Scientific Research Publishing

Open Journal of Acoustics, 2017, 7, 52-68
http://www.scirp.org/journal/oja
ISSN Online: 2162-5794
ISSN Print: 2162-5786

Virtual Pitch and Pitch Shifts in Church Bells

William A. Hibbert, Shahram Taherzadeh, David B. Sharp

School of Engineering and Innovation, Open University, Milton Keynes, UK
Email: d.sharp@open.ac.uk

How to cite this paper: Hibbert, W.A., Taherzadeh, S. and Sharp, D.B. (2017) Virtual Pitch and Pitch Shifts in Church Bells. *Open Journal of Acoustics*, **7**, 52-68.
https://doi.org/10.4236/oja.2017.73006

Received: August 3, 2017
Accepted: September 5, 2017
Published: September 8, 2017

Abstract

It is well established that musical sounds comprising multiple partials with frequencies approximately in the ratio of small integers give rise to a strong sensation of pitch even if the lowest or fundamental partial is missing—the so-called virtual pitch effect. Experiments on thirty test subjects demonstrate that this virtual pitch is shifted significantly by changes in the spacing of the constituent partials. The experiments measured pitch by comparison of sounds of similar timbre and were automated so that they could be performed remotely across the Internet. Analysis of the test sounds used shows that the pitch shifts are not predicted by Terhardt's classic model of virtual pitch. The test sounds used were modelled on the sounds of church bells, but a further experiment on seventeen test subjects showed that changes in partial amplitude only had a minor effect on the pitch shifts observed, and that a pitch shift was still observed when two of the lowest frequency partials were removed, so that the effects reported are of general interest.

Keywords

Virtual Pitch, Pitch Measurement, Church Bell, Missing Fundamental, Residue Pitch, Pitch Shift

1. Introduction

Pitch is the characteristic that allows a set of sounds to be ordered in a musical scale. Sometimes, the pitch of a sound corresponds to a single frequency component present in the sound (so-called spectral pitch). However, a sound comprising a harmonic series with partial frequencies $2f$, $3f$, $4f$ etc. gives a strong perception of pitch at frequency f even if no energy at that frequency is present in the sound. This effect is variously known as the missing fundamental, residue pitch or virtual pitch. Terhardt [1] [2] proposed an algorithm for the determination of pitch. A software implementation of this model due to Terhardt and Parncutt [3] allows spectral and virtual pitches to be predicted from partial frequencies and intensities.

The overview by Plack and Oxenham [4] of the work of various researchers shows that virtual pitch effects are most apparent in a restricted frequency range.

Experiments documented in [5] show that, in church bells at least, virtual pitch dominates pitch perception when $2f$ (the lowest partial physically present) lies in the approximate range of 500 Hz to 1500 Hz.

Shifts in the frequencies of the partials of a sound from exact harmonic ratios give rise to a shift in the virtual pitch. Experiments reported by Patterson and Wightman in [6] show that the size of the virtual pitch shift due to changes in the partial frequencies depends on the absolute position in the audible spectrum, being larger at lower frequencies. The existence of the virtual pitch effect and pitch shifts in piano tones is investigated in work by Järveläinen *et al.* [7] and Anderson and Strong [8]. Experimental evidence of virtual pitch shifts in bells at a single point in the audible spectrum due to changes in various partials is presented in [9] and [10].

The term pitch shift is used for more than one effect in the literature. In particular in Terhardt's work on virtual pitch [2], pitch shift is used to describe shifts in pitch due to partial amplitude and the close proximity of pairs of partials. However, the experiments discussed by Plomp [11] show no significant change in pitch shifts between two cases, one with all partials of equal amplitude, and the other with amplitude falling off as $1/n$ where n is the partial number. Meanwhile, work by Lin and Hartmann [12] shows that the virtual pitch effect survives when some of the partial frequencies are missing. In this paper, the term is used exclusively for shifts in virtual pitch due to changes in partial spacing.

The textbook definition of pitch given above mandates comparison by listeners as the only valid method of pitch estimation. Experiments show a variation in results between listeners—there is no absolute value for the pitch of a sound independent of a particular listener's ability to judge it. Various methods for pitch estimation are documented in the literature:

- comparison of complex tones against sine tones, either from tuning forks or electronically generated;
- the method of post vocalisation described by Terhardt and Seewann [13] in which test subjects are asked to sing or hum the pitch they perceive;
- comparison of complex tones against test tones with similar timbre.

This last method, used effectively by Järveläinen *et al.* [7] for work on piano tones, and also extensively validated by the author in [5], was used for the experiments described here.

Virtual Pitch in Church Bells

Church bells make an interesting case study for investigations of virtual pitch because perception of their pitch is dominated by virtual pitch effects. Terhardt documents investigations into virtual pitch in bells in [13]. An English translation of this paper appears in [5].

Figure 1 shows a typical spectrum of a bell sound (the old no. 9 bell from St Mary-le-Tower, Ipswich), comprising a mixture of inharmonic and near-harmonic partials. The identification of the significant partials follows the scheme described by Lehr in [14], except for the set of three partials labelled p1, p2 and p3. These three partials, not usually tuned or recorded by bell founders, have an effect on the timbre and therefore potentially on the pitch of the bell.

Table 1 shows the frequencies of the significant partials in the sound of the Ipswich bell together with the ratios between the partial frequencies. There are two series of partials with an approximately harmonic relationship (frequencies in the ratio 2:3:4), one based on the nominal and the other based on I-7, highlighted in **Table 1**. Experiments documented in [5] show that, for most bells, the series based on the nominal gives a dominant pitch sensation about an octave below the nominal. For low-pitched bells (nominals below 500 Hz), however, the higher series based on I-7 gives the dominant pitch sensation, and for high pitched bells (nominals above 1500 Hz), individual low partials dominate the pitch sensation as spectral pitches. For convenience in description, the nominal, superquint, octave nominal, I-7, I-8 etc. are commonly referred to as the "upper partials".

Due to the considerable importance of the nominal to the perceived pitch of bells, this partial is taken as the reference when considering the relationship

Table 1. Partial frequencies for Ipswich Old No. 9 bell, and ratios between the partial frequencies.

		Ratios between partials											
	Freq. (Hz)	Hum	Prime	Tierce	Quint	Nom.	p1	p2	p3	Superquint	Oct. Nom.	I-7	I-8
Hum	201.9	1.00											
Prime	323.7	1.60	1.00										
Tierce	424.2	2.10	1.31	1.00									
Quint	562.8	2.79	1.74	1.33	1.00								
Nominal	723.4	3.58	2.23	1.71	1.29	***1.00***							
p1	862.1	4.27	2.66	2.03	1.53	1.19	1.00						
p2	882.2	4.37	2.73	2.08	1.57	1.22	1.02	1.00					
p3	902.0	4.47	2.79	2.13	1.60	1.25	1.05	1.02	1.00				
Superquint	1091.4	5.41	3.37	2.57	1.94	***1.51***	1.27	1.24	1.21	1.00			
Oct. Nominal	1507.3	7.47	4.66	3.55	2.68	***2.08***	1.75	1.71	1.67	1.38	1.00		
I-7	1959.3	9.70	6.05	4.62	3.48	2.71	2.27	2.22	2.17	1.80	1.30	***1.00***	
I-8	2443.2	12.10	7.55	5.76	4.34	3.38	2.83	2.77	2.71	2.24	1.62	1.25	1.00
I-9	2948.3	14.60	9.11	6.95	5.24	4.08	3.42	3.34	3.27	2.70	1.96	***1.50***	1.21
I-10	3467.6	17.17	10.71	8.17	6.16	4.79	4.02	3.93	3.84	3.18	2.30	1.77	1.42
I-11	4000.8	19.82	12.36	9.43	7.11	5.53	4.64	4.54	4.44	3.67	2.65	***2.04***	1.64
I-12	4536.5	22.47	14.01	10.69	8.06	6.27	5.26	5.14	5.03	4.16	3.01	2.32	1.86
I-13	5077.1	25.15	15.68	11.97	9.02	7.02	5.89	5.76	5.63	4.65	3.37	2.59	2.08

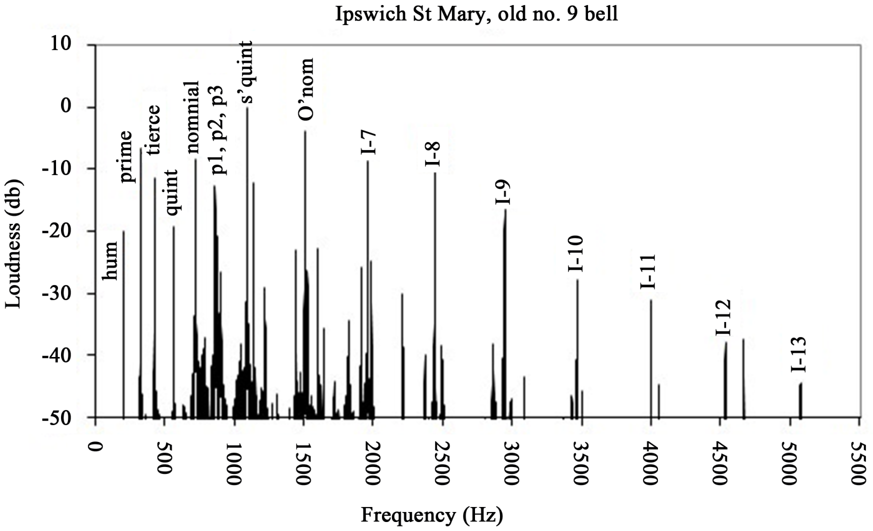

Figure 1. Typical spectrum of a bell sound.

between the frequencies of the various partials. The relationship between partial frequencies and the nominal frequency can be expressed either as a simple ratio or as a difference in cents (hundredths of a semitone) calculated using the following expression:

$$\text{cents} = \frac{1200}{\ln(2)} \ln\left(\frac{f_p}{f_n}\right) \tag{1}$$

where f_p is the partial frequency and f_n is the nominal frequency.

In bell sounds, the frequencies of the upper partials are either compressed together or stretched apart relative to the nominal, according to the shape or thickness of the bell. Informal observations suggest that these changes alter the perceived pitch of the bell by one quarter or more of a semitone. A simple relationship which allows the intervals of all the upper partials to the nominal to be represented by the interval between the octave nominal and the nominal is documented in [5]. The relationship can be visualised by imagining the spectrum of the bell to be drawn on a sheet of rubber. In thinner bells the rubber is stretched out more, moving the upper partials proportionally further apart.

The intervals of the other named partials to the nominal do not display this simple relationship.

2. Pitch Shift Experiment

An experiment was designed to quantify the shifts in perceived pitch due to changes in various partials. The experiment involved comparison of the pitches of test tones against reference tones with similar timbre. Both test and reference tones had partial frequencies and amplitudes typical of bells, derived from the average of measurements on a number of bell recordings. Nine test sets in all were constructed, with nominal frequencies spaced at 1/3 octave intervals (*i.e.* increasing in frequency by a factor of $2^{1/3}$ each time) from 315 Hz to 2000 Hz. To explore the effect of all the partial frequencies on the perceived pitch of a bell sound, in principle it is necessary to vary them all in relation to the nominal. However, because the intervals of all the upper partials to the nominal are related in bells, they must move together. This considerably simplifies the design of the experiment.

The independent variables chosen for the experiment, following some initial trials, were the intervals to the nominal of prime, tierce, quint, p1 to p3, and octave nominal (standing proxy for the intervals of all the upper partials). Intervals of each partial to the nominal occur over a range of values in bells seen in practice, and low, typical and high values of the interval were established for each partial. Values for the low and high values were chosen (based on examination of about 2000 bell recordings) to be extreme but realistic. To further simplify the experiment, all three of p1, p2 and p3 were assigned low, average or high values together even though, unlike the upper partials, no relationship exists between them. This gave five independent variables in total to be explored.

The amplitudes and intervals to the nominal (expressed as a frequency ratio) for the partials used in the tests are shown in **Table 2**.

A simple amplitude envelope for the partials was used, based again on measurements on partials from a number of bell recordings. The envelope used was $A/(1 + dt)$ where A is the initial amplitude (column labelled Amplitude in **Table 2**), t is time in seconds, and d is a decay constant set to 5 for all partials except the hum, for which it was set to 1. As will be demonstrated below, the virtual pitch shifts observed were relatively insensitive to amplitude.

Table 2. Intervals to nominal for virtual pitch tests.

Partial	Amplitude	Freq. ratio to nominal (low)	Freq. ratio to nominal (typical)	Freq. ratio to nominal (high)
Hum	3.14		0.25	
Prime	4.02	0.45	0.50	0.50
Tierce	6.78	0.59	0.60	0.63
Quint	0.82	0.71	0.77	0.84
Nominal	10		1.00	
p1	1.25	1.12	1.19	1.26
p2	2.43	1.19	1.28	1.37
p3	1.83	1.28	1.36	1.43
Superquint	6.88	1.44	1.48	1.51
Oct. Nominal	5.15	1.93	2.01	2.09
I-7	3.71	2.45	2.59	2.74
I-8	2.39	3.00	3.21	3.43
I-9	1.75	3.56	3.85	4.16
I-10	1.13	4.13	4.51	4.92
I-11	0.76	4.71	5.19	5.71

2.1. Experimental Design and Procedure

2.1.1. Test Sounds

Each of the nine test sets comprised 16 test sounds. All the test sounds in each set had the same nominal frequency, and the intervals of other partials were varied relative to this fixed nominal. The number of test sounds in each set was chosen to be 16 because this number of tests could be performed by each test subject in a reasonable period of time (no subject took longer than 45 minutes to complete a test set), while exploring the effect of four independent variables, each at either a high or low level; four variables each with two levels gives 16 combinations in total. To allow the effect of five independent variables to be explored, in different test sets either the quint or p1, p2 and p3 was varied, with the other held at the average value. These two variables were treated in this way because, although experience suggested that they did not affect virtual pitch, for completeness it was desired to explore their contribution.

The scheme used to explore the effect on virtual pitch of the five independent variables across the nine test sets is shown in Table 3.

The sixteen test sounds, using the case of nominal = 1000 Hz as an example, involved the combinations shown in Table 4. As will be explained below, the tests were actually presented in random order to the test subject.

2.1.2. Reference Sounds

To allow delivery of the tests across the internet, pre-created reference sounds were used in the tests. For each individual test, the subject was asked to select which of 16 reference sounds was closest in pitch to the test sound. The reference sounds for all 16 tests in a test set were the same. The reference sounds were built from the same partials as the test sounds, with the same amplitudes, but with their intervals to the nominal set to the typical values as shown in Table 2. The nominals of the reference sounds were spaced geometrically at equal intervals apart. The spread of nominals in the reference sounds for each test was

Table 3. Independent variables used in virtual pitch tests.

Test nominal freq. (Hz)	Prime	Tierce	Quint	p1, p2, p3	Upper partials
315	low or high	low or high	typical	typical	low or high
397	low or high	low or high	low or high	typical	low or high
500	low or high	low or high	typical	low or high	low or high
630	low or high	low or high	low or high	typical	low or high
794	low or high	low or high	typical	low or high	low or high
1000	low or high	low or high	low or high	typical	low or high
1260	low or high	low or high	typical	low or high	low or high
1587	low or high	low or high	low or high	typical	low or high
2000	low or high	low or high	typical	low or high	low or high

Table 4. Arrangement of independent variables in virtual pitch tests.

	Prime	Tierce	Quint	p1, p2, p3	Upper partials
test 1	low	low	low	typical	low
test 2	low	low	high	typical	low
test 3	high	low	low	typical	low
test 4	high	low	high	typical	low
test 5	low	high	low	typical	low
test 6	low	high	high	typical	low
test 7	high	high	low	typical	low
test 8	high	high	high	typical	low
test 9	low	low	low	typical	high
test 10	low	low	high	typical	high
test 11	high	low	low	typical	high
test 12	high	low	high	typical	high
test 13	low	high	low	typical	high
test 14	low	high	high	typical	high
test 15	high	high	low	typical	high
test 16	high	high	high	typical	high

established during preliminary trials to ensure that the range of pitches experienced by subjects lay within the range of reference sounds for that test. In the trials it was found that the pitch shifts, and hence the spread required in the reference nominals, was rather greater at lower nominal frequencies. The details of the reference sounds for each test set appear in Table 5.

2.1.3. Generation of the Test and Reference Sounds

The parameters for the test and reference sounds of all nine test sets were calculated in a spreadsheet and exported as csv files. These files of data were imported into bespoke software which generated each batch of test and reference sounds via additive synthesis of cosine waves.

The generated sounds were 16-bit with a sampling rate of 22,050 samples per second. The duration of all the sounds was set at 0.25 s. This short duration was intended to encourage pitch perception via virtual pitches rather than identification of individual partials (and is also typical of the spacing between bell sounds in English change ringing). During preliminary extensive trials of the tests, some tests with sounds 1 s long were conducted; the results were not significantly different from those at 0.25 s. In [13] Terhardt and Seewann report no significant difference in the pitch estimation ability of test subjects between sounds of 0.1 s and 3 s duration.

2.1.4. Execution of the Tests

Conduct of the tests and reporting of the results was automated in software so

Table 5. Reference nominals for each test run.

Test nominal freq. (Hz)	Lowest reference nominal freq. (Hz)	Highest reference nominal freq. (Hz)	Interval between lowest and highest reference nominal freq. (cents)
315	292.19	347.48	300
397	376.66	421.97	197
500	476.80	522.14	157
630	609.62	650.76	112.5
794	768.33	819.91	112.5
1000	968.03	1033.03	112.5
1260	1219.64	1301.53	112.5
1587	1553.37	1622.16	75
2000	1957.14	2043.79	75

that the tests could be carried out remotely across the internet. Test and reference sounds were played by the subject clicking on links on a webpage, using whatever sound reproduction equipment was attached to their PC. Each of the nine test sets was conducted separately (and not all subjects completed all nine sets). For each of the sixteen tests within a set, the subject was presented with links for the test sound and for the sixteen reference sounds, and asked to select which reference sound was closest in pitch (neither higher nor lower) to the test sound. Within each set of sixteen, the tests were displayed in a different random order on each test iteration. The subject was able to spend as long as required on each test but could not complete the set of sixteen tests without responding to them all. The result reported for each test was the nominal of the matched reference sound, which was taken to stand for the pitch the user had perceived. This was converted into a pitch shift in cents using Equation (1).

Extensive trials of the approach employed were carried out to validate the test procedure and documentation, analysis of the results, the effect on test results of quality of sound reproduction equipment and musical experience of the subject, and ranges of the various test parameters. The experiment design was proven to be robust against all these criteria, and in particular the quality of sound reproduction equipment was eliminated as a concern.

The test results were analysed using the standard analysis of variance (ANOVA) statistical technique [15]. This allowed the effect of each of the independent variables on pitch shift to be isolated from that of the others, and also allowed the statistical validity of the test results to be established.

2.2. Experimental Results

30 test subjects (including one of the authors) completed 116 test sets, a total of 1856 individual tests. Subjects were recruited through bell-ringing newsgroups with a request for people with "a PC or Mac with sound, a broadband connec-

tion and a musical ear". The statistical analysis of the results showed a small number of outliers (0.54% of the total) but all results were included in the analysis as the influence of the outliers on the test results was minimal.

The results were analysed using the standard ANOVA technique. Because the experiment design was orthogonal (*i.e.* all combinations of the independent variables were present; see **Table 4**) the ANOVA analysis allows two key sets of information to be obtained from the experiment results; 1) the effect of each independent variable individually on the pitch shift, averaged across all subjects and tests at a particular nominal frequency, and 2) the statistical significance of each of these effects.

As an example of the results, **Figure 2** shows the pitch shift due to upper partials experienced by 26 test subjects, for test tones at 1000 Hz. The quantity plotted is the interval in cents between the pitch observed with the upper partials at their low values, and pitch observed with the upper partials at their high values, averaged across all the tests at 1000 Hz. The order of the subjects is that in which they performed the test. The average shift for all subjects was 36 cents. Whether the spread in pitch shifts experienced was due to experimental variation or some characteristic of the subjects was not determined.

The statistical significance of each effect derived from the experiment results is expressed as a P-value. The P-value is the probability that a pitch shift of this size could have been observed (through chance variation) if in fact the independent variable did not give rise to a shift. A P-value of less than 0.05 indicates that the pitch shift is statistically significant at the 95% level.

Table 6 shows the pitch shifts in cents due to each of the independent variables in the nine test sets, together with the P-value derived from the statistical analysis. The interpretation of these results is as follows:

- the pitch shifts due to changes in the upper partials are both highly statistically significant, and musically significant. Pitch shifts range from 1/8 of a semitone at the highest frequency, through half a semitone in the middle range of frequencies, to well over a semitone at the lowest frequency;

Table 6. Virtual pitch test results.

Test nominal freq. (Hz)	Prime		Tierce		Quint		Middle partials		Upper partials	
	Effect (cents)	P-value	Effect (cents)	P-value	Effect (cents)	P-value	Effect (cents)	P-value	Effect (cents)	P-value
315	12.5	0.047	7.1	0.256					134.0	6.2×10^{-53}
397	10.1	0.057	9.6	0.070	1.1	0.841			71.4	2.0×10^{-27}
500	8.2	0.022	7.7	0.032			0.5	0.889	71.5	2.0×10^{-43}
630	6.6	0.019	9.3	0.001	0.7	0.793			43.2	2.0×10^{-31}
794	6.3	0.001	8.2	1.4×10^{-5}			1.8	0.330	42.2	1.1×10^{-56}
1000	6.2	3.9×10^{-6}	8.3	1.0×10^{-9}	−0.9	0.499			36.0	2.4×10^{-92}
1260	4.9	0.007	11.0	5.6×10^{-9}			0.3	0.853	31.6	7.6×10^{-43}
1587	6.3	0.001	5.9	0.001	0.3	0.849			18.8	1.2×10^{-18}
2000	0.4	0.806	8.1	1.3×10^{-6}			0.1	0.972	12.8	3.9×10^{-13}

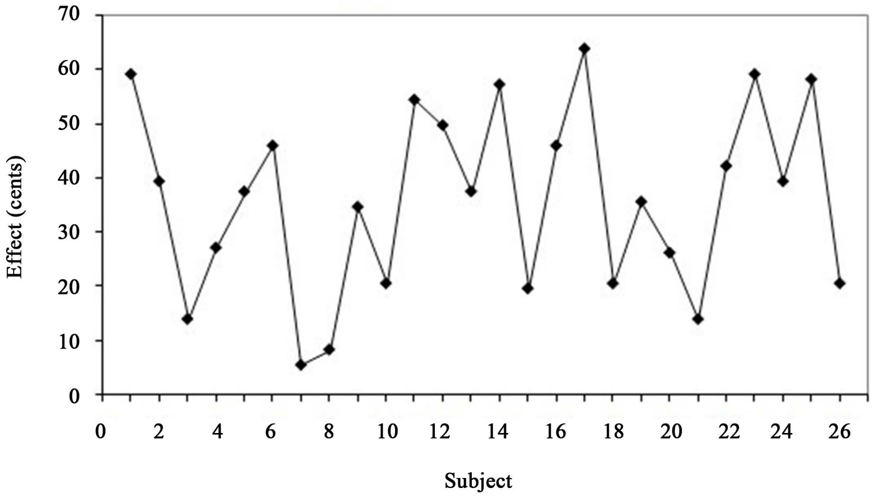

Figure 2. Pitch shift perceived by individual subjects due to change in upper partials, for test tones at 1000 Hz.

- pitch shifts due to the quint and the middle partials are neither statistically nor musically significant at any frequency;
- shifts due to prime and tierce are statistically significant at most frequencies, but musically quite small, generally a small fraction of the shifts due to the upper partials.

28 of the 30 test subjects gave an account of their musical experience, which ranged from near-professional experience to no experience at all. There was no correspondence between musical experience (or lack of it) and the size of the pitch shifts experienced by the test subjects.

As explained in Section 1.1, in bell sounds the spacing of the upper partials is uniquely parameterised by the interval of the octave nominal relative to the nominal. It is therefore convenient to show the upper partial pitch shift (in cents) for a one-cent shift in the octave nominal. The high and low frequencies of the octave nominal given in Table 2 are 1280 and 1140 cents above the nominal respectively, so the unit pitch shift can be calculated by dividing the upper partial effects in Table 6 by 1280 − 1140 = 140. The results, together with 95% confidence intervals, are as shown in Figure 3.

The dotted line is a regression assuming a simple log relationship (specifically $\ln(\text{pitch shift}) = -1.10012 \ln(f_n) + 6.12396$). No theoretical justification is proposed for this relationship; it just provides convenient shorthand allowing pitch shifts at an arbitrary frequency to be estimated. It is hypothesised that the very large pitch shifts at low nominal frequencies arise because, at this end of the spectrum, more subjects are hearing the pitch based on partial I-7 rather than the pitch based on the nominal.

A subsidiary experiment was conducted to explore the linearity or otherwise of the relationship between upper partial spacing and pitch shift. Test sounds were generated with the upper partial spacing in 16 steps from the low to the high values in Table 2, and all other partials set to their typical values. Figure 4

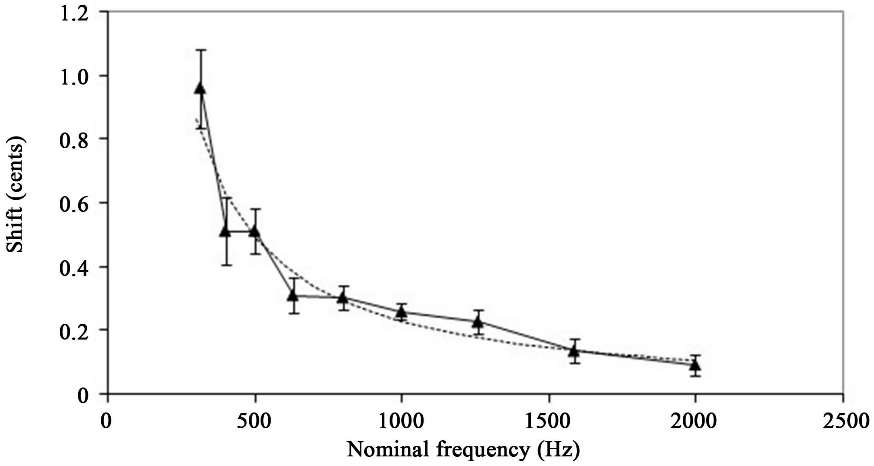

Figure 3. Average perceived pitch shift, per one-cent shift in octave nominal, due to change in upper partials.

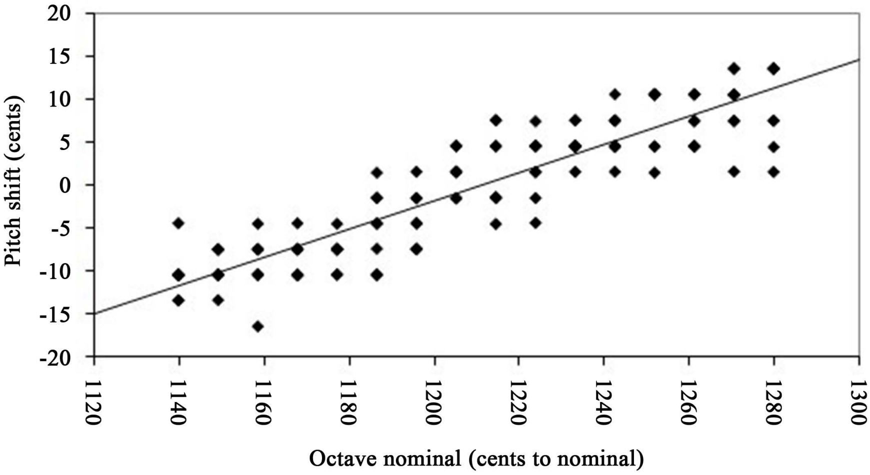

Figure 4. Dependence of perceived pitch shift on upper partial spacing.

shows the results of eight test runs by a single subject, together with a linear regression on the results. The interval of the octave nominal to the nominal is, as ever, used as proxy for the spacing of all the upper partials. The average pitch shift between the case with upper partials at their low and high values is 23.0 cents. There is nothing in these results to suggest that the pitch shift effect due to upper partial spacing is not linear over the range of spacings used in the experiments.

Additional work also documented in [5] used a model derived from the results of the main experiment above to investigate the tuning of peals of bells. Thick, heavy bells have compressed upper partials, and this experiment predicts that this will flatten their pitch. Comparison of the model against the partial frequencies of peals of bells known to have been tuned using virtual pitches rather than nominal frequencies, showed that the model successfully predicted the changes in bell pitch due to upper partial effects. The shifts found in practice can be up to 40 cents or 4/10 of a semitone, easily heard by a musical ear.

3. Effect of Partial Amplitude on Pitch Shift

A second experiment was constructed to investigate the effect of partial amplitude on the pitch shifts caused by changes in the upper partials. In this experiment, the amplitudes of superquint and octave nominal were varied from zero to 100% of the amplitudes in the previous experiment, while at the same time the frequencies of all the upper partials took on either the low or high values as defined in Table 2. All other partials took on their typical values. As before, the test subject was asked to match the pitch of the test tone against reference tones. The experiment was carried out at a single nominal frequency, 1000 Hz, chosen to be the middle of the range of nominals in the previous experiment.

3.1. Experimental Design and Procedure

To take advantage of the experimental procedure developed and proved in the previous experiment, this experiment followed a very similar approach and used the same software for delivery and execution.

The superquint and octave nominal amplitudes each took on four values; 0%, 33%, 67% and 100% of the amplitudes listed in Table 2. These values together with the low and high values for the upper partial frequencies give 32 combinations. For convenience, the combinations were divided into two sets of 16 tests to allow the established test procedure to be used. The two sets were constructed as shown in Table 7.

Table 7. Arrangement of independent variables in amplitude tests.

Test set 1				Test set 2			
	Superquint amplitude	Oct. Nom. amplitude	Upper partials		Superquint amplitude	Oct. Nom. amplitude	Upper partials
test 1	0%	0%	high	test 17	0%	0%	low
test 2	0%	33%	low	test 18	0%	33%	high
test 3	0%	67%	low	test 19	0%	67%	high
test 4	0%	100%	high	test 20	0%	100%	low
test 5	33%	0%	low	test 21	33%	0%	high
test 6	33%	33%	high	test 22	33%	33%	low
test 7	33%	67%	high	test 23	33%	67%	low
test 8	33%	100%	low	test 24	33%	100%	high
test 9	67%	0%	low	test 25	67%	0%	high
test 10	67%	33%	high	test 26	67%	33%	low
test 11	67%	67%	high	test 27	67%	67%	low
test 12	67%	100%	low	test 28	67%	100%	high
test 13	100%	0%	high	test 29	100%	0%	low
test 14	100%	33%	low	test 30	100%	33%	high
test 15	100%	67%	low	test 31	100%	67%	high
test 16	100%	100%	high	test 32	100%	100%	low

The arrangement of independent variables in each test set allowed statistically valid results to be derived if the subject only completed one of the two sets of tests. The reference sounds used were identical to those used in the previous experiment. All other aspects of the design and conduct of this experiment were identical to those previously described.

3.2. Experiment Results

Subjects were again recruited through bell-ringing newsgroups. In total, 17 test subjects undertook the tests, one of whom was one of the authors of this paper. One subject only completed one of the two sets of tests, so the average results of the remaining subjects were used to provide the missing results. In total, 528 individual tests were completed.

As in the previous experiment, all subjects experienced a shift in pitch due to the change in the upper partials. An analysis of variance showed that the pitch shift due to change in upper partials averaged across all test subjects and partial amplitudes was 15.8 cents with a P-value of 2.7×10^{-74}. As in the previous experiment, the shift was both highly statistically significant and musically significant.

The average values across all test subjects of the perceived pitch shift due to upper partial change were plotted against both superquint and octave nominal amplitude. The perceived pitch shifts at different superquint amplitudes are averaged across all octave nominal amplitudes, and vice versa. The plots appear as Figure 5 and Figure 6, with 96% confidence intervals marked.

These results suggest that pitch shifts due to changes in upper partials are relatively insensitive to the amplitude of the individual partials. The pitch shift averaged across all subjects with both superquint and octave nominal amplitudes at their maximum value is 21.1 cents. The pitch shift averaged across all subjects with both superquint and octave nominal absent is 8.8 cents. The shift with par-

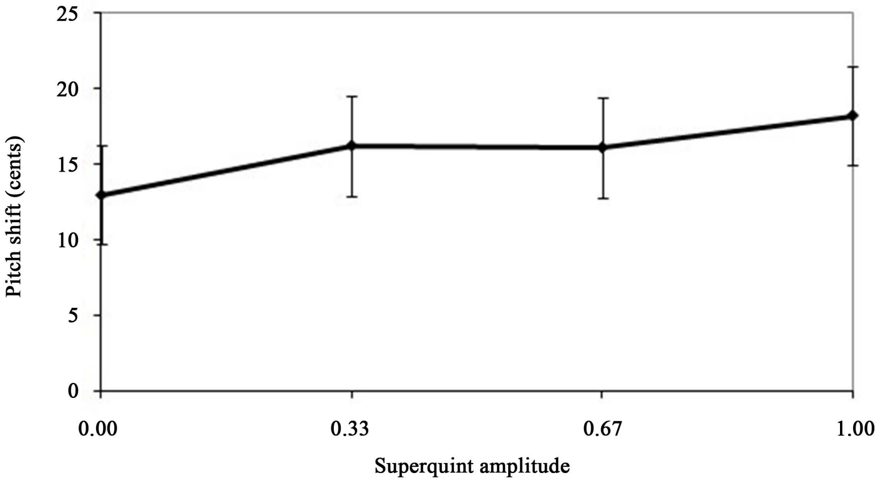

Figure 5. Average perceived pitch shift due to change in upper partials, plotted as a function of superquint amplitude.

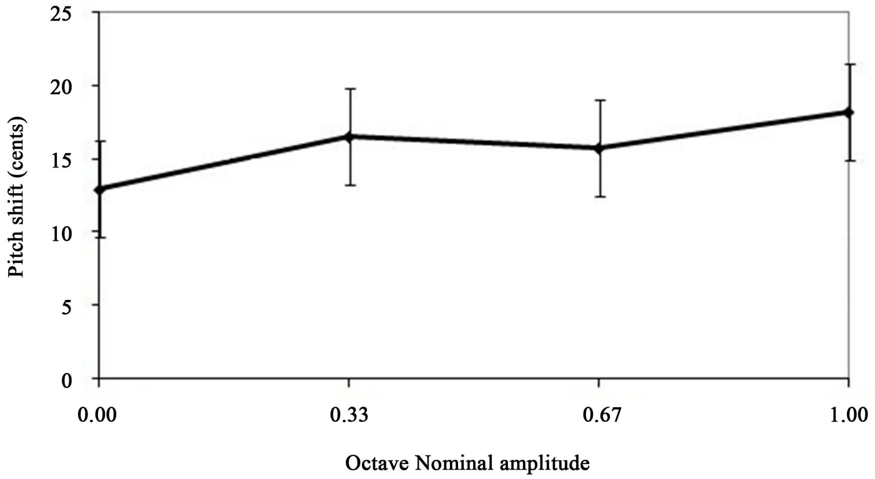

Figure 6. Average perceived pitch shift due to change in upper partials, plotted as a function of octave nominal amplitude.

tials at their maximum amplitude from this experiment is comparable with that shown in **Figure 4** of 23 cents but is rather less than the 36 cents seen in the first experiment. This difference is due to different groups of test subjects performing the experiments: **Figure 2** shows the considerable variation in pitch shift experienced by different subjects.

The complexity of the partial structure in bells obscures the simplicity of the effect being experienced in this experiment. Using the terminology of the introduction, the set of upper partials from the nominal upwards form an approximate harmonic series with partial frequencies close to $2f$, $3f$, $4f$, $5f$, $6f$ etc. The shift in perceived pitch due to changes in the spacing of these partials survives even if those at $3f$ and $4f$ are removed from the sound. These results imply that all the partials in a harmonic series, not just the lowest in frequency, contribute to the virtual pitch effect.

4. Comparison against Terhardt Virtual Pitch Model

Terhardt in [2] gives an algorithm for establishing perceived pitches, both spectral and virtual, from a set of partial frequencies and amplitudes. This algorithm was used in Terhardt's research on bell pitches described in [13]. In this work, the experimenters compared the predictions of the algorithm with pitch determination experiments on a number of bells. The predictions of the algorithm were a reasonable match for the experimental results as regards prediction of the note names of the pitches (*i.e.* to within half a semitone in either direction), but investigation of the detailed correspondence between measured and predicted pitches showed average deviations of between 15 and 40 cents. No attempt was made to explain the deviations as pitch shifts.

The Terhardt algorithm has been implemented in software (C program ptp2svp) by Terhardt and Parncutt [3]. This software implementation of the algorithm was used to investigate its predictions in respect of the first two experiments described in this paper. Partial frequencies and amplitudes for all 144 test sounds described in Section 2.1 (16 test tones at each of 9 nominal frequencies) were processed by the algorithm. The algorithm requires input of partial amplitudes as an absolute level in dB. As this absolute level was unknown for the test subjects, two sets of amplitudes were submitted, one with the amplitude of the loudest partial set to 85 dB (and the others scaled in proportion) and one with the amplitude of the loudest partial set to 60 dB. The results of the algorithm for both were very similar; only the results for 85 dB are reported here.

For the test sounds with nominal frequencies of 500 Hz and above, the algorithm predicted a single dominant pitch related to the nominal frequency. The effect on pitch of changes in individual partials was established by calculating contrasts and effects in the same manner as used in analysing the first experiment in this paper. The results obtained (quoted as pitch shifts in cents) are presented in Table 8.

Comparison with the experimental results given in Table 6 shows that the algorithm predicts no pitch shifts as a result of upper partial changes, whereas the experiments show that this is the dominant and significant effect. The prime and tierce effects predicted by the algorithm are at least an order of magnitude less than those seen in the experiments. The Terhardt algorithm therefore does not predict the pitch shifts seen in practice at these nominal frequencies.

The algorithmic results for test tones at nominal frequencies of 315 Hz and 397 Hz show multiple pitches with roughly equal perceived weight. This is in line with experience; due to the effect of the dominance region for virtual pitch, at lower nominal frequencies the higher harmonic series shown in Table 1 vies with the lower harmonic series as a generator of virtual pitch. The pitches predicted by the algorithm lie approximately two octaves below the nominal, one octave below the nominal, two octaves below partial I-7 (but not in every case at 397 Hz), and at the nominal frequency. Shifts in these pitches predicted by the algorithm due to changes in partials are shown in Table 9.

The pitch shifts seen in the experiment in Section 2 are an amalgamation of the various virtual pitches as predicted by the Terhardt algorithm, across the set

Table 8. Tehardt algorithm results.

Nominal freq. (Hz)	Prime effect (cents)	Tierce effect (cents)	Quint effect (cents)	p1, p2, p3 effect (cents)	Upper partial effect(cents)
500	0.455	0.065		0.000	0.000
630	0.450	0.040	0.000		0.000
794	0.395	0.015		0.000	0.000
1000	0.300	0.010	0.000		0.000
1260	0.195	0.005		0.000	0.000
1587	0.110	0.000	0.000		0.000
2000	0.050	0.000		0.000	0.000

Table 9. Tehardt algorithm results at lower frequencies.

Predicted pitch (315 Hz nominal)	Prime effect (cents)	Tierce effect (cents)	Upper partial effect (cents)	Predicted pitch (397 Hz nominal)	Prime effect (cents)	Tierce effect (cents)	Upper partial effect (cents)
Nominal/4	0.356	0.097	0.001	Nominal/4	0.424	0.085	0.000
Nominal/2	0.123	0.203	0.048	Nominal/2	0.131	0.251	0.029
Nominal	0.245	0.410	0.095	Nominal	0.264	0.504	0.056
I-7/4	0.330	5.225	23.685				

of test subjects. However, for both the 315 Hz and 397 Hz test sounds, the large upper partial effect seen in the experiment results is not predicted by the algorithm; nor are the significant prime and tierce effects.

The Terhardt algorithm was also applied to the test sounds used in the subsidiary experiment, the results of which are shown in **Figure 4**. The test tones used in this experiment had a nominal frequency of 1000 Hz, and as seen in **Figure 4**, an average pitch shift of 23.0 cents was seen between the test sounds with the closest and widest upper partials. However, the Terhardt algorithm predicted identical virtual pitches of 248.1 Hz and 505.4 Hz for all sixteen test sounds.

On the basis of these trials one can confidently say that the pitch shifts seen in practice are not predicted by the Terhardt algorithm.

5. Conclusions

The two experiments reported in this paper have shown that the perceived pitch of a complex tone due to the virtual pitch effect is shifted from the expected value (of half the lowest frequency in the harmonic series) if the frequencies of the partials in the harmonic series are stretched or compressed. The second experiment also shows that the size of this shift is only moderately affected by the amplitude of the next two higher frequency partials in the series. A pitch shift is still observed if these two partials are removed completely from the sound. The pitch shifts are large enough (at considerable fractions of a semitone) to be musically significant.

These findings have a number of consequences:

- although the complex tones used in the experiments were modelled on the sound of church bells, the relative insensitivity of pitch shifts to partial amplitude means that the principles, if not the exact numerical results, are of general applicability;
- many partials in the harmonic series, not just the lowest in frequency, contribute to the formation of the virtual pitch in the ear;
- the failure of Terhardt's well-established model of virtual pitch to predict the pitch shifts observed suggests that further investigation into pitch formation mechanisms in the ear is desirable.

References

[1] Terhardt, E., Stoll, G. and Seewann, M. (1982) Pitch of Complex Systems according to Virtual Pitch Theory: Tests, Examples, and Predictions. *Journal of the Acoustical Society of America*, **71**, 671-678. https://doi.org/10.1121/1.387543

[2] Terhardt, E., Stoll, G. and Seewann, M. (1982) Algorithm for Extraction of Pitch and Pitch Salience from Complex Tone Signals. *Journal of the Acoustical Society of America*, **71**, 679-688. https://doi.org/10.1121/1.387544

[3] Software Implementation of the Terhardt Virtual Pitch Algorithm; the C Code Was Originally Available from Richard Parncutt's Website (https://homepage.uni-graz.at/richard.parncutt/), but Can Now Be Found (Together with a Matlab Implementation) at https://github.com/IntersectAustralia/hcsvlab-psysound/tree/master/AudioAnalysers/%40VirtualPitch/private

[4] Plack, C.J. and Oxenham, A.J. (2005) The Psychophysics of Pitch. In: Plack, C.J., Fay, R.R., Oxenham, A.J. and Popper, A.N., Eds., *Pitch*: *Neural Coding and Perception*; *Springer Handbook of Auditory Research*, Vol. 24, Springer, New York, 7-55. https://doi.org/10.1007/0-387-28958-5_2

[5] Hibbert, W.A. (2008) The Quantification of Strike Pitch and Pitch Shifts in Church Bells. PhD Thesis, The Open University, UK. http://www.hibberts.co.uk/phd/phd.htm

[6] Patterson, R.D. and Wightman, F.L. (1976) Residue Pitch as a Function of Component Spacing. *Journal of the Acoustical Society of America*, **59**, 1450-1459. https://doi.org/10.1121/1.381034

[7] Järveläinen, H., Verma, T. and Välimäki, V. (2002) Perception and Adjustment of Pitch in Inharmonic String Instrument Tones. *Journal of New Music Research*, **31**, 311-319. https://doi.org/10.1076/jnmr.31.4.311.14167

[8] Anderson, B.E. and Strong, W.J. (2005) The Effect of Inharmonic Partials on Pitch of Piano Tones. *Journal of the Acoustical Society of America*, **117**, 3268-3272. https://doi.org/10.1121/1.1882963

[9] Eggen, J.H. and Houtsma, A.J.M. (1986) The Pitch Perception of Bell Sounds. *Instituutvoor Perceptie Onderzoek*, *Annual Progress Report*, **21**, 15-23.

[10] Greenhough, M. (1976) The Analysis and Synthesis of Bell Sounds. *Proceedings of the Institute of Acoustics*, **3**, 4.

[11] Plomp, R. (1967) Pitch of Complex Tones. *Journal of the Acoustical Society of America*, **41**, 1526-1533. https://doi.org/10.1121/1.1910515

[12] Lin, J-Y. and Hartmann, W.M. (1998) The Pitch of a Mistuned Harmonic: Evidence for a Template Model. *Journal of the Acoustical Society of America*, **103**, 2608-2617. https://doi.org/10.1121/1.422781

[13] Terhardt, E. and Seewann, M. (1984) Auditive und Objektive Bestimmung der Schlagtonhöhe von Historichen Kirchenglocken. *Acustica*, **54**, 129-144. (English Translation Included in [2] and at http://www.hibberts.co.uk/terhardt.htm)

[14] Lehr, A. (1986) Partial Groups in the Bell Sound. *Journal of the Acoustical Society of America*, **79**, 2000-2011. https://doi.org/10.1121/1.393208

[15] Walpole, R.E., Myers, R.H., Myers, S.L. and Keying, Y. (2002) Probability and Statistics for Engineers and Scientists. 7th Edition, Prentice Hall, New Jersey.