Title: Individual Differences in Perception of the Speech-to-Song Illusion are Linked to Musical Aptitude but not Musical Experience

Authors: Adam Tierney^{1†}, Aniruddh D. Patel^{2,3}, Kyle Jasmin¹, Mara Breen⁴

Affiliations

¹Department of Psychological Sciences, Birkbeck, University of London, London, UK

²Department of Psychology, Tufts University, Medford, MA

³Azrieli Program in Brain, Mind, & Consciousness, Canadian Institute for Advanced Research (CIFAR),
Toronto

⁴Department of Psychology, Mount Holyoke College, South Hadley, MA

†Corresponding author

Adam Tierney, Ph.D.

Birkbeck, University of London

Malet Street, London, WC1E 7HX

United Kingdom

Email: a.tierney@bbk.ac.uk

Phone: +44-020-7631-6368

Word count: 10346

Abstract

In the speech-to-song illusion certain spoken phrases are perceived as sung after repetition. Prior work has established that perception of the illusion is not limited to specialist listeners, such as individuals with musical training, but is perceived widely across the general population. However, little is known about whether there are stable individual differences in perception of the illusion, and if so, what factors underlie this variability. Here we assessed the strength of the song illusion across individuals using phrases that tend to be perceived as song when repeated, as well as phrases that continue to be perceived as speech when repeated, measuring the strength of the illusion as the rating difference between these two stimulus categories after repetition. Illusion strength varied widely, with differences in perception of the two categories ranging from 0% to 80% of the rating scale. Although variability in illusion strength was unrelated to degree of musical training, participants who perceived the illusion more strongly were proficient in several musical skills, including beat perception, tonality perception, and selective attention to pitch. This finding supports theoretical models of the speech-to-song illusion in which experience of the illusion is based on detection of musical characteristics latent in spoken phrases.

Public Significance Statement

In the speech-to-song illusion, certain spoken phrases begin, after repeated presentation, to sound as if they were sung. Here we investigated individual differences in perception of this illusion, finding wide variability across listeners that was reliably present across stimuli. We find that variability in perception of the song illusion is largely unrelated to demographic variables such as degree of musical training and language background. However, we also find that individuals who strongly perceive the illusion demonstrate high proficiency in a range of musical skills, suggesting that when perceiving the illusion listeners detect musical characteristics in spoken phrases.

1. Introduction

1.1. Sound-to-music illusions

Music often has acoustic patterns which distinguish it from speech or other sounds (Ding et al., 2017), and also elicits distinct neural responses compared to other sounds (Norman-Haignere et al. 2015; Zuk et al., 2020). Nevertheless, studies over the past decade have demonstrated that listeners sometimes report perceiving non-musical sounds (i.e., sounds not originally intended to be heard as music) as sounding like music. An important phenomenon in this regard is the speech-to-song illusion, first reported 20 years ago when Deutsch, Henthorn, & Lapidis (2011) described a phrase which was perceived as spoken when presented once but tended to be heard as sung after repetition. Building on this work, Simchy-Gross & Margulis (2018) and Rowland, Kasdan, & Poeppel (2019) demonstrated that repeated environmental sounds are rated as more musical after repetition, and Tierney, Patel, & Breen (2018a) found that pitch contours extracted from speech and resynthesized as complex tones increased in perceived musicality with repetition. However, not all non-musical stimuli sound musical after repetition, with some undergoing dramatic perceptual changes and others continuing to be heard as non-musical (Tierney, Patel, & Breen, 2018b). This raises the key question of why listeners perceive certain non-musical stimuli as music when they are played repeatedly.

1.2. Proposed mechanisms of the speech-to-song illusion

To explain the speech-to-song illusion (henceforth, “song illusion”), Deutsch et al. (2011) suggested that pitch salience is inhibited by default during speech perception, given that pitch plays a secondary role in communicating lexical information relative to other acoustic dimensions (at least in non-tonal languages). Based on this hypothesis, they predicted that an area anterolateral to primary auditory cortex which had previously been linked to pitch salience (Patterson, Uppenkamp, Johnsrude, & Griffiths, 2002; Penagos, Melcher, & Oxenham, 2004; Schneider et al., 2005) would show greater activation for stimuli which perceptually transform into song when repeated. This prediction was confirmed in an fMRI study by Tierney, Dick, & Sereno (2013), suggesting that

perception of the illusion is indeed linked to increased pitch salience. Deutsch et al. (2011) also hypothesized that perception of speech as song involves perceptual distortion of pitch, such that the perceived pitches of syllables are warped to fit a musical scale template. This hypothesis was supported by Vanden Bosch der Nederlanden, Hannon, & Snyder (2015a), who found that when participants perceived the song illusion in repeated spoken phrases, pitch changes which departed from the perceived melodic template were easier to detect than pitch changes which moved towards it.

The explanation for the song illusion advanced by Deutsch et al. (2011), however, does not specify why pitch perception should be disinhibited by stimulus repetition. A potential explanation was advanced by Castro, Mendoza, Tampke, & Vitevich (2018), who proposed that Node Structure Theory (MacKay, Wulf, Yin, & Abrams, 1993) could account for the song illusion. According to this theory, repeated activation of lexical nodes temporarily reduces their ability to become activated. Castro et al. (2018) suggest that while deactivation of lexical nodes diminishes the perception of speech, continued activation of syllable nodes leads to the emergence of a song percept. Supporting a role for lexical deactivation in the illusion, they find that words from denser phonological neighborhoods and words from an unfamiliar language are rated as more song-like after repetition. Similarly, Vitevitch, Ng, Hatley, & Castro (2020) report that words with a high phonological clustering coefficient (which measures the tendency for phonological neighbors of a word to also be neighbors of one another) are perceived as more song-like after repetition.

The Node Structure Theory explanation for the song illusion, however, struggles to explain why deactivation of lexical nodes along with continued activation of syllable nodes should lead to perception of song, rather than speech leached of semantic content (as in semantic satiation; Smith, 1984). The explanation advanced in Castro et al. (2018) is that since syllables are the unit of rhythmic structure in speech, and rhythm is also an important aspect of music, then syllable node activation without lexical node activation should lead to a song-like percept. According to this explanation, all spoken stimuli should transform into song when repeated, as lexical nodes are satiated and syllable

nodes continue to be active. Moreover, according to this explanation, the extent to which stimuli do or do not give rise to the illusion should be exclusively tied to phonological rather than acoustic or musical characteristics. Finally, this theory cannot explain increases in musicality with repetition of non-verbal stimuli.

These predictions, however, are not borne out by the literature. First, there is ample evidence that not all spoken phrases transform into song when repeated. Tierney et al. (2013), for example, assembled a corpus of "illusion" phrases, which listeners report transform into song when repeated, and "control" phrases, which continue to be perceived as speech when repeated. (These two types of phrases were matched for speakers, syllable rate, and duration.) The lack of transformation in the control phrases was replicated in Vanden Bosch der Nederlanden et al. (2015a), Tierney et al. (2018a), and Tierney et al. (2018b). Second, several studies have demonstrated that a range of acoustic and musical stimulus characteristics modulate the strength of the illusion. Tierney et al. (2013) showed that the illusion stimuli featured flatter pitch contours than the control stimuli. Falk, Rathcke, & Dalla Bella (2014) manipulated the stability of pitch contours within syllables and found that song percepts were more common for the more stable contours. Tierney et al. (2018b) found that higher song ratings for song illusion stimuli were linked to more isochronous rhythmic structure and a greater tendency for syllable pitches to follow musical melodic statistics. Rathcke, Falk, & Dalla Bella (in press) found that recordings with greater sonority led to earlier and stronger transformations to song. Third, differences across stimuli in perceived musicality after repetition were shown to be highly similar between spoken phrases and synthesized tone sequences with the same pitch contour and rhythm (Tierney et al., 2018a), suggesting that phonological neighborhood characteristics cannot be the primary factor determining the extent to which a spoken phrase sounds like song after repetition. Fourth, several studies have demonstrated that environmental sounds and non-verbal complex tone sequences increase in musicality after repetition (Simchy-Gross & Margulis, 2018; Tierney et al., 2018a; Rowland et al. 2019).

Thus Node Structure Theory alone is not a sufficient explanation of the song illusion. Given this observation, what theoretical framework could account for both the influence of phonological neighborhood and the influence of acoustic and musical characteristics on the perception of song in speech? One possibility is a hybrid of the accounts advanced in Deutsch et al. (2011) and Castro et al. (2018) in which deactivation of lexical nodes is *necessary but not sufficient* for perception of the illusion. According to this account, deactivation of lexical nodes frees up attention which can then be focused on acoustic characteristics of the stimuli, including pitch and timing patterns. If these melodic and rhythmic characteristics suggest that musical structure is present (e.g., an underlying tonal or beat pattern), then the stimulus will begin to sound like song. If not, the stimulus will continue to sound like speech.

1.3. Individual differences in perception of the speech-to-song illusion

While the above framework for explaining the song illusion merits further research and testing, such work will also need to address an important aspect of the illusion that has been neglected in prior research, namely individual differences in perception of the song illusion. Such differences are quite evident in classroom demonstrations and laboratory investigations of the illusion. Yet little is known about the extent of individual differences in the perception of the illusion across a broad sample of listeners or the factors driving this variability. In contrast, prior research has quantified individual differences in other auditory illusions. For example, perception of the ambiguous "Yanny/Laurel" recording varies such that roughly half of participants report each percept (Bosker, 2018; Gwilliams & Wallisch, 2020), perception of the McGurk illusion varies between 0% and 100% of trials (Cienkowski & Carney, 2002; Gurler, Doyle, Walker, Magnotti, & Beauchamp, 2015; Mallick, Magnotti, & Beauchamp, 2015; Brown et al., 2018), and perception of the octave illusion and of the missing fundamental varies widely across individuals (Deutsch, 1975; Craig, 1979; Davidson, Power, & Michie, 1987; Brennan & Stevens, 2002; Ladd et al., 2013). Thus, we predicted that there would be similarly robust and stable differences across participants in the

strength of the song illusion. The goal of the current work was to quantify these individual differences and to begin to explore predictors that help explain such variability.

To this end, we used the same set of illusion and control stimuli and testing methods employed in our previous research on the song illusion (Tierney et al., 2013; 2018b). Our first research question was whether this paradigm can elicit stable individual differences in the strength of perception of the illusion across individuals. To investigate this we recruited a large number of participants with a wide range of demographic characteristics (including age, language background, and degree of musical training). We then measured the degree of variability in perception of the illusion across this sample, using split-half reliability to assess the reliability of our measurement of individual differences in speech-to-song perception. We then investigated a number of possible factors driving individual differences in perception of the illusion, motivated by a range of theoretical explanations of the mechanisms underlying the illusion (surveyed above), as well as by existing research on individual differences in perception of other auditory illusions. Our first experiment focuses on potential demographic predictors of illusion strength, while our second experiment measures individual differences in auditory perception skill and musical aptitude. These experiments reveal the extent to which categorization of stimuli into domains (such as speech and music) differs across listeners and explore potential factors underlying variability in this domain-level categorization.

2. Experiment 1

2.1. Introduction: investigating demographic characteristics driving perception of the song illusion

Musical training is an obvious factor which could underly individual differences in the song illusion. If perception of the illusion requires detection of musical characteristics latent in non-musical stimuli, then musicians may be better able to detect these characteristics. Yet prior research has reported that musicians are somewhat less susceptible to auditory illusions in which acoustic details are subsumed by a top-down percept (Craig, 1979; Davidson et al., 1987; Brennan & Stevens, 2002; Pressnitzer, Graves, Chambers, de Gardelle, & Egré, 2018), perhaps because musicians are

better able to attend to low-level acoustic characteristics of stimuli. The existing evidence regarding musical experience and the speech-to-song illusion is inconsistent. Falk, Rathcke, & Dalla Bella (2014) found that participants with more years of musical training were no more likely to perceive the illusion, and in fact that their perception of the illusion was delayed relative to musically untrained participants. However, Vanden Bosch der Nederlanden et al. (2015b) found that musicians reported hearing all stimulus presentations as more musical than non-musicians, regardless of stimulus repetition or transposition. Given that Vanden Bosch der Nederlanden et al. (2015b) used only the original example from Deutsch et al. (2011) as a stimulus, this finding could indicate either that musicians are better able to detect the latent musical characteristics of this stimulus or that musicians are biased to perceive musicality regardless of stimulus characteristics.

Another possible factor underlying individual differences in the song illusion is age. As described above, Castro et al. (2018) speculated that one important precondition for perception of the illusion is deactivation of lexical nodes due to satiation. A similar explanation has also been advanced for the verbal transformation effect, in which repetition of a word eventually leads to an unstable percept that swaps between different words with different meanings (Warren & Gregory, 1958): MacKay et al. (1993) suggested that this illusion occurs when lexical nodes are satiated but phonological nodes remain activated. The verbal transformation effect is weaker in older participants (Warren & Warren, 1966; Pilotti & Khurshid, 2004; Pilotti, Simcox, Baldy, & Schauss, 2011), suggesting that satiation of lexical nodes due to repetition may be less rapid or less robust in older age. If so, then older participants may demonstrate less robust perception of the speech-to-song illusion as well. However, Mullin, Norkey, Kodwani, Vitevitch, & Castro (2021) reported no relationship between age and the strength of the speech-to-song effect, as perceived in the stimulus from Deutsch et al. (2011).

A third possible factor underlying individual differences in the song illusion is sex. There is some limited prior evidence for sex differences in the perception of auditory illusions. For example, Irwin, Whalen, & Fowler (2006) reported that female participants perceived the McGurk illusion

more strongly (i.e. greater influence of visual stimuli on auditory perception), but only when the visual stimuli were relatively brief. Similarly, female participants were more likely to report the illusory percept (i.e. "yanny") after hearing the Yanny/Laurel stimuli (Pressnitzer et al. 2018), a finding replicated by Gwilliams & Wallisch (2020). Given this prior evidence for sex differences in perception of auditory illusions, as well as evidence for sex differences in subcortical auditory processing (Krizman, Bonacina, & Kraus, 2019), we compared perception of the illusion between participants identifying as male versus as female.

Finally, there is some prior evidence that language background can affect perception of the speech-to-song illusion. Jaisin, Suphanchaimat, Candia, & Warren (2016) reported that native speakers of tonal languages did not perceive the speech-to-song illusion, although given the small sample sizes used (10 participants in each group) this finding is in need of replication. Castro et al. (2018) reported that repeated speech in an unfamiliar language was perceived as more musical than repeated speech in a familiar language. Finally, Margulis, Simchy-Gross, & Black (2015) reported that repeated speech in an unfamiliar language was perceived as particularly musical for languages that would be difficult for a participant to pronounce. Given these somewhat mixed prior findings, we investigated effects of language experience by determining both whether participants spoke English as their dominant language as well as whether participants were monolingual (only able to speak one language) or bilingual (able to speak at least two languages). More generally, given overall mixed findings regarding the roles of demographic variables in driving individual differences in perception of the song illusion, we did not have strong predictions about whether and how musical training, age, gender, and language background would be related to perception of the illusion, and so our analyses were exploratory.

2.2. Methods

2.2.1. Participants

One hundred seventy one (171) participants were initially recruited from a mixture of undergraduate psychology students and Mechanical Turk. Thirty-three (33) participants were removed from the dataset because they failed catch trials, not reporting an increase in song perception after stimuli switched from actual speech to actual song (see below for details). Responses from 138 remaining participants were analyzed. A power analysis conducted in GPower revealed that this sample size resulted in a power of 0.96 to detect a correlation of $r = 0.3$. Eighty-seven (87) participants identified as male, while 51 identified as female. (Participants were also given the additional options "Other" and "Prefer not to say", but no participants selected these options.) Participants reported a mean age of 29.96 (SD 10.37) years old (range 19-61). 112 participants reported that their dominant language was English, while 26 participants reported an alternate dominant language. These other dominant languages were Arabic ($n = 1$), Bulgarian ($n = 1$), French ($n = 1$), German ($n = 1$), Greek ($n = 2$), Italian ($n = 5$), Polish ($n = 3$), Portuguese ($n = 3$), Romanian ($n = 2$), Russian ($n = 1$), Slovak ($n = 2$), Spanish ($n = 3$), and Swedish ($n = 1$). Seventy-six (76) participants reported being bilingual (i.e. able to speak at least two languages), while 64 participants reported being monolingual. Participants reported a mean of 3.28 (SD 4.32) years of musical training (range 0-17). Study procedures were approved by the ethics board of the Department of Psychological Sciences at Birkbeck College.

2.2.2. Speech-to-song illusion experiment

Participants were presented with 48 spoken phrases drawn from the stimulus set first reported in Tierney et al. (2013). Prior studies (Tierney et al., 2013; 2018a; 2018b) have indicated that 24 of these stimuli tended to be heard as song after repetition (henceforth "illusion" stimuli), while 24 tended to be heard as speech after repetition (henceforth "control" stimuli). The mean length of stimuli was 5.5 (1.5) syllables and 1.33 (0.41) seconds. Stimuli were spoken by three different male talkers, represented in equal portions among illusion and control stimuli. These two stimulus categories did not differ meaningfully in syllable rate (illusion stimuli 5.13 syllables/second,

control stimuli 5.00 syllables/second) or duration (illusion stimuli 1.29 seconds, control stimuli 1.42 seconds).

In each trial, eight repetitions of each phrase were presented. After each repetition, participants pressed buttons labelled 1 through 10 to indicate whether the phrase sounded like speech or song, where 1 indicated completely speech-like, while 10 indicated completely song-like. Participants were given 2 seconds to respond to each phrase, after which time the program automatically went on to the next repetition. Any missing rating (due to a non-response) was replaced by the mean of the previous and following ratings. Missing first/last repetitions were replaced by the second/seventh repetitions.

In addition, four catch trials were presented. In each catch trial, a spoken phrase was presented during the first four repetitions, while a matched sung phrase (with the same words, roughly same rate, and similar pitch contour) was presented during the second four repetitions. Analysis was limited to data from those participants whose average rating of the last four repetitions exceeded that of the first four (by any amount) on these catch trials. As noted above, this exclusion criterion resulted in the removal of 33 out of 171 participants from the dataset, leaving 138 for analysis. The experiment lasted around 20 minutes.

To establish whether an overall speech-to-song illusion effect was present, a repeated measures ANOVA was conducted with two within-subjects factors, repetition (eight levels) and stimulus set (illusion versus control). Next, metrics summarizing illusion perception were compared to demographics characteristics. With eight ratings for each illusion and control stimulus, this is a high-dimensional dataset, potentially exacerbating multiple comparisons problems. To reduce the dimensionality of the data, therefore, we summarized each participant's responses in two ways for the purpose of investigating individual differences in perception of the illusion. First, the difference between the ratings of illusion and control stimuli after the eighth repetition was calculated and will be referred to as *illusion strength*. (In other words, for each participant we averaged their final rating for all illusion stimuli and all control stimuli, and then took the difference between these two mean

values to compute illusion strength.) Second the average rating across all repetitions and all stimuli was calculated and will be referred to as *response bias*.

Any variables not normally distributed according to a Jarque-Bera test (Jarque & Bera, 1980) were transformed prior to statistical analysis. A square root transformation was used for response bias, and age and years of musical training were converted to ranks. Pearson's correlations were used to assess the strength of the relationship between response bias and illusion strength versus age and years of musical training. Response bias and illusion strength were also compared between dominant English speakers and non-dominant English speakers, as well as between bilinguals and monolinguals, using unpaired t tests.

2.2.3. Available materials

Data from both Study 1 and Study 2 and test materials from Study 2 are available at <https://osf.io/hegxs/>. All stimuli from the speech-to-song illusion corpus are available at <https://osf.io/t4pqj/>.

2.3. Results

Overall, across participants the illusion was reliably experienced, as seen in the strong interaction between repetition and stimulus set ($F(7,959) = 229.6$, $p < 0.001$). There was, in addition, a main effect of repetition ($F(1,959) = 134.7$, $p < 0.001$) and of stimulus set ($F(1,137) = 408.3$, $p < 0.001$). However, there were very large individual differences across participants in the strength of the illusion, ranging from participants who reported no difference between illusion and control stimuli whatsoever, to participants who reported that the illusion stimuli sounded exactly like song when repeated while the control stimuli continued to sound exactly like speech. These individual differences are displayed in Figure 1, which plots average ratings of illusion and control stimuli across all eight repetitions from all 138 participants, sorted so that participants with larger ratings of illusion stimuli after the eighth repetition are closer to the top. (A color plot version of the same figure is presented in the supplementary information, see Figure S1. Additional plots showing individual differences for individual illusion and control stimuli are also available in supplementary

information, see Figures S2 and S3. Also see the supplementary information for sound examples of the illusion stimuli and control stimuli with the least and most variability in rating across participants, see Sound Examples S1 through S4.) In addition, Figure 2 displays the difference between the average rating of the last repetition and the average ratings of the first repetition in illusion and control stimuli, as well as the difference in average rating between illusion and control stimuli after the first repetition and the last repetition. The difference between the average last and first repetition rating for illusion stimuli ranged from -0.08 to 8 (mean 1.92, std 1.46), while this difference for control stimuli ranged from -1.17 to 4.92 (mean 0.33, std 0.86). The difference in average rating between illusion and control stimuli after the first repetition ranged from -0.83 to 4.08 (mean 1.33, std 1.02), while this difference after the last repetition ranged from -0.25 to 7.75 (mean 2.92, std 1.62).

To assess the reliability of this speech-to-song testing battery as a measure of individual differences in musicality perception, a Monte Carlo procedure was used. Across 100 iterations, the illusion and control stimuli were randomly divided into two halves. The average rating after the eighth repetition was then calculated separately for the illusion and control stimuli. The difference between ratings of the illusion and control stimuli was then calculated separately for each half of the trials (the *illusion strength* measure described above). The estimates from the two halves were correlated using a Pearson's correlation. The resulting correlation was corrected using the Spearman-Brown correction, to account for the fact that reliability would be expected to be slightly higher for a test with twice the number of trials. Finally, the median split half reliability across the 100 iterations was calculated. This procedure revealed a reliability of 0.88, suggesting that this testing battery provides a highly reliable measure of individual differences in the perception of musicality in spoken stimuli.

Figure 1: Waterfall plot displaying ratings of illusion and control stimuli across all 138 participants, sorted by the final rating of illusion stimuli. Each row shows the averaged ratings of a participant across all 24 stimuli in that category. The illusion stimuli matrix is sorted from highest to lowest final rating, and the control stimuli matrix shows each participant's corresponding average rating profile for control stimuli.

Figure 2: Top-left: histogram displaying the difference between the average rating after the last repetition minus the average rating after the first repetition for illusion stimuli. Bottom-left: histogram displaying similar data for control stimuli. Top-right: histogram displaying the difference between the average initial rating for illusion stimuli minus the average initial rating for control stimuli. Bottom-right: histogram displaying similar data for final ratings of illusion and control stimuli.

Years of musical training did not correlate with illusion strength ($r(136) = 0.13$, $p > 0.1$). However, years of musical training did correlate with response bias ($r(136) = 0.28$, $p < 0.001$), such that participants with a greater degree of musical training rated *all* stimuli (both illusion and control) as more song-like. Figure 3 displays scatterplots relating degree of musical training to both illusion strength and response bias, as well as average responses to each repetition of illusion and control stimuli in non-musicians and musicians. (For display purposes musicians were defined as having at

least six years of training, following Zhang et al. (2020), while non-musicians were defined as having no years of training.)

There was no correlation between age and illusion strength ($r(136) = -0.11$, $p > 0.1$). There was a small correlation between age and response bias ($r(136) = -0.20$, $p = 0.020$), such that older participants rated all stimuli as slightly less song-like. However, because this relationship would not survive post-hoc Bonferroni correction for multiple comparisons, this finding should be interpreted with caution. There was no difference in illusion strength ($t(136) = 1.87$, $p = 0.064$) between bilinguals (mean = 2.64, SD = 1.48) versus monolinguals (mean = 3.15, SD = 1.70). However, there was a significant difference in response bias tied to language experience ($t(136) = 5.01$, $p < 0.001$), with monolinguals showing a greater tendency to rate stimuli as song-like across both stimulus sets and all repetitions (mean = 4.10, std = 1.11) compared to bilinguals (mean = 3.18, std = 1.18). Figure 4 displays average responses to each repetition of illusion and control stimuli in bilinguals and monolinguals. There was no difference in illusion strength ($t(136) = 0.88$, $p = 0.381$) between participants whose dominant language was English (mean = 2.98, std = 1.69) compared to participants whose dominant language was not English (mean = 2.67, std = 1.29). Similarly, there was no difference in response bias ($t(136) = 1.49$, $p = 0.14$) between participants whose dominant language was English (mean = 3.62, std = 1.28) compared to participants whose dominant language was not English (mean = 3.96, std = 0.90). There was no difference in illusion strength ($t(136) = 0.71$, $p = 0.48$) between participants who identified as female (mean = 2.84, std = 1.43) versus participants who identified as male (mean = 3.05, std = 1.91). Similarly, there was no difference in response bias ($t(136) = 0.49$, $p = 0.63$) between participants who identified as female (mean = 3.63, std = 1.11) compared to participants who identified as male (mean = 3.79, std = 1.41).

Figure 4: Left: average responses to illusion and control stimuli in monolinguals. Right: average responses to illusion and control stimuli in bilinguals. Error bars indicate standard error of the mean.

2.4. Discussion

We find large individual differences in the magnitude of the song illusion across a diverse sample of participants drawn from the general population. There was widespread agreement across participants that the musicality of the control stimuli did not change with repetition. However, there was a high degree of variability across participants in the extent to which the illusion and control stimuli differed in musicality, and this variability increased as the stimuli were repeated: some participants reported large increases in musicality in the illusion stimuli with repetition, whereas others reported more modest increases. Importantly, our split-half reliability calculations showed a reliability of 0.88, suggesting that this paradigm can generate highly reliable measurements of individual differences in the perception of musicality in speech.

Given the high-dimensional nature of the data collected, musicality ratings were distilled into two summary measures. First, we calculated *response bias* by averaging across all musicality ratings regardless of repetition or stimulus category, forming an overall baseline measure of participants' tendency to perceive musicality in speech. Second, we calculated *illusion strength* by

averaging the final rating for all illusion stimuli and all control stimuli, and then taking the difference between these two mean values. These two measures related to demographic characteristics in very different ways. For example, we found that while degree of musical training did not relate to illusion strength, it did relate to response bias, such that participants with more musical training were more likely to produce high musicality ratings across all stimuli after all repetitions. This difference can be clearly seen in Figure 3, in which musicians' musicality ratings are higher than those of non-musicians even for control stimuli after the first repetition. This result replicates Vanden Bosch der Nederlanden et al.'s (2015b) finding that when musicians and non-musicians provided ratings of the original Deutsch et al. (2011) speech-to-song stimulus, musicians provided higher ratings overall, i.e. they heard all stimulus presentations more musically regardless of repetition or transposition. Here we replicate this finding in a larger number of participants across a much larger stimulus set. This musicianship effect, therefore, seems to be robust. One possible explanation is that musicians have sufficiently extensive knowledge of a wide variety of musical styles and examples that they can map any random sequence of pitches and durations onto a possible musical framework to some degree.¹

Turning to the effects of age, we did not find any differences between younger and older participants in illusion strength, and only a very small difference in response bias. Prior research has suggested that lexical node satiation decreases with age, as can be seen in decreased verbal transformation effects in older participants (Warren & Warren, 1966; Pilotti & Khurshid, 2004; Pilotti et al., 2011). This result suggests that individual differences in the readiness with which lexical nodes satiate is not a primary factor driving differences in perception of the song illusion in these stimuli. However, this conclusion may not generalize beyond the current stimulus set, and stimulus sets designed to elicit cross-stimulus differences in lexical satiation (by manipulating, for example, phonological neighborhood density; see Castro et al., 2018) could lead to larger age effects. We also found no differences in either illusion strength or response bias between participants who identified

¹ An extreme example of this ability can be seen when jazz accompaniment brings out the musical characteristics of natural speech, as in Henry Hey's orchestration of political speeches and interviews (<https://www.youtube.com/watch?v=9nlwwFZdXck>).

as male versus female, suggesting that perception of the speech-to-song illusion is relatively unaffected by gender.

In terms of the effect of language background, we found no differences in illusion strength between monolinguals and bilinguals, nor between dominant and non-dominant speakers of English. This result suggests that detecting musical characteristics in speech is not enhanced by exposure to multiple languages, and that lack of proficiency in a spoken language does not necessarily modulate perception of the speech-to-song illusion in that language. Initially this may seem to contradict prior research suggesting that language history can modulate the speech-to-song illusion, with song perception diminished in tone-language speakers (Jaisin et al., 2016) but enhanced for stimuli drawn from unfamiliar languages (Castro et al., 2018), especially difficult-to-pronounce languages (Margulis et al., 2015). Our findings, however, do not necessarily contradict the findings of Margulis et al. (2015) or Castro et al. (2018), as our participants were all relatively proficient speakers of English, even when English was not their dominant language. Instead, our findings suggest that language proficiency may not have an effect above a particular proficiency threshold. We do, however, find that monolinguals demonstrate increased response bias relative to bilinguals, i.e. they show a tendency to produce higher musicality ratings regardless of repetition or stimulus category. Regarding possible effects of tone language experience, this is a finding in need of replication, given that this result was reported in a study which included only ten participants per group (Jaisin et al., 2016). Our stimuli could be useful for studying effects of tone language experience on perception of the song illusion, provided the participants are bilingual in English.

Overall, we find that although our paradigm can reliably measure individual differences in the strength of the speech-to-song illusion, none of the demographic predictors we measured (including musical training) were significantly associated with illusion strength. What factors, then, might drive individual differences in illusion strength? Two possibilities are musical aptitude and the ability to direct attention to pitch information in speech. These possibilities were investigated in Experiment 2.

3. Experiment 2

3.1. Introduction: investigating musical aptitude and selective attention to pitch as drivers of the song illusion

One factor driving individual differences in perception of the song illusion could be musical aptitude, i.e., musical abilities which vary between individuals independently of any effect of musical training (Kragness et al., 2020). The extent to which speech stimuli sound musical after repetition has been linked to a variety of stimulus characteristics, including the presence of a steady beat and tonal structure (Tierney et al., 2018b) as well as the flatness of pitch within syllables (Tierney et al., 2013; Falk et al., 2014). Detection of these characteristics may be more successful in participants with greater musical aptitude and more precise auditory perception, leading to more robust perception of the speech-to-song illusion. Another possibility is that individual differences in the ability to direct attention to pitch information help determine the strength of the illusion. Prior research has indicated a link between perception of the illusion and increased pitch salience, as shown via increased activation in pitch-sensitive cortical areas (Tierney et al., 2013). Participants who can more readily direct attention to pitch information in speech when instructed to do so, therefore, may more robustly perceive the illusion.

To investigate these ideas, in Experiment 2 we examined whether illusion strength was linked to performance on six tests of auditory perception: two tests each of musical aptitude, psychoacoustics, and dimension-selective attention (Table 1). The musical aptitude tests were the Beat Alignment Test (Iversen & Patel, 2008), which asks participants to judge whether a series of tone pips is aligned with the beat or shifted away from the beat of clips of music, and the Tonality Alignment Test, a novel test constructed for the current study which asks participants to judge whether a sung melody is aligned with a tonal grid or misaligned (by being either compressed or expanded). The psychoacoustic tests were pitch and amplitude rise time discrimination (using parameters taken from Kachlicka et al., 2019), as we reasoned that participants with more accurate

auditory perception may be better able to detect subtle cues to musicality such as slight differences in pitch contour flatness or subtle amplitude envelope cues to beat location. The dimension-selective attention tests (first introduced in Jasmin, Sun, & Tierney, 2021) presented participants with pairs of spoken words which varied orthogonally in relative pitch and relative loudness. Participants were either asked to attend to pitch, indicating which of the two words was higher, or attend to loudness, indicating which of the two words was louder. We predicted that performance on the attention-to-pitch test, but not the attention-to-loudness test, would be linked to the robustness of perception of the speech-to-song illusion.

Test	Source
Frequency Discrimination	Kachlicka et al., 2019
Rise Time Discrimination	Kachlicka et al., 2019
Attention to Amplitude	Jasmin et al., 2021
Attention to Pitch	Jasmin et al., 2021
Beat Alignment Test	Iversen & Patel, 2008
Tonality Alignment Test	New

Table 1. Summary of measures included as possible predictors of illusion strength.

3.2. Methods

3.2.1. Participants

Ninety-five (95) participants were initially recruited from the online recruitment service Prolific. Because catch trials were not included in the speech-to-song illusion perception test in this experiment, participant inclusion was based on performance exceeding at least 55% across both dimension-selective attention tests, the Beat Alignment Test, and the Tonality Alignment Test. Eighty-one (81) participants passed these criteria. 54 of these identified as female, while 27 identified as male. A power analysis conducted in GPower revealed that this sample size resulted in a power of 0.80 to detect a correlation of $r = 0.3$. Participants reported a mean age of 30.8 years (SD 10.9, range 18-67). 79 of the participants reported that their dominant language was English, while 2 participants reported other dominant languages (1 Hindi and 1 Dutch). 67 participants reported being monolingual while 14 participants reported being bilingual (i.e. being able to speak at least two languages). Participants reported a mean of 5.30 years of musical training (std 7.71, range 0-36).

Participants were tested using the Gorilla platform for online testing (Anwyl-Irvine et al., 2020). The entire experiment lasted approximately 45 minutes. All participants completed the tasks in the same order, as follows: attention-to-amplitude, attention-to-pitch, song illusion experiment, Tonality Alignment Test, Beat Alignment Test, frequency discrimination, amplitude rise time discrimination. Participants were allowed to rest between tasks, but these breaks were optional rather than mandated. Study procedures were approved by the ethics board of the Department of Psychological Sciences at Birkbeck College.

3.2.2. Song Illusion Experiment

The song illusion experiment was conducted in a manner almost identical to that of Experiment 1, the only difference being that the catch trials were not included. Instead, the first four trials of the test consisted of practice items, i.e. 8 repetitions of phrases spoken by the same talkers included in the main song illusion dataset. (None of these four practice stimuli were drawn from the main stimulus set, and so cannot be defined as strictly being included among either the illusion or control stimuli. However, the first author, who assembled the original stimulus set (Tierney et al., 2013), perceived two of them as clearly transforming into song and two of them as continuing to be perceived as speech when repeated.) The inclusion of these items was meant to help participants get accustomed to rating the musicality of stimuli; the ratings of these stimuli were not analyzed.

3.2.3. Beat Alignment Test

This beat perception test used 22 stimuli drawn from Iversen & Patel (2008). These were instrumental musical excerpts from a broad variety of musical genres, including rock, classical, and jazz, onto which a sequence of 1 KHz 100 ms pure tones was superimposed. These tones were either aligned with the beat or shifted away by 25% of the inter-beat-interval. Participants were asked to indicate whether the tones were aligned with the beat or shifted away from the beat (by pressing buttons marked “on the beat” or “off the beat”). As soon as participants selected their response, the next stimulus was presented. A difference between this test and that reported in Iversen & Patel (2008) is that only a short excerpt of each stimulus was presented (long enough to contain seven

beats, average duration 3.81 seconds), to keep the experiment relatively brief. Task performance was summarized by calculating sensitivity (d'). Stimuli from this test are available at <https://osf.io/hegxs/>.

3.2.4. Tonality Alignment Test

To assess participants' ability to determine whether a stream of sung pitches was aligned with a tonal template, we used a new test devised for this study, the Tonality Alignment Test. Stimuli were drawn from the DSD100 (Liutkus et al., 2016), a dataset of 100 publicly available music recordings in which each component instrument can be downloaded separately. A single short portion of the vocal track from each of 22 different recordings was extracted. The duration of these musical passages ranged from 9.5 to 24.6 seconds (mean 15.7, SD 3.9). These 22 stimuli were then divided into three different sets that underwent three different types of pitch morphing in Praat (Boersma & Weenink, 2021). For 6 of the stimuli, the pitch intervals were expanded by 30% on a time-point-by-time-point basis. This was done by first extracting the fundamental frequency (F0) of each phrase using Praat (default settings, with a time step of 0.01 seconds), then comparing the ratio between the F0 value at each time point and the median F0 across all time points. The F0 of each time point was then adjusted to be equal to 1.3 times its original distance from the median F0, in semitones. (Previous perceptual research suggests that the semitone scale is more relevant to the perception of pitch in speech than the Hz scale [Nolan, 2003].) For 5 of the stimuli, F0 contours were contracted by 30% in a similar manner: the F0 of each time point was adjusted to be equal to 0.7 times its original distance from the median F0. For 11 of the stimuli, the F0 values were set to be equivalent to their original values. However, to ensure that any distortions due to the F0 manipulations were present across all stimuli, the "unaltered" stimuli were constructed by first expanding the F0 contours by 30% and then contracting them back to their original values.

Participants were told they would hear 22 melodies, some of which would be in tune, and some of which would not be in tune. They were asked to indicate whether the melody they heard was in tune (by pressing a button marked "in tune") or out of tune (by pressing a button marked

“out of tune”). Task performance was summarized by calculating sensitivity (d'). Stimuli from this test are available at <https://osf.io/hegxs/>.

3.2.5. Self-assessment of Music Perception Skills

Because the Tonality Alignment Test is a novel measure, it is of interest to establish its validity by comparing performance on this test to self-assessments of tonality perception ability. In addition, we were interested in comparing the validity of the Tonality Alignment Test to that of the Beat Alignment Test, an already established measure in the field. Participants were asked to indicate on a scale from 1 to 7 the extent to which they agreed with each of four statements, with 1 indicating “not at all” and 7 indicating “agree completely”. These statements were: 1) “I can tell when people sing or play out of tune”, 2) “When I sing I have no idea whether I’m in tune or not”, 3) “I can tell when people sing or play out of time with the beat”, and 4) “When I clap to music I have no idea whether I’m on the beat or not”. A composite measure of self-assessed tonality perception was calculated by averaging the response to question 1 with the inverted response to question 2. Similarly, a composite measure of self-assessed beat perception was calculated by averaging the response to question 3 with the inverted response to question 4.

3.2.6. Dimension-Selective Attention Tests

The ability to attend to individual acoustic dimensions in speech was assessed with a pair of tests first presented in Jasmin et al. (2021). Stimuli in this test are drawn from a pair of recordings of sentence fragments which are identical lexically but differ in the position of word emphasis: “STUDY music” and “study MUSIC”. These recordings were morphed onto one another using STRAIGHT (Kawahara & Irino, 2005) so that the extent to which individual acoustic dimensions resembled one versus the other recording could be precisely controlled, while all other acoustic characteristics were set to be constant across stimuli. For these tests, a 4 X 4 grid of stimuli was constructed in which pitch and amplitude varied in the extent to which they resembled the pattern in the initial-emphasis recording (“STUDY music”) versus the final-emphasis recording (“study MUSIC”). Specifically, the pitch levels used were 0%, 33%, 67%, and 100%, where 0% indicates patterns identical to the initial-

emphasis recording, 100% indicates patterns identical to the final-emphasis recording, and 33% and 67% indicate intermediate values. Similarly, the amplitude levels used were 0%, 33%, 67%, and 100%. To summarize, a set of 16 stimuli were constructed in which the extent to which pitch patterns versus amplitude patterns implied the existence of initial versus final word emphasis was independently varied (See Figure S4 for a schematic of the stimulus design).

On each trial, participants were presented with a single stimulus. The two tests involved presentation of stimuli drawn from the same set of 16 possible stimuli but differed in the instructions given to participants. For the Attention to Pitch test, they were asked to press a button to indicate whether the first or second word was higher in pitch, ignoring any differences in amplitude. For the Attention to Amplitude test, they were asked to press a button to indicate whether the first or second word was louder, ignoring any differences in pitch. Performance on each test was summarized as portion correct. For each test 40 trials were presented, with each trial pseudorandomly drawn from the set of 16 possible stimuli (a single randomization was performed for each test and used across all participants). This test has previously been shown to be sensitive to individual differences in dimensional salience: Mandarin speakers were found to display better performance on the Attention to Pitch test and worse performance on the Attention to Amplitude test, compared to native English speakers, suggesting that their experience speaking a tone language led to increased pitch salience across languages (Jasmin et al. 2021). Stimuli from these tests are available at <https://osf.io/hegxs/>.

3.2.7. Psychoacoustic Discrimination

Two tests were conducted to examine the precision of participants' perception of frequency and amplitude rise time. Each test consisted of a single 2-down 1-up adaptive staircase design. In each test, stimuli were drawn from a continuum of 101 stimuli, representing 101 different levels of the target acoustic continuum (i.e. either frequency or rise time). In each trial, participants were presented with 3 sounds, with either the first or the third being different from the other two, which

were identical. The two identical sounds always corresponded to level 1 of the target continuum. Participants were asked to indicate which of the three sounds was different by pressing either a button labelled “1” or a button labelled “3”. Initially, the target sound level (in other words, the level of the different sound) was set at 50. After every two correct responses, the target level decreased, becoming more similar to the comparison level, while after every incorrect response, the target level increased, becoming more dissimilar to the comparison level. The size of these level increases/decreases changed across the block, becoming smaller after each subsequent “reversal” or inflection point (in other words, two correct responses following a set of incorrect responses, or one incorrect response following a set of at least two correct responses). These step sizes were 10, 5, 2, 1, 1, 1, and 1 before the first through seventh reversals, respectively. The block continued until 75 trials were presented or participants reached the seventh reversal, whichever came first.

The stimulus continua for the two tests were constructed as follows. For the frequency discrimination test, the comparison stimulus was always presented at a fundamental frequency (F0) of 330 Hz, while the target stimulus ranged from 330.3 to 360 Hz in 100 equal linear steps. Frequency discrimination stimuli were 500-ms four-harmonic complex tones with equal amplitude across harmonics with a 0.015 linear amplitude ramp at the beginning and end to avoid perception of transient clicks. For the amplitude rise time test, the comparison stimulus was always presented with a linear amplitude ramp at the beginning of the stimulus of duration 15 ms, while the target stimulus ranged from a rise time of 17.8 to 300 ms in 100 equal linear steps. Rise time stimuli were 500-ms four-harmonic complex tones with equal amplitude across harmonics and an F0 of 330 Hz. Thresholds for each test were calculated as the mean of the target stimulus levels at the second through final reversals.

Any variables that were shown to be non-normally distributed according to a Jacque-Bera test were transformed prior to analysis. A rau transformation was used for the attention-to-pitch test and beat and tonality self-assessments, a log transformation was used for the frequency and rise time discrimination thresholds, and years of musical training were converted to ranks.

3.3. Results

As in Experiment 1, across participants the song illusion was very reliably experienced, as seen in the strong interaction between repetition and stimulus set ($F(7,560) = 85.9$, $p < 0.001$). There was once again a main effect of repetition ($F(1,560) = 53.9$, $p < 0.001$) and of stimulus set ($F(1,80) = 188.5$, $p < 0.001$).

To test the validity of the tonality perception and beat perception tests, performance on these measures was correlated with self-ratings of tonality perception (the ability to detect whether one's own singing and others' singing is in tune) as well as self-ratings of beat perception (the ability to detect whether one's own singing/playing and others' singing/playing is on the beat). Beat perception performance was correlated with beat perception self-rating ($r(79) = 0.44$, $p < 0.001$) but not tonality perception self-rating ($r(79) = 0.21$, $p > 0.05$). Tonality perception performance was correlated with both tonality perception self-rating ($r(79) = 0.40$, $p < 0.001$) and beat perception self-rating ($r(79) = 0.31$, $p < 0.01$). Figure 5 displays the relationship between tonality and beat perception performance and tonality and beat perception self-rating.

Figure 5. Upper left panel, relationship between tonality perception self-assessment and performance on tonality alignment test. Upper right panel, relationship between beat alignment perception self-assessment and performance on tonality alignment test. Lower left panel, relationship between tonality perception self-assessment and performance on beat alignment test. Lower right panel, relationship between beat perception self-assessment and performance on beat alignment test.

To investigate whether performance on the musical aptitude and auditory perception tests was associated with musical experience, Pearson correlations were used to investigate the relationship between musical experience and these measures. Years of musical training were weakly correlated with performance on the attention to pitch test ($r(79) = 0.26$, $p = 0.017$), but not with performance on the attention to amplitude test ($r(79) = -0.03$, $p > 0.1$). Years of musical training

were correlated with tonality perception performance ($r(79) = 0.30$, $p = 0.007$) and weakly correlated with beat perception performance ($r(79) = 0.22$, $p = 0.043$). Finally, years of training were not correlated with either discrimination of frequency ($r(79) = -0.07$, $p = 0.534$) or of amplitude rise time ($r(79) = -0.12$, $p = 0.275$). Note that the correlations between musical training and both beat perception and attention to pitch would not survive Bonferroni correction for multiple comparisons and should therefore be interpreted with caution. These weak-to-nonsignificant correlations underscore the distinction between musical aptitude and musical training as measures of musical ability (Kragness et al., 2020).

Pearson correlations were used to investigate the relationship between musical aptitude and experience and the strength of perception of the speech-to-song illusion, which was operationalised as the difference in the average ratings of illusion and control stimuli after the eighth repetition (i.e., illusion strength in Experiment 1). Years of musical training were not correlated with illusion strength, replicating a similar result in Experiment 1 ($r(79) = 0.150$, $p > 0.1$). Psychophysical discrimination thresholds were also not correlated with illusion strength, including frequency discrimination ($r(79) = 0.174$, $p > 0.1$) and amplitude rise time discrimination ($r(79) = -0.154$, $p > 0.1$). However, illusion strength was correlated with both perception of the spectral alignment of sung pitches with a tonal grid ($r(79) = 0.303$, $p < 0.01$) and perception of the temporal alignment of a tone sequence with the beat of a musical passage ($r(79) = 0.380$, $p < 0.001$). Illusion strength was also correlated with the ability to direct attention to pitch within spoken phrases ($r(79) = 0.279$, $p < 0.01$) but not with the ability to direct attention to amplitude within spoken phrases ($r(79) = 0.101$, $p > 0.1$). Figure 6 displays the relationship between illusion strength and musical aptitude. (Note that, although we report uncorrected p-values here, these relationships would survive Bonferroni correction for multiple comparisons.)

Figure 6. Left panel, relationship between attention to pitch and final illusion-control rating difference. Middle panel, relationship between tonality alignment perception and final illusion-control rating difference. Right panel, relationship between beat alignment perception and final illusion-control rating difference.

To examine whether these predictors explain independent variance in illusion strength, backwards linear regression was used, starting with a full model containing all predictors and removing predictors based on F probability until only significant predictors remained. The resulting model explained 22% of variance in illusion strength ($F(3,77) = 7.27, p < 0.001$). Predictors contained in the model included frequency discrimination threshold (standardized beta = 0.215, $t = 2.105, p < 0.05$), beat perception (standardized beta = 0.301, $t = 2.792, p < 0.01$) and tonality perception (standardized beta = 0.228, $t = 2.084, p < 0.05$). See Figure 7 for a depiction of predicted versus actual illusion strength values across participants. Note that the direction of the relationship between frequency discrimination is that individuals with less precise frequency discrimination perceive the illusion more strongly. Since this effect is only evident in the regression, while frequency discrimination and illusion strength do not correlate, we do not interpret this relationship further.

Figure 7. Scatterplot displaying predicted (x-axis) versus actual (y-axis) illusion strength values across participants.

3.4. Discussion

As in Experiment 1, we found that degree of musical training was not linked to robustness of perception of the speech-to-song illusion. However, we found that musical aptitude—including beat and tonality perception—as well as the ability to direct attention to the pitch of speech were linked to illusion strength. This result suggests that a key factor driving the song illusion is perceptual sensitivity to the degree of musical structure present in speech.

The strongest predictor of the strength of the speech-to-song illusion was the ability to judge whether a sequence of tones was temporally aligned with the beat of music. The presence of an isochronous beat is one of the most reliable characteristics distinguishing speech from music cross-culturally (Savage, Brown, Sakai, & Currie, 2015). Moreover, the presence of a steady pulse has also been shown to be a useful feature when constructing automatic classifiers of speech versus music (Scheirer & Slaney, 1997). Listeners may, therefore, be broadly aware that song tends to contain an isochronous beat while speech does not, and use the presence or absence of a musical beat as relevant evidence when categorizing a stimulus as speech versus song. Indeed, we have previously shown that within the stimulus set used in the current study, stimuli with more isochronous beats (as extracted via a computational model of beat perception) are perceived as more musical after repetition (Tierney et al., 2018b). Beat perception is not limited to musicians, but rather is a broadly

present skill in the general population: although neural entrainment to musical beats is enhanced by musical training, it is clearly present in non-musicians (Doelling & Poeppel, 2015). There is some evidence that beat perception emerges very early in development: infants coordinate their movements with the musical tempo based on musical pulse clarity (Zentner & Eerola, 2010) and can integrate auditory and somatosensory/proprioceptive information when deciding which notes carry musical beats (Phillips-Silver & Trainor, 2005). Nevertheless, although basic competence in beat perception is widespread, there are large individual differences in this ability (Müllensiefen, Gingras, Musil, & Stewart, 2014; Tranchant, Vuvan, & Peretz, 2016; Tranchant, Lagrois, Bellemare, Schultz, & Peretz, 2021), and impaired beat perception may limit individuals' ability to detect musical structure in speech and other non-musical stimuli.

The ability to assess whether pitch sequences align with scale structure was another significant predictor of the strength of the speech-to-song illusion. The use of discrete scales is another characteristic that distinguishes speech from music cross-culturally (Savage, Brown, Sakai, & Currie, 2015). Listeners may, as a result, take the presence of tonality as evidence that a stimulus should be categorized as song rather than speech. Tonality perception, like beat perception, is widespread in the general population: Western European musicians and non-musicians show a preference for scale structure in melodic sequences (Cross, Howell, & West, 1983). Our findings, however, suggest that not everyone is equally able to detect the presence of scale structure, and that individual differences in this ability help determine variability in the robustness of musicality perception in non-musical stimuli. The correlation between tonality perception and speech-to-song perception (as well as with self-assessment of tonality perception skills) demonstrates the validity of this novel measure, which may be useful in a variety of future studies of tonality perception, including investigations of the relationship between musical and language skills, studies of the effects of musical training, and research on the development of musical abilities.

We find that the ability to direct attention to pitch in speech is linked to the robustness of perception of the speech-to-song illusion. This relationship, however, does not reflect the influence

of general attentional skills, as the ability to direct attention to amplitude in speech was not linked to perception of the illusion. This finding aligns with the proposal by Deutsch et al. (2011) that perception of musical characteristics in speech is inhibited by default because pitch salience is down-regulated during speech perception, but that repetition of speech disinhibits attention to pitch, making possible the detection of musical characteristics in stimuli. A number of findings support the possibility that variation in other acoustic dimensions can interfere with pitch perception, possibly by diverting attention away from pitch. For example, Warrier & Zatorre (2002), Allen & Oxenham (2014), and Caruso & Balaban (2014) showed that the presence of variation in timbre can interfere with pitch perception, while Russo, Vuvan, & Thompson (2019) found that vowel content can interfere with relative pitch perception in speech stimuli. The idea that perceiving the speech-to-song illusion can lead to increased pitch salience is supported by the finding that pitch-sensitive cortical areas increase in activation when illusion stimuli are repeated (Tierney et al., 2013). This increase in salience may be greater in individuals who can more readily direct attention to pitch, leading to more robust perception of the illusion.

4. General discussion

Writing in 55 BCE, Cicero remarked in *De Oratore* that “even in speaking there may be a concealed kind of music.” Over 2,000 years later, the discovery of the song illusion provided compelling evidence for this claim (Deutsch et al., 2011). 20 years after that discovery, the current work on individual differences in the experience of the illusion suggests that Cicero’s phrase should be updated: “even in speaking there may be a concealed kind of music, *though only some can hear it.*” Specifically, our results suggest that the ability to focus on and extract several different kinds of auditory and musical cues is linked to the strength of the song illusion. These results align with prior work on other auditory illusions, which has generally shown that the salience of (and ability to perceive) different stimulus characteristics is linked to the influence these characteristics have on the final percept. In the “Yanny/Laurel” illusion, for example, whether individuals perceive “Yanny” versus “Laurel” is linked to their degree of prior exposure to lower versus higher auditory

frequencies in their environment (Gwilliams & Wallisch, 2020). Similarly, the magnitude of the McGurk effect correlates with lipreading skill (Strand, Cooperman, Rowe, & Simenstad, 2014; Brown et al., 2018), as well as the amount of time individuals spend fixating their gaze on a talker's mouth during audiovisual speech perception (Gurler et al., 2015). The song illusion, therefore, like other illusions, may involve the weighting of multiple potential cues to which individuals are differentially sensitive.

Our findings support theoretical models of the song illusion which suggest that perception of the illusion requires extraction of musical characteristics from the stimuli (Deutsch et al. 2011, Tierney et al. 2018b). The need to extract musical characteristics from the stimuli could partly explain the increase in musicality perception with repetition. Prior work has shown that perception of pitch interval size is surprisingly poor after a single presentation of a melody, and that melody repetition is necessary before participants achieve consistent interval perception (Deutsch, 1979). The idea that the increase in song illusion strength with repetition is driven by the time it takes to extract musical characteristics from the stimuli is supported by the finding that the increase in musicality with repetition is similarly sized when speech stimuli and synthesized complex tone stimuli with identical contours are presented (Tierney et al., 2018a).

Our findings constrain theories of the song illusion in significant ways. The finding that individual differences in perception of the illusion are linked to musical aptitude is not predicted by models which suggest that the illusion is driven by deactivation of lexical nodes and continued activation of syllable nodes (Castro et al., 2018). While our results indicate that Node Structure Theory does not provide a *sufficient* explanation of the song illusion, our results do not imply that deactivation of lexical nodes is *not* relevant to perception of the illusion. Our experiment was not designed to test the theory that individual differences in the speed or robustness of lexical satiation drive individual differences in perception of the illusion, an idea worth exploring in future research.

Here we find that musical aptitude predicts the strength of the song illusion, but musical training does not. This may seem somewhat contradictory, inasmuch as the main goal of musical

training is to boost musical skills. Indeed, we found that greater musical training was linked to better tonality and beat perception, as well as better ability to direct attention to pitch in speech. This result aligns with previous findings showing musicians demonstrate more precise mistuning perception (Hutchins, Roquet, & Peretz, 2012; Larrouy-Maestri, 2018) and stronger neural entrainment to musical beats (Doelling & Poeppel, 2015; but see Hickey et al., 2020). However, the lack of a relationship between musical training and illusion perception tracks with other recent work finding that cognitive/perceptual abilities are more strongly linked to musical aptitude than to degree of musical training (Swaminathan, Schellenberg, & Khalil, 2017; Swaminathan & Schellenberg, 2019). One possibility is that strong individual differences in musical aptitude exist prior to music training, and in fact help determine whether individuals begin (and stick with) musical training (Kragness, Swaminathan, Cirelli, & Schellenberg, 2020). Overall, our results suggest that individual differences in the perception of musicality in speech are tied to musical skills that may be relatively unaffected by musical training, possibly reflecting instead stable traits tied to genetic differences (Niarchou et al., 2021).

Although our finding of a relationship between musical aptitude and illusion strength represents an initial step towards understanding the factors driving individual differences in perception of the speech-to-song illusion, the results of our linear regression explained only 22% of the variance in illusion strength. It is likely, therefore, that other major factors driving individual differences in perception of the illusion remain to be uncovered. One possibility is that individuals who have a tendency to focus on the F0 rather than spectral information when assessing pitch contour may more robustly perceive the illusion. Research on perception of the missing fundamental illusion has shown that there exist large individual differences in the extent to which listeners rely on F0 versus spectral information when judging the interval between two notes (Schneider et al., 2005; Ladd et al., 2013). These divergent listening strategies may have consequences for the ability to extract musical information from sound sequences. Although pitch contour can be as easily extracted from inharmonic compared to harmonic sounds, the ability to extract F0 information

facilitates judgment of exact pitch intervals and tonality (McPherson & McDermott, 2018) as well as memory for pitch (McPherson & McDermott, 2020). As a result, spectrally-biased listeners may have difficulty detecting musical regularities latent in the speech-to-song illusion stimuli. Another possible factor driving individual differences in perception of the illusion is the extent to which individuals find music rewarding or absorbing: listeners with a greater hedonic or absorptive response to music may be more keen to seek out musical characteristics in non-musical stimuli. This hypothesis could be tested in future work using the Barcelona Musical Reward Questionnaire (Mas-Herrero, Marco-Pallares, Lorenzo-Seva, Zatorre, & Rodriguez-Fornells, 2013) or the Absorption in Music Scale (Sandstrom and Russo, 2013). Answers to specific questions within these scales may prove especially useful for helping predict susceptibility to the song illusion, e.g., “I like to find patterns in everyday sounds” from the Absorption in Music Scale.

In conclusion, we find that there are strong, reliable individual differences in the tendency to perceive certain repeated spoken phrases as sung. These differences seem to be largely independent of demographic characteristics, including language and musical experience, but are linked to individual differences in specific perceptual and musical skills. This finding suggests that individual differences in auditory abilities may strongly affect perceptual categorization, influencing not only categorization judgments within a domain (such as speech perception; Jasmin, Dick, Holt, & Tierney, 2020) but even perceptual categorization between domains (here, speech versus music). A deeper understanding of these individual differences would help researchers predict whether a given listener will experience a particular spoken phrase as sung when repeated. This in turn would provide a powerful tool for exploring the cognitive and brain mechanisms underlying selective neural responses to speech and music (Norman-Hagnere et al., 2015; Zuk et al., 2020), by using the same physical stimuli to elicit categorically different perceptual experiences.

5. References

- Allen, E., & Oxenham, A. (2014). Symmetric interactions and interference between pitch and timbre. *JASA*, *135*, 1371-1379.
- Anwyl-Irvine, A., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. (2020). Gorilla in our midst: an online behavioral experiment builder. *Behavior Research Methods*, *52*, 388-407.
- Boersma, P., & Weenink, D. (2021). Praat: doing phonetics by computer [Computer program]. Version 6.1.42, retrieved 15 April 2021 from <http://www.praat.org/>
- Bosker, H. (2018). Putting Laurel and Yanny in context. *JASA*, *144*, EL503-EL508.
- Brennan, D., & Stevens, C. (2002). Specialist musical training and the octave illusion: analytical listening and veridical perception by pipe organists. *Acta Psychologica*, *109*, 301-314.
- Brown, V., Hedayati, M., Zanger, A., Mayn, S., Ray, L., Dillman-Hasso, N., & Strand, J. (2018). What accounts for individual differences in susceptibility to the McGurk effect? *PLoS ONE*, *13*, e0207160.
- Caruso, V. C., & Balaban, E. (2014). Pitch and timbre interfere when both are parametrically varied. *PLoS ONE*, *9*(1), e87065.
- Castro, N., Mendoza, J., Tampke, E., & Vitevitch, M. (2018). An account of the speech-to-song illusion using Node Structure Theory. *PLoS ONE*, *13*, e0198656.
- Cicero, M. *On Oratory and Orators. With Notes Historical and Explanatory. A New Edition, Carefully Revised and Corrected. In Two Volumes 1808*. Translation by William Guthrie.

Cienkowski, K., & Carney, A. (2002). Auditory-visual speech perception and aging. *Ear and Hearing*, 23, 439-449.

Craig, J. (1979). The effect of musical training and cerebral asymmetries on perception of an auditory illusion. *Cortex*, 15, 671-677.

Cross, I., Howell, P., & West, R. (1983). Preferences for scale structure in melodic sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 444-460.

Davidson, B., Power, R., & Michie, P. (1987). The effects of familiarity and previous training on perception of an ambiguous musical figure. *Perception and Psychophysics*, 41, 601-608.

Deutsch, D. (1975). Musical illusions. *Scientific American*, 233, 92-105.

Deutsch, D. (1979). Octave generalization and the consolidation of melodic information. *Canadian Journal of Psychology*, 33, 201-205.

Deutsch, D., Henthorn, T., & Lapidis, R. (2011). Illusory transformation from speech to song. *JASA*, 129, 2245-2252.

Ding, N., Patel, A. D., Chen, L., Butler, H., Luo, C., & Poeppel, D. (2017). Temporal modulations in speech and music. *Neuroscience & Biobehavioral Reviews*, 81, 181-187.

Doelling, K., & Poeppel, D. (2015). Cortical entrainment to music and its modulation by expertise. *PNAS*, 112, E6233-E6242.

Falk, S., Rathcke, T., & Dalla Bella, S. (2014). When speech sounds like music. *Journal of Experimental Psychology: Human Perception and Performance*, 40, 1491-1506.

Gurler, D., Doyle, N., Walker, E., Magnotti, J., & Beauchamp, M. (2015). A link between individual differences in multisensory speech perception and eye movements. *Attention, Perception, and Psychophysics*, 77, 1333-1341.

Gwilliams, L., & Wallisch, P. (2020). Immediate ambiguity resolution in speech perception based on prior acoustic experience. *PsyArXiv*. doi:10.31234/osf.io/yptfr

Hickey, P., Merseal, H., Patel, A. D., & Race, E. (2020). Memory in time: Neural tracking of low-frequency rhythm dynamically modulates memory formation. *Neuroimage*, 213, 116693

Hutchins, S., Roquet, C., & Peretz, I. (2012). The vocal generosity effect: how bad can your singing be? *Music Perception*, 30, 147-159.

Iversen, J. R., & Patel, A. D. (2008). The Beat Alignment Test (BAT): Surveying beat processing abilities in the general population. In K. Miyazaki, M. Adachi, Y Hiraga, Y Nakajima, & M. Tsuzaki (Eds.), *Proceedings of the 10th International Conference on Music Perception & Cognition (ICMPC10)* (CD-ROM; pp. 465–468). Adelaide, Australia: Causal Productions.

Jaisin, K., Suphanchaimat, R., Candia, M., & Warren, J. (2016). The speech-to-song illusion is reduced in speakers of tonal (vs. non-tonal) languages. *Frontiers in Psychology*, 7, 662.

Jarque, C., & Bera, A. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, 6, 255-259.

Jasmin, K., Dick, F., Holt, L., & Tierney, A. (2020). Tailored perception: individuals' speech and music perception strategies fit their perceptual abilities." *Journal of Experimental Psychology: General*, 149, 914-934.

Jasmin, K., Sun, H., & Tierney, A. (2021). Effects of language experience on domain-general perceptual strategies. *Cognition*, 206, 104481.

Kachlicka, M., Saito, K., & Tierney, A. (2019). Successful second language learning is tied to robust domain-general auditory processing and stable neural representation of sound. *Brain and Language*, 192, 15-24.

Kawahara, H., & Irino, T. (2005). Underlying principles of a high-quality speech manipulation system STRAIGHT and its application to speech segregation. In P. Divenyi (Ed.), *Speech separation by humans and machines* (pp. 167–180). Boston, MA: Kluwer Academic Publishers.

Kragness, H., Swaminathan, S., Cirelli, L., & Schellenberg, G. (2020). Individual differences in musical ability are stable over time in childhood. *Developmental Science*. doi:10.1111/desc.13081

Ladd, R., Turnbull, R., Browne, C., Caldwell-Harris, C., Ganushchak, L., Swoboda, K., Woodfield, V., & Dediu, D. (2013). Patterns of individual differences in the perception of missing-fundamental tones. *Journal of Experimental Psychology: Human Perception and Performance*, 39, 1386-1397.

Larrouy-Maestri, P. (2018). I know it when I hear it: on listeners' perception of mistuning. *Music & Science*, 1, 2059204318784582.

Liutkus, A., Stöter, F., Rafii, Z., Kitamura, D., Rivet, B., Ito, N., Ono, N., & Fontecave, J. (2016). The 2016 signal separation evaluation campaign. In *International Conference on Latent Variable Analysis and Signal Separation*, pp. 323-332. Springer, Cham.

MacKay, D., Wulf, G., Yin, C., & Abrams, L. (1993). Relations between word perception and production: new theory and data on the verbal transformation effect. *Journal of Memory and Language*, 32, 624-646.

Mallick, D., Magnotti, J., & Beauchamp, M. (2015). Variability and stability in the McGurk effect: contributions of participants, stimuli, time, and response type. *Psychon Bull Rev*, 22, 1299-1307.

Margulis, E., Simchy-Gross, R., & Black, J. (2015). Pronunciation difficulty, temporal regularity, and the speech-to-song illusion. *Frontiers in Psychology*, 6, 48.

Mas-Herrero, E., Marco-Pallares, J., Lorenzo-Seva, U., Zatorre, R., & Rodriguez-Fornells, A. (2013). Individual differences in music reward experiences. *Music Perception*, 31, 118-138.

McPherson, M., & McDermott, J. (2018). Diversity in pitch perception revealed by task dependence. *Nature Human Behaviour*, 2, 52-66.

McPherson, M., & McDermott, J. (2020). Time-dependent discrimination advantages for harmonic sounds suggest efficient coding for memory. *PNAS*, 117, 32169-32180.

Müllensiefen, D., Gingras, B., Musil, J., & Stewart, L. (2014). The musicality of non-musicians: an index for assessing musical sophistication in the general population. *PLoS one*, 9(2), e89642.

Mullin, H., Norkey, E., Kodwani, A., Vitevitch, M., & Castro, N. (2021). Does age affect perception of the Speech-to-Song illusion? *PLoS one*, 16, e0250042.

Niarchou, M., Gustavson, D., Sathirapongsasuti, J., Anglada-Tort, M., Eising, E., Bell, E., McArthur, E., Straub, P., The 23andMe Research Team, McAuley, J., Capra, J., Ullén, F., Creanza, N., Mosing, M., Hinds, D., Davis, L., Jacoby, N., & Gordon, R. Unravelling the genetic architecture of musical rhythm: a large-scale genome-wide association study of beat synchronization.

bioRxiv 836197; doi: <https://doi.org/10.1101/8361>

Nolan, F. (2003). Intonational equivalence: an experimental evaluation of pitch scales. In *Proceedings of the 15th International Congress of Phonetic Sciences, Barcelona* (Vol. 39).

Norman-Haignere, S., Kanwisher, N. G., & McDermott, J. H. (2015). Distinct cortical pathways for music and speech revealed by hypothesis-free voxel decomposition. *Neuron*, 88(6), 1281-1296.

Patterson, R., Uppenkamp, S., Johnsrude, I., & Griffiths, T. (2002). The processing of temporal pitch and melody information in auditory cortex. *Neuron*, 36, 767-776.

Penagos, H., Melcher, J., & Oxenham, A. (2004). A neural representation of pitch salience in nonprimary human auditory cortex revealed with functional magnetic resonance imaging. *Journal of Neuroscience*, 24, 6810-6815.

Phillips-Silver, J., & Trainor, L. J. (2005). Feeling the beat: Movement influences infant rhythm perception. *Science*, 308, 1430–1430

Pilotti, M., & Khurshid, A. (2004). Semantic satiation effect in young and older adults. *Perceptual and Motor Skills*, 98, 999-1016.

Pilotti, M., Simcox, T., Baldy, J., & Schauss, F. (2011). Are verbal transformation sensitive to age differences and stimulus properties? *The American Journal of Psychology*, 124, 89-97.

Pressnitzer, D., Graves, J., Chambers, C., de Gardelle, V., & Egré, P. (2018). Auditory perception: Laurel and Yanny together at last. *Current Biology*, 28, R737-R759.

Rathcke, T., Falk, S., & Dalla Bella, S. (in press). Music to your ears: sentence sonority and listener background modulate the "speech-to-song illusion". *Music Perception*.

Rowland, J., Kasdan, A., & Poeppel, D. (2019). There is music in repetition: looped segments of speech and nonspeech induce the perception of music in a time-dependent manner. *Psychonomic Bulletin & Review*, 26, 583-590.

Russo, F., Vuvan, D., & Thompson, W. (2019). Vowel content influences relative pitch perception in vocal melodies. *Music Perception*, 37, 57-65.

Sandstrom, G. M., & Russo, F. A. (2011). Absorption in music: Development of a scale to identify individuals with strong emotional responses to music. *Psychology of Music*, 41(2), 216-228.

Savage, P., Brown, S., Sakai, E., & Currie, T. (2015). Statistical universals reveal the structures and functions of human music. *PNAS*, 112, 8987-8992.

Scheirer, E., & Slaney, M. (1997). Construction and evaluation of a robust multifeature speech/music discriminator. In *1997 IEEE International Conference in Acoustics, Speech, and Signal Processing* (pp. 1331– 1334).

Schneider, P., Sluming, V., Roberts, N., Scherg, M., Goebel, R., Specht, H., Dosch, H., Bleeck, S., Stippich, C., & Rupp, A. (2005). Structural and functional asymmetry of lateral Heschl's gyrus reflects pitch perception preference. *Nature Neuroscience*, *8*, 1241-1247.

Simchy-Gross, R., & Margulis, E. (2018). The sound-to-music illusion: repetition can musicalize nonspeech sounds. *Music and Science*, *1*, 1-6.

Smith, L. (1984). Semantic satiation affects category membership decision time but not lexical priming. *Memory and Cognition*, *12*, 483-488.

Strand, J., Cooperman, A., Rowe, J., & Simenstad, A. (2014). Individual differences in susceptibility to the McGurk Effect: links with lipreading and detecting audiovisual incongruity. *JSLHR*, *57*, 2322-2331.

Swaminathan, S., Schellenberg, G., & Khalil, S. (2017). Revisiting the association between music lessons and intelligence: training effects or music aptitude? *Intelligence*, *62*, 119-124.

Swaminathan, S., & Schellenberg, G. (2019). Musical ability, music training, and language ability in childhood. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *46*, 2340-2348.

Tierney, A., Dick, F., Deutsch, D., & Sereno, M. (2013). Speech versus song: multiple pitch-sensitive areas revealed by a naturally occurring musical illusion. *Cerebral Cortex*, *23*, 240-254.

Tierney, A., Patel, A., & Breen, M. (2018a). Repetition enhances the musicality of speech and tone stimuli to similar degrees. *Music Perception*, 35, 573-578.

Tierney, A., Patel, A., & Breen, M. (2018b). Acoustic foundations of the speech-to-song illusion. *Journal of Experimental Psychology: General*, 6, 888-904.

Toscano, J., & McMurray, B. (2010). Cue integration with categories: weighting acoustic cues in speech using unsupervised learning and distributional statistics. *Cognitive Science*, 34, 434-464.

Tranchant, P., Vuhan, D. T., & Peretz, I. (2016). Keeping the beat: A large sample study of bouncing and clapping to music. *PloS one*, 11(7), e0160178.

Tranchant, P., Lagrois, M. É., Bellemare, A., Schultz, B. G., & Peretz, I. (2021). Co-occurrence of Deficits in Beat Perception and Synchronization Supports Implication of Motor System in Beat Perception. *Music & Science*, 4, 2059204321991713.

Vanden Bosch der Nederlanden, C., Hannon, E., & Snyder, J. (2015a). Finding the music of speech: musical knowledge influences pitch processing in speech. *Cognition*, 143, 135-140.

Vanden Bosch der Nederlanden, C., Hannon, E., & Snyder, J. (2015b). Everyday musical experience is sufficient to perceive the speech-to-song illusion. *Journal of Experimental Psychology: General*, 2, e43-e49.

Vitevitch, M., Ng, J., Hatley, E., & Castro, N. (2020). Phonological but not semantic influences on the speech-to-song illusion. *Quarterly Journal of Experimental Psychology*. doi: 10.1177/1747021820969

Warren, R., & Gregory, R. (1958). An auditory analogue of the visual reversible figure. *The American Journal of Psychology*, 71, 612-613.

Warren, R., & Warren, R. (1966). A comparison of speech perception in childhood, maturity, and old age by means of the verbal transformation effect. *Journal of Verbal Learning and Verbal Behavior*, 5, 142-146.

Warner, C., & Zatorre, R. (2002). Influence of tonal context and timbral variation on perception of pitch. *Perception & Psychophysics*, 64, 198-207.

Zentner, M., & Eerola, T. (2010). Rhythmic engagement with music in infancy. *PNAS*, 107, 5768-5773.

Zuk, N. J., Teoh, E. S., & Lalor, E. C. (2020). EEG-based classification of natural sounds reveals specialized responses to speech and music. *NeuroImage*, 210, 116558.