Perceptual studies of violin body damping and vibrato

Claudia Fritz, Jim Woodhouse, Felicia P-H. Cheng, Ian Cross, Alan F. Blackwell, Brian C. J. Moore

► To cite this version:

Claudia Fritz, Jim Woodhouse, Felicia P-H. Cheng, Ian Cross, Alan F. Blackwell, et al.. Perceptual studies of violin body damping and vibrato. *Journal of the Acoustical Society of America*, 2010, 127 (1), pp.513-524. hal-00468019

HAL Id: hal-00468019

<https://hal.science/hal-00468019>

Submitted on 30 Mar 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Perceptual studies of violin body damping and vibrato

Claudia Fritz

Université Pierre et Marie Curie, UMR CNRS 7190, Institut Jean Le Rond d'Alembert, 4 place Jussieu, 75005 Paris, France^{a)}

Jim Woodhouse and Felicia P-H. Cheng

Dept. of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK

Ian Cross

Centre for Music and Science, Music Faculty, University of Cambridge, West Road, Cambridge CB3 9DP, UK

Alan F. Blackwell

Computer Laboratory, University of Cambridge, 15 JJ Thomson Avenue, Cambridge CB3 0FD, UK

Brian C. J. Moore

Dept. of Experimental Psychology, University of Cambridge, Downing Street, Cambridge CB2 3EB, UK

(Dated: October 2, 2009)

This work explored how the perception of violin notes is influenced by the magnitude of the applied vibrato and by the level of damping of the violin resonance modes. Damping influences the “peakiness” of the frequency response, and vibrato interacts with this peakiness by producing fluctuations in spectral content as well as in frequency and amplitude. Initially, it was shown that thresholds for detecting a change in vibrato amplitude were independent of body damping, and thresholds for detecting a change in body damping were independent of vibrato amplitude. A study of perceptual similarity using triadic comparison showed that vibrato amplitude and damping were largely perceived as independent dimensions. A series of listening tests was conducted employing synthesized, recorded or live performance to probe perceptual responses in terms of “liveliness” and preference. The results do not support the conclusion that “liveliness” results from the combination of the use of vibrato and a “peaky” violin response. Judgments based on listening to single notes showed inconsistent patterns for liveliness, while preferences were highest for damping that was slightly less than for a reference (real) violin. In contrast, judgments by players based on many notes showed preference for damping close to the reference value.

PACS numbers: 43.75.De, 43.66.Lj, 43.75.Cd, 43.66.Jh

Keywords: violin, vibrato, damping, perception

I. INTRODUCTION

Vibrato is the intentional modulation of the frequency and amplitude of a musical tone. It is an attribute of many sounds in the contemporary western musical world, and is generally linked to what might be termed “expressive performance” (Brown, 1988). It is employed in a wide range of musical contexts as one of the resources in the palette of expressive strategies available to a performer. It is typically applied at rates around 6 cyclic fluctuations per second, in line with the rates found for other expressive ornaments such as tremolos and trills (Moelants, 2004). Typical frequency excursions in vibrato are about 2% of the fundamental frequency, but they can range up to 4% (Prame, 1994, 1997). These frequency excursions have been shown to affect the rapidity with which listeners can make judgments of relative pitch (Yoo *et al.*, 1998), but they have a minimal effect on perceived pitch (van Besouw *et al.*, 1996). Vibrato has become an integral constituent of contemporary tech-

nique on several string and wind instruments (Moens-Haenen, 2009), and many performers employ it as part of the process of imparting desirable timbral qualities to the musical sounds.

Vibrato may help to define the timbre of a musical instrument, since the frequency modulation of the components causes them to move relative to the resonant frequencies that are characteristic of the instrument or voice, defining the center frequencies of the resonances more precisely (McAdams, 1989). McAdams and Gior-dano (2009) note that “vibrato may increase our ability to extract information relative to the resonance structure of the instrument”. The question of whether or not the use of vibrato has consistent effects that relate to the perceived timbre of musical sounds remains open.

For the particular case of the violin, Fletcher and Sanders (1967) were probably the first to suggest that fluctuations of spectral content due to vibrato are important for the perception of timbre. The frequencies of some harmonics will fall on positive slopes of the frequency response curve of the violin body, while others fall on negative slopes. The simple frequency modulation from the player’s finger movement is then converted, by the resonant body response “filter” of the instrument,

^{a)}Electronic address: claudia.fritz@upmc.fr

into a complex spectral modulation in which the amplitudes of different harmonics change in different ways.

A few years later Mathews and Kohut (1973) and Gorrill (1975) pioneered the methodology that will be exploited in this paper. Mathews and Kohut studied the effect of damping of the violin body resonances by creating violin sounds (a G major scale, performed with vibrato) using a near-silent electric violin and a set of analog filters: twenty resonant circuits were tuned to major resonance frequencies measured from a Stradivarius violin. Four values of damping were chosen, but the authors do not give clear details: one corresponded to a completely flat response (infinite damping), the second to a 10-dB peak-to-valley ratio in the response curve, and the remaining two corresponded to bigger ratios. They used their filtered stimuli to conduct informal listening tests using preference judgments. Their second value of damping, corresponding to a 10-dB peak-to-valley ratio, was preferred by their jury. They also reported that when the damping was too low the sound took on a “hollow” quality, which was found to be unpleasant.

McIntyre and Woodhouse (1974) suggested that this “hollow” quality might arise from transient excitation of the body resonances. The response to vibrato in the bowed-string input waveform does not consist simply of harmonics which track up and down the slopes of the response curve, as they would if the vibrato rate were extremely slow, because the typical rate of vibrato has a similar timescale to the free decay times of the body resonances. They found that simulations involving high-Q resonances but in which this transient effect was artificially omitted did not sound “hollow”, whereas with the transient effect the hollowness was heard.

More recent work has examined the extent to which the perceived qualities of vibrato tones derive from modulations of frequency or amplitude. Mellody and Wakefield (2000) conducted analyses of real violin sounds produced with vibrato, and resynthesized approximations to those tones, either co-varying both frequency and amplitude, as in the original sounds, or varying these parameters independently. They found that the absence of frequency modulation had little effect on perceptual judgments, while the absence of amplitude modulation resulted in large perceptual changes.

Alongside this scientific perspective there is strong anecdotal evidence based on the intuitions of violinists about their strategies in using vibrato and the timbral effects they expect it to achieve. The response of a violin to the use of vibrato is widely considered to be one of the key factors in the performer’s perception of the “responsiveness” of a particular instrument (Gough, 2005). Mathews and Kohut (1973), in the study described earlier, reported that when the frequency response was flat, the instrument seemed “unresponsive.” As a starting point for this study, we hypothesized that the desirable quality of “liveliness” or “responsiveness” in a violin may be connected with the interaction of vibrato with the “peaky” frequency response of the violin.

In another study by the present authors (Fritz *et al.*, 2008), experienced violinists were asked to arrange a collection of 61 words that were found to be commonly used

to describe violin timbre on a two-dimensional grid, so that words which were similar in meaning were close together and words with very different meanings were far apart. The results were analyzed using multidimensional scaling and led to a three-dimensional map, which showed that the word “lively” was considered similar to “alive,” “resonant,” “ringing” and “responsive,” and as opposite to “dead” and “dull”. Therefore, lively and responsive will be used interchangeably in this study.

The present study used a method analogous to that of Mathews and Kohut (1973) to explore the perceived quality and discriminability of violin timbre when varying vibrato amplitude and the resonance damping of the violin within the same experimental framework, and to probe the perceptual consequences in a series of studies. The authors were surprised by the results from every stage of the study. A succession of tests was designed, in each case with fairly clear expectations based on the intuitions of players, and those expectations were regularly not supported. This may indicate that the anecdotal evidence is wrong, but at least in some cases we feel that it points more strongly to the difficulty of designing tests which are focused enough to be quantitatively convincing without throwing out the “baby” of musical relevance with the “bathwater” of experimental control.

This is an important issue in its own right, underlying any study of musical psychoacoustics. The phrase “musical relevance” here covers two main aspects. First, very short sound samples tend to lose any musical quality to the listener, and repeated listening to similar sounds in a typical test erodes it still further. Second, there is the question of the realism or naturalness of the sounds used: if the test sounds are not close enough to the “training set” that a musical listener will have experienced from hearing real violin performances, it may be that finely-honed perceptions will not be able to operate in the way that is intended. Indeed, if some perceptions are of a categorical nature they may not be evoked at all by sounds of insufficient naturalness.

II. GENERAL SOUND SYNTHESIS METHOD

The methodology is based on creating “virtual violins”, as described in a previous study (Fritz *et al.*, 2007). The frequency response function of the violin is mimicked using a digital filter, and the output signal for listening tests is generated by applying this filter to an input signal representing the force exerted by the bowed string on the bridge of the violin. The main advantage of this approach lies in the fact that, once the violin response is represented in digital filter form, it becomes easy to make controlled variations of a kind which would be impossible to achieve by physical changes to a violin.

The input signal can be generated in three different ways, all used in different parts of this study. First, the bridge force may be recorded from a player using vibrato, on a violin whose bridge has been instrumented with piezoelectric force sensors. The same recording can then be used with many different digital filters, thus removing the influence of the player and their adaptation

to the instrument. Second, the bridge-force signal from an instrumented violin can be used in real time, passing the force signal to a digital filter system that generates the sounds the player hears. This approach works best if the violin body is essentially mute so that the only sound reaching the player's ears comes from the filtered signal. Third, for some tests the bridge force signal can be synthesized.

In the case of purely synthesized force signals, the approach exploits the fact that the usual Helmholtz motion of a bowed string produces a bridge force in the form of a sawtooth waveform, at least to a first approximation (e.g. Cremer, 1985). This waveform is easy to synthesize. To take account of the frequency modulation associated with vibrato, each successive period requires a slightly different length. The following formula is used to determine the k th period length:

$$T_k = T_0(1 + a_{vib} \sin(2\pi k f_v T_0) + r_k) \quad (1)$$

where T_0 is the period corresponding to the nominal frequency of the note, f_v is the vibrato frequency in Hz (typically 5 Hz), a_{vib} is the vibrato amplitude (typically 0.02) and r_k is a random number drawn from a uniform distribution between $-R$ and $+R$, where R determines the amount of randomness. In what follows, a_{vib} and R are both multiplied by 100 to express them as percentages. The random element was included in an attempt to increase the naturalness of the synthesized sounds. In the course of the series of experiments reported here, several further small changes were tried in the quest to improve naturalness, as will be described below.

A window function was applied to the entire synthesized input signal to give a smooth envelope resembling a détaché legato bowed note. For the first test, described in section III.A, this was a Hanning window, but in later tests this choice was changed in an effort to improve the naturalness: following Gough (2005) the first 600 ms of the signal was multiplied by $(1 - \exp(-t/T))$ and the last 100 ms by $\exp(-(t - 600)/T)$ with $T = 30$ ms. In these later tests, the whole signal was further processed to round off the ‘‘Helmholtz corners’’ of the ideal sawtooth waveform. This corner rounding was achieved by convolution with a Gaussian function whose width was initially chosen to be 3.5% of the period, a typical value from measured waveforms (see for example McIntyre *et al.*, 1981).

Filtering to represent the violin body was based on the bridge admittance frequency response of a good-quality modern violin made by David Rubio in 1992. The admittance was measured using a small impulse hammer and a laser vibrometer (see Fritz *et al.*, 2007, for details). The amplitude of the measured admittance is plotted as the solid line in Fig. 1. The dash-dot line in this figure shows a typical modification made to the response for the tests to be described: the modal damping factors have all been doubled. For comparison, the dashed line shows an approximation to one frequency response used by Mathews and Kohut (1973), the one with 10-dB peak-to-valley fluctuation, as preferred by their listening jury.

In order to make changes such as the damping modification shown in Fig. 1, the measured frequency response

FIG. 1. The input admittance of the original Rubio violin (solid line), a modified version of it (dash-dot line), where all the Q values have been divided by 2, and a resynthesized version of the frequency response function of Mathews and Kohut (1973, Fig. 4b) (dashed line). The original curve is shown as a calibrated measurement (re 1 m/s/N), and the other curves have been shifted down for clarity.

was first analyzed into modal contributions using standard pole-residue fitting procedures (e.g. Ewins, 2000), and then resynthesized from these parameters (with or without some modification being made first). The fitting procedure covered the frequency range up to 7000 Hz, and required 54 modes. The resynthesized response was used to construct a causal finite impulse response filter which was then applied to the chosen input signal. Damping is quantified throughout this work by the modal Q values: the Q value is the inverse of the damping factor, so high Q corresponds to low damping, and vice versa. The values of the fitted Q values for the Rubio violin are of some interest, and are plotted in Fig. 2. The estimated values for the study of Mathews and Kohut (1973) are also shown.

III. THRESHOLD AND INTERACTION STUDIES

A. Experiment 1: Discrimination of vibrato amplitude

In a preliminary study (Cheng, 2006), thresholds were measured for detecting a change in vibrato amplitude (a_{vib}) using reference sounds with or without vibrato ($a_{vib} = 0$). The sounds used were either the raw synthetic string signals or the same string signals filtered by the acoustical response of the Rubio violin, as described above. This gave a total of four different conditions.

1. Stimuli

All tests were based on the note G3 (fundamental frequency 196 Hz, the lowest note of the violin). String signals were synthesized using the method described above, using a sampling rate of 44100 Hz and 16-bit resolution.

FIG. 2. Comparison of the Q values used here (deduced by modal fitting of the response of the Rubio violin; solid line) and those used by Mathews and Kohut(1973). These correspond to a 10-dB peak-to-valley ratio (o), and two higher ratios ($\times$ and $+$)

The vibrato rate (f_v) was 5 Hz and the randomness in amplitude R was 1%. The stimuli were 2 s in duration. They were presented diotically in a relatively quiet environment, via Sennheiser HD580 headphones, chosen because of their diffuse-field response and low distortion.

2. Procedure

Thresholds were estimated using a three-alternative forced-choice (3AFC) procedure. Three sounds, two the same (the reference violin sound with a fixed vibrato amplitude of either 0 or 2%) and one different (with a greater vibrato amplitude), were played in a random order, and the participant was asked to choose which was different. The amount of modification was increased after a single incorrect response and decreased after three successive correct responses. The step size of these changes was initially a factor of 1.414 (relatively large, for fast convergence toward the threshold region). After two turnpoints (changes from decreasing to increasing vibrato amplitude of the test sound and vice versa), the step size was reduced to a factor of 1.189. Eight turnpoints were obtained and threshold was taken as the geometric mean of the values of the amount of modification at the last six turnpoints. Participants were given visual feedback during the experiment and were given some practice by performing the test twice. The thresholds shown here are those obtained for the second run.

3. Participants

The participants in this study were four experienced string players and seven other musicians (UK Grade 8). All subjects reported having normal hearing, although this was not checked. The two tests based on a reference sound with no vibrato used all 11 participants, while the

two involving a reference sound with 2% vibrato amplitude used the four string players and five of the other musicians.

4. Results

The results did not differ for the two types of musicians, and so the results were averaged over all participants. When the reference sound had no vibrato, the threshold for detecting the vibrato was 0.5% for the raw string signal and 0.8% for the ‘Rubio’-filtered signal. When the reference sound had a vibrato amplitude of 2%, the thresholds for detection of a change were 2.9% and 3.2%, for the unfiltered and filtered cases, respectively. Contrary to our expectations based on earlier findings (Mathews and Kohut, 1973; Meyer, 1992; Gough, 2005) thresholds were not lower for the filtered than for the unfiltered signals. Thus, the fluctuations in amplitude of individual harmonics produced by passing the vibrato string signal through the synthesized violin body did not lead to enhanced detection or discrimination of the vibrato. However, listeners may still be sensitive to fluctuations in amplitude induced by vibrato. Indeed, as described earlier, such fluctuations might be used to infer properties of violins, such as the degree of damping. To explore this, it was decided to carry out the inverse test to that described above: the threshold was measured for detecting a change of the Q values of the violin response, using an input signal with no vibrato and with two fixed amplitudes of vibrato.

B. Experiment 2: Effect of vibrato amplitude on the discrimination of damping

1. Stimuli

The stimuli (again the bottom violin note G3) were synthesized as before. Some details were changed from those used in experiment 1 in an attempt to improve the naturalness of the synthesized sound. The randomness in amplitude was decreased to 0.2% and the vibrato rate was increased to 6 Hz. The whole signal was “corner-rounded” as described in section II. The duration of the sounds was shortened to 700 ms to enhance the effectiveness of echoic memory (Darwin *et al.*, 1972). Three amplitudes of vibrato were used: 0, 1 and 2%. Each raw signal was filtered, either using the input admittance of the Rubio violin (reference sound), or with a modified version of that admittance (test sound) resulting from a multiplication of all Q values by a factor whose value was varied to determine the threshold for discrimination.

2. Procedure

Thresholds for discriminating the reference and test sounds were estimated using the 3AFC procedure described earlier. This time, the quantity being varied was the scaling factor applied to all the modal Q values (an

upwards shift in all cases). The initial scaling factor was 2.8. The step size in the factor was 1.189 until two turn-points had occurred and 1.091 thereafter.

3. Participants

In this and all subsequent experiments the participants were experienced violinists (UK grade 8), who practiced regularly. They were paid for their participation. For this experiment there were 14 participants whose hearing was checked to be normal (defined here and below as audiometric thresholds below 15 dB HL at the standard audiometric frequencies).

4. Results

A one-way within-subjects analysis of variance (ANOVA) showed no significant difference in threshold for the three vibrato amplitudes, including the case with no vibrato: in all cases the threshold Q scaling factor was 1.4. In other words, a 40% reduction in damping was required for “threshold”. Again, the result seems surprising. One might have expected that an input signal with vibrato would lead to enhanced sensitivity to changes in damping, since the frequency modulation produced by the vibrato would cause fluctuations in spectrum which might provide a cue related to the Q values. A possible explanation for the lack of effect of vibrato amplitude involves informational masking (Neff and Green, 1987): the auditory stimulus is more complex in the presence of vibrato, with more “irrelevant” information, and this may make the task of discriminating a change in damping more difficult, even though more information is being presented to the auditory system.

C. Experiment 3: The effect of vibrato amplitude and level of damping on perceptual dissimilarity

Experiments 1 and 2 showed, somewhat surprisingly, that the detection and discrimination of vibrato amplitude was not affected by the presence or absence of resonances in the frequency response, and that the detection of changes in Q value was not influenced by the amount of vibrato in the input signals. However, both of these experiments involved discrimination of stimuli varying along a single dimension. It could be argued that the results are not relevant to the type of judgments made when assessing differences in quality between virtual “instruments” with different amounts of vibrato. Violinists often describe such quality differences in terms of “liveliness” or “responsiveness”. Accordingly, an experiment was designed to investigate how both vibrato amplitude and damping influence listeners’ judgment of differences between synthesized violin sounds.

The hypothesis to be tested was that perceptual similarity or dissimilarity may correlate with the richness or complexity of spectral fluctuations induced by the resonant body response when vibrato is applied. An intu-

itive aspect of a more “responsive” instrument is that it does not require a very large vibrato amplitude to evoke the level of spectral fluctuations for the desired richness. Many players might therefore expect to be able (to a degree) to compensate for inadequate responsiveness of an instrument by increasing the vibrato amplitude. Following this line of reasoning, perhaps the perceived degree of richness depends on the interaction of damping and vibrato amplitude. This idea was tested by obtaining ratings of dissimilarity for pairs of sounds which differed along two dimensions, amount of damping and vibrato amplitude.

1. Stimuli

Sound files were synthesized corresponding to all combinations of two different parameter variations: three values of the vibrato amplitude (1, 2 and 3%) and three sets of Q values (original, divided by 2 and multiplied by 2); for a detailed rationale of these values, see the account of Experiment 7 below. The synthesis details were the same as for section III.B, except that this time some reverberation was added to further simulate natural listening conditions, as typically experienced with recorded sound. We used an industry-standard digital audio processing package (Digidesign Pro Tools), using a “small room” option, with a decay time of 1.42 s and diffusion of 87%.

2. Procedure and participants

The method of triadic comparisons was used (Wickelmaier and Ellermeier, 2007). Participants were presented with triads of sounds, each of which could be heard individually as often as desired, by clicking a button on the computer screen. For each triad of sounds, participants had to specify the most similar pair and the most dissimilar pair. Each participant listened to 84 different triads (all possible triadic combinations of the nine pairs of parameter values) plus the repetition of thirteen of them to check consistency. The 14 participants reported having normal hearing (which was not checked).

3. Results

Responses were generally consistent across repeated stimuli; chance responses would have resulted in 16% of second responses being the same as initial responses, whereas our subjects performed at 56.5%

A simple rating scale for dissimilarity was used to process the results: 2 points were allocated to the most dissimilar pair, 1 to the intermediate pair, and 0 to the most similar pair. By adding all participants’ dissimilarity points for all triads, a dissimilarity matrix was constructed. This was then analyzed with the multi-dimensional scaling (MDS) algorithm ALSCAL available in SPSS. The two-dimensional map shown in Fig. 3 was obtained with an S-stress value of 0.03.

FIG. 3. (Color online) Perceptual space obtained using MDS for stimuli that were varied in two ways: (1) by changing the Q values of the violin resonances (Q values divided by 2 ($Q/2$), original Q values (Q) and Q values multiplied by 2 ($Q \times 2$); (2) by changing the vibrato amplitude (a_{vib} equal to 1, 2 and 3%).

The two dimensions correspond approximately to the two physical parameters which were varied: vibrato increases along dimension 2 and the Q value along dimension 1. Q also maps partly onto dimension 2, but the overall shape indicates that the two physical parameters have largely independent perceptual effects. High Q values combined with a low amplitude of vibrato (e.g. $Q \times 2$, $a_{vib} = 1\%$) cannot be perceptually substituted for low Q values combined with a large amplitude of vibrato (e.g. $Q/2$, $a_{vib} = 3\%$); these combinations of parameters are widely separated in the perceptual space. Contrary to our initial hypothesis, the interaction of vibrato and Q in determining perceived similarity is not large. There is some interaction, however, as evidenced by the parallelogram form in Fig. 3: the top left and bottom right corners were found closer than the top-right and bottom-left corners.

Dimension 2 appears to be mainly related to the vibrato amplitude while dimension 1 appears to be mainly related to the Q values, and may reflect a perceptual dimension that is related to spectral smoothness versus spectral unevenness in the acoustic properties of the violin sound. The perceptual effect of a modification of the Q values was approximately uniform on a logarithmic scale, with the map location for Q lying in the middle of the locations for $Q/2$ and $Q \times 2$ along dimension 1. The perceptual effect of an increase of the vibrato amplitude did not fit well to either a linear or a logarithmic interpretation: the distance between 1% and 2% is approximately twice the distance between 2% and 3%.

IV. PERCEPTUAL CORRELATES OF VIOLIN VIBRATO

A. Experiments 4a and 4b: Effects of damping and vibrato amplitude on judgments of liveliness

1. Stimuli

Experiment 3 revealed the perceptual dimensions associated with change in damping and vibrato amplitude, but did not establish whether these were associated with changes in the quality of “liveliness” often described by violinists; one might expect that a decrease of the vibrato amplitude and/or the Q values would reduce the perceived liveliness. A linked pair of experiments was performed to probe this question.

For both experiments, sounds were synthesized exactly as for experiment 3. To investigate fully the influence of vibrato, cases were included with a very small amount of vibrato and with no vibrato at all. For experiment 4a, the range of the vibrato amplitudes was therefore increased relative to those used for experiment 3: the values 0, 0.5, 1, 2 and 3% were used. The Q values were the same as for experiment 3, giving a corpus of 15 sounds. The sounds were equalized in loudness, using the methodology described in section II.D of Fritz *et al.* (2007), to minimize the effect of loudness on perceptual differences.

Experiment 4b was similar, except that the damping now had 5 values, with the Q values being either normal or multiplied/divided by 2 and 4, and the number of values of vibrato amplitude was reduced to 3 (1, 2 and 3%) to make the duration of the test reasonable.

2. Procedure and participants

Participants first listened to the whole corpus of sounds and were then asked to rate the liveliness of each sound — presented in random order — on a scale from 0 to 10. Each sound was presented twice, to check consistency. As described below, consistency was very good. Data were therefore averaged for each participant.

There were 17 participants for the first experiment and 11 for the second. Their hearing was checked to be normal.

3. Results

The results of experiment 4a were as follows. First, a reliability analysis was performed with SPSS (using a two-way mixed-effects model), which gave an Intra-class Correlation Coefficient (ICC) of 0.9 with $p < 0.001$ for average measures. The correlation being very good, we can consider the averaged liveliness ratings to be meaningful. They are given in Table I.

A within-subjects ANOVA was performed, with Q and vibrato amplitude as factors. The influence of Q was significant [$F(1.3, 19.9) = 4.3, \epsilon = 0.7, p = 0.02$], as was that of vibrato amplitude [$F(2.3, 33.8) = 34.7, \epsilon = 0.6, p < 0.001$]. The interaction was also significant

TABLE I. Average liveliness ratings for the five values of vibrato amplitude and the three values of damping for experiment 4a.

Q value	Vibrato Amplitude				
	0%	0.5%	1%	2%	3%
$Q/2$	2.6	5.1	6.7	6.6	6.8
Q	2.2	4.6	6.8	6.0	5.1
$Q \times 2$	1.3	4.5	6.5	5.3	4.6

$[F(5.4, 80.9) = 2.5, \epsilon = 0.7, p = 0.03]$. While the larger values of vibrato amplitude were rated the liveliest for $Q/2$, this was not the case for Q or $Q \times 2$, for which ratings were highest for the middle vibrato amplitude. This effect was most pronounced for $Q \times 2$. A significant quadratic trend was found for vibrato amplitude $[F(1, 15) = 44.5, p < 0.001]$: liveliness increased when the amplitude increased from 0 to 1% and then stayed constant or decreased for higher values. This pattern of results can be explained by subjective reports of the participants that a large amount of vibrato makes the sound appear artificial and unpleasant, especially when combined with low damping. Thus, if “liveliness” is considered as a positive quality, such sounds may be rated as less lively.

There was a significant linear trend for the Q value $[F(1, 15) = 5.0, p = 0.04]$: liveliness decreased when the Q value increased, although this effect was small for intermediate values of vibrato amplitude. This effect is the opposite of what would be expected from the claim of Mathews and Kohut (1973) referred to in the introduction.

For experiment 4b, a reliability analysis was again performed. The ICC was equal to 0.5, with $p < 0.02$. Since the ICC was much lower than for experiment 4a, we calculated Pearson’s bivariate correlations to assess the extent to which the pattern of results was similar across participants and the extent to which each participant’s results were consistent across repetitions. Five of the eleven participants showed no significant correlations with ratings of the other subjects, nor within their own ratings across repetitions. These five participants were considered as unreliable. The ratings for the remaining six showed positive correlations with each other and within their own ratings.

When the results for the five unreliable participants were removed from the analysis, the ICC increased to 0.9, with $p < 0.001$. The mean liveliness ratings for the remaining six participants are given in Table II. The mean ratings decreased with increasing vibrato amplitude and with increasing Q value.

A within-subjects ANOVA was performed on the data for the six reliable participants, with factors vibrato amplitude and Q value. The effects of both vibrato amplitude $[F(2, 10) = 36.1, p < 0.001]$ and Q $[F(4, 20) = 50.4, p < 0.001]$ were significant. The surprising linear trend with damping that was found for experiment 4a was also found here $[F(1, 5) = 100, p < 0.001]$, liveliness again decreasing as Q increased. The decrease in liveliness with increasing vibrato amplitude was confirmed

TABLE II. Average liveliness ratings for the three values of vibrato amplitude and the five values of damping for experiment 4b, for the six participants whose results were consistent.

Q value	Vibrato Amplitude		
	1%	2%	3%
$Q/4$	9.4	8.7	6.7
$Q/2$	8.7	6.3	5.5
Q	7.3	5.1	2.7
$Q \times 2$	6.3	3.0	1.8
$Q \times 4$	5.3	3.0	2.3

by a significant linear trend $[F(1, 5) = 38.2, p = 0.002]$, which is consistent with results of experiment 4a for vibrato amplitudes of 1% or more. However, it should be noted that some rating values in identical conditions differed markedly between experiments 4a and 4b. For example, for the “standard” (middle) amount of damping, and 3% vibrato amplitude, the mean rating was 5.1 for experiment 4a and 2.7 for experiment 4b. This may reflect individual differences across participants and/or an influence of the range of conditions presented on judgments for any specific condition.

The fact that five of 11 participants in experiment 4b gave unreliable results is an indication that the liveliness scale may not be entirely appropriate for characterizing perceptual differences among this particular set of sounds. The fact that the synthesized sounds were sometimes unrealistic (especially for large vibrato amplitudes) may have contributed to the unreliability of listener judgments. It was therefore decided to repeat part of the experiment, but using a recorded string signal instead of a synthesized sawtooth signal and using only a single representative value of vibrato amplitude.

B. Experiment 5: Liveliness judgments using recorded rather than synthesized string signal

1. Stimuli and procedure

A recording was made of a C sharp (277 Hz) played on the G string by an experienced violinist, instructed to use vibrato as in normal expressive playing. This was then taken as a typical signal. Vibrato amplitude was measured to be 2%. The Q values of the violin acoustical response were varied over the same five levels as in experiment 4b. The experiment was conducted in the same way as experiment 4b, except that participants were additionally asked to assess how much they liked each sound on a 0 to 10 scale. There were 12 participants, whose hearing was checked to be normal.

2. Results

The reliability analyses gave an ICC of -0.3 for liveliness ($p = 0.56$) and of 0.88 for preference ($p < 0.001$). Therefore, averaging the liveliness ratings across participants would not be meaningful, as there was no consis-

tent pattern of the results across participants. With the recorded string signal as input, but no variation in vibrato amplitude, “liveliness” seems even less appropriate as a quality for listeners to use for rating variations in damping, at least in the context of this kind of single-note test. Lack of anything resembling musical context may be important here, and this is explored later.

In contrast to the liveliness ratings, the ratings for preference proved to be similar across participants and the averaged ratings are presented in Table III.

TABLE III. Average liveliness (* not reliable because of low ICC) and preference ratings for the five values of damping, using a recorded string signal as input.

	Q value				
	Q/4	Q/2	Q	Q×2	Q×4
Liveliness *	5.5	5.2	5.2	4.9	4.5
Preference	2.8	3.5	5.3	6.5	5.3

A one-way within-subjects ANOVA on the preference ratings showed a significant effect of the amount of damping [$F(1.9, 21.3) = 8.2, \epsilon = 0.5, p = 0.002$], with a preference for somewhat higher Q than for the measured violin used as the baseline in this study (which is typical of conventional violins in this regard). The trend was cubic [$F(1, 11) = 11.4, p = 0.006$]: too much damping (Q/4) was definitely not liked but the extreme case with very low damping (Q×4) was reported as making the sound hollow and metallic (echoing comments noted by Mathews and Kohut (1973)). This trend makes more intuitive sense than the linear trend for liveliness obtained in experiments 4a and 4b.

The inconsistent use of the liveliness scale in this test prompted the question of whether “liveliness” is interpreted by participants as a quality of each individual sound, or whether listeners interpret it as a more global property of an instrument, in terms of how it responds to changes in, for example, the amount of vibrato applied. A first attempt to address this possibility led to another pair of experiments.

C. Experiments 6 and 7: Liveliness and preference for sounds grouped by “violin”

1. Experiment 6: stimuli, procedure and participants

The same synthesized stimuli as in experiment 4b were used (section IV.A), but this time they were not presented in an inter-mixed fashion but grouped by amount of damping, resulting in three sounds, differing in the amount of vibrato, for each of five “violins”. Participants were asked to listen to the three sounds for each violin, and then to assess the liveliness of that violin and indicate how much they liked it. There were seven participants, whose hearing was checked to be normal.

TABLE IV. Average liveliness (* not reliable because of low ICC) and preference ratings for the five “violins”, each corresponding to a different amount of damping.

	Q value				
	Q/4	Q/2	Q	Q×2	Q×4
Liveliness *	4.4	4.3	4.5	5.1	4.6
Preference	2.3	3.0	5.9	6.3	4.4

2. Results for experiment 6

The degree of agreement among the participants was extremely low for the liveliness ratings (ICC=-6.1, $p = 0.97$), whereas it was high (ICC = 0.92, $p < 0.001$) for the preferences. The mean ratings are shown in Table IV. Liveliness ratings did not show a clear dependence on damping. The results for preference are very similar to those of experiment 5. Again, there was a significant effect of damping on preference [$F(4, 24) = 12.8, p < 0.001$], with a significant cubic trend [$F(1, 6) = 24.3, p = 0.03$]: the highest mean rating was given for half of the damping (Q×2) of the original violin.

3. Experiment 7: stimuli, procedure and participants

The results so far suggest that “liveliness” is not a suitable word for characterizing the perceptual effects of a change in the modal damping, at least within the constraints of the experiments reported here. Hence we designed an experiment which assessed preferences for vibrato amplitude and body damping without using any specific verbal description of the quality that was being judged.

It is also fair to say that none of the synthesized sounds used in the experiment so far really achieved satisfactory realism. Experiment 7 used a synthesized input signal, but it incorporated some further small adjustments to the synthesis details in a continuing effort to improve naturalness. The randomness in amplitude was increased to 0.6% and the percentage of the period used in the Gaussian filtering was increased to 4.5%. Furthermore, the vibrato amplitude was not constant throughout the duration of the note (Schoonderwaldt and Friberg, 2001): after a period of 100 ms where the vibrato amplitude was small and constant (called the delay), the vibrato amplitude was linearly increased over 400 ms (called the attack), then linearly increased with a lower slope or remained constant during the sustained part (400 ms), and finally decreased linearly to zero over 500 ms. The amounts of vibrato at the end of the delay, attack and sustain phases are denoted a_d , a_a and a_s , respectively.

The note was a C4 sharp (273 Hz), chosen because a measured string signal was available for this note. The spectral envelope of the synthesized input signal was adjusted from the very regular form of an ideal sawtooth to approximately match the rather more irregular spectral envelope of the recorded sound, using FFT-based equalization.

Four vibrato envelopes were used:

1. $a_d = a_a = a_s = 0.02\%$ (i.e. almost no vibrato)
2. $a_d = 0.2\%$, $a_a = a_s = 0.5\%$
3. $a_d = 0.2\%$, $a_a = a_s = 1\%$
4. $a_d = 0.2\%$, $a_a = a_s = 2\%$

The numbers 1-4 are used as labels for the amount of vibrato. There were five “violins”, corresponding to the original Q values and Q values divided by 4 and 2, and multiplied by 1.5 and 2. These asymmetric ranges were motivated by the constraints imposed by real materials: damping could feasibly be increased by a factor of 4 by the violin maker or even by the player (by the way that they hold the violin against the chin and body, or with the left hand grip), but the damping cannot be reduced by the same factor, given the conventional wooden construction of the violin body. The larger factor was used in earlier experiments to create large effects, but for this test, in the interests of naturalness of the sounds, the range of Q values was restricted.

The participants were asked to choose, for each “violin”, the amount of vibrato which was the most natural for them. The five selected sounds were then re-presented in a second stage, and the participants were asked to rate how much they liked them, on a 0 to 10 scale. The 10 participants reported normal hearing.

4. Results for experiment 7

The degree of agreement among the participants for the first stage of the experiment was moderately good ($ICC = 0.63$, $p=0.046$), so average results are shown in Fig. 4. A one-way within-subjects ANOVA showed no significant effect of the amount of damping [$F(1.2, 11.2) = 2.7$, $\epsilon = 0.31$, $p = 0.12$] indicating that, when the frequency response of the violin is relatively flat (low Q), naturalness is not increased by the use of more vibrato. This is consistent with the conclusion drawn from the results of experiment 3, that damping and amount of vibrato are largely independent in their perceptual effects.

FIG. 4. Average number of the preferred amount of vibrato for the five “violins”, each corresponding to a different amount of damping. The error bars represent $\pm$ one standard deviation.

There is evidence that the discrepancy between judgments could represent a difference in strategy. In particular, when high Q values are considered, half of the participants judged $a_{vib} = 0.5\%$ to be most natural, and half judged $a_{vib} = 2\%$ as most natural. But very few judged $a_{vib} = 0\%$ or $a_{vib} = 1\%$ as most natural.

For the second stage of the experiment, which involved judgments of preference, the agreement of results across participants was much higher ($ICC = 0.88$, $p < 0.001$). The average preference score for each amount of damping is shown in Fig. 5. A one-way within-subjects ANOVA showed a significant effect of the amount of damping [$F(4, 32) = 7.7$, $p < 0.001$] with a significant linear trend [$F(1, 8) = 11.3$, $p = 0.01$], showing that the participants preferred higher Q values.

FIG. 5. Average preference score (on a scale from 0 to 10) for each violin (each corresponding to a different amount of damping) obtained using the preferred amount of vibrato. The error bars represent $\pm$ one standard deviation.

This preference for somewhat higher Q values than for the original response is consistent with the results of experiments 5 and 6. However, it seems not to be consistent with the results of Mathews and Kohut (1973). The discrepancy may derive from the fact that their study involved more extensive playing rather than single-note comparisons. Perhaps the reason for their participants not liking a peak-to-valley ratio larger than 10 dB was partly associated with unevenness of the violin notes across the scale, some being very loud and some being very soft. This was not a factor in the experiments described so far, as only a single note was used in each experiment. A further experiment was therefore carried out that involved real performance.

D. Experiment 8: Real-time playing on an electric violin

Violinists may need to play a violin to judge it reliably. Tests were therefore conducted with an electric violin, which allowed players to assess each virtual violin under fairly natural conditions.

1. Methodology

Instead of doing the filtering off-line using pre-synthesized/measured string input signals, the filtering was done in real time, using the output signal from a Harley Benton electric violin. The filtering was realized with Signal Wizard hardware, developed by Gaydecki (2009). The filtered sound was played through a stereo pair of high-quality loudspeakers (ATC SCM100ASL), in a recording studio. The player was facing the loudspeakers and two metres away from both. The sound level was adjusted for each subject in order to provide a comfortable level, while being sufficient to mask the direct sound from the electric violin. The experiment took place in a recording studio.

This methodology allows a player to test different violin sounds using the same physical violin. Also, it allows violinists to play a wide variety of effects, which increases the validity of the results with respect to violin performance, in contrast to the passive listening tests reported above.

2. Test Procedure and participants

There were six synthetic violins with a range of Q values; the original, divided by 4 and 2, multiplied by 2 and either by 1.5 or by 4. This experiment was in fact carried out simultaneously with the listening experiments, so the first six subjects were tested using $Q \times 4$. Later, $Q \times 4$ was dropped and the remaining nine subjects were tested using $Q \times 1.5$ instead. The 15 violinists reported normal hearing. Two of them were also violin makers of a high standard. The test was divided into two stages. In the first stage, violinists were instructed to play whatever they wanted, but without vibrato. In the second stage, they were encouraged to play with vibrato. In each stage, they were presented with the six violins in order of ascending Q value. After a first experience with each violin, to give an idea of the range of variation, they had to give a score to each violin (on a 0–10 scale) for both their preference (i.e. how much they liked it), and liveliness/responsiveness (i.e. how the violin responded to what they wanted to get from it).

3. Results

The agreement between participants was very good. The ICC was 0.76 ($p = 0.01$) for preference and 0.91 ($p < 0.001$) for liveliness/responsiveness. A within-subjects ANOVA showed no significant difference between ratings for the first stage (playing without vibrato) and the second stage (playing with vibrato), for either preference and liveliness. This can perhaps be explained by the comments of several players: vibrato is used to make the sound musically more interesting, but it does not change the quality of the violin. Average results for the two stages are presented in Table V.

While the liveliness increased significantly and progressively with decreasing damping, preference showed

TABLE V. Average liveliness and preference ratings for the six virtual violins played through the electric violin, each corresponding to a different amount of damping.

	Q value					
	$Q/4$	$Q/2$	Q	$Q \times 1.5$	$Q \times 2$	$Q \times 4$
Liveliness	4.6	5.4	5.5	5.9	6.4	6.8
Preference	5.4	6.1	5.8	6.0	4.4	2.5

a broad peak around the original damping. It seems that low damping increases liveliness, but is not preferred overall. This is a different pattern than was seen in Tables III and IV from the single-note tests, where slightly lower than normal damping was preferred. This may mean that there can be “too much liveliness,” or more likely that judgments of pleasantness are influenced by perceptual factors such as unevenness, or increased awareness of the body impulse response during the initial transient, especially with vigorous bowing techniques such as *martelé* (Woodhouse, 1983).

V. DISCUSSION AND CONCLUSIONS

The literature of violin acoustics has to date concentrated predominantly on physics, but most of the key questions have a perceptual dimension. We have attempted, through a series of tests, to shed light on one aspect of why a violinist or listener prefers one instrument over another. There are several threads running through the work: specifically, the perceptual effect of violin vibrato, the perceptual effect of changing body damping, and the interaction (or lack of it) between the two. More generally, we raise the issue of what is needed to make a synthesized sound “natural,” and the much bigger issue of how to design controlled empirical tests to access high-level musical perceptions and judgments using terms that are recognizable and meaningful to performers and instrument makers.

Any study which is to be relevant to the concerns of players and makers must take notice of their informal evidence, beliefs and experience. This study took as its starting point two observations. First, western classical violinists habitually use vibrato, and the interaction of vibrato with the “peaky” frequency response of the violin undoubtedly produces clear perceptual effects (Gough, 2005; Curtin and Rossing, in press). Second, among the verbal descriptions very commonly used by violinists to describe differences between instruments is a desirable quality of “liveliness,” or “responsiveness,” two terms which appear to be regarded as closely related (Fritz *et al.*, 2008). It seemed a promising hypothesis that these two things were connected, i.e. that “liveliness” would result from the combination of the use of vibrato and a peaky violin response, but any direct link between them has proved remarkably elusive to demonstrate.

The overall picture revealed by the tests reported in this paper is complicated. In order to probe the perceptual correlates of changes in vibrato amplitude and of damping of the violin body resonances, several experi-

ments were conducted. The simplest of these measured just-noticeable differences in vibrato amplitude and in violin body damping. Even these first results gave surprises: experiment 1 showed that the threshold for detection of a change in vibrato amplitude was unaffected by the amount of damping, and, conversely, experiment 2 showed that the threshold for detection of a change in damping was unaffected by the presence or amount of vibrato.

Experiment 3 investigated the combined perceptual effects of vibrato amplitude and body damping, based on judgments of similarity and dissimilarity. The results showed that, to a large extent, the two parameters are perceived as independent dimensions. However there was some systematic interaction which seems intuitively plausible: low damping with small vibrato amplitude was judged somewhat similar to high damping with high vibrato amplitude, when compared to the opposite pair of extreme combinations.

A subsequent series of listening tests examined how vibrato and body damping affect the perceived “liveliness” of, and general preference for, the sound. For those tests in which the amplitude of vibrato needed to be varied, the only way to achieve appropriate input without additional uncontrolled factors was by synthesis. However, all synthesized sounds achieved so far have suffered from a lack of naturalness, despite successive attempts to improve this aspect. This artificial quality may have compromised the ability of listeners to make judgments reflecting real musical contexts. Recorded signals taken from an actual performance are much closer to a musical context, but offer less flexibility in test design. Live playing using an electric violin and real-time filter system is better again, but is only suitable for certain kinds of test and also brings a host of new factors into play.

In summary, the results show that, when listening to single notes, participants find it difficult to make judgments of liveliness, and in most tests the word was not used in a consistent way across participants. This was especially true in the later tests, when participants were also asked to judge overall preference for the sounds. Judgments of preference were more consistent than judgments of liveliness. Only for the tests with live playing on the electric violin were participants able to judge both liveliness and preference consistently. The results of this experiment showed a trend which the authors had been expecting at the outset: lower damping produces progressively greater judged liveliness. However, no direct link with vibrato was found. Also, preference was broadly centered around the damping values for the original, reference violin. In the earlier tests with single synthesized notes, the liveliness ratings which were sufficiently consistent produced unexpected patterns. Sounds with no vibrato were judged less lively than sounds with some vibrato, but as vibrato amplitude increased the liveliness ratings flattened off or even decreased. More surprisingly, liveliness was judged to decrease, not to increase, as the damping was reduced. This all seems to suggest that liveliness is a quality more relevant to a player than to a non-playing listener, or at least that it is given a different interpretation in the two contexts.

Despite the fact that vibrato clearly influences the perception of musical timbre, we have been unable to confirm a clear relationship between the timbral properties of the instrument itself, and those produced by use of vibrato. It has previously been suggested that vibrato might either accentuate or compensate for resonance characteristics of stringed instruments. We did find evidence that varying the resonant behavior of an instrument can sometimes improve its musical properties. But the property of “liveliness” or “responsiveness”, often used to describe the capacity of an instrument for expressive sound production, appears to be largely independent of the specific expressive technique of vibrato. The popularly described relationship between vibrato and violin resonance may arise indirectly from the fact that both factors influence the perceived timbre in any given musical context.

A consistent trend from the tests involving preference judgments of single notes was a slight preference for damping somewhat lower than for the reference case (i.e. slightly higher Q values). However, this trend was not found in the final experiment, with the electric violin. In that experiment, the preferences were roughly symmetrical around the reference case. These different results are not necessarily in contradiction. They probably indicate a genuine tension, whereby some aspects of violin sound (such as timbre of an individual prominent note) benefit from lower damping, while other aspects (such as evenness) benefit from higher damping. This may echo a debate in the violin-making world concerning the selection of wood for instruments. Luthiers frequently express a preference for wood with a “good ring”, suggesting low damping, but there is also a persistent belief that some instruments improve with age. In some cases the wood in old instruments has degenerated to a chalky texture, suggesting a high density of micro-cracks, and consequently high damping.

What makes for a perception of “naturalness” in a synthesized violin sound? Even for single short notes with vibrato, it has proved surprisingly hard to produce sounds which do not evince an immediate sensation of artificiality. One would guess that the explanation has something to do with irregularities of various kinds, on different timescales: cycle by cycle, correlated variations over several cycles, and variations at timescales relevant to the player’s moment-to-moment input. There are few published studies on this question: the work of Schumacher (1992) is a notable exception. This question of realistic synthesis is a worthy subject of study in its own right. There is some knowledge and experience within the world of synthesis for the purposes of musical performance, but it is important to note a philosophical difference. If the aim is simply to make a good sound, there are no rules about what can and cannot be tried. But if the aim is to understand key features of actual played musical notes on conventional acoustic instruments, the ingredients of the synthesis should all have some basis in physics or in the physiological limitations on, for example, bow control and vibrato production.

Finally, there are issues of general methodological significance. In this paper the authors have attempted to design tests which have internal empirical and experi-

mental validity but which also retain at least some external validity through relevance to musical practice. This has proved challenging, and some of the results seem contradictory, but they should not be dismissed on those grounds. The experience of observing skilled musicians testing instruments suggests very strongly that violinists can make consistent judgments about subtle qualities of instruments. The fact that simple tests do not readily yield the knowledge of how they do it is a spur to further work, not an indication of a dead end. High-level musical perceptions are presumably of a comparable order of complexity to, for example, face or voice recognition, and they may take a comparable scientific effort to begin to understand.

In conclusion, while there are relationships between perceived timbre, Q values and vibrato, the primary function of vibrato when applied in real performance does not appear to be that of enhancing the perception of the properties of a violin body (damping). In conjunction with the fact that the use of vibrato, even on the violin, appears to have varied significantly over the last four hundred years of western musical history (see e.g. Moens-Haenen, 2009), this finding could be taken to suggest that the principal purpose of the use of vibrato in contemporary musical performance is to enhance expressivity.

VI. ACKNOWLEDGMENTS

The authors are grateful to Prof P. Gaydecki for making the Signal Wizard system available for this study and to the Leverhulme Trust for having funded this research. We thank an anonymous reviewer for helpful comments.

Brown, C. (1988). "Bowing styles, vibrato and portamento in nineteenth-century violin playing", *J. Roy. Music. Assn.* 97–128.

Cheng, F.P.-H. (2006). "Physics and perception of stringed instrument vibrato", MPhil thesis, University of Cambridge, UK.

Cremer, L. (1985). *The Physics of the Violin* (MIT Press, Boston, MA).

Curtin, J. and Rossing, T. (2009, in press). *Violins In Science of String Instruments*, edited by T. Rossing, (Springer, New York).

Darwin, C. J., Turvey, M., and Crowder, R. G. (1972). "An auditory analogue of the Sperling partial report procedure: Evidence for brief auditory storage", *Cognitive Psychol.* 3, 255–267.

Ewins, D. (2000). *Modal Testing: Theory, Practice and Application* (Research Studies Press Ltd, Baldock, England).

Fletcher, H. and Sanders, L. (1967). "Quality of violin vibrato tones", *J. Acoust. Soc. Am.* 41, 1534–1544.

Fritz, C., Blackwell, A.F., Cross, I., Moore, B. C. J., and Woodhouse, J. (2008). "Investigating English Violin Timbre Descriptors", in *Proceedings of the 10th International Conference on Music Perception and Cognition* (Sapporo, Japan).

Fritz, C., Cross, I., Woodhouse, J., and Moore, B. C. J. (2007). "Perceptual thresholds for detecting modifications applied to the acoustical properties of a violin", *J. Acoust. Soc. Am.* 122, 3640–3650.

Gaydecki, P. (2009). "Signal wizard 2", <http://www.signalwizardsystems.com/>, last accessed 1 July 2009.

Gorrill, S. (1975). "A viola with electronically synthesised resonances", *Catgut Acoust. Soc. Newsletter* 11–13.

Gough, G. E. (2005). "Measurement, modelling and synthesis of violin vibrato sounds", *Acustica - Acta Acustica* 91, 229–240.

Mathews, M. V. and Kohut, J. (1973). "Electronic simulation of violin resonances", *J. Acoust. Soc. Am.* 53, 1620–1626.

McAdams, S. (1989). "Segregation of concurrent sounds. 1: Effects of frequency modulation coherence", *J. Acoust. Soc. Am.* 86, 2146–2159.

McAdams, S. and Giordano, B. L. (2009). "The perception of musical timbre", in *Oxford Handbook of Music Psychology*, edited by S. Hallam, I. Cross, and M. Thaut, 72–80 (Oxford University Press, Oxford).

McIntyre, M.E., Schumacher, R. T., and Woodhouse, J. (1981). "Aperiodicity in bowed-string motion", *Acustica* 49, 13–32.

McIntyre, M.E. and Woodhouse, J. (1974). "Towards a psychoacoustically realistic violin physics", *Catgut Acoust. Soc. Newsletter* 22, 18–19.

Mellody, M. and Wakefield, G. (2000). "The time-frequency characteristic of violin vibrato: modal distribution analysis and synthesis", *J. Acoust. Soc. Am.* 107, 598–611.

Meyer, J. (1992). "Zur klanglichen Wirkung des Streicher-Vibratos (On the tonal effect of string vibrato)", *Acustica* 76, 283–291.

Moelants, D. (2004). "The timing of tremolo, trills and vibrato by string instrument players", in *Proceedings of the 8th International Conference on Music Perception and Cognition*, edited by S. Lipscomb, R. Ashley, and P. Webster (Northwestern University, Evanston).

Moens-Haenen, G. (2009). "Vibrato", In *Grove Music Online*. Oxford Music Online. <http://www.oxfordmusiconline.com/subscriber/article/grove/music/29287>, last accessed 1 Jul. 2009.

Neff, D. L. and Green, D. M. (1987). "Masking produced by spectral uncertainty with multi-component maskers", *Percept. Psychophys.* 41, 409–415.

Prame, E. (1994). "Measurements of the vibrato rate of ten singers", *J. Acoust. Soc. Am.* 96, 1979–1984.

Prame, E. (1997). "Vibrato extent and intonation in professional western lyric singing", *J. Acoust. Soc. Am.* 102, 616–621.

Schoonderwaldt, E. and Friberg, A. (2001). "Towards a rule-based model for violin vibrato", in *Proceedings of the Workshop on Current Research Directions in Computer Music* (Barcelona, Spain).

Schumacher, R. (1992). "Analysis of aperiodicities in nearly periodic waveforms", *J. Acoust. Soc. Am.* 91, 438–451.

van Besouw, R. M., Breteton, J. S., and Howard, D. M. (1996). "Range of tuning for tones with and without vibrato", *Music Percept.* 26, 145–155.

Wickelmaier, F. and Ellermeier, W. (2007). "Deriving auditory features from triadic comparisons", *Percept. Psychophys.* 69, 287–297.

Woodhouse, J. (1983). "On recognising violins: starting transients and the precedence effect", *Catgut Acoust. Soc. Newsletter* 22–24.

Yoo, L., Sullivan, D., Moore, S., and Fujinaga, I. (1998). "The effect of vibrato on the response time in determining the pitch relationship of violin tones", in *Proceedings of the 5th International Conference on Music Perception and Cognition*, edited by S. W. Yi (Seoul, Korea).