Is Vibrato Acoustically Necessary for Operatic Singing?

A thesis submitted in partial fulfillment of the requirements for the Bachelor of Science degree in Physics
from the College of William and Mary

by

Kevin Green

Advisor: Dr. Keith Griffioen

Senior Research Coordinator: Gina L. Hoatson

Date: May 10, 2017

Abstract

In opera, modern singers employ singing techniques that allow them to sing over large orchestras and to fill a large hall. In particular, they sing in such a way as to take advantage of formants, a resonance region with fixed amplification of upper harmonics of a fundamental pitch. During the Romantic period in music, roughly 1790 - 1910, major changes were ocurring that affected the singer's ability to be heard in a concert setting: concert halls began to be built with larger halls while composers began to write music requiring significantly more instruments, so that singers had to project into a larger space over a larger sound background. One of the developments of this time was the increased use of vibrato in singing. In this document I discuss formants to illuminate their importance and report an experiment I conducted to see whether vibrato has a significant effect on the ability of a listener to detect a pitch, ultimately concluding that a pitch with straight tone is on average more easily detected than the same pitch with vibrato tone played on the same white noise background by 8.86 ±1.14 seconds.

Contents

1 Introduction

1.1 Motivation

Having followed opera and having studied singing privately, it is of great interest to me how opera singers are able to sing how they do. In listening to operatic singers, certain charasteristics stick out, namely a steady use of vibrato not found in other forms of singing. Vibrato is defined as "periodic changes in the pitch of a musical tone" but also includes "periodic changes in amplitude" in singing vibrato ([1], 106). There exists a debate in the musical world about the use of vibrato in operatic singing with one side claiming that the use of vibrato should be required and the other that vibrato should be reserved for ornamentation; there are also various arguments in between these two sides. This debate is where I formed my question, in the hopes of providing some defense for one side over the other: does vibrato help to distinguish the human voice among an orchestra?

1.2 History

Over the course of the 19th century, opera singers experienced a shift in the demands that were placed on them. Opera houses began to be constructed with larger auditoriums as patronage shifted from the aristocracy to the box office, resulting in an interest in providing space for more people. Simultaneously, orchestras began to experience an increase in size as composers began to incorporate more instruments in their music, some in numbers over 100 instruments, and so greater amplification became a mandatory part of operatic singing. In order to compete with these orchestras, operatic singers take advantage of formants that appear naturally in the human voice. In this paper I don't discuss the ways in which singers attempted to achieve this as there is little record about particularities of a technique and mostly just self-reporting of perceived effects of technique. In doing whatever it was that they did, singers eventually came to some kinesthetic understanding of how to sing over an orchestra into a larger space; simultaneously, the steady use of vibrato in their singing began to emerge as a side-effect of their attempts to sing louder. It's important to mention that a sung pitch is never truly without vibrato and what I refer to in describing a "straight tone" is a tone perceived as being without vibrato. I discuss in the following sections some possible reasons for this, as well as my findings on formants and their usefulness to this topic.

1.3 Formants

A formant is a broad frequency range within which harmonics of a fundamental pitch are amplified; people have three resonating cavities that allow for multiple formants to develop: the larynx, mouth, and nasal cavity. The main determinant of a formant is the resonating space that the sound passes through: in the singer's case this means the mouth cavity, which a singer can modify by thinking of producing different vowel sounds, shifting the formants that are produced to center around different frequencies. Indeed, the difference between vowels is a shift in the arrangement of formants in the overtone series. The effect of different formants results in a difference in perceived timbre, which is the characteristic of a pitched sound that allows it to be distinguished from another identical pitch. In this sense, it is easy for the lay person to understand formants intuitively: they need only listen to different vowels.

When a sound is produced, it produces a particular pitch, as well as a group of pitches that lie at higher frequencies from it. This lowest-frequency pitch is the fundamental pitch, and the higher-frequency pitches are referred to as overtones. In producing a pitch, a particular instrument will, by its shape, material, length, size of resonating body, etc, have a distinct formant (or, in some cases such as the voice, multiple) that make more prominent certain overtones, and, in doing so, create a timbre associable with the instrument. An interesting feature of the human voice is that each voice will have variations in the formants that they possess, allowing for different people to be understood to have a distinct sound to them.

The more precise manner in which these formants are created is to change several different aspects of their 'vocal instrument,' adjusting the shape of their mouth or positioning their tongue differently in the mouth, elongating and shortening the larynx; modifying the resonating cavities in order to produce an environment in which certain formants are created and thereby producing the different vowels that we are familiar with. One of the typical complaints of operatic singing is that it can be difficult to understand the singer even if they are singing in the listerner's native language; one of the reasons for this is in the singer's decision to modify their vowels to produce a certain sound; in essence, they are picking and choosing certain positions for their resonating cavities to sing with certain formants that have been decided to be favorable - though, for the sake of the resulting sound rather than to create a particular frequency spectrum.

2 Experiments

2.1 Overview of Semester

Over the course of the Fall semester, my primary goal was to try to understand formants and to work on coming up for an experiment to perform in the Spring. I focused my attention to basic/introductory texts on acoustics, such as Helmholtz's *On the Sensations of Tone* [2], Berg & Stork's *The Physics of Sound* [1], and Roederer's *The Physics and Psychophysics of Music* [3]. I was also able to find information regarding formants in these books and was able to use the computer program Audacity to look at real formants with the help of a professional soprano.

The data I have collected comes from a session with Marje Palmieri, a professional singer in Falls Church, Virginia, who performed different vocalises with slight modifications, allowing me to look at the different effects they had in the overtone series she produced. While I was able to come up with many different tracks, for the purpose of this paper I decided to include just a few to show what effects occurred given the modifications she employed.

2.1.1 Exploring formants

The test subject was instructed to perform various vocalises in order to explore the formants of the vowels: /ɑ/, /ɛ/, /i/, /o/, /u/. In employing her technqiue, the test subject determines that at the pitch located at approximately 555 Hz ($C^{\#}5$), she begins to modify her vowels to favor rounder vowels, which produce a sound she describes as "fuller, with more shimmer." In lay terms, a rounded vowel is one which is produced with the lips in a circular shape; there are many different degrees of roundness and unroundness in linguistics and that particular research was beyond the scope of my project, so I just report the two as a dichotomy.

The test subject was asked to sing one of two pitches with certain constraints, attempting to control for the multiple variables that could affect the sound produced. The first two plots depict the same pitch with her natural vibrato at a similar "loudness," something which would be difficult if not impossible to do exactly. What has changed between the two is the vowel being sung, and the two plots provide a direct image of the change in the spectral envelope that occurred.

The following plots are of the pitch of A4 (~440 Hz) in different vowel configurations. The true value of the pitch is reported as the first peak.

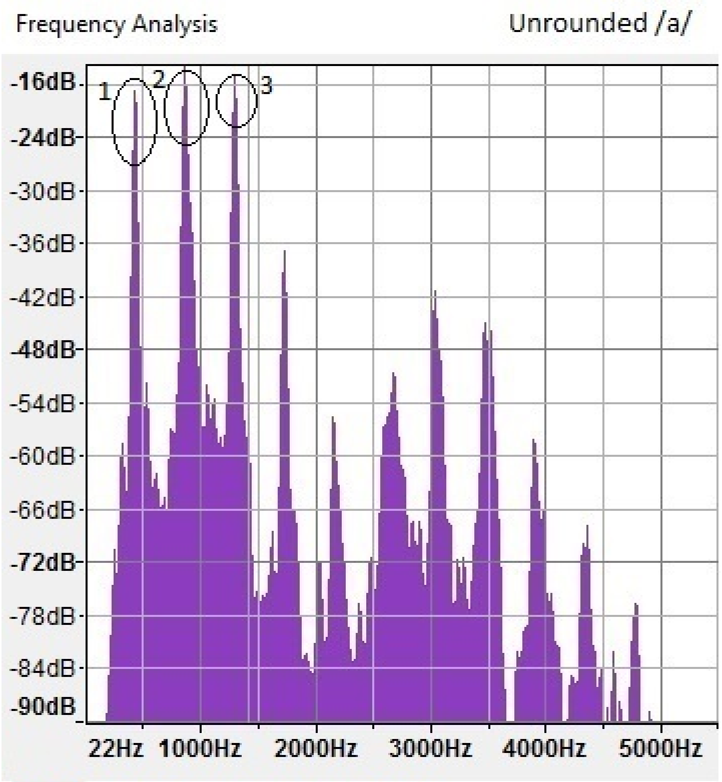

Figure 1. Operatic singer singing unrounded /a/
Frequency analysis of a sung tone on /a/ showing the formants amplifying certain overtones that lie within them.

In figure 1 above, the peaks are: (1) 450 Hz (A4); (2) 850 Hz (A5); 1300 Hz (E6).

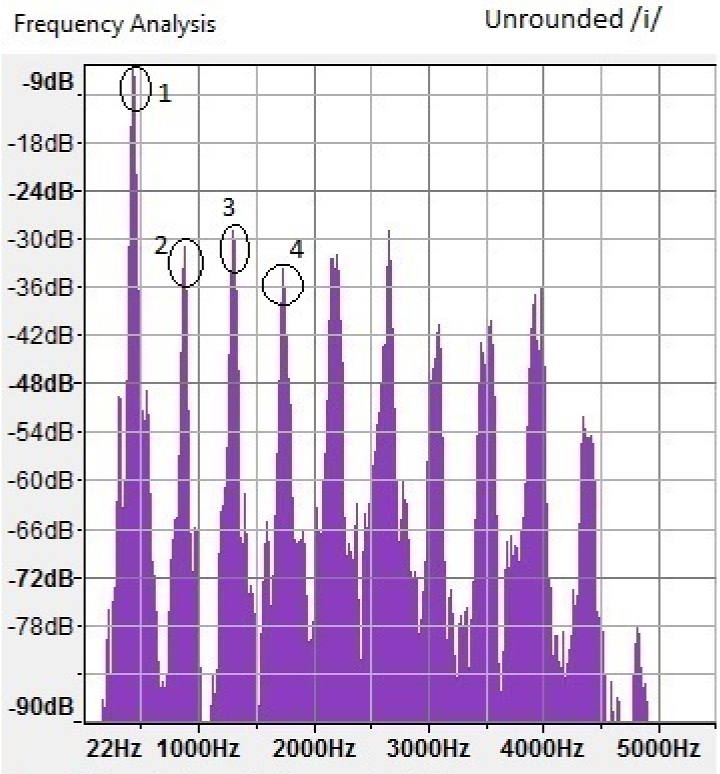

Figure 2. Operatic singer singing unrounded /i/
Frequency analysis of a sung tone on /i/ showing the formants amplifying certain overtones that lie within them.

In figure 2 above, the pitches circled are: (1) 444 Hz (A4); (2) 890 Hz (A5); (3) 1339 Hz (E6); (4) 1750 Hz (A6).

The second plot contains less acoustical noise, which does have some impact on the sounds, but the relative peak intensities of the two are so different that they can be appreciated approximately, at least visually. The dramatic change perhaps explains the ease with which the two vowels can be distinguished from each other. The musical implications are a little more subjective, but, given the far greater amplification of the first two overtones beyond the fundamental in the first plot indicates a greater ability to be heard. The second plot, however, does show more consistency among the overtones which would provide an easier frequency spectrum for the listener to process and understand as a specific sound.

In general, in modern classical singing there is the thought that the voice has several different regions, each separated by a *passaggio*, a break in the voice at a particular pitch. In these different regions, singers shift their vowels to approximate a more preferred vowel. In these two plots, the test subject is below her self-reported passaggio of C$^{\#}$5; she indicates that below this passaggio she attempts to bring her vowels closer to the /i/ sound (like the ee in greet), and that, above, she attempts to bring her vowels closer to a /œ/ sound (like the German ö which unfortunately is unused in English). In these two plots we can start to see why one would prefer an /i/ sound: while it's true that the first plot there is greater amplification of the first two overtones above the fundamental, the others are of so much lower intensity that the sound produced seems to have a less dynamic timbre than the /i/ vowel, where the collection of many of the first several overtones all sound at relatively similar amplitudes. There must be an aesthetic effect from this, too, but that unfortunately isn't part of the purview of this paper.

The following four plots are of the pitch at approximately 580 Hz (D5). The first two have the singer alternating between an unrounded /ɑ/ and the more rounded vowel attempting to approximate /œ/; the second two plots also change in the same way, with the shift being from /i/ to the vowel rounded to approximate a /ʏ/.

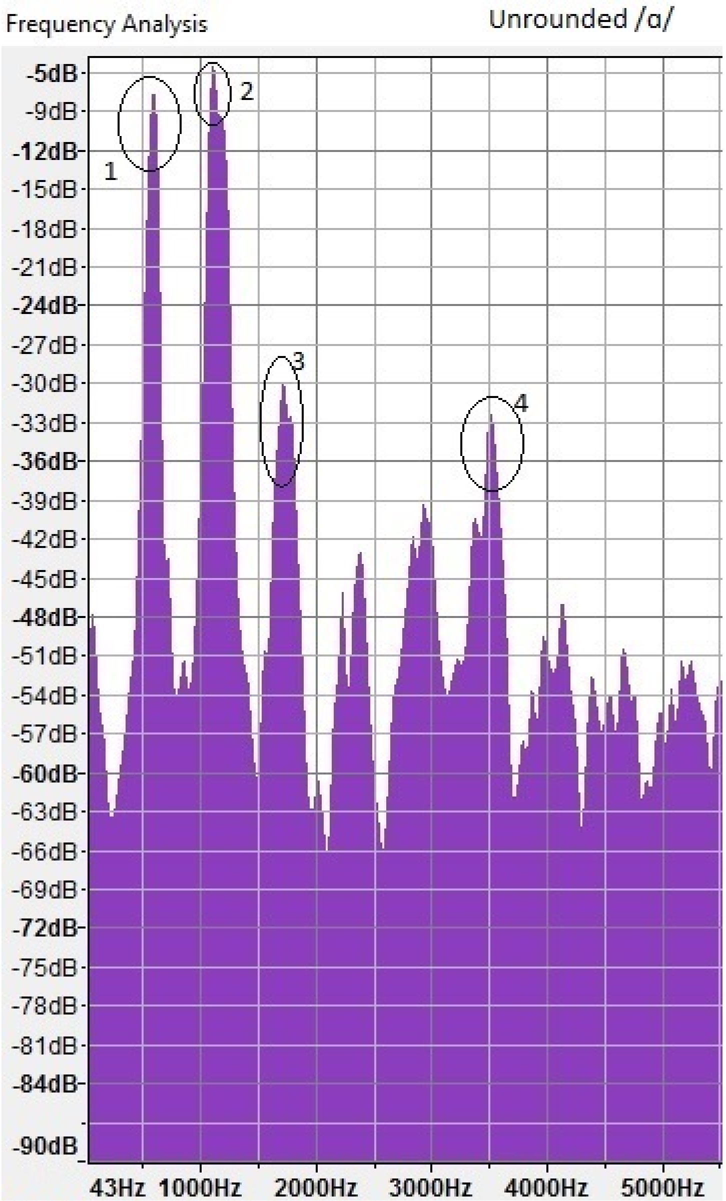

Figure 3. Operatic singer singing unrounded /ɑ/
Frequency analysis of a sung tone on /ɑ/ showing the formants amplifying certain overtones that lie within them

In figure 3 above, the circled pitches are: (1) 602 Hz (D5); (2) 1140 Hz (D6); (3) 1717 Hz (A6); (4) 3517 Hz (A7).

Figure 4 plots what a hearer would expect to still be the same vowel, an /ɑ/, but modified so that /ɑ/ would approach the sound /œ/:

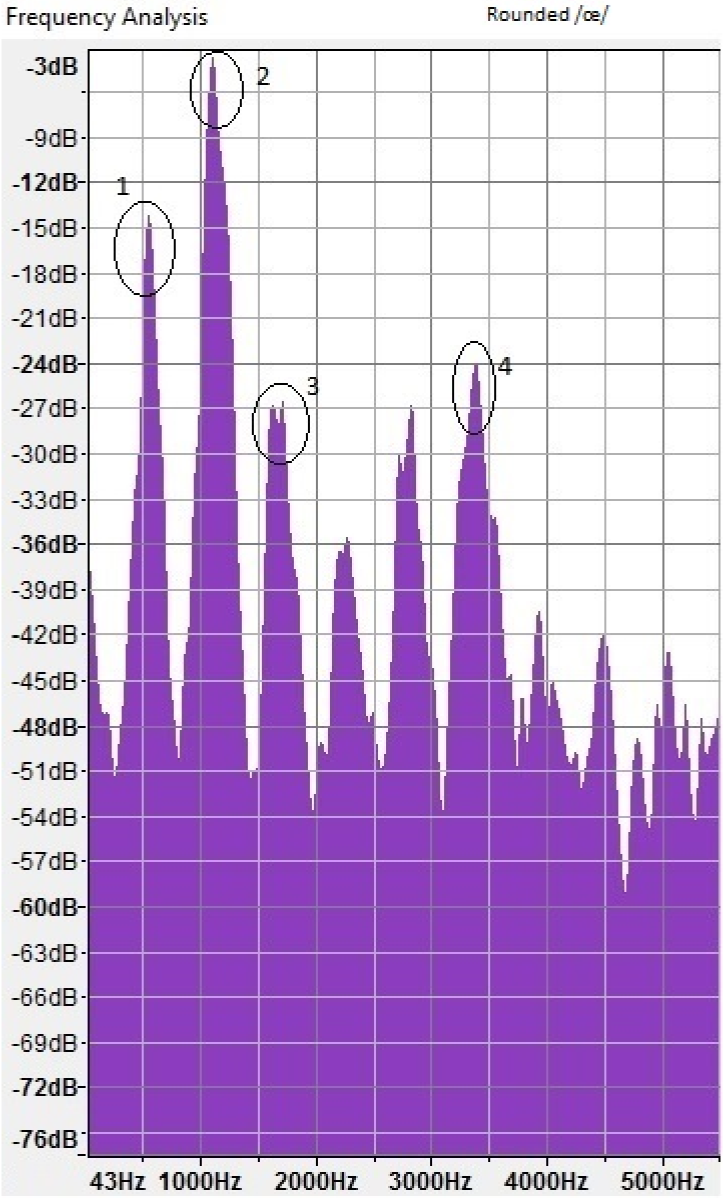

Figure 4. Operatic singer singing a rounded /œ/
Frequency analysis of a sung /œ/ showing the formants amplifying certain overtones lying within them.

The pitches indicated in figure 4 are: (1) 570 Hz (D5); (2) 1180 Hz (D6); (3) 1730 Hz (A6); (4) 3400 Hz (A$^\flat$7/A7). Pitch four is on the cusp between the A$^\flat$ and the A, so I have reported both rather than try to allocate it to one or the other.

In figures 3 and 4, the biggest difference is actually quite similar to the previous comparison: one has a higher amplification, while the other has a more even distribution. The unrounded vowel has a more prominent first overtone, but the overtones between the tones marked (3) and (4) are smaller, whereas in the rounded vowel they are more prominent if only slightly.

Finally, I have two plots of the same pitch but sung on the vowel /i/ and /i/ -> /ʏ/ (modified):

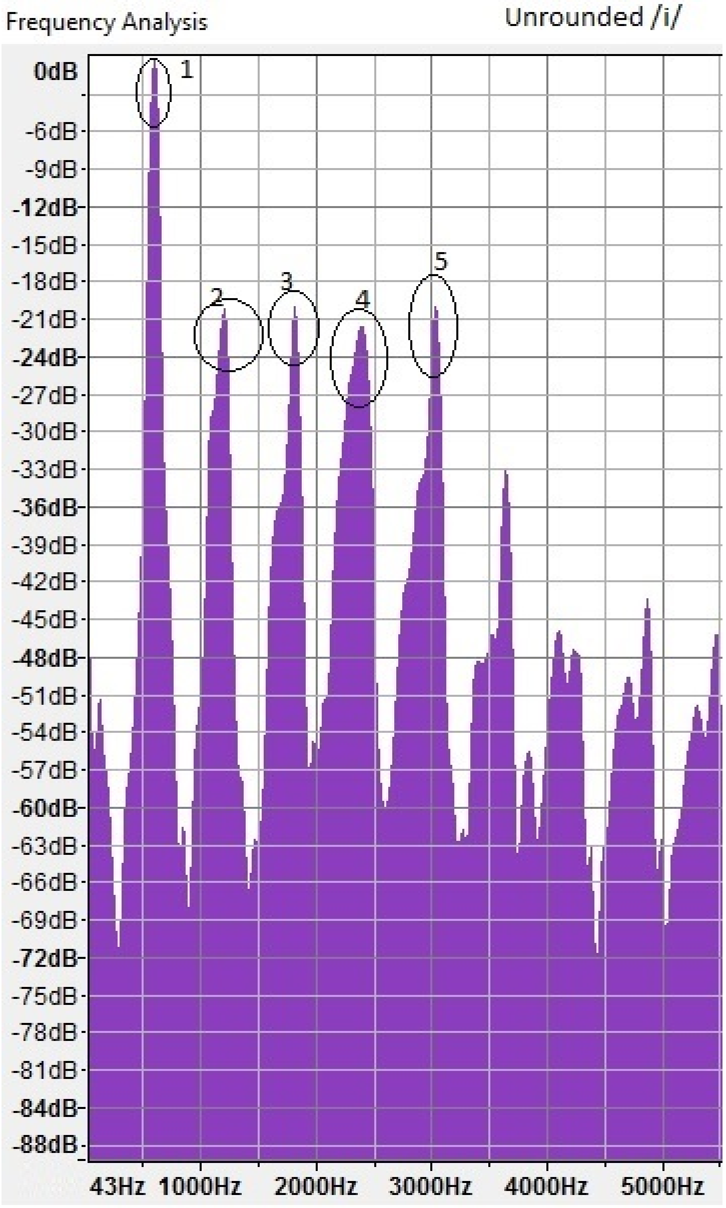

Figure 5. Operatic singer singing an unrounded /i/
Frequency analysis of a sung /i/ showing the formants amplifying certain overtones lying within them.

The pitches in figure 5 are: (1) 600 Hz (D5); (2) 1180 Hz (D6); (3) 1800 Hz (A6); (4) 2404 Hz (D7); (5) 3040 Hz (F#7).

The next plot features the modification of /i/ that shifts it towards the rounded /ʏ/:

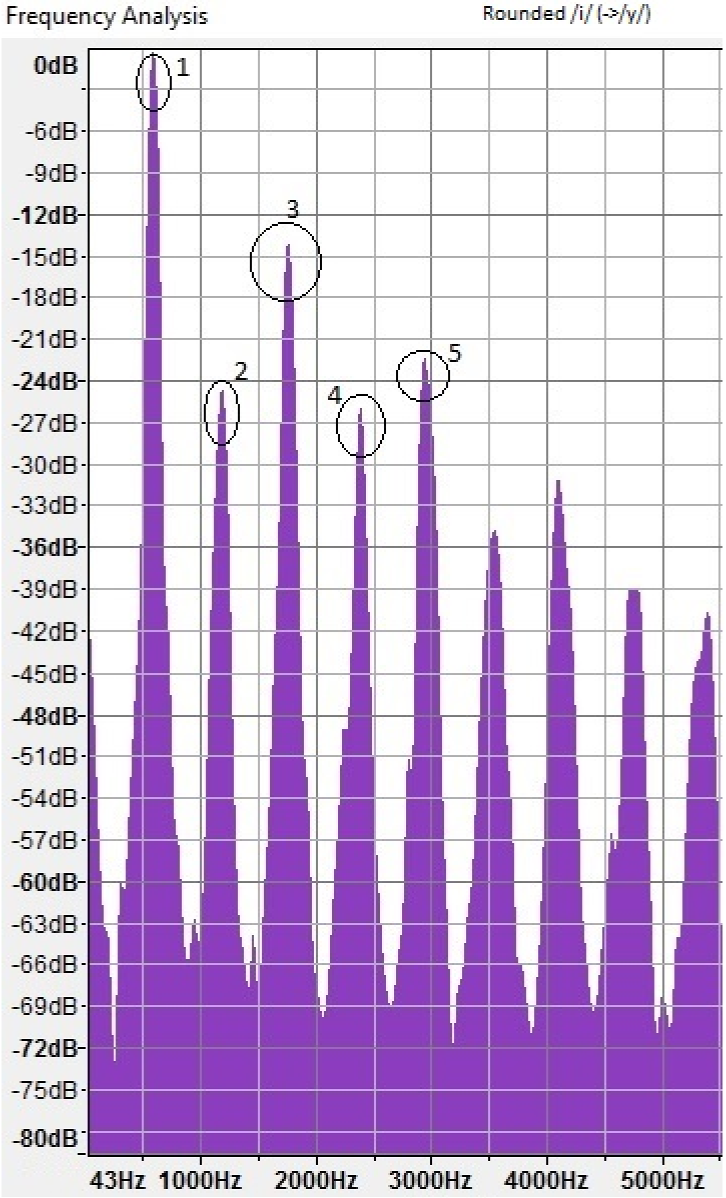

Figure 6. Operatic singer singing a rounded /i/ (-> /y/)
Frequency analysis of a sung /y/ showing the formants amplifying certain overtones that lie within them. The /y/ is a result of an /i/ being rounded.

With pitches: (1) 598 Hz (D5); (2) 1150 Hz (D6); (3) 1700 Hz (A6); (4) 2380 Hz (D7); (5) 2900 Hz (F#7).

Interestingly enough, the difference between figures 5 and 6 run counter to the vowel modification pattern employed by the singers might imply; in this case, the unmodified vowel is the more consistent, so that the peaks from the second overtone to the fifth are of similar intensities, while the rounded vowel has more variation among the peak amplitudes and a pronounced peak at the third overtone. What this indicates to me is that there must be an aesthetic quality that is produced by these changes not explainable acoustically.

2.2 Vibrato experiment

Pitch Detection through White Noise

2.2.1 Goal

Determine whether there is any appreciable difference in the recognition of a pitch that has vibrato vs. a pitch that does not have vibrato through a constant white noise background.

2.2.2 Process

2.2.3 Experiment overview

In this experiment, I had test subjects give a time in seconds when they were able to detect a pitch through white noise. The white noise was played at a constant amplitude while the amplitude of the pitch slowly grew from an amplitude of zero to an amplitude one half that of the white noise over 45 seconds. For each pitch provided there was one with no frequency or amplitude modulation to represent a **straight** tone; the other had a combination of frequency and amplitude modulation to represent a **vibrato** tone. Because, as explained before, the voice has simultaneous periodic changes in both frequency and amplitude, I created the vibrato tone incorporating both; these are explained in greater detail in section 2.2.4. The test subjects were given blank sheets of paper and prompted to simply number from one to eight before the experiment began. When I collected the papers, I wrote on them the group number and then numbered them.

I played a total of 3 pitches across 6 tracks. Each pitch had two tracks associated with it: one with a straight tone and a second with vibrato tone. Each track is labeled as belonging to a particular pitch by the letter associated with it (A,B,C), and the presence of modulation will be denoted by a subscript, where a subscript 1 will represent straight tone and a subscript 2 will represent vibrato tone.

The process for performing the experiment was kept simple. The test subjects were asked to number their papers and to mark next to each number the time at which they heard the pitch in a given track; I projected a stopwatch onto the front of the room and instructed them to use that time. This also allowed for some consistency in the direction they faced throughout the experiment. In determining what volume to use, I was able to set the speakers to their default setting and selected the speaker volume of 50 on the computer that I used. This was consistent among all the groups.

When the experiment began, I played track A_1, gave a short delay, played track A_2, gave the same delay, then repeated the two tracks in reverse order. I have denoted this section “Set 1.” For the second section, “Set 2,” I played all four tracks in the B and C groups in an arbitrarily selected order that was different between groups. I noted how these were different so as to be able to organize them in my tables. The tracks were all 45 seconds long, so there was a total of 6 minutes of audio for the subjects to listen to; tallying the breaks within the experiment, the total runtime of the experiment (post-script reading and questions) was approximately 7 minutes. The subjects were asked to stay to give feedback and comments after the experiments were performed.

2.2.4 Audio files

All pitches are relative to $A_4 = 440$ Hz. Each track is 45 seconds.

For the pitches labeled as having vibrato, the details of the vibrato are as follows:

Tremolo (amplitude modulation)

depth: 20%

frequency: 4 Hz

Vibrato (frequency modulation)

speed: 4%

depth: 20%

Note that the vibrato includes both of these modifications simultaneously.

The files are organized as:

File ($\mathbf{A_1}$)

Pitch: C_5= 525 Hz; tenor high C

Modulation: Straight

File ($\mathbf{A_2}$)

Pitch: C_5= 525 Hz; tenor high C

Modulation: Vibrato

File ($\mathbf{B_1}$)

Pitch: C_6= 1050 Hz; soprano high C

Modulation: Straight

File ($\mathbf{B_2}$)

Pitch: C_6= 1050 Hz; soprano high C

Modulation: Vibrato

File ($\mathbf{C_1}$)

Pitch: C_7= 2100 Hz; first in overtone series for C_6

Modulation: Straight

File ($\mathbf{C_2}$)

Pitch: C_7= 2100 Hz; first in overtone series for C_6

Modulation: Vibrato

2.2.5 Data

The data was collected from the students as described above. I collected the papers and organized the data into tables.

Data collected from experiment is :

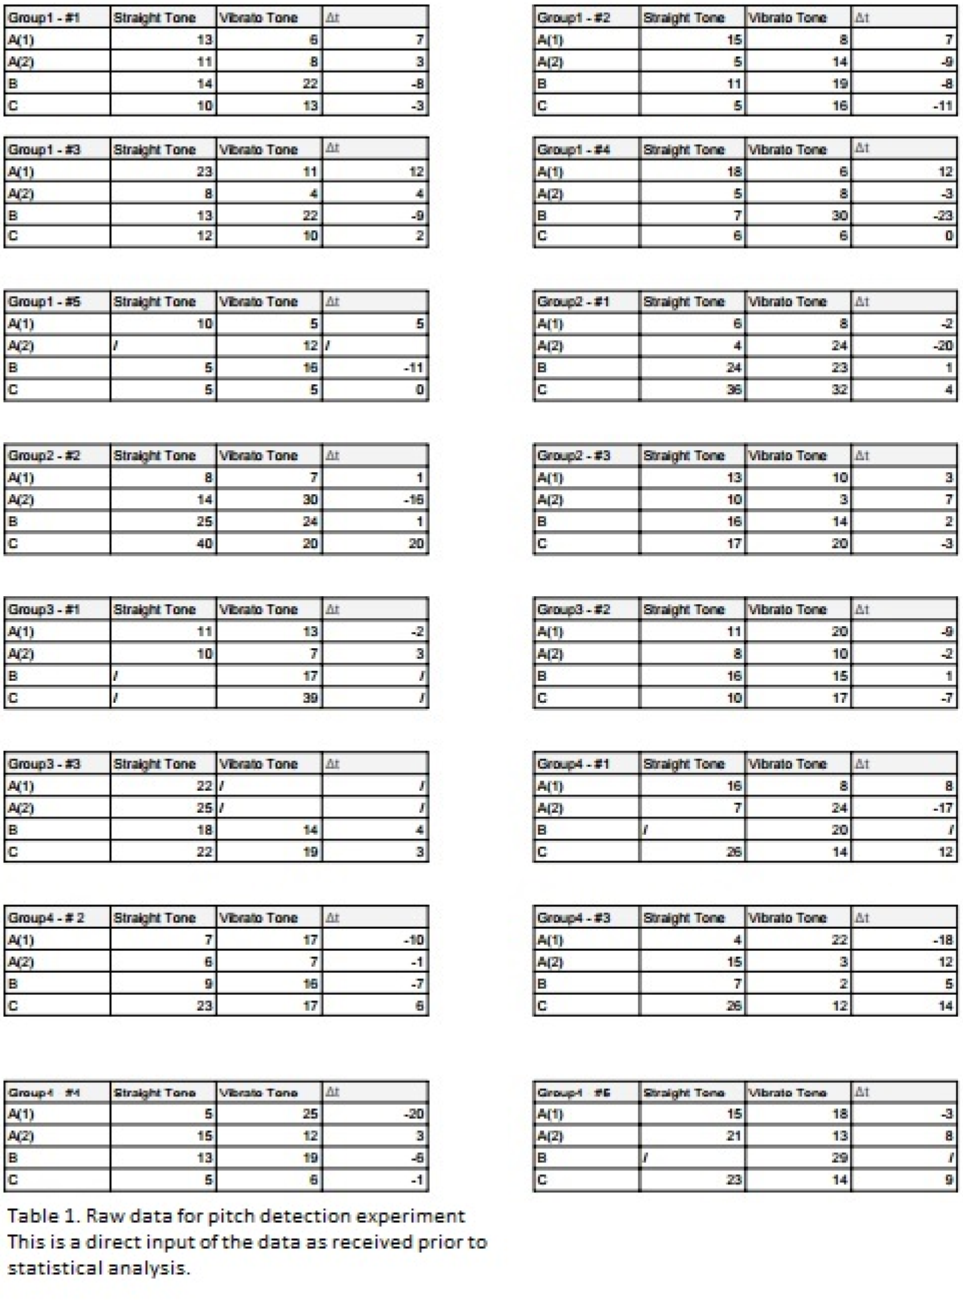

Group1 - #1	Straight Tone	Vibrato Tone	Δt
A(1)	13	6	7
A(2)	11	8	3
B	14	22	-8
C	10	13	-3

Group1 - #2	Straight Tone	Vibrato Tone	Δt
A(1)	15	8	7
A(2)	5	14	-9
B	11	19	-8
C	5	16	-11

Group1 - #3	Straight Tone	Vibrato Tone	Δt
A(1)	23	11	12
A(2)	8	4	4
B	13	22	-9
C	12	10	2

Group1 - #4	Straight Tone	Vibrato Tone	Δt
A(1)	18	6	12
A(2)	5	8	-3
B	7	30	-23
C	6	6	0

Group1 - #5	Straight Tone	Vibrato Tone	Δt
A(1)	10	5	5
A(2)	/	12	/
B	5	16	-11
C	5	5	0

Group2 - #1	Straight Tone	Vibrato Tone	Δt
A(1)	6	8	-2
A(2)	4	24	-20
B	24	23	1
C	36	32	4

Group2 - #2	Straight Tone	Vibrato Tone	Δt
A(1)	8	7	1
A(2)	14	30	-16
B	25	24	1
C	40	20	20

Group2 - #3	Straight Tone	Vibrato Tone	Δt
A(1)	13	10	3
A(2)	10	3	7
B	16	14	2
C	17	20	-3

Group3 - #1	Straight Tone	Vibrato Tone	Δt
A(1)	11	13	-2
A(2)	10	7	3
B	/	17	/
C	/	39	/

Group3 - #2	Straight Tone	Vibrato Tone	Δt
A(1)	11	20	-9
A(2)	8	10	-2
B	16	15	1
C	10	17	-7

Group3 - #3	Straight Tone	Vibrato Tone	Δt
A(1)	22	/	/
A(2)	25	/	/
B	18	14	4
C	22	19	3

Group4 - #1	Straight Tone	Vibrato Tone	Δt
A(1)	16	8	8
A(2)	7	24	-17
B	/	20	/
C	26	14	12

Group4 - # 2	Straight Tone	Vibrato Tone	Δt
A(1)	7	17	-10
A(2)	6	7	-1
B	9	16	-7
C	23	17	6

Group4 - #3	Straight Tone	Vibrato Tone	Δt
A(1)	4	22	-18
A(2)	15	3	12
B	7	2	5
C	26	12	14

Group4 #4	Straight Tone	Vibrato Tone	Δt
A(1)	5	25	-20
A(2)	15	12	3
B	13	19	-6
C	5	6	-1

Group4 #6	Straight Tone	Vibrato Tone	Δt
A(1)	15	18	-3
A(2)	21	13	8
B	/	29	/
C	23	14	9

Table 1. Raw data for pitch detection experiment
This is a direct input of the data as received prior to statistical analysis.

The data is organized by group number and participant number in the top left corner. A(1) refers to the first time that the participant heard the A pair of tones, and A(2) refers to the second time; they heard the B and C pairs once each. Δt is always (Straight tone) - (Vibrato tone) across the row. In doing this, I define a positive value as representing an event when the vibrato tone was heard at a faster time than the straight tone in that pair, and so the negative value refers to the opposite. In events where the participant was unable to hear the pitch in a particular track, they were instructed to write a slash by that number.

The data from Table 1 was, in this form, rather unusable. It seemed random and I couldn't really piece together any conclusions, so I reorganized the data in two particular ways so as to be able to see how some of my expectations resolved with it. This is shown in Table 2.

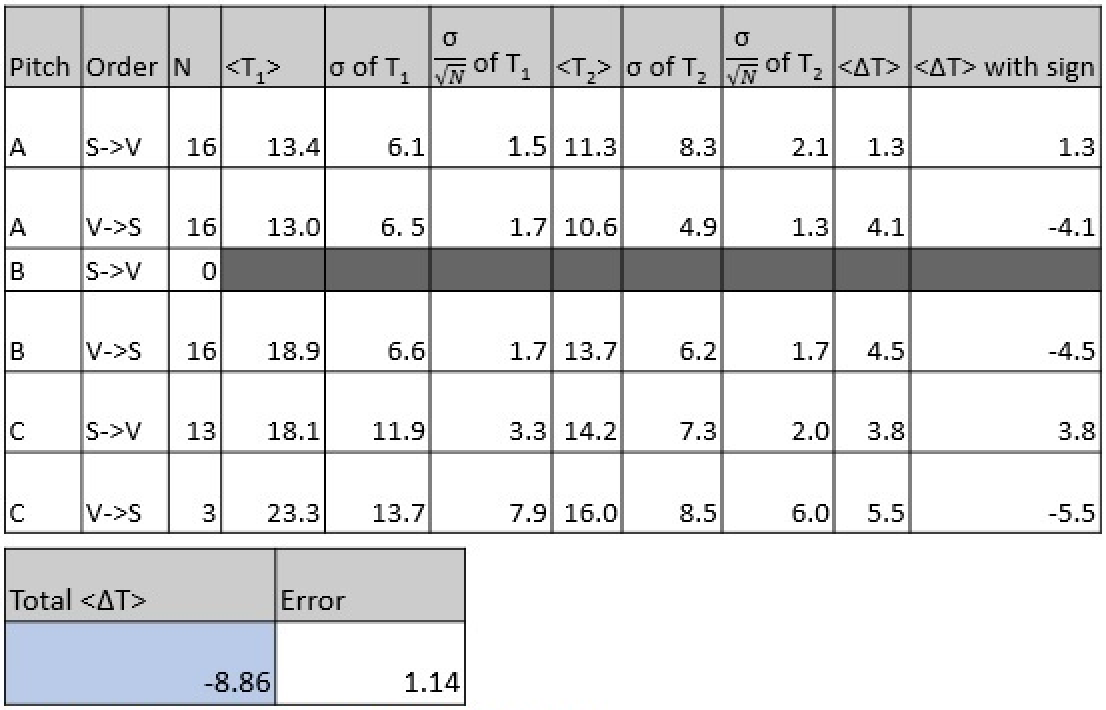

Pitch	Order	N	$<T_1>$	σ of T_1	$\frac{\sigma}{\sqrt{N}}$ of T_1	$<T_2>$	σ of T_2	$\frac{\sigma}{\sqrt{N}}$ of T_2	$<\Delta T>$	$<\Delta T>$ with sign
A	S->V	16	13.4	6.1	1.5	11.3	8.3	2.1	1.3	1.3
A	V->S	16	13.0	6. 5	1.7	10.6	4.9	1.3	4.1	-4.1
B	S->V	0								
B	V->S	16	18.9	6.6	1.7	13.7	6.2	1.7	4.5	-4.5
C	S->V	13	18.1	11.9	3.3	14.2	7.3	2.0	3.8	3.8
C	V->S	3	23.3	13.7	7.9	16.0	8.5	6.0	5.5	-5.5

Total $<\Delta T>$	Error
-8.86	1.14

Table 2. Statistical analysis of pitch detection experiment data
The table is split into the two forms of analysis done: first a collation of the data by each individual pitch (Row A, Column T1 represents the straight tone time of detection), the standard deviation, and standard error. The lower portion of the table provides the conclusion drawn from the experiment.

For each pitch I separated the data into two based on the order in which the tracks were played, marked as S->V if the $_1$(no vibrato; "Straight tone") track was played before the $_2$(Vibrato) track and vice versa. Interesting to notice is that the B pair, by random chance, were never heard with the straight tone played before the vibrato tone; whether this would have had any effect is hard to determine. The $<T_1>$ is the average time for the first track played in that row, so if it is S->V, it is the time for the straight tone, and if V->S then it is the time for the vibrato tone; I corrected for this in the $<\Delta T>$ column wherein, if the order was V->S, I applied a negative sign to flip the order for the sake of being consistent in analyzing the data. In the event that one pitch in a pair wasn't heard, I did not use the other pitch from the pair in these $<\Delta T>$ slots. I then found the standard deviation for the times, then the standard error. For the standard deviation I used Google Sheets's STDEV() function and then calculated the standard error by dividing it by $\sqrt{N}$ as indicated in the table.

From this group of data I found one of the more interesting results of this experiment. That the pitch had vibrato or straight tone seemed significantly less important than the order in which they were heard; whichever pitch was heard second was consistently heard quicker. It's important to note that the two pitches in a pair were not necessarily heard simultaneously except for the A pair; in some cases, the B pair could have been heard with both tracks from the C pair between them, in other cases they were heard simultaneously. It's an easy conclusion to draw, then, that the order in which the pitches are heard in this particular experiment is significant in the recognition of the pitch. It's also worth noticing that the values of the $<\Delta$T$>$ column are greater when the straight tone came second, indicating that the straight tone is in fact more easily heard than the vibrato tone through the white noise background, which is further shown by the Total $<\Delta$T$>$ figure, which was calculated by going through the entirety of the data and calculating the overall average. Doing this also allowed me to come up with the error in a like manner as in the previous figures.

The result, then, is that, on average, the pitch with straight tone was heard 8.86 ±1.14 seconds sooner; a rather solid result. I discuss my conclusions in the final section of this paper.

2.2.6 Participant Comments

I instructed the test subjects to simply mark with a forward slash any track in which they could not detect the pitch. Most of the comments that I received had to do with their insecurity in determining whether they were actually hearing the pitch or not. Some common issues were that they couldn't tell the pitch apart from the white noise clearly enough, even though they could sense some sort of difference. I asked these subjects what they did and, while they varied in how they responded to it, each individual was consistent with what they did themselves. They all noted the time at which they thought they heard the pitch and waited until they were certain; some individuals reported the first time, others the second time, and one individual reported the average time. My sense is that the important factor in this is that they were all consistent which they indicated they were.

Interestingly, there was a pretty even spit among test subjects between ones who were able to tell that there was a change in the frequency/amplitude modulation, and those who could not. When asked whether they perceived a difference in their ability to notice the pitch given the difference of vibrato and straight tone, the test subjects indicated that they were able to note the difference as they continued to hear the pitch, but rarely at first hearing the pitch. The common comment regarding this was that the pitch with vibrato seemed harder to distinguish from the white noise because it contained a vibrato of its own.

Another interesting question I received was whether or not the amplitude changes between the different tracks were the same. I assured the student that I had kept the rate of amplification the same, which was met with surprise by others in the group. What's interesting to me about these things is that it provides more information about the psychology going on in the test subjects, especially with regards to their own insecurity with hearing the pitch. I am not sure at this time how one could better preempt these effects in order to make this experiment more consistent across large groups.

3 Conclusion

In undertaking this experiment, I had an intuitive belief that vibrato made the detection of a pitch easier among a large background, but this data seems to indicate the opposite; a rather startling result. In particular, a recurring comment from participants was that the pitch with vibrato was harder to detect because of the perceived vibrato within the white noise itself; the lack of vibrato would presumably then make the straight tone a more distinguished sound within the noise.

My question becomes, then, what is the purpose of vibrato in operatic singing? An intuitive response would be that it has to do with its physical effect: in using vibrato, the power the singer is applying to the vocal cords is distributed over a larger frequency, presumably allowing for greater power to be applied than if the vocal cords were trying to sound a single pitch. This would indicate that vibrato could simply be a by-product of a singer's attempt to sing louder while protecting their vocal cords. There is an intersection here between physics and physiology that would be of great interest to me.

Another, perhaps less satisfying answer, is that there might simply be some aesthetic reasoning behind the use of vibrato. The subjectivity of this makes it difficult for me to envision pursuing, but I suspect that it could have some mixture of physics, psychophysics, and history that could be used to determine some answer.

The conclusion that I drew from this experiment was not as broad as I might have hoped, but it is not too difficult to connect it to the real world parallel of an operatic setting: in an orchestra, there is a mixture of instruments which utilize vibrato consistently throughout their playing and others which do not. This data seems to indicate to me that if the voice is heard at a volume that is lower than a particular instrument, the use of vibrato in the voice would be detrimental if the instrument were one being played with a straight tone; the data does not directly point to the situation in which both use vibrato and I suspect that in this operatic situation the amplification of overtones caused by the formants would play a much larger part in allowing the voice to be distinguished.

In expanding upon this research I would be inclined to try to replicate the experiment using as close to a realistic setting as possible: having an orchestra create the background and having an operatic singer performing alongside it applying a straight tone in one trial and a vibrato in another trial. This could then be used similarly to my pitch detection experiment to ascertain whether the use of vibrato has any effect on the listener's ability to distinguish the voice from the orchestra. A primary difficulty in this would be amplitude regulation, particularly between the singer and the orchestra; presumably one could take a step back and have the two perform in isolation and combine the tracks later, adjusting for the amplitudes relative to each other so that they are as close to identical. Regardless of how close to realistic this would be, a significant result could just as easily be drawn from this experiment with this small adjustment.

Despite the many complications that arose in my doing this experiment and which would arise in further experiments in this, I can at least firmly state that vibrato does not by itself have any advantage over straight tone in being detected over a white noise background and, in fact, straight tone is the preferable configuration for a pitch in this environment by a significant amount.

References

[1] Richard E. Berg and David G. Stork. *The Physics of Sound*. Pearson; 3 edition, August 27, 2004.

[2] Hermann Helmholtz. *On the Sensations of Tones*. Dover Publications; 2nd ed. edition, June 1, 1954.

[3] Juan G. Roederer. *The Physics and Psychophysics of Music: An Introduction*. Springer-Verlag New York, January 1995.