Explaining the Railsback stretch in terms of the inharmonicity of piano tones and sensory dissonance

N. Giordano[a)]
Department of Physics, Auburn University, 246 Sciences Center Classroom, Auburn, Alabama 36849, USA

(Received 11 March 2015; revised 28 August 2015; accepted 2 September 2015; published online 23 October 2015)

The perceptual results of Plomp and Levelt [J. Acoust. Soc. Am. **38**, 548–560 (1965)] for the sensory dissonance of a pair of pure tones are used to estimate the dissonance of pairs of piano tones. By using the spectra of tones measured for a real piano, the effect of the inharmonicity of the tones is included. This leads to a prediction for how the tuning of this piano should deviate from an ideal equal tempered scale so as to give the smallest sensory dissonance and hence give the most pleasing tuning. The results agree with the well known "Railsback stretch," the average tuning curve produced by skilled piano technicians. The authors' analysis thus gives a quantitative explanation of the magnitude of the Railsback stretch in terms of the human perception of dissonance.
 [http://dx.doi.org/10.1121/1.4931439]

[TRM] Pages: 2359–2366

I. INTRODUCTION

It is well known that the notes of a well tuned piano do not follow an ideal equal tempered scale. Instead, the octaves are "stretched"; that is, the frequencies of the fundamental components of piano tones that would differ by precisely a factor of 2 in the ideal case are separated by a slightly greater amount. This stretched tuning was noted many years ago by Railsback.[1,2] He also studied the tuning curves (the deviation from an equal tempered tuning) of a collection of pianos tuned by expert technicians and found that all follow approximately the same stretched tuning. It is now widely accepted that this effect is caused by the inharmonicity of real piano strings. This inharmonicity is an important factor in piano design; it has thus been studied in great detail and there is an accurate theory of the inharmonicity of piano strings that has been verified through a number of quantitative studies (e.g., Refs. 3 and 4).

The fact that all well tuned pianos appear to follow a sort of universal tuning curve suggests that this tuning is related to a fundamental aspect of human perception, i.e., that the human perception of dissonance (which we will also refer to as "sensory dissonance" or simply "dissonance") is a key factor in determining how a piano is tuned. In some cases, the dissonance of two tones is easy to understand and predict. For example, for two pure tones the dissonance is minimized when the tones have the same frequency, i.e., when there is no beating. The notion that less beating results in lower dissonance also explains the perceived dissonance of two complex tones in some simple cases, such as complex tones that each contains only two components (which are called partials in the case of piano tones). As we will explain, such cases apply quite well to the treble region of the piano, so this simple notion about sensory dissonance explains the piano tuning curve in the treble, a fact that has been known for some time.[3]

However, piano tones in the bass region contain many partials with significant amplitudes. When considering the dissonance of two complex tones of this kind, it is not possible to achieve a "zero beat" condition for all pairs of partials. Instead, there will always be beating involving many partials, and one would expect minimum dissonance to be achieved when some (presumably complicated) function of the beat rates of many partials is minimized. This complicated function is a product of the human sensory system and we do not currently have a theory for this function based on fundamental physics or physiology. A complete understanding of the piano tuning curve that applies throughout the bass and treble thus requires a model or at least a description of sensory dissonance for complex tones.

Interestingly, such a description of sensory dissonance does exist[5] but until now it has not been applied to the problem of explaining the Railsback stretch. That is the purpose of this paper. We will see that our analysis, which combines what is known about sensory dissonance with what we know about the inharmonicity of real piano tones, leads quantitatively to the Railsback tuning curve. Our result also serves as a test of how the perceptual dissonance of pure tones can be used to determine the dissonance of complex tones that consist of many inharmonic components.[6,7]

II. BACKGROUND

Piano tones are produced by vibrating strings. The vibrations of an ideal, flexible string held rigidly at both ends give a complex tone whose components are perfectly harmonic. That is, if the frequency of the fundamental component is f_1^0, the frequency of the nth harmonic is precisely

$$f_n = n f_1^0, \tag{1}$$

[a)]Electronic mail: njg0003@auburn.edu

where n is an integer. The harmonic spectra produced by ideal strings and a number of other vibrating objects suggest why the octave plays an important role in musical scales. Two tones described by Eq. (1) and whose fundamental frequencies differ by a factor of exactly 2 have many overlapping harmonics. The notion that tones whose components overlap in this way are perceived as "pleasing," i.e., consonant, has been suggested many times over the past centuries.[8]

Real piano strings are not ideal; they have some stiffness, which causes deviations from the ideal harmonic spectrum in Eq. (1) (see, for example, Refs. 3 and 9). The components of real piano tones are therefore called partials, to distinguish them from ideal harmonics. The stiffness of real piano strings is small, and in that limit the partial frequencies are given to lowest order by[4,10]

$$f_n \approx n f_1 (1 + \alpha n^2), \tag{2}$$

where α is proportional to the Young's modulus of the string and the frequency f_1 of the first partial is shifted slightly from the value f_1^0 found for the ideal string. In practice there are corrections to Eq. (2) that depend on the boundary conditions of the string, the soundboard impedance, and other factors. However, these corrections do not change the fact that the spacing between neighboring partials is not constant but generally increases with n, as predicted by Eq. (2). Indeed, our analysis will not rely on or even use the functional form in Eq. (2), as we will make use of the values of f_n measured for the tones of a real piano. The key point is that the partials of two piano tones whose fundamental frequencies differ by precisely a factor of 2 (a perfect octave) will have few, if any, partials that overlap exactly. Indeed, the second partial of the lower tone will overlap with the fundamental component of the upper tone only if their fundamental frequencies are spaced by slightly more than a factor of 2, hence giving a stretched octave.

When a skilled technician tunes a piano he or she presumably has the goal of making it sound as "good" as possible. This judgment may be based simply on listening, or rely on what a commercial tuning device says is the desired tuning, or some combination, but it is reasonable to assume that this judgment is ultimately based on human perception. It is also reasonable to assume that this "best" tuning is designed to minimize the dissonance of the most important musical intervals, and it is widely believed that the process of minimizing the dissonance amounts to adjusting the octave spacing so that the partials of notes an octave apart overlap as closely as possible. However, it is not clear what the term "as closely as possible" means when one is considering tones with many partials as is the case for real piano tones, especially in the bass region of the scale.

It is useful at this point to consider the spectra of a few piano tones, which will help in understanding how pairs of partials can combine to increase or decrease the sensory dissonance, and why the problem is much simpler for notes in the treble as compared to the bass. Figure 1 shows the spectra for the notes A5 and A6 of a piano; in this notation, A4 has a fundamental frequency of 440 Hz. (The piano and

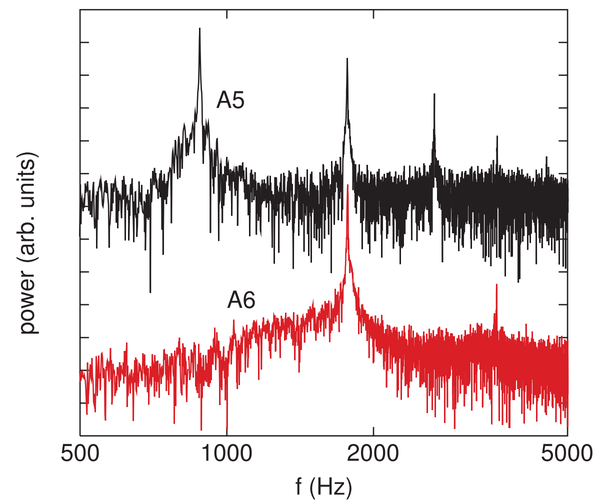

FIG. 1. Measured spectra of notes A5 and A6. Note that both scales are logarithmic and that the spectra have been displaced vertically for clarity.

details of the measurement will be described below.) These spectra illustrate a key property of notes in the treble; the intensities of the different partials fall off very rapidly with partial number. For example, the fourth partial of A5 is about 3 orders of magnitude smaller than the fundamental (i.e., the first partial), while for A6 the second partial is about 4 orders of magnitude smaller than its fundamental component. This strongly suggests that when considering possible beating between partials of A5 and A6, only the second partial of A5 and the fundamental of A6 need to be considered; the combined intensities of the fourth partial of A5 and the second partial of A6 are likely far too small for their beating to be discernible. The second partial of A5 and the fundamental of A6 will thus determine the sensory dissonance of these two notes.

The situation is completely different in the bass. Figure 2 shows spectra for notes A1 and A2, which show that the partials vary in intensity in a complex manner as the partial number increases. We must then expect the beating of many

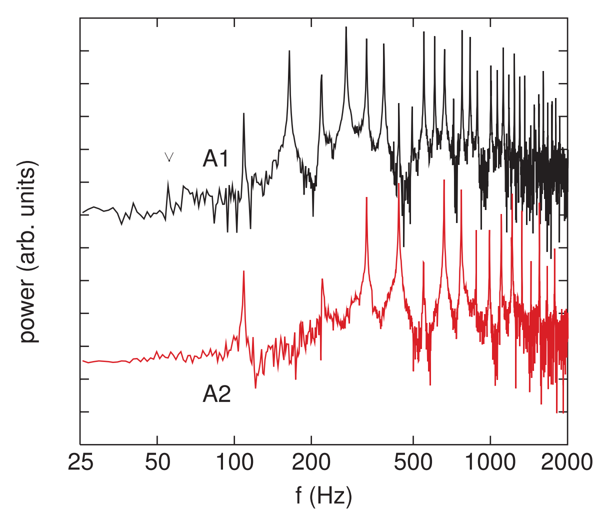

FIG. 2. Measured spectra of notes A1 and A2. Note that the vertical scale is logarithmic and that the spectra have been displaced vertically for clarity. The vertical arrow indicates the small peak at the fundamental frequency of A1.

pairs of partials to be important when considering the sensory dissonance of these notes and others in the bass.

These insights on contributions to sensory dissonance have been recognized for some time. For example, Schuck and Young[9] used measurements of the spectra of real piano tones, similar to those shown in Figs. 1 and 2 to explore how the overlap of partials produces a stretched tuning. They did not attempt to calculate dissonance quantitatively, but instead used only the observed frequencies of a single pair of partials (one partial from each note) to estimate the optimum stretch. For example, if one assumes that the lowest partials are most important one can estimate the octave stretch by requiring that the second partial of each note be equal to the fundamental frequency of the note an octave higher. Schuck and Young found that this assumption gives a good account of the stretch in the treble (as we have suggested based on the spectra in Fig. 1) but greatly underestimates the amount of stretch found in the bass. They also showed that matching pairs of higher partials of notes separated by either an octave or by two octaves leads to different predictions for the degree of stretch, but none of their matching algorithms, which were based solely on the frequencies of a single pair of partials and ignored their amplitudes, was able to account quantitatively for the tuning curves observed in practice. For example, matching the lowest partials of notes spaced by two octaves led to an underestimation of the octave stretch in the bass while overestimating it in the treble. Schuck and Young were thus not able to quantitatively account for the magnitude of the Railsback stretch in the bass.

More recently, there have been various analyses of stretched tuning, in which purely algorithmic approaches have been used to determine the best tuning.[11,12] Specifically, these approaches have used algorithms that minimize an effective entropy or similar constructs, in which these functions have no relation to what is known about human sensory dissonance.

Surprisingly, the author knows of no attempt to explain the Railsback stretch quantitatively in terms of what is known about the human perception of dissonance. The purpose of this paper is to take a step forward from the analysis of Schuck and Young in which established descriptions and models of human perception are used to evaluate the sensory dissonance of real piano tones and thereby predict the Railsback stretch. Our approach allows us to account for the many partials of real piano tones and their amplitudes, using the spectra of tones measured for a particular piano when estimating the sensory dissonance of two tones. This dissonance is estimated quantitatively using the results of the perceptual studies of the dissonance of pure tones by Plomp and Levelt as described in Sec. III.[5] The results are compared with the actual tuning of the piano under study and with the Railsback tuning curve. Overall good agreement is found. We also explore the sensitivity of this agreement to plausible alterations of the dissonance model.

III. MODEL OF TONAL DISSONANCE

This analysis described in this paper makes use of relations concerning perceptual dissonance that are based on the perceptual studies of Plomp and Levelt[5] (see also Kameoka and Kuriyagawa[13,14]). Plomp and Levelt used listening tests to determine how the dissonance of two pure tones with frequencies f_1 and f_2 depends on the difference in frequency of the tones. They showed that the perceived dissonance depends on both the frequency difference $f_2 - f_1$ and the lower of the two frequencies (f_1). Figure 3 shows a convenient parameterization of the Plomp and Levelt results due to Sethares,[6] the functional form of which will be given below. The dissonance is zero when the frequencies are equal, $f_2/f_1 = 1$. The dissonance exhibits a maximum when $f_2 - f_1$ equals a value Δf_0 that depends on the frequency of the lower tone [Fig. 3(b)]. The frequency difference Δf_0 is approximately equal to one-quarter of the critical band as identified in other studies (e.g., Ref. 15). The dissonance then decreases smoothly to zero as the frequency separation is increased beyond Δf_0.

The dissonance curves in Fig. 3 are noteworthy for their lack of structure. One might have expected to find dips in the dissonance (and hence maxima in the consonance) for musically pleasing intervals such as the octave ($f_2/f_1 = 2$), perfect fifth ($f_2/f_1 = 3/2$), and so on, but that is not the case. The

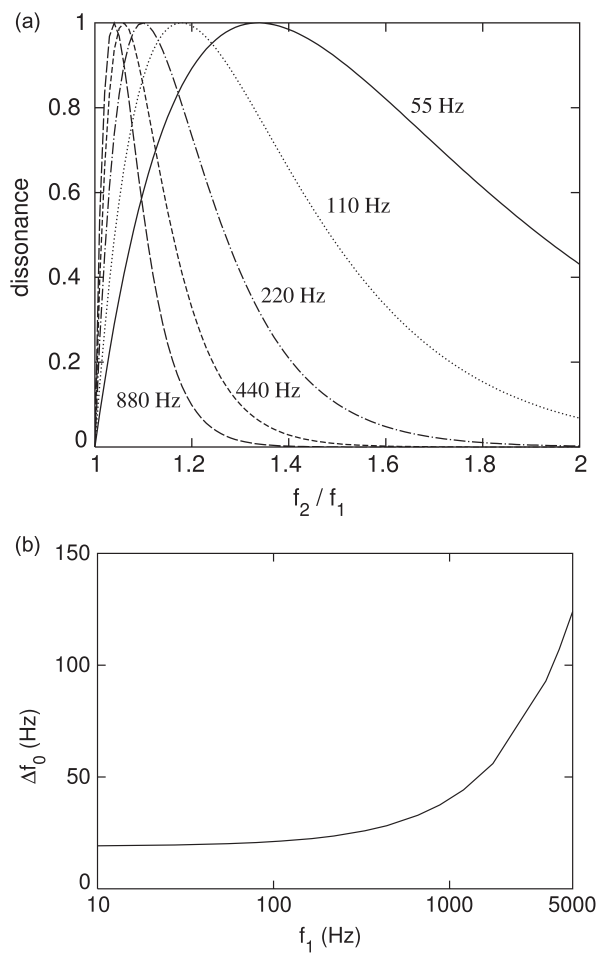

FIG. 3. (a) Perceived dissonance of two pure tones with frequencies f_1 and f_2 for different values of the frequency f_1 of the lower tone (given by the label on each curve). (b) Frequency Δf_0 at which the dissonance in a maximum as a function of the frequency of the lower tone f_1. These curves are computed using parameterizations of the results of Plomp and Levelt (Ref. 5) as given by Sethares (Ref. 6).

reduction of the dissonance of such intervals appears when one considers complex tones with harmonic spectra. Plomp and Levelt (and many others) have applied their result for pure tones to compute the dissonance of complex tones whose components form a harmonic series [Eq. (1)]. For such tones the total dissonance is assumed to be the sum of the pairwise dissonances of each of the components of both complex tones. One now has the possibility for the dissonance to be minimized by having some of the harmonics of the two tones overlap. For complex tones with harmonic components, such overlap does indeed produce minima in the dissonance for tones that form various musical intervals,[5] thus providing an "explanation" for the mathematical structure of the 12 tone musical scale. In this picture, the intervals which minimize the dissonance (i.e., maximize the consonance) depend on the spectra of the associated complex tones. Harmonic tones lead to the standard consonant intervals of an octave and perfect fifth, etc., but other types of complex tones can give very different results (see, e.g., Refs. 16 and 17).

In this paper we describe a similar analysis using the complex tones measured for a real piano. Our main goal is to determine if the inharmonicity of complex tones found in a real piano can be combined with a model of dissonance based on perceptual studies to quantitatively predict the stretched tuning actually found for that piano. Our analysis uses results for the partial frequencies and amplitudes measured for a real piano, and is not limited to the lowest one or two partials but explicitly includes the contributions to the dissonance from all significant partials.

IV. DISSONANCE OF COMPLEX TONES

Figure 3(a) shows the results of Plomp and Levelt[5] for the dissonance of two pure tones with frequencies f_1 and f_2. These curves are a parameterization derived by Sethares,[6] according to which the dissonance is proportional to

$$d_2(f_1,f_2) = (e^{-b_1 s|f_2-f_1|} - e^{-b_2 s|f_2-f_1|})\,. \tag{3}$$

Here

$$s = s^\star/[s_1 \min(f_1,f_2) + s_2], \tag{4}$$

$\min(x,y)$ is a function equal to the lesser of its two arguments, and $b_1 = 3.5$, $b_2 = 5.75$, $s^* = 0.24$, $s_1 = 0.021$, and $s_2 = 19$. These relations were derived from studies involving pure tones, and to apply them to real musical tones we must consider how the dissonances of two or more pairs of pure tones combine to give a total dissonance. This question is extremely important, since we have seen in Fig. 2 that piano tones can contain a large number of partials and for the bass notes the strongest partials may not be the ones with small n. Various authors have used the function $d_2()$ in Eq. (3) (or its equivalent) to compute the dissonance of complex tones. Plomp and Levelt applied it to complex tones consisting of six harmonics [with frequencies given by Eq. (1) with $n = 1, 2, \ldots, 6$] but did not specify the relative amplitudes of the harmonics and appeared to simply add the dissonance functions for each of the components of the two complex tones. Kameoka and Kuriyagawa[14] gave a rather extensive discussion of the issues involved in combining the dissonances of different dyads. In the end, they seem to have added the dissonance functions $d_2()$ for each pair of components weighted approximately by the product of the amplitudes of the two components in each dyad. Sethares[6] initially took a similar approach and assumed that the dissonance of two complex tones is given by

$$D_{\text{total}} = \frac{1}{2}\sum_{i=1}^{n_1}\sum_{j=1}^{n_2} B_{i,j} d_2(f_{1,i}, f_{2,j})\,, \tag{5}$$

with

$$B_{i,j} = a_{1,i} a_{2,j}\,. \tag{6}$$

Here the ith partial of tone 1 has frequency $f_{1,i}$ and amplitude $a_{1,i}$, and we assume that only n_1 partials have significant amplitudes, etc., for tone 2. (Note that here we omit the "self-dissonance" of a tone, which is included by some authors but is not important for our work since we will wish to find the condition in which the dissonance of two tones is a minimum. The self-dissonance may certainly be important when judging the quality of a single complex tone.[18]) We will term D_{total} with $B_{i,j}$ given by Eq. (6) the "amplitude product" model. Sethares[6] applied this model to complex tones consisting of seven harmonics with amplitudes that fall at a rate of 0.88 from one harmonic to the next, and showed, in agreement with the results of Plomp and Levelt, that this yields to consonances at the well-known and expected intervals, including octaves and perfect fifths.

In later work, Sethares[7] suggested that instead of Eq. (6) the dissonance should be proportional to the loudness of the weaker component so that

$$B_{i,j} = \min(\ell_{1,i}, \ell_{2,j})\,, \tag{7}$$

where $\ell_{n,i}$ is the loudness of the ith partial of tone n. In our applications of Eq. (7) in this paper we will assume that loudness is defined in the usual way [$\ell \sim 2^{\text{SPL}/10}$ where the sound pressure level (SPL) is related to the power by SPL $\sim \log(P)$].[7,10] We thus ignore the dependence of perceived loudness on frequency (the so-called Fletcher-Munson curves); we will consider this below in Sec. VI D. We will refer to D_{total} with $B_{i,j}$ given by Eq. (7) as the "minimum loudness" model. One argument in favor of a model based on loudness rather than partial amplitudes is that human hearing judges the loudness of a sound on a logarithmic scale, as accounted for in the definition of $\ell_{n,i}$. We will consider both the amplitude product and the minimum loudness model in our analysis. So far as we know, there have been no experiments that have studied which, if either, of these models is the appropriate way to measure the relative dissonance of pairs of pure tones.

V. METHOD

This analysis described in this paper makes use of a Steinway model M piano that belongs to the author and had recently been tuned by a professional technician. No special instructions were given to the technician. Recordings were made of every note A, C, E, and G, ranging from A0 to C8. These notes were chosen so as to provide multiple notes spread over each octave and give the option of studying the dissonance of common intervals and chords (a problem left for the future). Each note was recorded at an *mf* level, and all of the strings for each note were allowed to vibrate. Notes A0–E1 employed a single wound string, G1–C3 had two wound strings each, and E3–C7 employed three non-wound strings each. Each note was recorded using a PCB Piezotronics (Depew, NY) model 130P10 microphone with a sampling rate of 44.1 kHz. Spectra were calculated with a standard fast Fourier transform (FFT) algorithm using data for approximately 0.75 s taken just after the attack portion of the tone. The frequencies of each partial were obtained by fitting the spectral peaks to either Gaussian or sinc^2 lineshapes,[19] which were found to give essentially identical results for the partial frequencies. The frequencies derived in this way and for sound data of this duration were sufficiently precise for our purposes since the dissonance according to Eq. (3) varies slowly compared to the frequency resolution of these FFTs. More sophisticated signal processing methods were thus not required. For the same reason, the small difference in the frequencies of unison strings[20] was not important.[21]

This analysis yielded the frequencies and amplitudes of all partials with significant strength for each note, which were used in computing the dissonance functions of various pairs of notes. For most notes the spectra also yielded directly the fundamental frequency of each note and hence the Railsback stretch function for this piano as tuned by the technician. For the notes below G1 the fundamental component was difficult to discern from the noise; in those cases the fundamental frequency was estimated using the frequencies of the lowest partials and extrapolating f_n/n as $n \to 1$.

VI. RESULTS

We first consider the actual tuning curve of the piano under study; that is, the tuning as set by the professional tuner and presumably judged by him (according to his tuning device or simply by ear) to be the best and presumably most consonant tuning possible for our piano. We will find that it agrees well with the Railsback result. We then consider the measured spectra of individual tones and use the models of dissonance described in Sec. IV to calculate the most consonant tuning curve predicted by the models.

A. Tuning curve from the fundamental frequencies

The tuning curve for our piano as derived from the measured fundamental frequencies is shown by the solid symbols in Fig. 4. Here we plot the deviation Δf of the fundamental frequency f_1 of each note from an ideal equal tempered scale with A4 ($f_1 = 440$ Hz) as the reference value. This figure also shows the Railsback curve (solid line) and the result of Schuck and Young for their particular piano (dotted line). The Railsback curve is smooth, since it is the average for many pianos. The Schuck and Young result is not smooth, as it is the result for a single piano; similar variations from note to note have been reported by other workers (see especially the results for a number of pianos by Koenig[22]). The Schuck-Young result in Fig. 4 shows typical note-to-note fluctuations from the Railsback curve, to be expected for any particular piano. It is seen that the tuning curve for our piano has note-to-note variations from the Railsback curve that are similar to those found by Schuck and Young. We thus conclude that the results obtained by our professional tuner are quite similar to those obtained by many other tuners in a similar context. This is indeed a normal well-tuned piano.

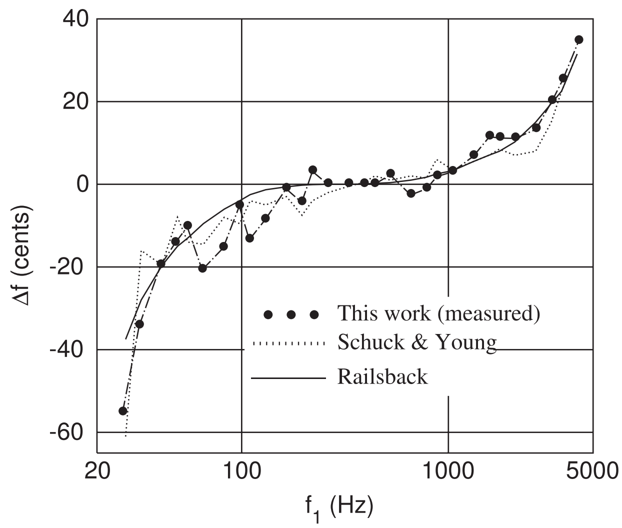

FIG. 4. Solid symbols: Deviation Δf of the fundamental frequencies of notes of the piano under study from an equal tempered scale as a function of the fundamental frequency f_1 for all notes A, C, E, and G. Solid curve: Railsback tuning curve. Dotted curve: Tuning curve for the piano studied by Schuck and Young (Ref. 9).

B. Understanding the general shape of the Railsback curve

One can gain a qualitative understanding of the shape of the Railsback curve from a few details of the spectra of piano tones. In the middle of the piano range the inharmonicity is relatively small, as can be seen from either the spectra of notes in this region or the theoretical value of α in Eq. (2). As one moves toward higher notes in the treble the inharmonicity increases, increasing the amount of stretch per octave and causing the upward curvature of the Railsback curve in that region. In the bass the inharmonicity parameter α is smaller than found in the upper treble but the number of significant partials increases as one moves toward lower notes. Indeed, for the lowest notes 20 or more partials may be significant (see Fig. 2 and below) while in the extreme treble only two might have been important (Fig. 1). Since the inharmonicity increases with partial number n, the amount of stretch per octave is largest in the bass causing the large downward curvature of the Railsback curve in that region. These qualitative observations will be confirmed in the quantitative analysis presented in Secs. VI C and VI D.

C. Finding the least dissonant tuning curve

To apply a model of dissonance to calculate the optimum tuning curve we select a pair of notes that for perfect tuning are separated by precisely an octave, for example, A0 and A1. If the fundamental frequencies of the notes are f_1(A0) and f_1(A1), then in the ideal case $f_1(A1)/f_1(A0) = 2$. We then consider shifting the pitch of the upper note by shifting its fundamental frequency by a small amount df_1. At the same time, we shift the frequencies of each partial of this note by $n \times df_1$ where n is the partial number.[23–26] Since the amplitudes of the partials are known from measurements of the spectrum of each note, the dissonance function can be evaluated as a function of df_1. The value of df_1 that minimizes the dissonance function then gives a prediction for the amount Δf the octave should be stretched for optimum tuning. As a final point, in all of our calculations we assume that the two notes have equal total power. This is accomplished by normalizing for each note the sum of the power contained in each of the partials for that note.

Before applying the models of dissonance discussed in Sec. IV, we first consider one of the approaches taken by Schuck and Young, in which they computed the stretch by considering notes separated by an octave and requiring that the frequency of the second partial of the lower member of the octave coincide with the fundamental frequency of the higher member. This is a very simple dissonance function that ignores the amplitudes of different partials; we will refer to this as the "2/1 model." This approach yields the predicted stretch of each octave, from which one can deduce the tuning curve by adding the stretch predicted for each octave as one moves away from the notes in the octave near A4. The results for the 2/1 model for our piano are shown in Fig. 5 where it is compared to the measured tuning curve. Note that in applying the 2/1 model we have assumed that the notes C4, E4, G4, and A4 all have $\Delta f = 0$ and hence all lie on the horizontal axis (which is consistent with the results in Fig. 4). The stretch values for other octaves are then obtained by adding the stretch values octave by octave. The results of the 2/1 model applied to our piano are consistent with those found by Schuck and Young. Namely, for notes in the middle range to the extreme treble the 2/1 model is consistent with both the Railsback curve (Fig. 4) and with the observed tuning. However, there are substantial deviations in the bass, as the model predicts a much smaller stretch than gives the best tuning.

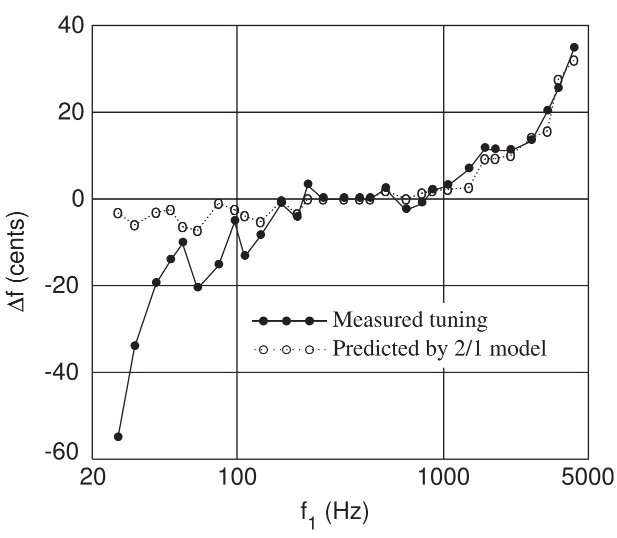

FIG. 5. Open symbols and dotted curve: Predictions of the 2/1 model for our piano. Solid symbols and solid line: Actual (measured) tuning of the piano.

The reasons for the agreement and disagreement with the 2/1 model of dissonance are important for critically understanding the predictions of other dissonance models. As we saw in Fig. 1, notes in the extreme treble have only a few significant partials, and their amplitudes decrease very quickly with partial number n. The dissonance is therefore dominated by the lowest partial(s) of the lower member of each octave and the 2/1 model, which accounts for only the very lowest partials, gives an acceptable result. In this range we would expect virtually *any* plausible model of dissonance to give a result identical to that of the 2/1 model and thus account for the observed tuning curve. However, as one moves down through the midrange and into the bass, the notes have many partials with significant amplitudes. For example, for the note A1 of our piano one must go to partials above $n = 27$ to reach the point at which all higher partials have amplitudes less than 3% of that of the strongest partial. Moreover, the strongest partials for this note are of order $n = 7$ and $n = 14$, so an accurate description of dissonance must include at least these strong partials. For these reasons, the 2/1 model is expected to fail in the bass, as is indeed found.

We have discussed the 2/1 model here to illustrate two crucial points. First, any reasonable model of dissonance will almost certainly work in the treble. Second, the bass notes provide the real test of models of dissonance.

We next consider the amplitude product and minimum loudness dissonance models. We have computed the tuning curves with these models following the procedure of considering notes separated by an octave and calculating the dissonance as the frequencies of the partials of one of the notes is shifted, as was done in applying the 2/1 model. The results are shown in Fig. 6. For comparison this figure also shows the measured tuning curve. Above a fundamental frequency of about $f_1 \approx 110$ Hz, corresponding to note A2, the two models agree well with each other and with the measured stretch. The fact that the two models agree in the treble is expected since, as already discussed, the notes in the treble have only a few significant partials and the amplitude falls rapidly with partial number. For the lowest notes the amplitude product model seems to be in slightly better agreement with the actual tuning curve. This is perhaps suggestive evidence in favor of the amplitude product model, but in our opinion both models provide an acceptable quantitative account of the tuning curve.

Application of the amplitude product and minimum loudness models confirms our earlier claims about the need to include a large number of partials to account for the dissonance of two tones in the bass. For example, when computing the dissonance of the note pairs A0–A1, and A1–A2, it is necessary to include at least 16 partials of the lower note and 8 of the higher member of the pair to reach the asymptotic result, i.e., the stretch found by including all significant partials. Including fewer partials gives a significantly smaller predicted stretch. By comparison, for the note pair A2–A3 the

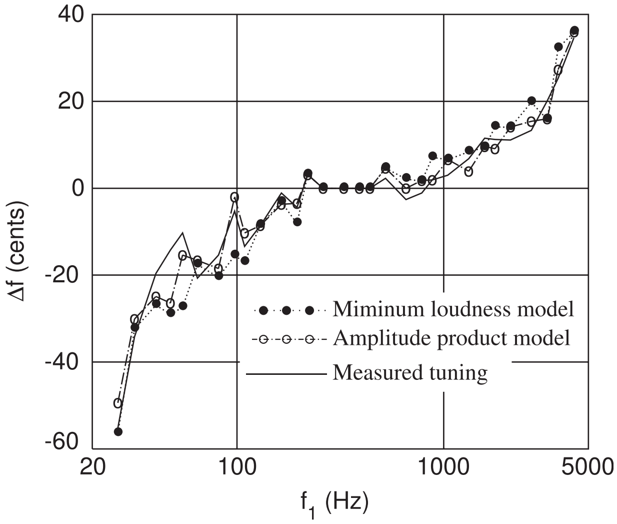

FIG. 6. Solid symbols and dotted curve: Tuning curve predicted by the minimum loudness model. Open symbols and dotted-dashed curve: Tuning curve predicted by the amplitude product model. Solid curve: Actual (measured) tuning of the piano.

asymptotic value for the stretch is reached with only six partials from the lower note and three from the upper one.

It is also interesting to consider why the two models give similar results. One might have thought that the very different factors $B_{i,j}$ in Eqs. (6) and (7) would give rather different results, especially since the loudness of a pure tone depends logarithmically on its amplitude. The answer to this question seems to contain two parts. First, when comparing two piano tones an octave apart with equal total powers, the pairs of partials that are close in frequency tend to have similar amplitudes. Hence if one member of such a pair of partials is small then both are usually small, so the amplitude product model and minimum loudness model tend to give most weight to the same sets of partials. Second, one might wonder why the logarithmic dependence in Eq. (7) gives a result similar to Eq. (6). The reason is not entirely clear to the author, but may be due to that fact that the loudness of a pure tone is approximately proportional to $p^{0.3}$ where p is the power in that tone. This is not far (arithmetically) from the relation between amplitude and power.

D. Consideration of other models

We have seen that the amplitude product model and the minimum loudness model both give a good quantitative account of the Railsback tuning curve. In this section we consider how sensitive this result is to certain aspects of the models.

One obvious extension of the dissonance models pointed out by Sethares[7] would be to include the dependence of perceived loudness on frequency and amplitude as captured by the so-called Fletcher-Munson curves.[10,25] We have performed a full calculation with the minimum loudness model in which the Fletcher-Munson results are included when computing the effective loudness. This had no significant effect on the results. The reason for this can be traced to the fact that the lowest partials of the bass notes all have small values of ℓ so suppressing their contributions to the dissonance further (by applying the Fletcher-Munson correction) has a negligible effect on the predicted dissonance.

We have also considered the effect of changes to the Plomp-Levelt dissonance relation, which is the basis for Eq. (6). There are several parameters that determine the precise shape of this function; one of the most important is the frequency difference at which the dissonance is a maximum, Δf_0, which is proportional to s in Eq. (4). We have therefore considered the tuning curves predicted for different values of s given by

$$s = \beta \times s^{\star}/[s_1 \min(f_1, f_2) + s_2]\,. \tag{8}$$

The value of the factor β in Eq. (8) determines how much the dissonance maximum Δf_0 is shifted from its value in Fig. 3.

We have calculated the tuning curves that minimize the dissonance for values of β in the range 0.1 to 10, corresponding to either decreasing or increasing Δf_0 by a factor of 10. (Recall that Δf_0 is proportional to a physiological parameter, the critical band.) We found that increasing Δf_0 by as much as a factor of 10 produced very little change in the predicted tuning. However, decreasing Δf_0 made a significant difference, as shown in Fig. 7 which gives results for $\beta = 0.5$ and 0.1.

With $\beta = 0.5$ the tuning curve in the extreme bass is displaced slightly below both the result calculated with $\beta = 1.0$ and the measured tuning curve. With $\beta = 0.1$ the predicted tuning curve is displaced much farther and for the lowest notes there is even no dissonance minimum for any reasonable value of the stretch. This result thus gives an upper limit on acceptable values of Δf_0; decreasing this parameter by more than about a factor of 2 leads to predicted tuning curves which disagree with the Railsback result. This provides a test of one aspect of the dissonance function.

VII. SUMMARY AND CONCLUSIONS

The results of perceptual studies of the dissonance of pure tone by Plomp and Levelt together with the measured

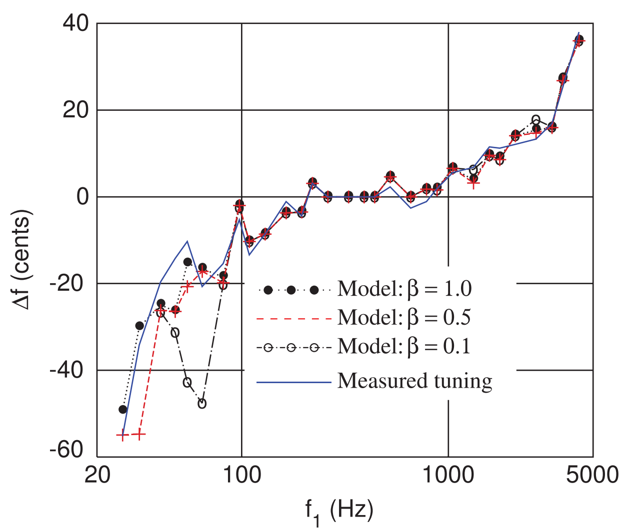

FIG. 7. Solid curve: Actual tuning of our piano. The symbols give the tuning curves predicted by the amplitude product model with the parameter s given by Eq. (8) with $\beta = 1.0$ (filled circles), $\beta = 0.5$ (filled squares), and $\beta = 0.1$ (open circles).

inharmonic spectra of real piano tones have been used to predict the optimum tuning of a piano. This prediction agrees well with the way pianos are found to be tuned in practice, the Railsback stretch. This result is further confirmation that the inharmonicity of piano tones is indeed responsible for the stretched tuning of a piano. This is not a surprising result; this inharmonicity has long been assumed to be responsible for the Railsback stretch. However, to our knowledge this is the first *quantitative* calculation that uses what is known about the human perception of sensory dissonance to correctly predict the magnitude of the stretch. Since our analysis relies on a model of dissonance, the agreement between the predicted and observed stretch serves as a check on that model.

The model of dissonance we have used is based on the perceptual studies by Plomp and Levelt using pure tones. The extension of their results to complex tones requires an assumption as to how the dissonance of two pure tones depends on their amplitudes. Our analysis has considered two proposals of Sethares, one based on the product of the amplitudes and another based on loudness of the weakest tone, and both are consistent with the Railsback result. Further work with tones of greatly differing amplitudes will be required to distinguish which of these models gives the best description of sensory dissonance.

ACKNOWLEDGMENTS

The author thanks David E. Ferrell for some very illuminating discussions and critical comments.

[1]O. L. Railsback, "Scale temperament as applied to piano tuning," J. Acoust. Soc. Am. **9**, 274 (1938).
[2]O. L. Railsback, "A study of the tuning of pianos," J. Acoust. Soc. Am. **10**, 86 (1938).
[3]R. W. Young, "Inharmonicity of plain wire piano strings," J. Acoust. Soc. Am. **24**, 267–273 (1952).
[4]N. Fletcher, "Normal vibration frequencies of a stiff piano string," J. Acoust. Soc. Am. **36**, 203–209 (1964).
[5]R. Plomp and W. J. M. Levelt, "Tonal consonance and critical bandwidth," J. Acoust. Soc. Am. **38**, 548–560 (1965).
[6]W. A. Sethares, "Local consonance and the relationship between timbre and scale," J. Acoust. Soc. Am. **94**, 1218–1228 (1993).
[7]W. A. Sethares, *Tuning, Timbre, Spectrum, Scale* (Springer, London, 2005).
[8]S. Dostrovsky, "Early vibration theory: Physics and music in the seventeenth century," Arch. Hist. Exact Sci. **14**, 169–218 (1975).
[9]O. H. Schuck and R. W. Young, "Observations on the vibrations of piano strings," J. Acoust. Soc. Am. **15**, 1–11 (1943).
[10]N. H. Fletcher and T. D. Rossing, *The Physics of Musical Instruments* (Springer-Verlag, New York, 1991), Chaps. 2 and 12.
[11]F. Rigaud, B. David, and L. Daudet, "A parametric model and estimation techniques for the inharmonicity and tuning of the piano," J. Acoust. Soc. Am. **133**, 3107–3118 (2013).
[12]H. Hinrichsen, "Entropy-based tuning of musical instruments," Rev. Bras. Ens. Fís. **34**, 2301 (2012).
[13]A. Kameoka and M. Kuriyagawa, "Consonance theory part I: Consonance of dyads," J. Acoust. Soc. Am. **45**, 1451–1459 (1969).
[14]A. Kameoka and M. Kuriyagawa, "Consonance theory part II: Consonance of complex tones and its calculation method," J. Acoust. Soc. Am. **45**, 1460–1469 (1969).
[15]E. Zwicker, B. Flottorp, and S. S. Stevens, "Critical band width in loudness summation," J. Acoust. Soc. Am. **29**, 548–557 (1957).
[16]J. R. Pierce, "Attaining consonance in scales," J. Acoust. Soc. Am. **40**, 249 (1966).
[17]W. A. Sethares, "Adaptive tunings for musical scales," J. Acoust. Soc. Am. **96**, 10–18 (1994).
[18]D. E. Farrell, private communication (2015).
[19]R. N. Bracewell, *The Fourier Transformation and its Applications*, 2nd ed. (McGraw-Hill, New York, 1978), Chaps. 2 and 4.
[20]R. E. Kirk, "Tuning preferences for piano unison groups," J. Acoust. Soc. Am. **31**, 1644–1648 (1959).
[21]The reason for this seems to be that for a well-tuned piano, the difference in frequency between unison strings is much smaller than frequency differences associated with the stretch tuning.
[22]D. M. Koenig, *Spectral Analysis of Musical Sounds with Emphasis on the Piano* (Oxford University Press, Oxford, 2015), Chap. 12.
[23]This approach for shifting the frequency of a partial is valid for all normal partials. It will not apply to longitudinal modes or so-called phantom partials (Refs. 26 and 27) which are known to be present in bass tones. Careful examination of the spectra, as in Fig. 2, allowed us to distinguish the longitudinal modes and phantom partials, and we found that they were all sufficiently weak and small enough in number that they had no effect on the calculated dissonance.
[24]This approach for shifting the partial frequencies should be accurate to order α in Eq. (2). Since α was typically of order 5×10^{-5}, this approximation should be adequate.
[25]H. Fletcher and W. A. Munson, "Loudness, its definition, measurement and calculation," J. Acoust. Soc. Am. **5**, 82–108 (1933).
[26]J. H. A. Conklin, "Design and tone in the mechanoacoustic piano. Part III. Piano strings and scale design," J. Acoust. Soc. Am. **100**, 1286–1298 (1996).
[27]J. H. A. Conklin, "Piano strings and 'phantom' partials," J. Acoust. Soc. Am. **102**, 659 (1997).