came largely from 5 of the subjects. To check that these 5 subjects were not distorting the pattern of response to the other conditions, they were omitted from a reanalysis of the data. However, their omission did not substantially change the main pattern of results for the other conditions, although it did slightly reduce the number of continuous judgments made to the rich sounds.

DISCUSSION

The main experiment clearly shows that whether a particular set of frequency components are heard as continuous is influenced by the auditory object that those frequency components belong to. The lower three harmonics were heard as continuous much more in the mixed condition, where they could segregate from the upper harmonics by virtue of a difference in f_0 , than in the rich condition, where all six harmonics shared a common fundamental. Continuity judgments are, therefore, influenced by simultaneous grouping, a conclusion that sits comfortably with the experimental evidence (Bregman et al., 1999) that continuity requires the (simple) sounds on either side of the interrupting noise to be groupable into the same stream.

Although the perception of continuity requires some criterion to be met that is primarily cochlear in nature (the masking condition), the experiment shows that decisions about continuity are made about auditory objects: The masking criterion must, presumably, be satisfied for each of the object's constituent frequencies.

This experiment required considerable training of the subjects in order to familiarize them with the different sounds used in the experiment. This training, together with the specific contexts of the different parts of the experiment, may have had an influence on how the listeners perceived the sounds. The listeners might, for example, have been more likely in the later stages of the experiment to perceptually decompose the rich sound into the dull and thin sounds (and so, to think that they were hearing the mixture sound, instead of the rich) by virtue of being exposed to the dull, thin, and mixed sounds earlier in the experimental procedure or earlier in a particular block of trials. The experiment may thus be underestimating the consistency of the effect of grouping on continuity judgments.

The instructions to the listeners encouraged them to listen for the dull sound even when they were actually just hearing a block of rich sounds (they were told that they were hearing both rich and mixed sounds in both the solely rich and the solely mixed blocks). If they heard a sound that actually was rich decompose into dull plus thin (which is the definition of what they knew as a mixed sound), the instructions encouraged them just to make a continuity judgment on this dull sound. The instructions thus biased the listeners toward a decomposed percept

and, thus, potentially toward more continuity judgments in both the rich and the mixed blocks. The subjects were not asked to make explicit decisions in the main experiment about the category of sound (rich, mixed, etc.) that they heard on each trial. Although this would have provided useful information, it would have made an already complex task impracticable.

The conclusion that simultaneous auditory grouping precedes decisions about continuity has implications for the relationship between simultaneous auditory grouping and attention (Carlyon, Cusack, Foxton, & Robertson, 2001; Cusack, Deeks, Aikman, & Carlyon, 2004), since it has been claimed that the continuity illusion can occur outside the focus of attention. Micheyl et al. (2003) asked their subjects to watch a silent movie and to ignore the sounds that were played to them. These sounds consisted of sequential pairs of pure tones with either silent or noise-filled gaps. The noise was either of an appropriate frequency range to produce the illusion of continuity in the tone or of a different frequency range. Their carefully controlled study showed that mismatch negativity was smaller in conditions in which both the deviant and the standard stimuli were heard as either continuous or not continuous than when one was heard as continuous and the other not. Similarity in terms of continuity thus influenced mismatch negativity, despite their subjects' attending to a visual input. Their pure tones did not, of course, require any simultaneous grouping, so it is an experimental question whether the continuity effects reported in the present article would also be reflected in mismatch negativity changes when the subjects' attention was diverted elsewhere. If they were, there would then be a *prima facie* case that simultaneous grouping also takes place outside the primary focus of attention, in contrast to the claim that the buildup of sequential grouping requires attention (Carlyon et al., 2001; Cusack et al., 2004).

REFERENCES

- BREGMAN, A. S. (1990). *Auditory scene analysis: The perceptual organization of sound*. Cambridge, MA: MIT Press, Bradford Books.
- BREGMAN, A. S., COLANTONIO, C., & AHAD, P. A. (1999). Is a common grouping mechanism involved in the phenomena of illusory continuity and stream segregation? *Perception & Psychophysics*, **61**, 195-205.
- CARLYON, R. P., CUSACK, R., FOXTON, J. M., & ROBERTSON, I. H. (2001). Effects of attention and unilateral neglect on auditory stream segregation. *Journal of Experimental Psychology: Human Perception & Performance*, **27**, 115-127.
- CARLYON, R. P., DEEKS, J., NORRIS, D., & BUTTERFIELD, S. (2002). The continuity illusion and vowel identification. *Acta Acustica United With Acustica*, **88**, 408-415.
- CIOCCA, V., & BREGMAN, A. S. (1987). Perceived continuity of gliding and steady-state tones through interrupting noise. *Perception & Psychophysics*, **42**, 476-484.
- CUSACK, R., DEEKS, J., AIKMAN, G., & CARLYON, R. P. (2004). Effects of location, frequency region, and time course of selective attention on auditory scene analysis. *Journal of Experimental Psychology: Human Perception & Performance*, **30**, 643-656.
- HOUTGAST, T. (1972). Psychophysical evidence for lateral inhibition in hearing. *Journal of the Acoustical Society of America*, **51**, 1885-1894.
- HOWARD-JONES, P. A., & ROSEN, S. (1993a). The perception of speech in fluctuating noise. *Acustica*, **78**, 258-272.

- HOWARD-JONES, P. A., & ROSEN, S. (1993b). Unmodulated glimpsing in "checkerboard" noise. *Journal of the Acoustical Society of America*, **93**, 2915-2922.
- LAYTON, B. (1975). Differential effects of two nonspeech sounds on phonemic restoration. *Bulletin of the Psychonomic Society*, **6**, 487-490.
- MICHEYL, C., CARLYON, R. P., SHYROV, Y., HAUKE, O., DODSON, T., & PULLVERMÜLLER, F. (2003). The neurophysiological basis of the auditory continuity illusion: A mismatch negativity study. *Journal of Cognitive Neuroscience*, **15**, 747-758.
- PLOMP, R. (1982, July). *Continuity effects in the perception of sounds*. Paper presented at Conference on the Psychoacoustics of Music, Jablonna, Poland.
- PLOMP, R. (2002). *The intelligent ear: On the nature of sound perception*. London: Erlbaum.
- SAMUEL, A. G. (1981). The role of bottom-up confirmation in the phonemic restoration illusion. *Journal of Experimental Psychology: Human Perception & Performance*, **7**, 1124-1131.
- THURLOW, W. R. (1957). An auditory figure-ground effect. *American Journal of Psychology*, **70**, 653-654.
- VICARIO, G. (1960). L'effetto tunnel acustico. *Rivista di Psicologia*, **54**, 41-52.
- VOGTEN, L. L. (1978). Low-level pure-tone masking: A comparison of "tuning curves" obtained with simultaneous and forward masking. *Journal of the Acoustical Society of America*, **63**, 1520-1527.
- WARREN, R. M. (1970). Perceptual restoration of missing phonemes. *Science*, **167**, 392-393.
- WARREN, R. M. (1982). *Auditory perception: A new synthesis*. New York: Pergamon.

NOTE

1. Because of cochlear nonlinearities (Houtgast, 1972; Vogten, 1978), this condition was slightly different from one in which the tone would have been masked by the noise, had the tone actually been present during the noise (Warren, 1982). I am grateful to Brian Moore for drawing this distinction to my attention.

(Manuscript received October 12, 2004;
revision accepted for publication February 23, 2005.)