When Poems Don't Meet our Expectations:

Effects of Slant Rhyme on Duration

Magda Andrews-Hoke

Advisor: Natalie Weber

LING 491: The Senior Essay

April 24, 2019

Acknowledgements

My sincerest thanks to my advisor Natalie Weber for all of their guidance and encouragement, to Raffaella Zanuttini and Hadas Kotek for their additional guidance, to Chris Geissler for his assistance with R and his kindness, to Jason Shaw for his advice concerning statistics, to Benjamin Glaser for guiding my interest in poetry and first introducing me to Breen (2018), to Ryan Bennett for introducing me to phonology and encouraging me to explore its overlap with my poetic interests, to my fellow linguistics majors James Lin, Will Merrill, Anelisa Fergus, Jay Pittman, Jisu Sheen, Noah Macey, and Rose Mintzer-Sweeney for their continued support and good cheer, to my friends Ryan Howzell, Kristine Chung, Raquel Sequeira, and many others for listening to me talk about rhyme and for their love and support, and to my mom, Agatha, my dad, Frank, and my brother, Levi for absolutely everything.

Contents

> Marry, I cannot show it in rhyme; I have tried: I can find out no rhyme to 'lady' but 'baby,' an innocent rhyme; for 'scorn,' 'horn,' a hard rhyme; for, 'school,' 'fool,' a babbling rhyme; very ominous endings: no, I was not born under a rhyming planet.
>
> – *Much Ado About Nothing*, Act V, Scene 2

1 Abstract

What do we know when we know how to read poetry aloud? This, like much of our ability to use language, is a form of inherent linguistic knowledge, often unexamined by those who use it. This thesis builds on the work of Breen (2018), which examines how readers audibly mark their internal knowledge of poetic structure. This thesis focuses particularly on how readers might realize rhyming structure by altering word durations. Breen (2018) finds that in rhyming couplets, the second of the two rhyme words is noticeably shorter in duration than the first. In section 2 the findings of Breen (2018) are expanded upon, after which a linguistic definition of rhyme is proposed. I conclude section 2 by hypothesizing that the rhyme word shortening effect that Breen finds may be due to the predictability of the second rhyme word, as predictability is known to decrease duration.

In section 3, the experiment that was used to test this hypothesis is described. In this experiment, participants read rhyming couplets aloud, some couplets containing exact rhyme, and some containing slant (non-exact) rhyme. In slant-rhymed couplets, the second rhyme word is unpredictable, and therefore, if the stated hypothesis is correct, the more unpredictable the second rhyme word is, the longer it should be, relative to an exact rhyme. The experiment first attempts to recreate Breen's initial finding of shorter word duration in the second rhyme word, and then investigates how the unpredictability of slant rhyme affects duration. In sections 4 and 5, the results of the experiment and the unanswered questions it leaves are presented.

Ultimately, this thesis seeks to explore the way in which poetic form falls into the category of English speakers' linguistic knowledge, as well as how predictability and unpredictability play a role in the writing and reading of poetry.

2 Introduction

2.1 Breen (2018) and rhyme shortening

This thesis finds its original inspiration in the work of Breen (2018). Breen recorded 18 participants reading the children's book *The Cat in the Hat*, and analyzed both word duration and inter-onset interval (the time between the beginning of one word's onset and the next's) to determine the levels of metrical hierarchy readers are cognizant of and verbally marking while reading aloud. One of Breen's finding – the one that is of particular relevance here – is that **readers tended to pronounce the second of two rhyming words with a shorter duration than the first.** Breen's finding provokes some interesting questions about how readers recognize and articulate rhyming structure.

Before delving into this finding and the unanswered questions associated with it, a brief overview of Breen's experiment and methodology is in order. *The Cat in the Hat* is written in anapestic tetrameter rhyming couplets – lines with four feet, each with a weak-weak-**strong** stress pattern:

(1) A: "Put me **down**!" said the **fish**.
This is **no** fun at **all**!
B: Put me **down**!" said the **fish**.
"I do **NOT** wish to **fall**." (Dr. Seuss 1957)

Every rhyme is exact and is monosyllabic, such as [ɑl] 'all' ∼ [fɑl] 'fall'. Breen follows Fabb and Halle (2008) in the use of a metrical grid to describe the metrical poetic structure. Fabb and Halle (2008), in turn, follow Hayes (1995) in their use of metrical stress theory, applying its methods to metrical poetry.

Metrical stress theory, in its most basic form, describes the way in which stresses are distributed across words, feet, phrases, and other linguistics domains. Hayes (1995) lays out several basic principles that are foundational for metrical stress theory. Three of these principles that are relevant to Breen's use of metrical stress theory are **culminativity, rhythmic distribution, and hierarchy**. **Culminativity** is the notion that stress is culminative, or that each word or phrase must contain only one strongest stress-bearing syllable. In poetry, culminativity also applies. Each metrical foot of, say, iambic pentameter has one strongest stress, although this does not necessarily align with the natural stress of the line's words. **Rhythmic distribution** describes the way in which stress is always distributed in a rhythmically alternating pattern. Stress clash (two stresses right next to each other) is remarkably dis-preferred cross-linguistically, while an alternation of stressed and unstressed syllables is common (such as can be found in many poetic meters). **Hierarchy** describes the way in which stress comes in multiple degrees. Languages usually have primary, secondary, and sometimes tertiary stresses, which are of hierarchically different degrees.

When applying these concepts to a metrical understanding of poetry, the application is not always directly stress-related. For example, a culminative stress does not always rest on the syllable pronounced with the most "stress-connoting" features (higher pitch, longer length, etc.). In (1), hierarchically tertiary stress occurs on the word "down" (both times), "no," and "NOT." Secondary stresses occur on both instances of "fish." Primary stresses occur on both "all" and "fall". These degrees of stress are not dependent on the phonological character of the individual words, but rather on their position in the larger metrical structure. While the usual use of metrical stress theory depends heavily on the native stress system of a language, the poetic use of metrical stress theory focuses on abstract meters, which create stress hierarchies not innate to the individual words in question. Whether these degrees of stress correspond directly to audible differences in pronunciation in terms of word duration and inter-onset interval is the question that Breen then takes up.

Within metrical stress theory, the use of metrical grids is common. Metrical grids are used to show the rhythmic distribution of different degrees of stress. For example, in the word *Mississippi*, there are four syllables and two stresses, one primary and one secondary. The metrical grid of the word [mɪ.sə.ˈsɪ.pi] is as follows:

(2)			x	
	x		x	
	x	x	x	x
	mɪ	sə	sɪ	pi

The height of the columns of x-marks represents the degree of stress on each syllable. Every syllable contains some kind of stress (or at least is the "head" of its own constituent: itself). The first and third syllables project to a higher level of stress, so they receive a second x-marking. The third syllable contains the culminative stress of the word, and therefore is the only syllable that projects to the third level.

Breen describes five different metric levels in *The Cat in the Hat*. She lays them out in a table similar to the metrical grid in (2), which can be seen below in Table 1.

In Table 1, the projecting levels are related to metrical "heads." Within one line of *The Cat in the Hat*, there are four anapests (a foot of the form unstressed-unstressed-**stressed**), two half lines, and one full line. The final words of these, since anapests and rhymes are right-headed (the right edge of the foot and line is most prominent), are projected to the next metrical level. In the metrical grid in Table 1, Breen differentiates five metrical levels: **1)** the word, **2)** the (anapestic) foot , **3)** the half-line (two feet), **4)** the line, and **5)** the couplet. Each asterisk marks words that are heads at the indicated metrical level. Parentheses mark those heads that will project to the next metrical level. Again, the length of the columns indicates the level of stress, although here the columns are oriented downward.

Table 1: Metrical levels, adapted from Breen (2018)

Metric Level	Put	me	down	said	the	fish	this	is	no	fun	at	all
1	*	*	*)	*	*	*)	*	*	*)	*	*	*)
2			*			*)			*			*)
3						*						*)
4												*
5												
Metric Level	Put	me	down	said	the	fish	I	do	not	wish	to	fall
1	*	*	*)	*	*	*)	*	*	*)	*	*	*)
2			*			*)			*			*)
3						*						*)
4												*)
5												*

As can be seen in this Table 1, the metrical levels indicate words, anapests, half-lines, full lines, and rhymed couplets. Breen finds that participants mark these different metrical levels using different phonetic indicators, including word duration and inter-onset interval. Inter-onset interval was used to indicate all five different metrical levels, meaning that participants took greater pauses between rhymed couplets than they did between lines, and greater pauses between lines than between half-lines, etc. Word duration generally only marked three levels of metrical hierarchy; heads on the third metrical level (the half-line level) were pronounced with longer duration than heads on the anapest or word level. However, heads on the half-line levels were not significantly different in length from those on the fourth metrical level (the line level).

Interestingly, though, Breen finds that the second rhyme word (the rhyme target) in a rhyming couplet was pronounced with a shorter duration than the first rhyme word (the rhyme prime). This means that in the rhyming lines in Table 1 above, [fɑl] 'fall' had a shorter duration than [ɑl] 'all', its rhyme prime. Why readers would pronounce a rhyme target with shorter duration than a rhyme prime is unclear. On the first three levels of metrical hierarchy, participants appeared to use *increased* word duration to emphasize the metrical heads. And yet, on the fifth and final metrical level, Breen finds that *decreased* word duration is used to mark the final and ultimate head of the poetic phrase.

This thesis will take up the question of *why* rhyme targets are shorter in duration than rhyme primes. In section 2.2, predictability will be examined as one possible explanation for this rhyme word shortening effect, but, first, section 2.2.1 will discuss the compelling reasons for taking up this question at all.

2.2 Reasons for rhyme word shortening

2.2.1 Does phrase-final shortening exist?

Something that makes Breen's finding of rhyme word shortening particularly interesting is the large body of research on phrase-final lengthening. Phrase-final lengthening is one of several effects that occur towards the end of an utterance, including declination (the lowering of the pitch over the course of the utterance) and creaky voice (Ladd 1984). Phrase-final lengthening is often called "pre-boundary lengthening" or "pre-pausal lenghtening" (Beckman and Edwards 1990; Klatt 1975; Cooper and Paccia-Cooper 1980). In this phenomenon, words at the right edge (i.e. at the end) of a syntactic, prosodic, or some other constituent are lengthened, compared to what their duration would be in initial or phrase-medial position. Word lengthening, therefore, is being used to mark either the end of a phrase, the beginning of a pause, or both.

The rhyme word shortening effect that Breen finds stands sharply in contrast to this. It is, in fact, essentially the reverse of phrase-final lengthening, given that durational shortening seems to be being used to mark the end of a poetic phrase. This phenomenon, if Breen's description is accurate, seems to be unique to the poetic domain of language.

If we are to account for the rhyme word shortening effect, we will have to look towards other explanatory factors than the basic intonational patterns of speech, which ostensibly stand opposed to this effect. As further possibilities are discussed, it should be kept in mind how particular this pattern seems to be to poetic utterances. Poetry seems to make use of the natural grammar of language in surprising ways, and this particular effect of phrase-final shortening is no exception.

2.2.2 Predictability and word duration

One possible reason for the rhyme word shortening effect that Breen observes is the predictability of rhyme words, which is not a common feature of natural speech. The rhyme target in a rhyming pair is highly predictable phonologically. As we are reading a rhyming poem, we already know almost everything about the rhyme target long before we have read it or said it out loud. Certainly in a book like *The Cat in the Hat*, where the rhymes are all exact and monosyllabic, predictability is extremely high.

It is well attested in a broad range of linguistic contexts that increased predictability leads to reduction in word duration, greater degrees of coarticulation, and other effects that generally decrease the distinctiveness of the phonetic signal (Blevins 2005; Hall 2018). This makes sense. The more obvious it is to both speaker and listener what word or phrase is about to be said, the less clearly articulated that word or phrase needs to be.

Hall (2018) distinguishes between two types of predictability. To begin with, some words are more predictable by virtue of their **frequency**: they simply occur often in the general corpus of speech. These words tend to be closed class words, such as adpositions, determiners, and pronouns. Additionally, Hall (2018) describes the way in which words can become predictable by virtue of their **context**. This type of predictability has little to do with the frequency of the actual words, and more to do with the phrasal, syntactic, and conversational context around them. For example, if someone says,"He's dead. He really kicked the bucket," the word *bucket* is highly predictable. This is not because the word *bucket* is so frequently used, but because once someone has said, "He's dead. He really kicked the [   ]," the likelihood that the next word will be *bucket* is extremely high. Both types of predictability create similar phonetic effects, although words whose average predictability is lower have longer word duration regardless of context. The effects of predictability that Hall (2018) points out are quicker articulation, less vowel dispersal (all vowels are more centralized), and shorter duration. In spite of these effects, listeners are more likely to recognize predictable words quickly and accurately, even when reduced and when the phonetic signal is very noisy.

In rhyming pairs, we are looking at a case largely of contextual predictability. Regardless of whether a high frequency word like "you" or a low frequency word like "symphony" is being rhymed, the metrical structure of a poem and our expectations of rhyme make the contextual predictability the overriding factor. Word frequency certainly still plays a role, as we will rarely expect the word "symphony" when a much more predictable word could take its place in the rhyme, such as "see." In *The Cat in the Hat*, the rhyme words are made predictable by their context in the metrical rhyming structure of the book and by their frequency; Dr. Suess attempted to pick high frequency words that children would encounter regularly in their schooling and reading.

As Breen (2018) shows, the second rhyme word, which is always highly predictable, has a shorter duration than the first. This leads us to our first hypothesis.

(3) **Hypothesis 1:** In a rhyming pair, the rhyme target tends to have a shorter duration than the rhyme prime *due to its predictability.*

This hypothesis will be expanded upon in section 2.3, and an experimental design will be proposed to test it in section 3. Before delving into this hypothesis and the ways in which it can be tested, it is important to have a foundational understanding of what rhyme is, and how it can be thought about in more concrete terms. Section 2.3 will take up this task of better defining rhyme.

2.3 Towards a definition of rhyme

2.3.1 Exact rhyme

In order to investigate any question related to rhyme, it is important to have a basic definition of rhyme and to understand the phonological theories that are most applicable to it. In English, the genre of poetry has traditionally necessitated the use of end rhyme, in which lines of poetry always end in rhyming words (as opposed to internal rhyme, where rhyme words are embedded inside lines). End rhyme can be found in couplets, quatrains, sonnets, church hymns, and many other forms throughout the range of English poetry. *Exact* end rhyme (as opposed to non-exact or slant end rhyme, which will be discussed in section 2.3.2) is characterized by *complete identity* between the vowel and all the following consonants in the two rhyme words. For example, Alexander Pope's poem "Epistle to Dr. Arbuthnot," published in 1735, makes use almost exclusively of exact end rhyme:

(4) Shut, shut the door, good John! fatigu'd, I said,
Tie up the knocker, say I'm sick, I'm dead. (Pope 1735)

In this case, the rhymes come at the end of lines of iambic pentameter – lines with five feet of alternating unstressed and stressed syllables – and both are monosyllables. Complete identity exists between the vowels and final consonants of the two rhyme words, [sɛd] 'said' and [dɛd] 'dead'.

In order to better define exact end rhyme and its near companion, slant (non-exact) end rhyme, it will be crucial to understand syllable structure. Generally, linguists have segmented the syllable into three distinct components: the onset, the nucleus, and the coda. The onset is made up of any consonants that stand before the nucleus and are contained in that syllable, the nucleus is typically a vowel and is the sonority peak of the syllable, and the coda is made up of any consonants that follow the nucleus and are contained in that syllable.

Besides including these three components parts, syllables (σ) also have an internal hierarchical structure (articulated syllable structure) which will be quite relevant to the discussion of rhyme. The nucleus and coda are grouped into a constituent called the "rime" (distinguished by spelling from the more general topic of this paper) in the following manner (Halle and Vergnaud 1980; Ladd 1984; Durand 1990):

(5)

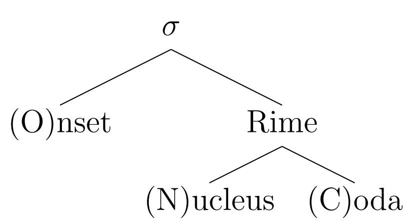

Two words are considered a rhyming pair when their rimes are identical or sufficiently identical. The onsets are independent of any rhyming constraints and do not need to be similar for the two words to rhyme. Exact rhyme can then be defined simply as a pair of words in which the rimes are exactly the same, while the onsets need not be. In fact, it has generally been considered bad form in the history of English poetry to rhyme homophonous words, whether they are merely homophones (*pear/pair*), or are in fact homographs (*fair/fair*) (Encyclopaedia Brittanica 2017). There seems to be a preference for different onsets in rhyming pairs, although this is not, strictly speaking, *essential* to the definition of rhyme described here. Instead, rhyme is only truly concerned with identity between the rimes of two words.

As an example of this, we can break down the rhyme words from (4) into the syllable structure in (5) and see that what makes these words rhyme is the identity of their rimes:

(6) 'said' ∼ 'dead' → s ɛ d ∼ d ɛ d

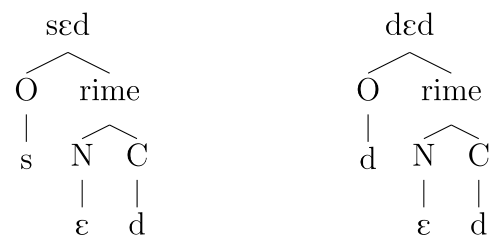

What really makes these words rhyme is that the rimes have identical *features*. The nuclei of the two words are identical in height, backness, rounding, and all other features pertinent to vowels. Likewise, the codas are identical in place of articulation, voicing, manner, and all other features pertinent to consonants. The onsets of these words, on the other hand, have no featural relationship that is contributing the rhyme. The words [sɛd] 'said' and [ɹɛd] 'red' would be an equally good rhyme, so we can see that the features of the onset are not related in any meaningful way.

In section 2.3.2, this syllable- and feature-based analysis of rhyme will be extended to include slant rhyme and the importance of featural identity will be expanded upon.

2.3.2 Slant rhyme

What happens if the rimes of two words are not featurally identical? Can these words still rhyme without fitting into the restrictive definition of rhyme as total identity between rimes? Yes, I would argue, but it is a matter of degree. The more featurally different two words' rimes are, the less well they rhyme. As speakers and readers of English, we can intuit that there is more "rhyminess" between a pair like [tɹuθ] 'truth' ∼ [pɹuf] 'proof' than between [tɹuθ] 'truth' ∼ [læmp] 'lamp'. This comes down to the featural similarity between [tɹuθ] ‘truth’ and [pɹuf] ‘proof’, which, although their rimes are not identical, is still significant. Their nuclei are identical, both containing the back, high, rounded, tense vowel [u]. Their codas are quite similar, as both are voiceless fricatives, but they have slightly different places of articulation: [θ] is (inter)dental, while [f] is labiodental. The rimes of [tɹuθ] ‘truth’ and [læmp] ‘lamp’, on the other hand, have so many featural differences that it is barely worth enumerating them all. It is fair to say the two words do not rhyme at all.

I will use the term “exact rhyme” to describe the case in which two words’ rimes are featurally identical. I will use the term “slant rhyme” to describe all cases in which two words’ rimes are not identical, but are featurally similar enough to suggest rhyme. This is, of course, a rather subjective matter. Where is the line between a slant rhyme and a non-rhyme? This is not entirely clear, and therefore slant rhyme could be said to exist on a continuum, moving from exact rhyme, to one featural difference, two featural differences, three featural differences, and so on, down to total featural dissimilarity. In the work of Andrews-Hoke (2017), which analyzed the types of rhyme used by Emily Dickinson, I found this model of a featural similarity gradient to be useful in describing slant rhyme. Most of Dickinson’s slant rhyme is what might be called “vowel rhyme,” a form of slant rhyme in which the vowels of codaless syllables rhyme but are not identical in features. For example, Dickinson uses the rhyme:

(7) I think that Earth feels <u>so</u>
to folks in Heaven – <u>now</u> – (Dickinson ?b)

This rhyme can be diagrammed just like (6), making it clear that the rhyme [so͡ʊ] ‘so’ ∼ [na͡ʊ] ‘now’ does not meet the definition of exact rhyme, since the rimes are not featurally identical:

(8) [so͡ʊ] ‘so’ ∼ [na͡ʊ] ‘now’

According to most featural definitions, the first half of the diphthongs [o͡ʊ] and [a͡ʊ] differ by two features: height and backness. (They also differ in features that are redundant in English, like rounding.) This slant rhyme, therefore, is slightly more featurally distant than [tɹuθ] 'truth' and [pɹuf] 'proof', which differ by only one feature, but the pair [so͡ʊ] 'so' ∼ [na͡ʊ] 'now' is still squarely in the realm of slant rhymes. Andrews-Hoke (2017) presented the following table, which shows the distribution of Dickinson's use of different types of rhyme. Dickinson uses exact rhyme most often, and then is most likely to use slant rhymes that change only one feature, then two features, then three, etc. In this table, only featural changes in vowels are counted. Dickinson very rarely changes the features of coda consonants.

Table 2: Rhyme continuum

Rhyme: Features:	Exact 0	Slant 1	Slant 2	Slant 3	Slant 4
	91 (72%) [bi] ∼ [si]	20 (16%) [so͡ʊ] ∼ [tu]	7 (5.5%) [ɹɛd] ∼ [hɪd]	7 (5.5%) [bi] ∼ [la͡ɪ]	1 (<1%) [stɑɹ] ∼ [fiɹ]

Although Dickinson largely uses exact rhyme, she uses a fair number of slant rhymes that differ by only one feature. From there, the numbers decrease. Of course, these are the rhyming habits of *one* poet, who uses one particular type of rhyme ("vowel rhyme") far more than any other, but these trends suggest larger tendencies about the perception (and production) of rhyme: we perceive rhyme words as more and more similar, and therefore are more willing to accept them as rhymed pairs, when there are fewer featural differences between their rimes.

Slant rhyme interacts in a particularly interesting way with Breen's finding of rhyme target shortening. If, as was hypothesized in (3), the predictability of rhymes accounts for their decreased duration, then the unpredictability of slant rhymes should throw a wrench in this idea. Slant rhymes are, by their very nature, unpredictable. Some phonological feature that we expect, as readers, to remain identical, has changed. Not only is slant rhyme unpredictable, but it is phonologically unpredictable in a contextually predictable environment, giving it an extra edge of surprise for the reader.

If Breen's word shortening effect is related to rhyme's predictability, then we should find a different effect when readers encounter slant rhyme. Given slant rhyme's unpredictability, we should find that slant rhymes shorten less than full rhymes do in the same poetic conditions.

A second prediction goes along with this one. Namely, that the number of featural differences in a slant rhyme will also have an effect on word duration. The closer a slant rhyme is to an exact rhyme (just one featural change, for example) the more it will shorten. As the number of featural changes in a slant rhyme increases, the word duration will also increase. Hypothesis 1 can therefore be expanded to include this definition of rhyme predictability as featural similarity:

(9) **Hypothesis 1*:** In a rhyming pair, the rhyme target tends to have a shorter duration than the rhyme prime due to its predictability, which is highest when there are zero featural differences between the two rhyme words, and becomes lower with each additional featural difference between the two rhyme words.

This expanded hypothesis leads to a prediction about how slant rhyme will interact with Breen's word shortening effect:

(10) **Prediction:** If predictability drives word shortening, then the more featurally similar a rhyme pair is, the more the rhyme target should shorten. The less featurally similar a rhyme pair is, the less the rhyme target should shorten.

This prediction will be tested in an experiment which will be described in section 3. Before an experiment is proposed, however, section 2.3.3 will discuss a basic rhyme typology, which will contribute to the construction of the experiment and will round out the understanding of rhyme described thus far.

2.3.3 A Digression into a typology of rhyme

Rhyme comes in several flavors, and several of these varieties fall under the umbrella of slant rhyme. Exact rhyme, as already mentioned, occurs when two words rimes are entirely featurally identical. Slant rhyme can be divided into three distinct categories (Small 2010). As can be seen in Table 3, one of these is **consonance**, a type of rhyme in which the coda consonants of two rhyme words are featurally identical, but there is featural change in the nucleus. Dickinson's "vowel rhyme" could be considered a sub-type of consonance, since she alters the features of the nucleus, and not the coda (since there is no coda). The second type of slant rhyme is **assonance**, in which the nucleus remains identical, but the features of coda consonants are altered. The third type of slant rhyme is a combination of **assonance and consonance**. In other words, features of both the nucleus and coda are changed.

Table 3: Rhyme types

	Nucleus identical	Nucleus non-identical
Coda identical	Exact rhyme *time* ∼ *lime* [ta͡ɪm] ∼ [la͡ɪm]	Consonance *lad* ∼ *odd* [læd] ∼ [ɑd]
Coda non-identical	Assonance *rib* ∼ *tip* [ɹɪb] ∼ [tɪp]	Assonance + Consonance *time* ∼ *ran* [ta͡ɪm] ∼ [ræn]

Although all three types of rhyme are, of course, the same in some sense, in that they all fall under the category of slant rhyme, they each have unique featural qualities. There is a reason Dickinson uses "vowel rhyme" so freely – the vowel space is continuous and all vowels have some degree of perceptual similarity that consonants do not necessarily have.[1] On the other hand, many poets use almost exclusively assonance (this is common in hip-hop, for example) (Walser 1995; Tsujimura and David 2010). This makes sense also from a perceptual angle: if two nuclei are featurally identical, this will be the most perceptually salient thing about the two words, since vowels are louder and longer than consonants (Johnson 2003). This allows consonantal feature differences to fade into the background.

This typology of rhyme will be revisited in section 3, when an experimental procedure is proposed. In the proposed experiment, only assonance rhyme will be used, rather than consonance or consonance + assonance rhyme. This choice was made largely because consonance requires vowel features to be changed, and the continuity of the vowel space (and the extensive difference in vowels between individual speakers) makes this approach somewhat more difficult to control.

2.4 Hypotheses and Predictions

I will now distill the discussions of the previous sections, and describe the experiment that was run with the goal of shedding some light on these matters of slant rhyme predictability and word duration.

[1]The understanding of rhyme and slant rhyme as conditioned by *perceptual* similarity, rather than featural similarity, will be returned to in section 5.1. It complicates the featural approach somewhat.

Breen (2018) finds evidence of a word shortening effect in rhyming quatrains in *The Cat in the Hat*, wherein the second rhyme word of a rhyming pair is shortened in duration relative to the first. One possible reason for this is that rhyme words are predictable. All of the phonological material in the second word's rime is easily predictable after reading the first half of the quatrain. Thus, it makes sense that the predictable rhyme word would be shortened, since highly predictable words in highly predictable contexts are always shorter (and less distinctly articulated) than unpredictable words in unpredictable contexts. If predictability is a factor in the shortening of rhyme words, then slant rhyme should produce a different word shortening effect than the exact rhyme found in *The Cat in the Hat*. Slant rhyme is inherently unpredictable, and therefore should not be shortened in the same way that exact rhyme is.

I ran an experiment in which participants were asked to read poems aloud. These poems contained couplets that either were exactly rhymed or slant rhymed, with one or two featural differences. Exact rhymes were included to reconfirm Breen's finding of a rhyme shortening effect. If predictability plays a large role in this effect, then slant rhyme targets should be produced with a longer duration than exact rhyme targets, since predictability is lower. If the analysis of slant rhyme presented in section 2.3.2 is accurate, then slant rhymes that differ by one feature should be shorter than those that differ by two features. In section 3 I will further describe how the experiment was run, what data were collected, and how those data were analyzed.

3 Methods

In sections 3.1, 3.2, and 3.3, the design of the experiment, how it was run, and what data was collected will be described. Section 3.1 will describe the way in which the rhyming couplets were written and arranged that were then read aloud by participants. This was one of the most important steps in the experimental process, since the types of slant rhyme used determined what types of questions could be asked about the nature of slant rhyme. Section 3.1.2 will also describe the construction of the poems more generally, since the ordering of the rhyming couplets mattered to participants' reading of them. Section 3.2 will then describe the experimental procedure during which participants actually read the poems aloud and were recorded. Finally, section 3.3 will address what data were collected, and how those data were organized and analyzed. In section 4, the results of this analysis will be discussed.

3.1 Writing poems

While Breen's experiment was run using an existing and known poetic text, I decided to construct a new set of texts to specifically address the issue of slant rhyme predictability and word duration. The particular construction of these poems is important for the experiment, since as many confounding variables as possible were excluded from the poems and only a few particular features were variables in the study.

Rhyming couplets were decided upon as the ideal structure for the poem. As can be seen in (4), Pope's "Epistle to Dr. Arbuthnot," and in (7), Dickinson's "I'm wife—," rhyming couplets are a commonly used form in English poetry. Quatrains could also have been used:

(11) I died for beauty, but was scarce
Adjusted in the <u>tomb</u>,
When one who died for truth was lain
In an adjoining <u>room</u>. (Dickinson ?a)

However, couplets were selected because the rhyme words would be closer together. Metrically, the couplets used were iambic tetrameter, with some hopefully unobtrusive metrical variation. Tetrameter contains four feet, in this case, iambic feet, which are right-headed (unstressed-stressed). An example of an iambic tetrameter rhyming couplet is:

(12) When all the bells began to ring,
The choir then began to sing.

The choice to use tetrameter couplets will be discussed further in section 5, since the form is not hugely common across English poetry. Generally speaking, rhymed couplets are usually pentameter, while rhymed tetrameter is found in quatrains. Therefore, the use of rhymed tetrameter puts the rhyme words unusually close together for English poetry. Iambic tetrameter is known for its sing-songy quality, so it may have a particularly noticeable metrical pattern, more so than some other meters.

3.1.1 Choosing Consonants

This experiment focused on only one type of slant rhyme. Instead of interrogating all types of slant rhyme, this experiment focused on assonance, in which there is always identity between the nuclei, and non-identity between the codas. Assonance could include any of the following rhymes: [hɪm] 'him' ∼ [tɪn] 'tin', [mæs] 'mass' ∼ [bæt] 'bat', [ɹɪb] 'rib' ∼ [lɪp] 'lip', [kɪŋ] 'king' ∼ [bɪg] 'big'. These codas differ by individual features – place, manner, voicing, and nasality respectively. In the construction of these poems, the number of features under consideration was limited as well. In the poems constructed for this experiment, only the features place of articulation and nasality were changed.

Nasality was decided upon as a relevant feature because nasals are the most confusable segments, regardless of place of articulation (Weber and Smits 2017). Weber & Smits' analysis, which will be described more fully in section 5.1, expands and challenges the featural analysis of slant rhyme. While, in section 2.3.2, an analysis of slant rhyme has been laid out that is based on featural similarity, Weber & Smits' analysis suggests that features cannot account entirely for the "rhymability" of segments. Instead, some sort of phonetic confusability plays a role. Under this analysis, the slant rhyme [hɪm] 'him' ∼ [tɪn] 'tin' will be better (more "rhyme-y") than the slant rhyme [ɹɪb] 'rib' ∼ [hɪd] 'hid', even though the same exact feature (place of articulation) is being altered. Again, this analysis will be expanded upon in section 5.1, after the featural approach has been given its full due experimentally.

The place feature was chosen as the only feature that can interact with nasality without fundamentally altering nasality. Any other featural change (manner or voicing, say) would result in the loss of nasality, since nasals do not have more than one possible manner or voicing quality. Therefore, place and nasality were chosen since they interact well without innately altering each other. The combination of place and nasality features allows the rhymes [ɹɪb] 'rib' ∼ [hɪd] 'hid' and [hɪm] 'him' ∼ [tɪn] 'tin' to be compared. If these rhymes behave in the same way, since both contain the same single featural change, this would support the understanding of slant rhyme laid out in section 2.3.2. If they do not behave in the same way, then the nasal "confusability" that Weber and Smits (2017) describe and which will be discussed more in section 5.1 may be at play.

Voiceless stops were excluded also because voicing is known to create a length contrast in itself (Delattre, Liberman, and Cooper 1955). Because word and rhyme duration are the measurements of study, the phonetic difference in nucleus duration before voiced and voiceless codas ([ɹɪb] 'rib' ∼ [ɹɪp] 'rip') might be too much of a confounding factor in itself.

Using the features nasality and place, 18 rhyming couplets were created. Nasality is a binary feature, meaning a segment is either nasal or non-nasal. Place, however, is not a binary feature and three different places of articulation were used in these couplets: labial, alveolar, and velar. The two features were used to create a set of exact rhymes and slant rhymes. Only six consonants were used as codas, all of which differed by the features in question: *b, d, g, m, n, ŋ*.

Using the six consonants above as codas, the 18 couplets were created, in which the rhyme words ended in 18 different pairs of coda consonants:

(13)

	b	**d**	**g**	**m**	**n**	**ŋ**
b	**b - b**	b - d	b - g	b - m	b - n	b - ŋ
d		**d - d**	d - g	d - m	d - n	d - ŋ
g			**g - g**	g - m	g - n	g - ŋ
m				**m - m**	m - n	m - ŋ
n					**n - n**	n - ŋ
ŋ						**ŋ - ŋ**

As can be seen in (13), all consonant pairs in the upper right of this table will be mirrored in the bottom left, so it is unnecessary to write them out. Counting every cell in the upper right, there are 21 distinct pairs of consonants. Several of these pairs were not used (*g-m, b-n,* and *d-ŋ*). This was an oversight in the design of the experiment, but also should not have left too much of a gap in the data, since all of the possible nasality and place combinations are represented by the pairs *b-ŋ, d-m* and *g-n*, which were used.

This left a total of 18 couplets to write, using the 18 coda consonant pairs from the table above. Six couplets were exact rhyming pairs, and twelve were slant rhymed, differing either in nasality, place of articulation, or both. Of the twelve slant rhyme couplets, three differed in nasality, six differed in place of articulation (three stop pairs and three nasal pairs), and three differed in both place and nasality.

This ultimately breaks down into four distinct categories of rhyme: 1) exact rhymes 2) slant rhymes in which codas differed only by place of articulation 3) slant rhymes in which codas differed only by nasality 4) slant rhymes in which codas differed by place and nasality. Table 4 shows the number of couplets in each rhyme type and Table 5 shows the individual rhyme words in each rhyme type category.

Table 4: Number of couplets of each rhyme type

Exact	**Slant 1** (place)	**Slant 2** (nasality)	**Slant 3** (place + nasality)
6 couplets (3 nasal, 3 oral)	6 couplets (3 nasal, 3 oral)	3 couplets (nasal + oral)	3 couplets (nasal + oral)

Table 5: All rhyme pairs, categorized by rhyme type

Exact	**Slant 1** (place)	**Slant 2** (nasality)	**Slant 3** (place + nasality)
[kʊb] ∼ [tʊb]	[kæb] ∼ [wæg]	[hɛn] ∼ [bɹɛd]	[hʊm] ∼ [mʊd]
[ta͡ɪm] ∼ [la͡ɪm]	[wɛb] ∼ [bɛd]	[tʊŋ] ∼ [ɹʊg]	[ɛg] ∼ [tɛn]
[bæd] ∼ [mæd]	[lɪd] ∼ [wɪg]	[kɹɪb] ∼ [d͡ʒɪm]	[jʊŋ] ∼ [pʊb]
[lun] ∼ [spun]	[klæm] ∼ [mæn]		
[hʊg] ∼ [bʊg]	[kɪŋ] ∼ [swɪm]		
[sɪŋ] ∼ [ɹɪŋ]	[tɪn] ∼ [wɪŋ]		

The rhyme words were all monosyllabic and were chosen carefully so as not to have nasality affect vowel quality too much. For example, a rhyme like [kɹɪb] 'crib' ∼ [d͡ʒɪm] 'gym' was chosen over a rhyme like [kɹæb] 'crab' ∼ [d͡ʒæm] 'jam', because the vowel quality of [æ] is affected greatly by nasality, to the point where speakers do not always have the intuition that [kɹæb] 'crab' and [d͡ʒæm] 'jam' have the same vowel. This intuition is supported by the phonetics of nasalization, since low vowels (like [æ]) are known to have some amount of passive nasalization, i.e. some secondary resonance occurs in the nasal cavity even without the presence of a nasal consonant nearby. This passive nasalization combined with the active nasalization in a word like [d͡ʒæm] 'jam' can lead to the perceptual re-evaluation of low vowels, like [æ] (Johnson 2003). Therefore, it makes sense that the vowels in [kɹæb] 'crab' and [d͡ʒæm] 'jam' might be interpreted as different enough vowels that they do not rhyme exactly .

3.1.2 Constructing Couplets and Poems

As described above, each couplet was iambic tetrameter. The couplets were written so as not to contain sharp enjambment between lines. Poetic enjambment occurs when one line transitions to another without a terminating punctuation mark, often cutting through a phrase or clause (Literary Devices 2019). An example of the type of enjambment that was avoided can be seen in (14).

(14) Elise walked to the corner. In
the street there was a silver pin.

Instead, each line was written to be a natural prosodic unit, meaning that often the two lines were separate clauses, as in:

(15) There was a little lion cub,
who climbed into an empty tub.

Once the 18 couplets were constructed, they were randomly ordered using an online randomization generator.[2] This ensured that the exact rhyme and slant rhyme couplets were randomly distributed throughout the poem. Two different random orders were created (called Poem 1 and Poem 2), so that the locational effect of a particular couplet in the poem would be minimized. In other words, it is possible that early couplets would be lengthened, as the reader was still adjusting to the cadence of the poem, while later couplets would be rushed through and shortened, or some similar effect. Creating two different random orders would minimize potential effects like these.

After 18 couplets were written, 18 additional couplets were written in which the two rhyme words were reversed in their order, creating Poem 1R(everse) and Poem 2R(everse):

(16) The spider walked around her web,
And ate some flies, and went to bed.

The spider ate and went to bed,
And woke up early in her web.

This theoretically minimized any effects of word frequency. Words that are more frequent, as discussed in section 2.2.2, are usually pronounced with shorter durations across different linguistics contexts (Hall 2018). If, for example, [bɛd] 'bed' were more frequent than [wɛb] 'web' and therefore pronounced with a shorter duration on average,

[2] The online randomization generator used can be found at www.random.org/lists/

the reversal of the couplets should neutralize this factor. In (16), in other words, [bɛd] ‘bed’ should be shorter in the first couplet and [wɛb] ‘web’ should be shorter in the second couplet. Regardless of the frequency of either rhyme word, the second rhyme word is expected to be shorter in duration, as Breen found and I hypothesized.

Again, what is being protected against through this reversal process is the artificial creation of (or obscuring of) the word shortening effect. Always having one word come first or second could inadvertently create the word shortening effect that Breen finds, due to the frequency predictability effects Hall (2018) describes.

Couplet reversal also had the added benefit of distributing experimenter bias when it came to extracting rime durations in the data analysis phase. Any unrealized biases in marking the beginnings or ends of the rime were distributed evenly across both first and second rhyme words. These decisions and potential biases will be discussed more in section 3.3.

After creating all 18 couplets and their reversed versions, two “framing” couplets were created. These couplets were designed to begin and end the poem, and both rhyme exactly:

(17) The girl was wise of heart and mind.
She spoke these truths for all mankind:

[the 18 experimental couplets randomly ordered]

Here are the sayings of the girl.
Go on and give her truths a whirl.

The initial couplet of (17) was included because its exact rhyme was necessary to prime the reader to expect rhyme throughout the poem. Some participants encountered a slant rhymed couplet immediately after this first framing couplet, so it was important that they expect rhyming couplets before being faced with slant rhyme. The final couplet of (17) was included so that any poem-final effects on word duration would be minimized. When a reader can see that these are the final words of the entire poem, they might change their pace, word length, and/or pitch to express this.

As was discussed in section 2.2.1, such phrase- or utterance-final effects have been observed in other linguistic contexts, and may very well apply to the case of poem-final utterance. As speakers reach the end of any linguistic utterance, they often will decline in pitch and have a tendency to lengthen words (Ladd 1984; Beckman and Edwards 1990). By using the final framing couplet, that potential effect on word duration could be minimized.

In total, each poem was made up of 20 couplets. Two of these were the framing couplets, and the middle 18 were the experimental couplets, randomly ordered. Two different random orders were used and for each random order, a "reverse" poem was created, using the couplets with the rhyme words reversed. This process yielded four poems in total:

Table 6: Four different poem orders

	Random Order 1	Random Order 2
Normal couplets	Poem 1	Poem 2
Reverse couplets	Poem 1R	Poem 2R

Each of the four poems was read by four or five participants. The recording procedure will be described further in section 3.2. Poem 1 and Poem 1R can be found in Appendix A. All couplets and their reversed forms can be read in full in the poems in the appendix.

As can be seen from the couplets in (15)–(17), these poems do not entirely "make sense." To some degree, these are nonsense poems made up of randomly ordered nonsense couplets. In this experiment I largely ignored semantic and syntactic factors and tried to make the couplets relatively syntactically easy to parse. The focus is on the meter and the rhymes.

3.2 Reading poems

Recordings were made in a sound booth, using a Logitech USB microphone. Recordings were made directly into Praat on a Mac laptop at a sampling frequency of 44,100 Hz and saved as .wav files. A total of 19 participants were recorded. Each participant was first asked to fill out a brief survey, which asked a few questions about demographic information, such as age and gender, as well as about the participant's linguistic background. All participants were native speakers of English, and most spoke a standard American

English dialect. A few spoke other dialects, including African-American English, British English, and South African English.

After filling out this survey, participants were given a poem. The poem was printed on two sheets of paper in size 16 font so that it was easily legible and was laid out on the table in front of the participant, so that they could see the entirety of the poem. There were four different poems (see Table 6) and four or five participants read each poem. (For one participant, one poem was one couplet shorter than the others, due to error, but this oversight was corrected for the following participants.) Participants were asked to read the poem silently to themselves, so that they would be familiar with the language before reading aloud. After reading silently, participants were then asked to read through the poem aloud while being recorded. All participants read and were recorded twice. (One file was lost due to investigator error.) Participants were told to read at whatever felt like a normal pace and with whatever felt like a normal cadence. If they made an error while reading, participants were asked to begin that particular couplet again and continue.

After 19 participants read, the resulting data was 37 recordings of poems, each 18 couplets in length (plus two framing couplets, which were not measured for word duration). Again, one file was lost and one participant's recordings were only 19 couplets long, due to error.

3.3 Sorting the Data

Once all of the recordings were gathered, the process of organizing and analyzing the the data began. Each .wav file was opened in Praat and a TextGrid was used to mark the word duration of rhyme words. Each TextGrid had two important tiers: Word and Rime.

In the Word tier, the durations of all rhyme words were marked. A diligent attempt was made to mark exactly from the beginning of each rhyme word's onset to the end of the coda. Sometimes, the exact beginning of an onset was difficult to determine, since it was phonetically similar to the coda of the previous word. For example, one couplet read: "We hunted with a silver spoon, but did not find a swimming loon." When constructing the couplets, ease of segmentation was not a major priority, which resulted in hard-to-segment phrases like [swɪmɪŋ lun], which contains a transition between [ŋ] and [l] that is difficult to detect in either spectrogram or waveform, as can be seen in Figure 1.

As can be seen in Figure 1, the division of [ŋ] from [l], marked by the vertical line to the left of the loon_B notation, cannot exactly be determined by eye from the spectrogram. A mix of careful listening and subtle spectrogram clues (like the nondescript rise in the second formant right after the segmentation line) were used to make these more difficult choices.

In the Rime tier, only the rime of each rhyme word was marked off. Since rimes are the parts of the words that actually *rhyme*, this seemed like a relevant distinction to make. Additionally, some rhyme words had complex onsets, like [klæm] 'clam', while others did not, like [mæn] 'man'. Therefore, the duration of the rimes might contain Breen's word shortening effect, even if the durations of the entire rhyme words did not, due to differences in entire word duration being complicated by onset length. Again, it was sometimes difficult to determine where the onset of a word ended and where the nucleus began, particularly when onsets were liquids or glides. (For example, [j], [w], [l], and [ɹ] were especially difficult to segment.) An example of one of these onset to nucleus transition that was difficult to segment can be seen in Figure 2.

When words began with plosives, the word selection was marked at the beginning of the burst (Figure 3), and when a word ended in a plosive, if there was a final burst, the end of the burst was marked (Figure 4).

In each case, the selection was made from when the sound wave crossed the zero axis. Not every .wav file collected was marked and included in the final data set. However, an equal number of recordings of each of the four poems was used. All of the marked recordings were read by participants who spoke Standard American English dialects, although there were several different dialects spoken in the original collected data.

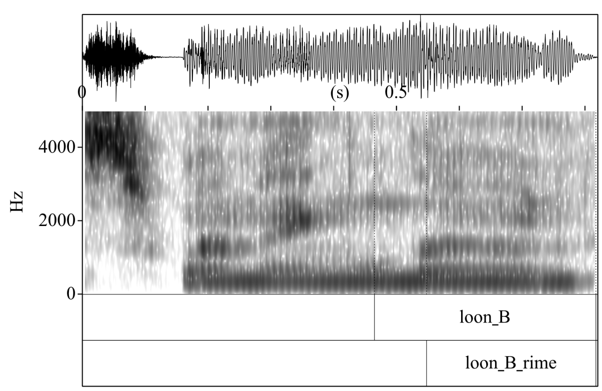

Figure 1: [swɪmɪŋ lun] "swimming loon"

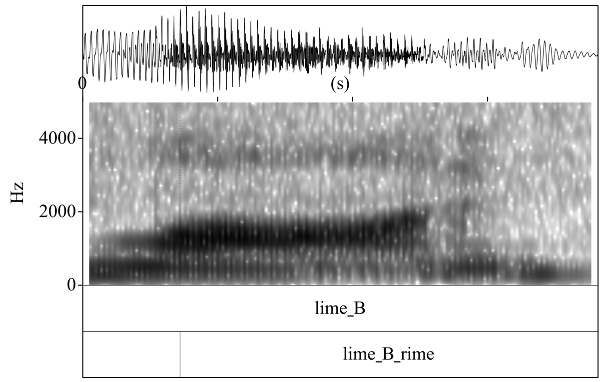

Figure 2: [la͡ɪm] “lime”

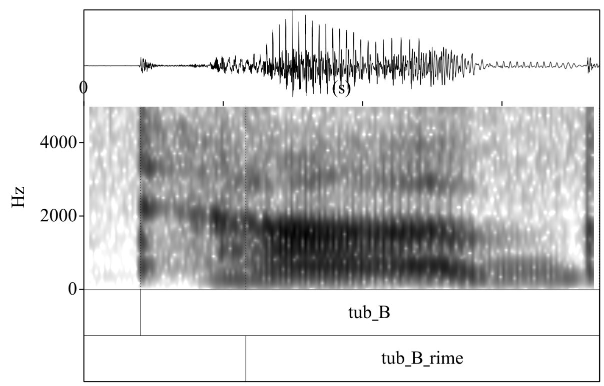

Figure 3: [tʊb] “tub”

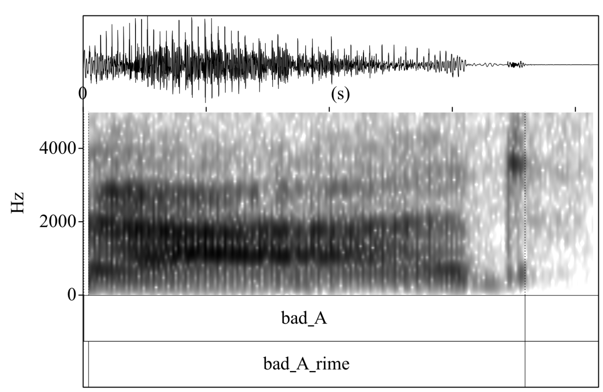

Figure 4: [bæd] “bad”

3.3.1 Extracting Durations

After segmenting a substantial number of the recordings on the Word and Rime tiers, the durations of the rhyme words and rimes were extracted from Praat using a Praat script. This script was found online, and was called "calculate segment durations."[3] For every sound file and related TextGrid, the script was run to extract the duration of every rhyme word and of every rime. All of the extracted durations were compiled into two Microsoft Excel documents, one containing word durations and one containing rime durations. In section 4, only rime durations will be used, since some rhyme words had complex onsets while others did not, a factor that can be ignored when focusing only on rimes. The durations of a total of 177 rhyme pairs were extracted, with 57 exact rhyme pairs and 120 slant rhyme pairs. Of the slant rhyme tokens, 60 were of the Slant 1 variety and 30 each were of the Slant 2 and Slant 3 varieties (see Table 5). (There should have been 60 tokens of slant rhyme, but the exclusion of one couplet accidentally resulted in a few lost tokens.)

Ultimately, 177 rhyme pairs falls far below the amount of data actually collected. The use of a forced aligner could have increased the amount of data available, since the main limiting factor was the time it took to hand-mark the word and rime durations in each .wav file. In any further analysis of this data, a forced aligner would ideally be used to increase the number of tokens under consideration. However, the amount of data analyzed should still give an accurate account of the trends and patterns within the different rhyme type groups.

4 Results

After the experiment was run, resulting pieces of data included the duration of the first rhyme word's rime, the duration of the second rhyme word's rime, and, importantly, the difference between the two durations. The duration difference was found by subtracting the duration of what I will call Rhyme 2 (the duration in milliseconds of the second rhyme word's rime) from Rhyme 1 (the duration in milliseconds of the first rhyme word's rime). If the value of the duration difference is positive, that means that Rhyme 2 is shorter than Rhyme 1. If the value of the duration difference is negative, then Rhyme 2 was longer than Rhyme 1. From the hypothesis previously discussed, the duration difference

[3] www.mv.helsinki.fi/home/lennes/praat-scripts/public/calculate_segment_durations.praat

is expected to be positive, which would confirm Breen's rhyme word shortening effect.

Looking at all of the tokens of exact rhyme that were collected, Breen's finding was indeed confirmed. This can be seen in Figure 5 below, in which all tokens of duration difference (Rhyme 2 duration – Rhyme 1 duration) are shown in a histrogram.

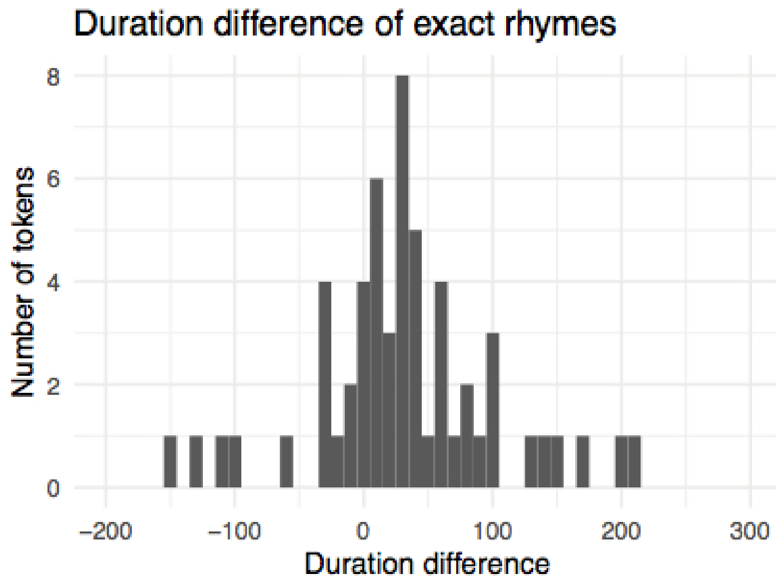

Figure 5: Histogram showing distribution of duration difference of all exact rhyme tokens

As can be seen, Figure 5 is essentially a normal curve, with the center of the curve set above zero. Simply from looking at this histrogram, Breen's finding appears to be confirmed, since the average duration difference certainly is above zero. Again, this is only the exact rhymes, so Figure 5 has yet to demonstrate anything concerning the status of slant rhymes.

Breen's word shortening effect is further confirmed by running a t-test to compare the means of Rhyme 1 duration and Rhyme 2 duration in all exact rhymes. After running a two-tailed t-test, using the durations in milliseconds of Rhyme 1 and Rhyme 2, it can be determined that there is a statistically significant difference in the means of the two groups, with the higher mean being found among the Rhyme 1 durations. This confirms Breen's finding of a word shortening effect in exact rhymes.

(18) t-test results comparing Rhyme 1 and Rhyme 2 durations for exact rhymes

	Mean (ms)	Standard Dev.	p-value (<5%?)
Rhyme 1 duration	287.83	55.15	.0015
Rhyme 2 duration	256.22	65.03	”

Although this is simply the confirmation of Breen's finding, it should not be forgotten that this finding is quite interesting. Given the general linguistic evidence of phrase-final (or pre-pausal) lengthening (Liberman and Prince 1977; Beckman and Edwards 1990), this tendency to shorten towards the end of poetic phrases is a reversal of the normal linguistic tendencies. This fact will be returned to in greater depth in section 5.

Breen's finding can be extended to include slant rhymes. Overall, slant rhymes (of the Slant 1, 2, and 3 groups) exhibit the same shortening effect. The following *t*-test shows that the average duration of Rhyme 2 is shorter than that of Rhyme 1 to a statistically significant extent in all slant rhyme groups.

(19) *t*-test results comparing Rhyme 1 and Rhyme 2 durations for all slant rhymes

	Mean (ms)	Standard Dev.	p-value (<5%?)
Rhyme 1 duration	254.72	52.99	0.00001
Rhyme 2 duration	227.25	46.51	”

(19) shows that the word shortening effect is not exclusively reserved for exact rhymes. It remains to be shown, however, whether there is a noticeable difference in the amount of word shortening between the exact and slant rhyme groups. To revisit the hypothesis briefly, slant rhyme's unpredictability is expected to have a lengthening effect on otherwise shortened rhyme words. The more featurally different a Rhyme 2 is from its Rhyme 1, the more lengthened Rhyme 2 is expected to be compared to an exact rhyme, by virtue of its featural unpredictability. We now know that, on average, Rhyme 2 is shorter than Rhyme 1 in exact *and* slant rhymes, but is Rhyme 2 slightly longer in slant rhymes than in exact rhymes?

In Figure 6, the density of duration difference is shown, divided up by rhyme type. This gives a general sense of where the average duration difference lies for different rhyme types. As a reminder, Exact rhyme contains zero featural differences, Slant 1 and Slant 2 both have one featural difference, and Slant 3 has two featural differences.

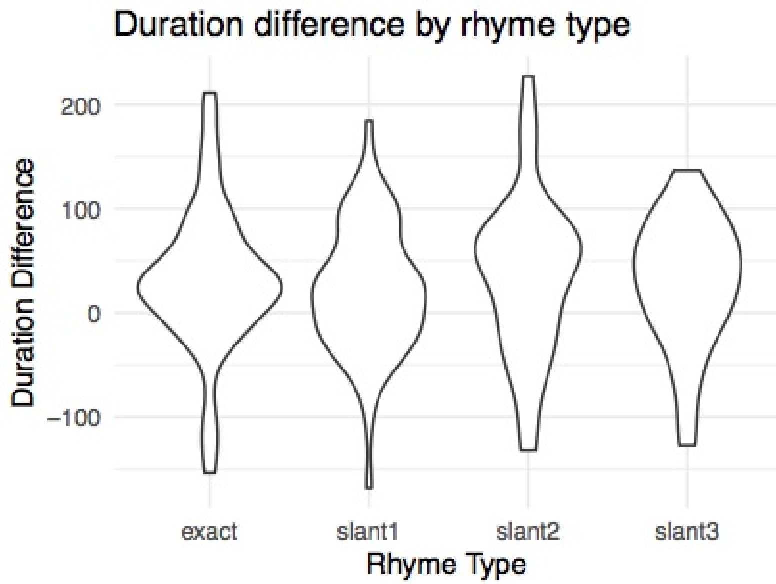

Figure 6: Density of duration difference separated by rhyme type

If the hypothesis were true, we would see a descending average duration difference across these four rhyme types. (Slant 1 and Slant 2 should theoretically be the same, since they both have one featural difference.) This trend, however, cannot be seen in (6). If anything, the average duration difference rises slightly between the Exact and Slant 3 groups, which would mean that, on average, a slant rhyme with two featural changes is shorter than an exact rhyme. This would be the opposite of the predicted result.

Indeed this finding (or lack of a finding) is confirmed by a further t-test. Using a two-tailed t-test, the means of the duration differences of Exact rhymes and *all* Slant rhyme groups were compared. In order to confirm the hypothesis previously stated, there would need to be evidence that the Slant rhyme groups had a lower average duration difference to a statistically significant extent. This is not confirmed by the t-test in (20).

(20) t-test results comparing Exact rhyme duration difference and Slant rhyme duration difference

	Mean (ms)	Standard Dev.	p-value (<5%?)
Exact rhyme duration diff	31.62	70.78	0.60
Slant rhyme duration diff	24.82	63.23	"

This *t*-test does not confirm the hypothesis presented in section 2.4, since the p-value is far above that needed for the difference in means to be statistically significant. Instead, this suggests that the featural identity of a rhyme, whether it be an exact rhyme or a slant rhyme of any of the varieties tested, does not interact with word shortening in the way predicted. Although there is an overall effect of rhyme word shortening in both exact and slant rhymes, this effect does not appear to be gradient, and certainly is not gradient based on individual featural differences in the way predicted.

Although the hypothesis has not been confirmed, the negation of it leads towards many other interesting questions which will be further discussed in section 5.

5 Discussion

The confirmation of Breen's finding, and the extension of it to include slant rhyme, is worth pausing over. As mentioned, this sort of phrase-final shortening is the opposite of what occurs in other realms of linguistic utterance (i.e. phrase-final lengthening). In combination with the negation of the hypothesis concerning featural slant rhyme, the rhyme word shortening phenomenon remains unaccounted for. If, as is shown in the *t*-test in (20), there is indeed no significant connection between number of featural differences in a slant rhyme and the strength of the shortening effect, then predictability cannot be the main factor accounting for shortening, at least not in the way expected.

Something about the understanding of predictability laid out in section 2.2.2 must be incorrect. Theoretically, as has been described, a higher number of featural differences should lead to a lower predictability, which should in turn lead to a longer rhyme word duration. This, however, is not the case, which leaves two possibilities: either features do not contribute to word predictability in the way expected *or* predictability does not affect word duration (and thus, duration difference) in the way expected, or both. If number of featural difference in fact does not add to or detract from the predictability of a rhyme word, then some other type of predictability might be at play. On the other hand, perhaps the featural account of slant rhyme *does* have something to say about the overall predictability of a rhyme – a predictability which might emerge in other contexts or other experiments – but the predictability of individual words, in fact, has nothing to do with the shortening effect.

If predictability does not drive the word shortening effect in the way expected, **what does drive the word shortening effect?** In the rest of this section, shortcomings or possible expansions of the current experimental design will be discussed, in hopes of presenting more ways to potentially investigate this question.

One addition to the existing experimental structure that could shed some light on what might be driving the shortening effect would be the addition of non-rhymes. Although various exact and slant rhymes were used, no complete non-rhyme words with included, which means that a crucial category of comparison is absent from the results. There were not couplets in the experiment that could be feasibly categorized as non-rhymes, like, for example:

(21) There was a little lion cub,
who climbed into an empty box.

The words [kʊb] and [bɑks] are quite different phonologically, and yet both are still monosyllables, which allows them to fit into the metrical structure, if not the rhyme scheme. Would [bɑks], in this context, shorten in the same way that a rhyme or slant rhyme would? If so, the entire foundation of this experiment – in which the word shortening effect is suspected to be attributed to rhyme in some fashion – might need to be reexamined. Rather than rhyme, the word shortening effect could be related to the overall metrical structure in which any word finds itself. For example, the following "poem," found in a Yale University dining hall, contains many strong metrical effects, and yet no rhyme. Would this, too, be subject to the word shortening effect, either at the end of every line or at the end of the stanza as a whole, despite containg no rhyme?

(22) We are the Farmers
Who work the Land
And grow the Crops
To feed the Cows
That give the Milk
For you to Enjoy!

Although a form like this would need to be researched more thoroughly, as readers of poetry and speakers of English, we can intuit that different metrical structures can have strong effects on our reading speed, pitch (including metrical "sing-songiness"), and other vocal qualities. Throughout this paper, the word shortening effect has been often referred to as the "rhyme word shortening effect," as if the effect must have something to due with rhyme, but it is plausible that rhyme words are neither the cause nor the only "victim" of the shortening effect. It is possible, in fact, that the word shortening effect is entirely a product of some other poetic force, perhaps of the metrical structure in which a word finds itself. Another way to approach this possibility experimentally would be to use the same rhyme word sets used in this experiment and laid out in Table 5, but to embed them in different metrical structures (and non-metrical structures) and run a similar experimental design:

(23) Where is the lioness mama and where is her <u>cub?</u>
There is her baby – he's climbing into an old <u>tub</u>!

(24) There was a little lion <u>cub</u>,
who climbed over a big old tin <u>tub</u>.

(25) There's a small lion <u>cub</u>
climbing into a <u>tub</u>.

Instead of asking about the effect of rhyme type, this type of experiment would investigate different meter (and non-meter) types, which conceivably could be responsible for the word shortening effect observed – an effect which seems independent of rhyme type, but nonetheless is unique to poetic language.

Indeed, rhyme and meter are often described as the two features that make traditional English poetry what it is, formally speaking. If the identity of rhymes cannot be found to account for the linguistic abnormality of phrase-final shortening, then meter is the next candidate to be investigated, since the word shortening effect seems to be unique to poetry, and has only been found thus far in highly metrical poetic texts (*The Cat in the Hat* and this experiment's iambic tetrameter couplets). Metrical structure could be said to be a sort of predictability as well, although a more abstract contextual predictability, since meter primes readers to know how many strong beats there will be in a line and where they will fall. The effect of metrical predictability, rather than rhyme word predictability, could easily be the next step of inquiry on this topic.

As it stands, the metrical form used in this experiment (iambic tetrameter couplets) is not the most natural English poetic form anyway. It was relatively arbitrarily chosen (although it is similar to *The Cat in the Hat*'s anapestic tetrameter), but is not widely in use in much of English poetry. The results found using this form might not be as widely applicable to the broader body of English poetry as wished, since it is not hugely common.

The problem with iambic tetrameter couplets – which might account for its unpopularity across English poetry – is that rhyme words come only four beats (four feet) apart, which is quite close. As a result, pentameter (five-footed) rhyming couplets and tetrameter rhyming quatrains are preferred. Again, Pope's "Epistle to Dr. Arbuthnot," which is in pentameter rhyming couplets:

(26) Shut, shut the door, good John! fatigu'd, I <u>said</u>,
Tie up the knocker, say I'm sick, I'm <u>dead</u>.

Tetrameter rhyming quatrains are also preferred over tetrameter rhyming couplets. An example of this can be found in Thomas Hardy's "Channel Firing:"

(27) That night your great guns, una<u>wares</u>,
Shook all our coffins as we **lay**,
And broke the chancel window-<u>squares</u>,
We thought it was the Judgment-**day** . . . (Hardy 1914)

In (26) and (27), the rhymes are five and eight beats apart respectively. The few rare cases of tetrameter rhyming couplets in English poetry are regarded as sing-songy, probably because the rhymes are so closely spaced. For example, Joyce Kilmer's poem "Trees" is in tetrameter rhyming couplets. "Trees" is also the poem that inspired Columbia University's Alfred Joyce Kilmer Memorial Bad Poetry Contest, which speaks to the general disrepute of the form.

(28) I think that I shall never <u>see</u>
A poem lovely as a <u>tree</u>.
. . .
Poems are made by fools like <u>me</u>,
But only God can make a <u>tree</u>. (Kilmer 1913)

Taking these things into consideration, it may have been more representative of the poetic contexts in which people usually encounter rhyme to have used pentameter rhyming couplets or tetrameter rhyming quatrains. Nonetheless, this consideration takes a back seat to the previous discussion of the inclusion non-rhymes. If a different experiment were undertaken in which a more natural English metrical form were used, it would have to also include non-rhyme words.

Another important possibility, which has not been addressed thus far, is the possibility that first rhyme words are being lengthened, rather than second rhyme words being shortened, creating a positive duration difference by a separate avenue.

Why would the first rhyme word be lengthened? There are several reasons to follow this line of thought. First of all, in the data recorded for the experiment proposed here, it can be seen across speakers that often there is no audible pause between the first and second lines of a couplet. As Breen's findings show, however, there is a longer inter-onset interval between the last word of one line and the first word of the next than between any words within the lines. This suggests that the poetic line break is enforcing some audible separation between the two lines, but speakers also try to bridge the line break by connecting the two line vocally. These suppositions are largely being made by eye when observing the collected data, but it certainly seems that investigating the possibility of word lengthening would be a valuable next step.

The effects of word frequency and lexical neighborhood (Marian 2017) were also not considered in the results section, which could have added something to the analysis, even though predictability's role in the word shortening effect is uncertain.

To close out the discussion of the results, section 5.1 will present an alternate approach to rhyme that both expands on an challenges the featural understanding of rhyme laid out in section 2.3. The featural understanding of rhyme, while useful, does not contain ultimate explanatory power, and so any future work on the topic of rhyme should consider alternative understandings as well.

5.1 Rhyme as perceptual similarity

To conclude, I will present another approach to rhyme which could be pursued in future research on these topics. This approach sees rhyme as a function of perceptual similarity or "confusability," rather than featural similarity. To what degree does the featural model of rhyme and slant rhyme discussed so far actually describe how listeners/readers understand rhyme? It is hard to say. Do speakers of English really find [tɹuθ] 'truth'

~ [pɹuf] 'proof' to be a "better" rhyme than [so͡ʊ] 'so' ~ [na͡ʊ] 'nau' due to the lower number of featural changes? Again, this study does not fully answer these questions.

On the surface level, it seems that an increased number of featural changes between rimes *does* decrease the acceptability of a rhyme. This is shown in the work of Weber and Smits (2017), who ran an experiment to determine the confusability of various segments in American English. The confusability of two segments, Weber and Smits (2017) find, is closely related to featural similarity, which supports the approach to slant rhyme already described. In the following table, from Weber & Smits, the confusion rates of different consonants are laid out. Listeners were given an actual input (in the lefthand column) and were asked to report what segment they heard. The segments participants reported hearing are in the upper row, and the rate of confusion of the input for a certain output is in each cell.

(29)

	b	d	m	n
b	**27.3**	5.8	6.5	1.0
d	5.8	**28.8**	1.9	12.7
m	2.7	1.0	**60.0**	9.0
n	0.4	2.9	12.5	**68.8**

(Weber and Smits 2017)

As can be seen in (29), consonants that are more featurally similar are more often confused for each other. For example, [b] and [d], which differ by one feature (place of articulation), are each confused for the other 5.8% of the time. [b] is confused for [m], which differs only in the nasality feature from [b], 6.5% of the time. However, [b] is confused with [n] only 1.0% of the time. This makes sense, since [b] and [n] differ by two features: place of articulation and nasality.

The notion of confusability is similar to "rhymability." In other words, the more confusable two segments are, the easier it is to rhyme them. Thus, to the extent that Weber and Smits find that more featural similarity is equal to higher confusability/rhymability, this supports the featural model of slant rhyme laid out in previous sections.

Weber & Smits' rates of confusability, however, also present some difficulties for the featural approach. Although number of featural differences plays an important role in rhymability, it also seems that *which* features are present matters. In (29), [b] and [d] are confused for each other 5.8% of the time, but [m] was confused for [n] 9.0% of the time and [n] confused for [m] 12.5% of the time. This suggests that the nasal consonants are simply more confusable than the oral stops. Although the pairs [b] & [d] and [m] & [n] both differ only in place of articulation, [m] and [n] are somehow "more rhymable." In other words, a rhyme like [dɪm] 'dim' ∼ [tɪn] 'tin' should be more acceptable than a rhyme like [wɛb] 'web' ∼ [ɹɛd] 'red', even though the exact same featural changes are made between the rimes in each rhyme pair.

This asymmetry in rhymability between oral stops and nasal must be due to some other factor than the featural differences described before. One plausible explanation for the higher rhymability of nasals is perceptual. It is well-attested cross-linguistically that nasals are simply less phonetically distinguishable from each other than oral stops, fricatives, and other consonants are. All nasals are characterized phonetically by "nasal murmur" (a very low frequency periodic murmur) and have few other distinctive qualities at higher frequencies (Ladefoged and Maddieson 1996). Therefore, the primary way that speakers distinguish place of articulation in nasals is by the formant transitions of the preceding and/or following vowels (Johnson 2003). This is not true of oral stops, which tend to have phonetically and perceptually distinctive bursts. Therefore, by this line of thinking, it is important, when considering slant rhyme, to not rely only on the featural similarity of two sounds, but also on the perceptual similarity, independent of feature specifications.

Considering other approaches to rhyme, slant rhyme, and word shortening can add to our understanding of poetic speech, and determining the ways in which speakers and readers of poetry realize these phenomena can create a deeper understanding of what knowledge a linguistic art like poetry is tapping in to. Any future study in this area would do well to consider some of the previously discussed shortcomings of this experiment, as well as the expansions and alternate approaches proposed here.

6 Conclusion

In this thesis, I have expanded experimentally on the work of Breen (2018), following from her finding that rhyme targets are shorter in duration than rhyme primes. The role of rhyme predictability, and specifically featural predictability, in this shortening effect was then tested.

The results of this experiment confirmed Breen's original hypothesis, and expanded its scope to include slant rhymes, but did not confirm the subsequent hypothesis concerning the effect of the featural predictability of slant rhyme on word duration. Breen's original finding, however, and its confirmation point to the unique application of linguistic rules in the realm of poetry. Although phrase-final lengthening is found in most linguistic contexts, in poetic contexts, phrase-final shortening seems to exist. This suggests that poetry, as a medium, not only plays consciously with language, rearranging syntax and heightening sonic qualities, but also plays with our grammar at a deeper level, manipulating our unconscious hold on linguistic structures to create new and playful grammars within us.

A Appendix: Example Poems (Poem 1 and Poem 1R)

The girl was wise of heart and mind.
She spoke these truths for all mankind:

I think if he would eat a clam,
Then he would be a better man.

The grandmother began to hum,
Although the boy was eating mud.

I meant to give the man a hug.
Instead I ate the little bug.

She had a rooster and a hen.
For dinner she would bake them bread.

The cat stuck out its little tongue,
And sat upon the Persian rug.

Whenever I become a king,
I'll always take a daily swim.

For breakfast he would eat an egg,
Then count to five or maybe ten.

He drove the yellow taxi cab.
His little dog began to wag.

The spider walked around its web,
And ate some flies, and went to bed.

We hoped that we would find the loon.
We hunted for it with a spoon.

If children don't keep track of time,
They will not get their evening lime.

There was a bird made out of tin.
He cried because he had one wing.

The choir soon began to sing,
And all the bells began to ring.

The dog is being very bad,
And that makes me a little mad.

There was a little lion cub
Who climbed into an empty tub.

The jar was covered by a lid,
On top of which there was a wig.

The baby cried inside the crib,
but mom and dad were at the gym.

Although the boy was very young,
He spent his Sundays at the pub.

Here are the sayings of the girl.
Go on and give her truths a whirl.

The girl was wise of heart and mind.
She spoke these truths for all mankind:

I think he'd be a better man,
If he would only eat a clam.

The little boy was eating mud,
While grandmother began to hum.

I ate the little silver bug,
and did not give the man a hug.

For dinner, she would bake some bread,
And feed her rooster and her hen.

The cat sat on the Persian rug,
and licked the milkshake with its tongue.

I'll always take a daily swim,
Whenever I become a king.

At breakfast he would count to ten,
Then start to eat his scrambled egg.

The little dog began to wag,
And drove a yellow taxi cab.

The spider ate and went to bed,
And woke up early in her web.

We hunted with a silver spoon,
But did not find a swimming loon.

They will not get their evening lime,
if children don't keep track of time.

The metal bird had just one wing.
His body was made out of tin.

When all the bells began to ring,
The choir then began to sing.

The dog makes me a little mad,
because he's being very bad.

Into the empty metal tub,
There climbed a little lion cub.

Inside the jar there was a wig,
On top of which there was a lid.

Mom and dad were at the gym.
The baby cried inside its crib.

He spent his Sundays at the pub,
although the boy was very young.

Here are the sayings of the girl.
Go on and give her truths a whirl.

References

Andrews-Hoke, Magda. 2017. An OT Approach to Emily Dickinson's Slant Rhyme .

Beckman, Mary E., and Jan Edwards. 1990. Lengthenings and shortenings and the nature of prosodic constituency. In *Papers in Laboratory Phonology: Volume 1, Between the Grammar and Physics of Speech*, 152–178. Cambridge University Press.

Blevins, Juliette. 2005. The Role of Phonological Predictability in Sound Change: Privileged Reduction in Oceanic Reduplicated Substrings. *Oceanic Linguistics* 44:517–526.

Breen, Mara. 2018. Effects of metric hierarchy and rhyme predictability on word duration in The Cat in the Hat. *Cognition* 174:71–81.

Cooper, William E., and Jeanne Paccia-Cooper. 1980. *Syntax and Speech*. Cognitive Science Series. Harvard University Press.

Delattre, Pierre C., Alvin M. Liberman, and Franklin S. Cooper. 1955. Acoustic loci and transitional cues for consonants. *The Journal of the Acoustical Society of America* 27:769–773.

Dickinson, Emily. ?a. I died for beauty.

Dickinson, Emily. ?b. I'm "wife"—I've finished that—.

Dr. Seuss. 1957. *The Cat in the Hat*. Penguin Random House LLC.

Durand, Jacques. 1990. *Generative and non-linear phonology*. Harlow: Longman.

Encyclopaedia Brittanica. 2017. Rime riche. *Encyclopaedia Brittanica* .

Fabb, Nigel, and Morris Halle. 2008. *Meter in Poetry*. Cambridge: Cambridge University Press.

Hall, Kathleen. 2018. The Role of Predictability in Shaping Phonological Patterns. *Linguistics Vanguard* .

Halle, Morris, and Jean-Roger Vergnaud. 1980. Three dimensional phonology Journal of Linguistic Research:83–105.

Hardy, Thomas. 1914. Channel Firing.

Hayes, Bruce. 1995. *Metrical Stress Theory*. Chicago: University of Chicago Press.

Johnson, Keith. 2003. *Acoustic and Auditory Phonetics*. Blackwell Publishing.

Kilmer, Joyce. 1913. Trees.

Klatt, Dennis H. 1975. Voice Onset Time, Frication, and Aspiration in Word-Initial Consonant Clusters. *Journal of Speech and Hearing Research* 18:686–706.

Ladd, D. Robert. 1984. Declination: A Review and Some Hypotheses. *Phonology Yearbook* 1:53–74.

Ladefoged, Peter, and Ian Maddieson. 1996. *The sounds of the world's languages*. Cambridge, MA: Blackwell Publishing.

Liberman, Mark, and Alan Prince. 1977. On Stress and Linguistic Rhythm. *Linguistic Inquiry* 8:249–336.

Literary Devices. 2019. Enjambment. *Literary Devices* .

Marian, Viorica. 2017. Orthographic and Phonological Neighborhood Databases across Multiple Languages. *Writ Lang Lit* 20:2–27.

Pope, Alexander. 1735. Epistle to Dr. Arbuthnot.

Small, Judy Jo. 2010. *Positive as Sound: Emily Dickinsons's Rhyme*. Athens: University of Georgia.

Tsujimura, Natsuko, and Stuart David. 2010. Rhyme and the Reinterpretation of Hip Hop in Japan 157–171.

Walser, Robert. 1995. Rhythm, Rhyme, and Rhetoric in the Music of Public Enemy. *Ethnomusicology* 39:193–217.

Weber, Andrea, and Roel Smits. 2017. Consonant And Vowel Confusion Patterns By American English Listeners. *15th ICPhS Barcelona* .