Music-Perception
Winter 1990, Vol. 8, No. 2, 203-214

The Tone Height of Multiharmonic Sounds

ROY D. PATTERSON
MRC Applied Psychology Unit, Cambridge, England

Multiharmonic tones containing resolved harmonics (1-7) of a fundamental, unresolved harmonics (8-24) of a fundamental, or both resolved and unresolved harmonics (1-24) were presented to listeners at each of six octaves ranging from 32 to 1024 Hz. There also were modified versions of each stimulus in which alternate harmonics were either attenuated or phase shifted. Listeners were asked to judge the tone height of the sounds when presented in a musical context. The results show that all three manipulations can affect tone height and that stimuli with the same period produce tone-height perceptions that differ by more than an octave!

Introduction

In the traditional representation, musical pitch is a helix—a single valued function existing in three-dimensional space, much like a stretched coil spring (see Ueda & Ohgushi, 1987 for a review). The pitch helix has a circular dimension, referred to as tone chroma, and a vertical dimension referred to as tone height, and pitch rises one octave per revolution of the circular dimension. This representation of pitch is supported by the perception of pure tones; in this case the only continuous path from a note to its octave is along the helix through all of the intervening tone-chroma values. For any given chroma value, say C, tone height is a discrete variable, namely the set of octaves CO, Cl, C2, . . . C8.

In the case of multiharmonic tones, it is possible to move continuously from a note toward its octave without changing chroma. For example, consider the complex sound composed of all harmonics of 100 Hz starting in cosine phase (i.e., a pulse train) and the perceptual change that occurs when the odd harmonics are attenuated as a group. As the attenuation increases, tone height rises smoothly from the original octave (100 Hz) to the final octave (200 Hz), and there is no change in tone chroma during this transition. This indicates that the domain of pitch for musical notes is more like the surface of the space enclosed by the pitch helix rather than a helical wire.

There are at least three ways of modifying a pulse train and producing sets of waves that all have the same tone chroma but that sound "higher" or "lower" than the starting pulse train. This article presents an experiment in which listeners judged the tone height of a set of these sounds presented over a range of six octaves.

Modifying Tone Height Without Changing the Fundamental

The starting point for all of the manipulations in the current experiment was the multiharmonic tone whose long-term power spectrum is shown in Figure 1. The tone is a modified pulse train. Specifically, it is a set of 24 harmonics of a fundamental that ranged from 32 to 1024 Hz in octave steps. The level of the components was reduced 1.5 dB/octave in an effort to balance the loudness across different conditions. In one condition, all of the harmonics started at their maximum value; these stimuli are referred to as cosine-phase sounds. In a second condition, the starting phase of each component was randomized; these stimuli are referred to as random-phase sounds. Cosine-phase and random-phase sounds with the same power spectrum have the same tone height.

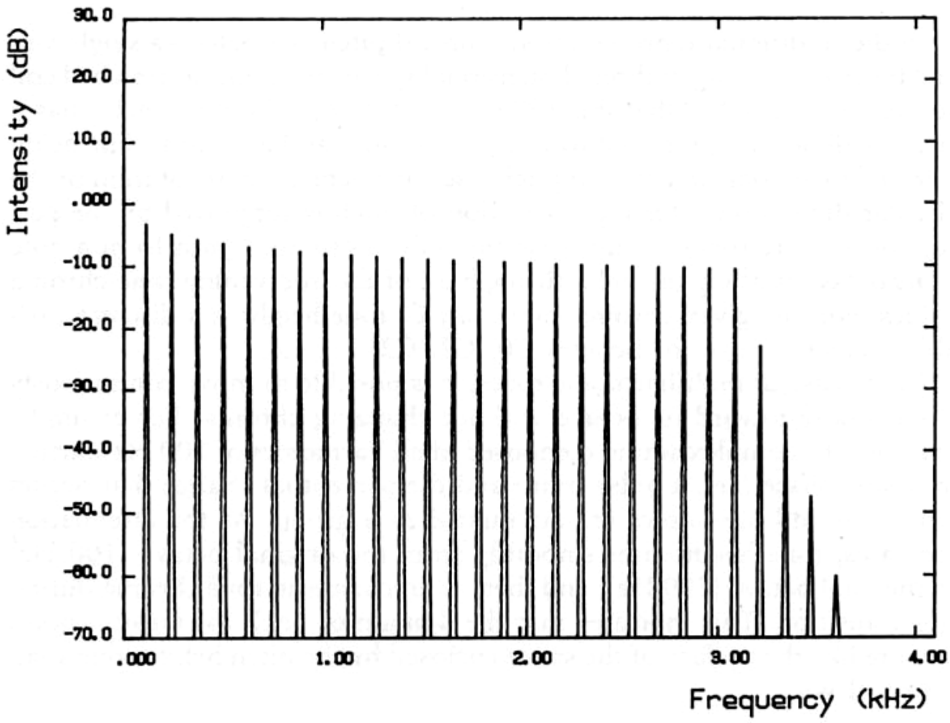

Fig. 1.The power spectrum of the cosine-phrase sound with both lower and upper harmonics. All of the experimental sounds are modified versions of this cosine-phrase sound.

ATTENUATION OF THE ODD OR EVEN HARMONICS

The first method of altering tone height involves reducing the amplitude of either the odd harmonics or the even harmonics of the cosine-phase sound. The effect of removing the odd harmonics has already been described: It doubles the component spacing and the fundamental of the harmonic series, and so it increases tone height by an octave. As their tone height is entirely predictable, the even-harmonic sounds are not considered further. When the even harmonics are attenuated, leaving the odd harmonics stronger, the situation is quite different. In the extreme, like the previous manipulation, the attenuation doubles the component spacing in the spectrum. Unlike the previous manipulation, however, it does not double the fundamental, which remains fixed at 100 Hz for all attenuations. The tone height is not immediately predictable and so the experiment contrasted the cosine-phase and odd-harmonic sounds, referred to as odd-cosine-phase sounds.

SHIFTING THE PHASE OF THE ODD OR EVEN HARMONICS

The second method of altering tone height involves shifting the phase of all of the odd harmonics of the cosine-phase stimulus, or all of the even harmonics of the cosine-phase stimulus, without changing their amplitudes. This produces what are referred to as alternating-phase sounds. All of the stimuli in the set have the identical period and the identical long-term power spectrum. When the phase change is sufficiently large (40° for a moderate-level stimulus with a fundamental of 125 Hz), it produces a change in the timbre of the sound. Patterson (1987) studied this phenomenon and, contrary to the predictions of place theory, the effect exists over a wide range of stimulus conditions that are relevant to music. All though it is not discussed in that paper, the alternating-phase stimuli sound "higher" than the corresponding cosine-phase stimuli. Accordingly, the odd-alternating-phase and even-alternating phase stimuli were compared with each other and the cosine-phase/random-phase sounds. In each case the phase shift was 90 degrees, the condition that produces the largest effect.

REDUCING THE LOWER OR UPPER HARMONICS BY FILTERING

The final method of manipulating tone height involves filtering either the lower or upper harmonics of the sounds. When the lower harmonics are attenuated, the resulting sound is perceived to be "higher" than the original. In one case, the unresolved harmonics were attenuated, leaving harmonics 1—7; in the other case, the resolved harmonics were attenuated, leaving harmonics 8—24. The former are referred to as lower-harmonic sounds and the latter as upper-harmonic sounds.

In order to avoid the perception of edge tones, the lower-harmonic and upper-harmonic sounds were not square filtered; rather, in the lower- harmonic sounds, the harmonics above the seventh were progressively attenuated 12 dB/harmonic, and similarly, in the upper-harmonic sounds, the harmonics below the eighth were progressively attenuated 12 dB/harmonic. This procedure also reduces the effect of any cubic difference tones in the upper-harmonic sounds.

BACKGROUND

The alternating-phase stimuli were designed to study phase perception (Patterson, 1987), and they were originally conceived to be wideband extensions of the three-component QFM stimuli used by Ritsma and Engel (1964) to study residue pitch. Ritsma and Engel noted that there was some ambiguity concerning the perceived octave of these sounds, as did Moore (1977). Although it was not reported at the time, the listeners in Patterson's experiments said that the alternating-phase stimuli sounded "higher" than the cosine-phase stimuli did and that this was the cue they used to perform the discrimination.

Flanagan and Guttman (1960a, 1960b) presented a series of experiments that included cosine-phase and odd-cosine-phase sounds and showed that "octave errors" occurred when the odd-cosine-phase stimuli had low fundamentals (below 150 Hz). Guttman and Flanagan (1964) and Rosenberg (1965) showed that the presence of octave errors is influenced by highpass filtering the stimuli. Patterson (1973) used odd-cosine-phase sounds when extending residue pitch measurements to include sounds with low, resolved harmonics. He noted that the listeners made octave errors in response to these sounds and reported that the tone height of an odd- cosine-phase sound with fundamental /p, seems to fall between those of cosine-phase sounds with fundamentals /o and 2/o.

In general, however, there has been little research on the perception of tone height, as distinct from octave errors. This may be in part due to the traditional view that tone height is discrete for any given chroma value; that is, for the note C there are only eight points on the tone height dimension—Cl, C2, ... C8. For instruments with a fixed timbre, this conception is essentially correct. However, the fact that one can progressively attenuate the odd harmonics of a cosine-phase sound and move continuously from an initial tone height to its octave suggests that tone height is a continuous dimension when timbre is allowed to vary and that tone height should be measurable with standard psychological techniques.

Method

PROCEDURE

When periodic sounds contain resolved harmonics, the auditory system has the option of two modes of analysis, typically referred to as the analytic mode and the synthetic mode. In the analytic mode the focus is on individual, resolved components and changes in those components from sound to sound. In the synthetic mode, the system focuses on the complete auditory image produced by the set of components as a whole. The latter mode is more common in music perception and in the natural environment generally.

In order to promote the synthetic mode of perception, the individual trials of the experiment were constructed to give the impression of a tonal melody at the end of a musical phrase. Rhythmically, this subphrase was a triplet followed by three half notes as shown in Figure 2a. Melodically, each trial was restricted to the tonic plus the note above it and below it on the diatonic scale; the figure shows the case when the tonic is C. The triplet that initiated the trial was chosen at random from the notes ti and re with the restriction that the triplet not be entirely ti's or entirely re's. The three half notes were all do's. The trial-to-trial variation is illustrated by the changes in the triplet of Figure 2b.The fast-moving triplet encourages synthetic listening and provides information about the instrument and the octave for the current trial. The three do's provide long clear notes (500 ms) for the listeners to make their octave judgments. On each trial the instrument/ octave combination was chosen at random (without replacement) from the full set of sounds. There were four replications in the experiment, and there were four listeners with normal hearing.

The task was simply to listen to the melody on each trial and write down a number between one and six to indicate the octave of the three final do's. All of the listeners had musical experience, but none was a professional. All of the listeners found the task easy to perform, and they all produced the same pattern of results. For brevity, the results will be restricted to average data.

DESIGN

The filtering process that separated unfiltered stimuli into lower-harmonic or upper- harmonic stimuli was applied to all four of the stimulus types (cosine-phase, odd-cosine- phase, odd-alternating-phase and even-alternating-phase) and so there were 12 stimulus conditions in the experiment. For convenience, the stimulus types are referred to as "instruments." For comparison, a thirteenth instrument, in the form of a sine tone, was included. Each of the instruments was presented at each of the six octaves four times.

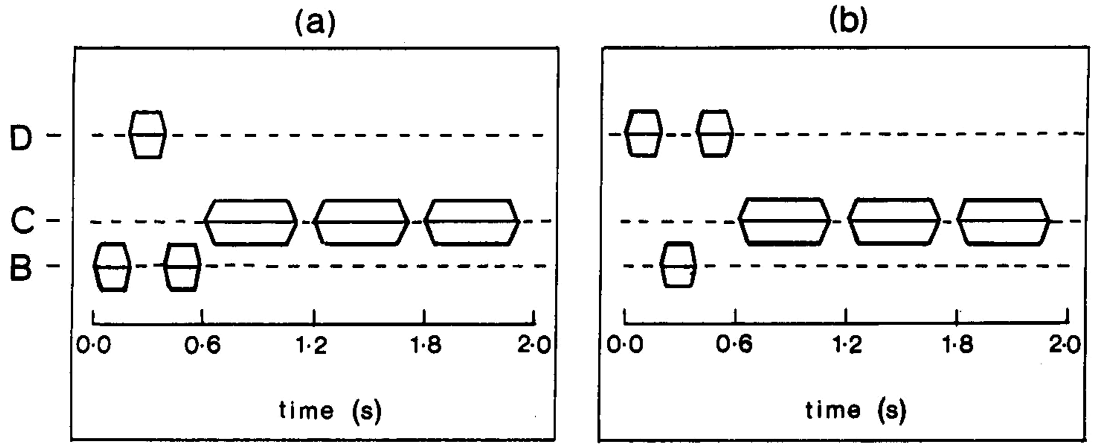

Fig. 2.A schematic diagram of the melody in two trials of the experiment. The melody is a triplet followed by three held notes; the triplet varies from trial to trial (a and b)

Results and Discussion

The data take the form of "confusion-matrix" plots, with the physical octave (i.e., the period of the wave) on the abscissa and the average response of the four listeners on the ordinate (Figure 3). If the listeners were invariably correct in identifying the physical octave, the data would all fall on the central, dashed, diagonal line running from (1,1) to (6,6). The upper dashed diagonal shows the position of responses an octave above the physical octave and the lower dashed line shows the position of responses an octave below the physical octave.

AVERAGE RESPONSE AS A FUNCTION OF PHYSICAL OCTAVE

The average of all the responses at each physical octave is represented by the solid line with diamonds in Figure 3. it shows that the octave of the sound is readily identifiable and that the primary determinant of tone height is the period of the wave or, alternatively, the fundamental of the harmonic series. The average data fall above the diagonal at the lowest octaves and below the diagonal at the highest octaves. These departures from the diagonal are largely attributable to the fact that the listeners' responses were restricted to the range 1-6. This response scale was used because the listeners knew the experimental design and the stimulus range.In retrospect, it would probably have been better to instruct them to use the range 0—7. Although the response-range limit reduces the slope of the data slightly, it does not prevent us from observing the effects of the main manipulations.

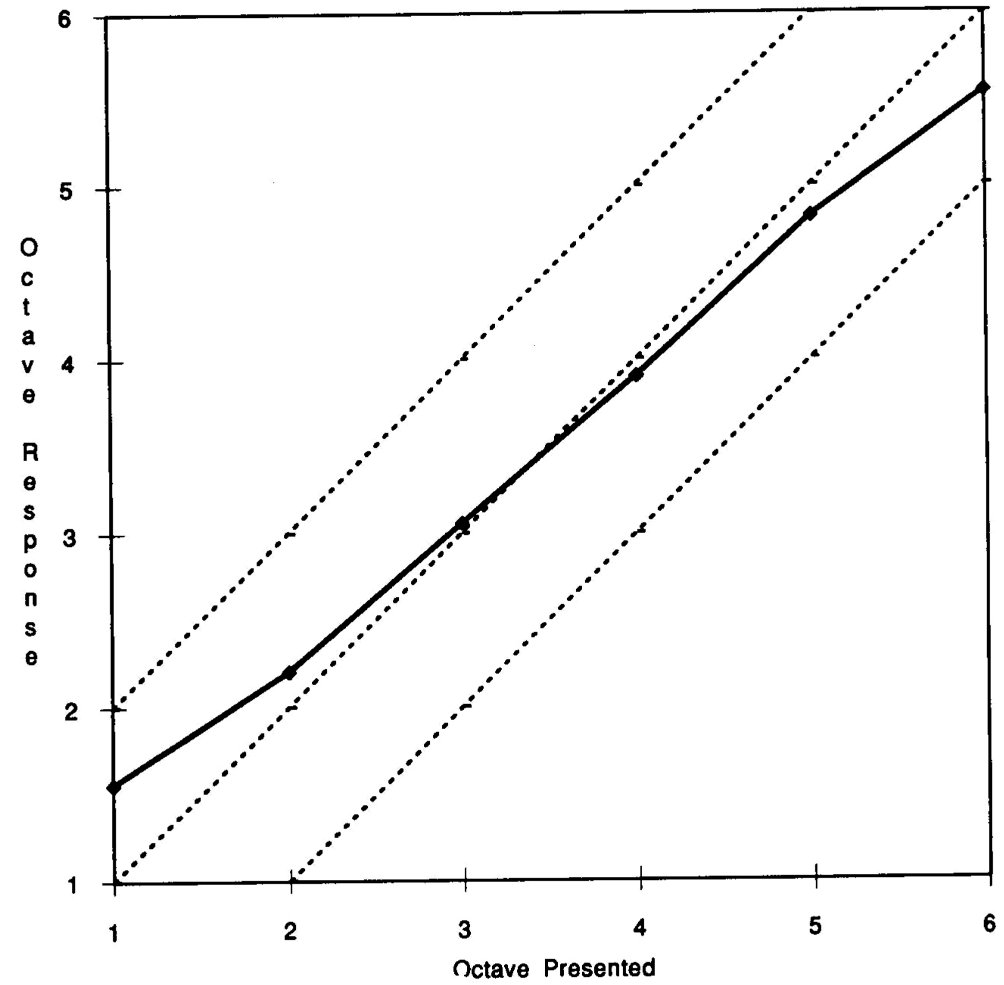

Fig. 3. The average of all of the responses from all of the listeners at each physical octave (diamonds). The central dashed line shows the position of responses that correspond to the physical octave.

THE EFFECT OF ATTENUATING THE UPPER OR LOWER HARMONICS

The mean data of Figure 3 are separated into average responses for the unfiltered stimuli (squares), the lower-harmonic stimuli (triangles), and the upper-harmonic stimuli (diamonds) in Figure 4. The data show that the spectral composition of the sound has a strong effect on tone height—the upper harmonics on their own are perceived to be an octave above their physical octave when the period is longest (octave 1) and still half an octave above their physical octave when the period is shortest (octave 6). Note, however, that it is not the spectral center of gravity that determines tone height; removing harmonics 8-24 from the stimulus that has all harmonics does not change tone height. This indicates that tone height is largely determined by the resolved harmonics when they are present.

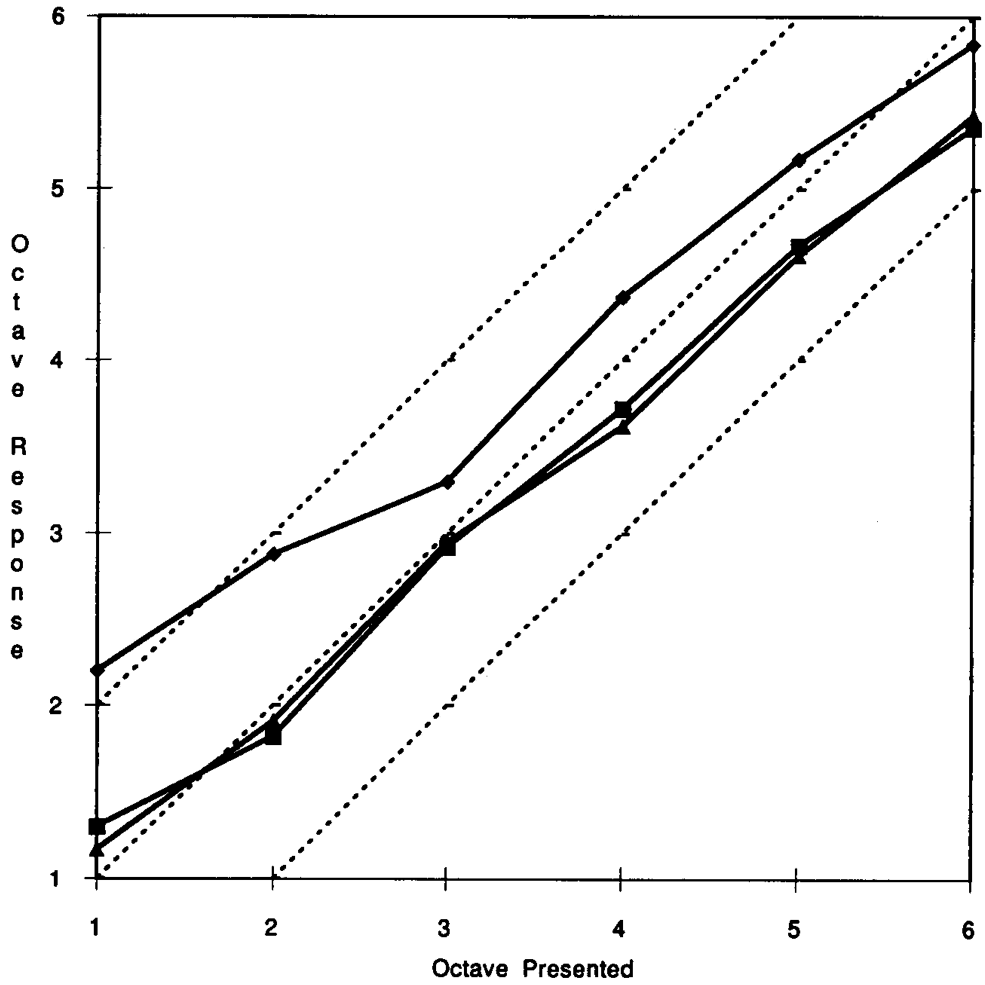

Fig. 4. Average data for sounds with both lower and upper harmonics (squares), lower harmonics alone (triangles), and upper harmonics alone (diamonds). On their own, the upper harmonics lead to elevated tone-height judgments; with the lower harmonics present, they have no influence on tone height.

THE EFFECT OF ATTENUATING THE EVEN HARMONICS

The data for the odd-cosine-phase sounds on their own are presented in Figure 5, with the unfiltered, lower-harmonic, and upper-harmonic stimuli presented separately as in Figure 4. The unfiltered data (squares) and the upper-harmonic data (diamonds) are similar to the corresponding average data in Figure 4. The lower-harmonic data in Figure 5 are lower in the lower octaves than the average lower-harmonic data in Figure 4 (triangles). The data for the sinusoid, on its own, are very similar to those for the odd-cosine-phase lower-harmonic stimuli, which supports the interpretation that the low harmonics carry the most weight in the determination of tone height.

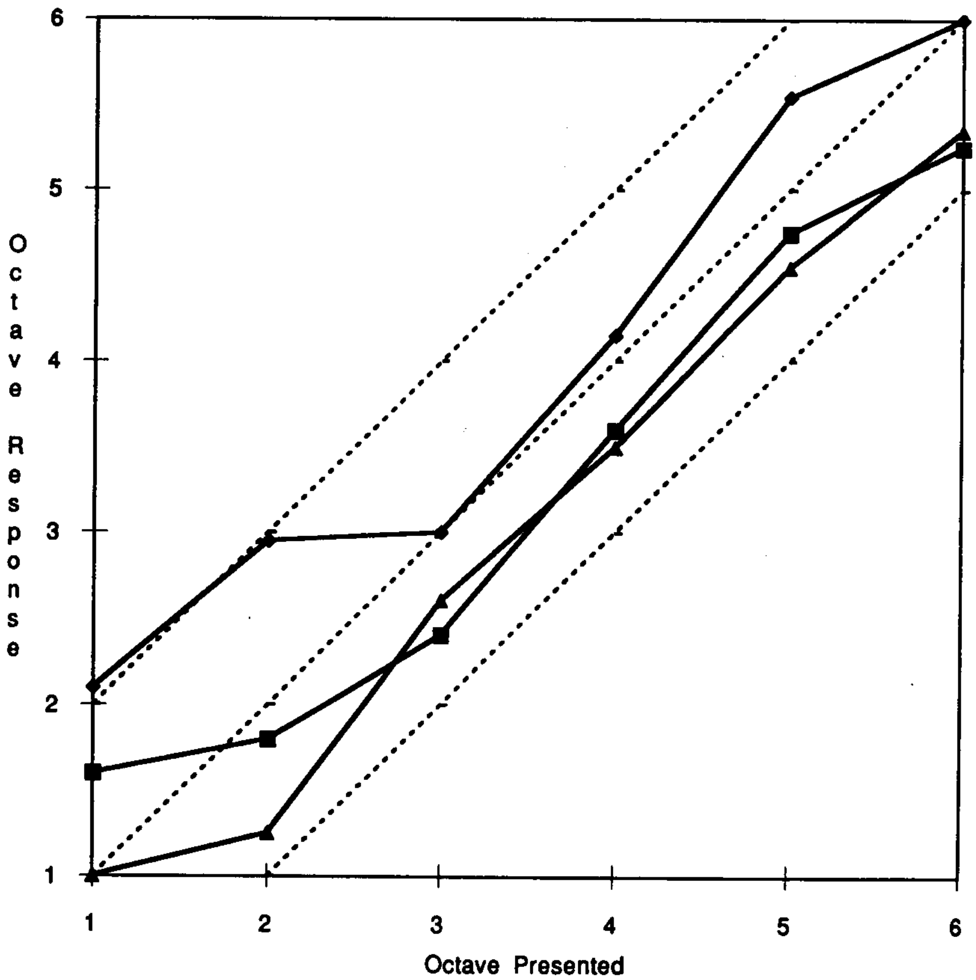

Fig. 5. Average data for odd-cosine-phase sounds with both lower and upper harmonics (squares), lower harmonics on their own (triangles), and upper harmonics on their own (diamonds). The lower harmonics on their own produce the lowest tone-height judgments; as low as a single sinusoid at the fundamental frequency.

THE EFFECT OF PHASE SHIFTING THE ODD OR EVEN HARMONICS

The tone height results for the odd-alternating-phase sounds were very similar to those for the even-alternating-phase sounds. Indeed, odd- alternating-phase and even-alternating-phase sounds are difficult to discriminate. Random-phase sounds are distinguishable from cosine-phase sounds (Patterson, 1987), but the tone-height judgments were essentially the same. Accordingly, the odd- and even-alternating-phase sounds were averaged, and the cosine-phase and random-phase sounds were averaged, to provide the best opportunity for observing the effect of phase shifting the odd or even harmonics. Furthermore, in these average data, the unfiltered responses and the lower-harmonic responses were very similar at all octaves, and so they were averaged in each case. The resulting comparison between sounds with a fixed phase shift (alternating-phase) and those with either no shift (cosine-phase) or a random shift (random-phase) is shown in Figure 6.

The lower pair of curves present the average unfiltered and lower-harmonic data for the alternating-phase sounds (open squares) and the cosine-phase/random-phase sounds (filled squares). They show that there is no effect of phase when the low harmonics are present. The upper pair of curves present the upper-harmonic data, and they show that when the phase of the odd or even harmonics is shifted 90° (open diamonds), there is a substantial increase in tone-height responses over and above that produced by attenuating the resolved harmonics (filled diamonds).

Conclusions

The tone height of multiharmonic tones is readily apparent to listeners, and the primary determinant of tone height is the period of the wave or, alternatively, the fundamental of the harmonic series.

When the resolved harmonics (1-7) are all present, tone height is fixed. Introducing many high harmonics does not raise tone height, nor does shifting the phase of the components.

Removing the even resolved harmonics (leaving 1, 3, 5, and 7) reduces tone height in the lower octaves. A sine tone on its own leads to similarly low tone-height judgments. This indicates that low harmonics are essential for low-octave perceptions.

When the resolved harmonics are removed, tone height rises about half an octave, and when the phase of alternate harmonics is shifted 90°, there is a further increase in tone height of about one quarter of an octave.

All of the effects are largest at low octaves and smallest at high octaves.[1]

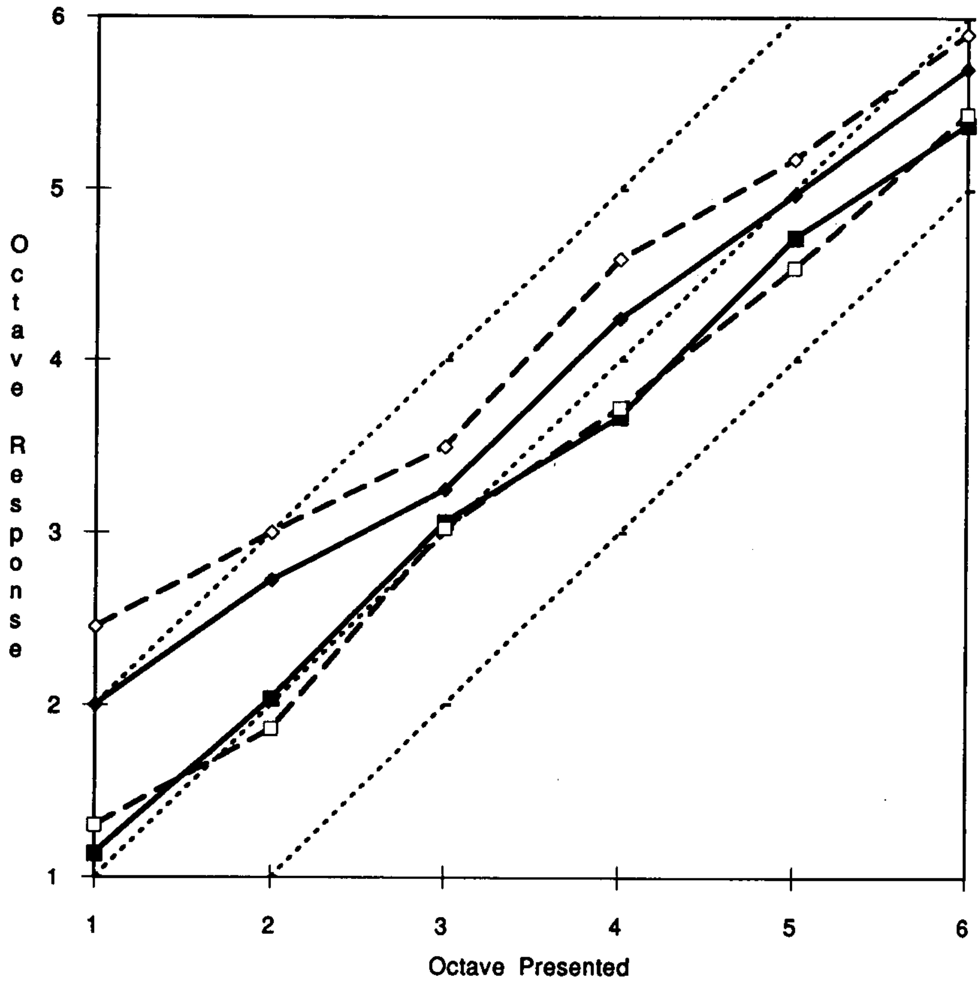

Fig. 6. Average data for alternating-phase sounds (open squares) and cosine-phase/random-phase sounds (filled squares) with both lower and upper harmonics or lower harmonics on their own and with upper harmonics on their own (open and filled diamonds, respectively). Removing the lower harmonics elevates tone height (filled diamonds), and shifting the phase of alternate harmonics elevates it further (open diamonds).

---

[1] 1. This paper was presented at the First International Conference on Music Perception and Cognition, Kyoto, Japan, October 1989.

2. The author thanks Robert Milroy, Ken Robinson, and Paul Cosgrove, who participated in the experiments and assisted with the data analysis.

References

Flanagan, J. L., & Guttman, N. On the pitch of periodic pulses. Journal of the Acoustical Society of America, 1960a, 32, 1308-1319.

Flanagan, J. L., & Guttman, N. Pitch of periodic pulses without fundamental component. Journal of the Acoustical Society of America, 1960b, 32, 1319-1328.

Guttman, N., & Flanagan, J. L. Pitch of high-pass filtered pulse trains. Journal of the Acoustical Society of America, 1964, 36, 757-765.

Moore, B. C. J. Effects of relative phase of the components on the pitch of three-component complex tones. In E. F. Evans & J. P. Wilson (Eds.), Psychophysics and physiology of hearing. London: Academic Press, 1977, pp. 349-358.

Patterson, R. D. The effects of relative phase and the number of components on residue pitch. Journal of the Acoustical Society of America, 1973, 53(6), 1565-1572.

Patterson, R. D. A pulse ribbon model of monaural phase perception. Journal of the Acoustical Society of America, 1987, 82(5), 1560-1586.

Ritsma, R. J., & Engel, F. L. Pitch of frequency modulated signals. Journal of the Acoustical Society of America, 1964, 36, 1637-1655.

Rosenberg, A. E. Effect of masking on the pitch of periodic pulses. Journal of the Acoustical Society of America, 1965, 38, 747-758.

Ueda, K., & Ohgushi, K. Perceptual components of pitch: Spatial representation using a multidimensional seeing technique. Journal of the Acoustical Society of America, 1987, 82(4), 1193-1200.