

A Segment-specific Metric for Quantifying Participation in Harmony

Avery Ozburn

University of British Columbia

1 Introduction

There is a long history of phonological research into quantifying and analyzing lexically gradient co-occurrence patterns between segments. This research has most frequently dealt with dissimilatory patterns (i.e. OCP effects) among consonants (e.g. Pierrehumbert 1993, Frisch et al. 2004), though there has also been similar work on assimilatory patterns (e.g. Rose & King 2007, Brown 2008, Arsenault 2012). However, most of this work has focused primarily on consonant co-occurrences. With respect to lexically gradient vowel co-occurrence patterns, specifically vowel harmony, the existing work focuses more on modelling variability in heteromorphemic contexts (e.g. Hayes & Londe 2006 on Hungarian), gradient carry-over of categorical harmony into other domains like compounding (e.g. Martin 2007), and quantifying overall trends towards harmony in the vowel system as a whole (e.g. Harrison et al. 2004).

Despite all of this research on segmental co-occurrence patterns, there is little research into gradient morpheme-internal vowel harmony patterns. Moreover, there remains no systematic way of calculating the gradient degree to which a segment participates in a harmony system. The standard Observed/Expected (O/E) values give a measure of the behaviour of each individual pair of segments, but a comprehensive picture of how harmonic a segment is, across its co-occurrences with all segments, is still lacking. Such a measure of overall segment-specific gradient harmony participation is a useful concept because it could illuminate gradient language-internal and cross-linguistic patterns in harmony participation that are not apparent from more categorical descriptions or entirely clear from O/E values.

In this paper, I adopt the statistical concept of *relative risk* as a measure of participation in harmony. I compute both O/E values and the relative risk measure for vowels in corpora of three languages with front/back harmony, to show that relative risk can capture the intuitive notion of how much a vowel participates in harmony. I then consider the implications of the results, given what is known about categorical trends of participation in front/back harmony systems in other languages. This measure does in fact provide new insights into harmony participation in these languages.

2 Participation in harmony

The question of which vowels participate in harmony is a long-standing one in the literature. For example, it is well-known that non-low, front, unrounded vowels are neutral to front/back harmony in Finnish and Hungarian, and this pattern extends to a variety of other languages (e.g. Kiparsky & Pajusalu 2003). Specifically, [i] is a common neutral vowel, and neutrality of lower front unrounded vowels like [e] typologically implies neutrality of [i] (e.g. Anderson 1980, Benus 2005, Finley 2015). There are clear typological trends in which vowels tend to be neutral to a given type of harmony, and moreover these patterns often have a phonetic basis (e.g. Archangeli & Pulleyblank 1994; Ozburn in press). For example, Beddor et al. (2001) argue that the tendency of [i] and [e] to be neutral in front/back harmony systems results from smaller effects of vowel-to-vowel coarticulation on them. Similarly, Benus & Gafos (2007) argue that neutrality of [i] to front/back harmony can be explained by its acoustic/perceptual stability under coarticulatory displacement along the front-back dimension; lower vowels have less of this stability. This property is used to explain the *height effect* of neutrality in Hungarian, in which, for example, [i] is more

* Thank you to Paul Tupper for discussion of the quantification measure, and to Gunnar Hansson, Douglas Pulleyblank, and the attendees of AMP 2018 for feedback on this work.

neutral than [e:], even though [e:] is still frequently neutral (e.g. Hayes & Londe 2006). As such, participation and neutrality are not clear-cut categories in Hungarian. This type of intermediate neutrality suggests that participation is more complex than categorical descriptions would suggest.

While non-participation in harmony is commonly linked to inventory structure, with neutral vowels being the ones that lack harmonic counterparts, this cannot be the entire story; at least some cross-linguistic patterns of participation are independent of contrast. For example, in Mayak (Nilotic; Andersen 1999), low vowels are generally neutral to ATR harmony, but contrastive for ATR; this corresponds to a general cross-linguistic trend for low vowels to be neutral to ATR harmony (Archangeli & Pulleyblank 1994). Moreover, vowels that lack counterparts may nonetheless participate in harmony by *re-pairing* with vowels that differ in an additional feature value (Baković 2000). For example, Maasai (Nilotic) has no [+ATR] low vowel, but its [-ATR] low vowel [a] alternates with the mid vowel [o] for tongue root harmony in certain contexts (e.g. Baković 2000). Such patterns provide clear evidence that whatever drives the typological tendencies in participation goes beyond mere contrast and inventory structure.

3 Gradient harmony and quantification

As previously mentioned, it is well-known in the literature that harmonic co-occurrence restrictions can be gradient; this occurs for instance in laryngeal consonant harmony restrictions in the Ethio-Semitic languages Chaha and Amharic (e.g. Rose & King 2007). The typical measure used to quantify gradient co-occurrence restrictions is observed over expected (O/E) values (e.g. Pierrehumbert 1993; Frisch et al. 2004; Coetzee & Pater 2008).¹ O/E indicates, for each pair of segments, how over- or under-represented the co-occurrence of those segments is. This means that O/E values, as traditionally used, can only express the participation of a specific pair in co-occurrence restrictions. For example, O/E can say whether [i] followed by [a] is over- or under-represented, but does say whether [i] in general tends to follow harmonic co-occurrence restrictions, as compared to other vowels.²

Beyond O/E, harmony has also been quantified overall for a language (Harrison et al. 2004). This value expresses the extent to which the percentage of harmonic co-occurrences in a corpus exceeds chance, given the relative proportions of each harmonic category in the corpus. This measure requires stipulating which vowels are neutral, and is an overall measure, not segment-specific.

Missing from the gradient quantification of harmony is a segment-specific measure of participation. Such a measure could provide interesting new insights because of the cross-linguistic patterns of participation described in Section 2. Given that these typological trends can exist independently of inventory contrast, and given that harmony can be gradient, we might expect the categorical trends in participation/neutrality to also manifest in a gradient way. For example, given that [i] is a common neutral vowel in front/back harmony, we might predict that [i] should participate gradiently less in front/back harmony even in languages in which it is not categorically a neutral vowel. However, there is no existing way of measuring this. In particular, it is insufficient to measure participation with the percentage of the segment's co-occurrences that are harmonic, because that fails to take into account the relative frequency of the two harmonic categories; if one category is more frequent, then the likelihood of segments in that category being in harmonic co-occurrences is increased.

4 Relative risk measure

I propose to adopt relative risk as a way of computing segment-specific harmony participation.

¹ Note that other work, in particular Graff & Jaeger (2009) and Wilson & Obdeyn (2009), has pointed out problems with the O/E metric and argued for the rejection of its use.

² It is worth noting, however, that it would be possible to get a general measure of participation using O/E, by constructing a table that groups together the vowels not under consideration by category. For example, for [i] in a front/back system, such a table would have rows/columns for [i], non-[i] front vowels, and back vowels. Performing O/E on such a table could give two general measures for [i], one of how over- or under-represented it is with front vowels, and one of how over- or under-represented it is with back vowels. Such a measure is slightly different from the one pursued in this paper, and is worth considering further in future work. One reason for using the measure in this paper, over the O/E method, is because it combines both harmonic and disharmonic behaviour into a single measure, rather than having two separate O/E values.

Relative risk (RR) is a standard statistical measure for count data, used to determine the relative probability of an ‘event’ in one condition versus another; it answers whether the ‘event’ is more likely in one condition, and if so by how much (Agresti 2013). It has previously been used in linguistics as a measure for speech errors (Tupper & Alderete in prep). The general formula for RR is as in (1), where $P(E|C_x)$ is the probability of event E given condition C_x .

$$(1) \quad RR = P(E|C_1)/P(E|C_2)$$

Since this measure is a ratio of probabilities in different conditions, it corrects for the relative frequency of C_1 and C_2 .

Interpreting the ‘event’ as the segment of interest, and the conditions as harmonic versus non-harmonic contexts, RR provides a measure of how much more frequent a segment is in harmonic contexts, and therefore how harmonic the segment is. In cases where all relevant counts are known, as in a corpus study of vowel co-occurrence patterns, the formula in (1) can be simplified to one that references only counts, rather than probabilities. The RR for vowel V from category C_1 is given in (2).

$$(2) \quad RR(V) = \frac{(\# V \text{ in context } C_1) \times (\# \text{ anything in context } C_2)}{(\# V \text{ in context } C_2) \times (\# \text{ anything in context } C_1)}$$

This formula can also be re-written as in (3), as the ratio of ratios. These two formulas are equivalent.

$$(3) \quad RR(V) = \frac{\frac{(\# V \text{ in context } C_1)}{(\# V \text{ in context } C_2)}}{\frac{(\# \text{ anything in context } C_1)}{(\# \text{ anything in context } C_2)}}$$

Note that if the vowel V belongs to category C_2 , then the position of C_1 and C_2 in the formula will be reversed. Thus, this measure effectively multiplies the ratio of V in harmonic to disharmonic contexts (i.e. $(\# V \text{ in context } C_1) / (\# V \text{ in context } C_2)$) by a constant for each category: if V is in category C_1 , the constant is as in (4); if V is in category C_2 , then the constant is the inverse of this, as in (5).

$$(4) \quad \text{constant for V in category } C_1 = \frac{(\# \text{ anything in context } C_2)}{(\# \text{ anything in context } C_1)}$$

$$(5) \quad \text{constant for V in category } C_2 = \frac{(\# \text{ anything in context } C_1)}{(\# \text{ anything in context } C_2)}$$

If V is in category C_1 , then the context C_1 also includes V. Thus, in order to avoid double-counting occurrences, it is simpler to treat this formula as directional, computed twice for each vowel: once for all co-occurrences involving V as the first of the two vowels (‘V1’), and once for co-occurrences where V is the second of the two vowels (‘V2’).³ In the former case, the ‘context’ refers to the second vowel, while in the latter case, it refers to the first vowel. This approach has the additional advantage of determining whether there are directional asymmetries in harmony participation, which we might expect to be the case, for instance, if harmony applies strictly in one direction triggered by only a single value of the feature.

One way of thinking of $RR(V)$, as shown in (3), is that it represents the ratio of occurrences of V in harmonic contexts to disharmonic contexts, relative to the ratio of occurrences of V’s harmonic versus disharmonic contexts in general. As such, the higher $RR(V)$ is, the more harmonic V is; a high $RR(V)$ indicates that V occurs more in harmonic than disharmonic contexts, relative to what we would expect given the relative frequency of those contexts. For the purposes of this study, I will use $RR(V)$ as a relative measure, comparing it across vowels to establish what counts as high and low values.

To illustrate how $RR(V)$ works, I will briefly consider several toy examples. The relevant factors to

³ Note that ‘V1’ and ‘V2’ are labels for relative position, and they do not necessarily reflect absolute position within the word. For example, if considering the co-occurrence between the third and fourth vowels within a given word, ‘V1’ would be the third vowel, as it is the first of the vowels under consideration.

consider are whether V tends to occur in harmonic contexts or is relatively evenly distributed between contexts, and whether or not the distribution of contexts (i.e. harmonic categories) is skewed, and if so in which direction. Table 1 gives toy counts and $RR(V)$ values for each of the six resulting contexts to consider.

	Symmetric category distribution	Skewed category distribution (to category of V)	Skewed category distribution (away from category of V)
V generally harmonic	# V in context $C_1 = 50$ # V in context $C_2 = 10$ # any in context $C_1 = 250$ # any in context $C_2 = 250$ $RR(V) = 5$	# V in context $C_1 = 50$ # V in context $C_2 = 10$ # any in context $C_1 = 400$ # any in context $C_2 = 100$ $RR(V) = 1.25$	# V in context $C_1 = 50$ # V in context $C_2 = 10$ # any in context $C_1 = 100$ # any in context $C_2 = 400$ $RR(V) = 20$
V relatively evenly distributed between contexts	# V in context $C_1 = 30$ # V in context $C_2 = 30$ # any in context $C_1 = 250$ # any in context $C_2 = 250$ $RR(V) = 1$	# V in context $C_1 = 30$ # V in context $C_2 = 30$ # any in context $C_1 = 400$ # any in context $C_2 = 100$ $RR(V) = 0.25$	# V in context $C_1 = 30$ # V in context $C_2 = 30$ # any in context $C_1 = 100$ # any in context $C_2 = 400$ $RR(V) = 4$

Table 1: Values for toy examples, with vowel V in category C_1

As is evident, both the distribution of V and the distribution of categories affects the $RR(V)$ value. If both V and the contexts are evenly distributed, then $RR(V) = 1$. This would also be true if both V and the contexts are skewed in exactly the same way (e.g. if V occurs 5 times more often in its harmonic context, and its harmonic context occurs 5 times more often than the other context). If V co-occurs with other vowels in its harmonic category less frequently than we would expect based on the relative frequency of the categories, as is the case if V is evenly distributed but the category distribution is skewed towards the category of V , then $RR(V) < 1$. If V occurs in its harmonic category more often than we would expect based on the distribution of categories, then $RR(V) > 1$, and the degree to which $RR(V)$ exceeds 1 depends on how much more harmonic V is than we would expect. For example, if V is generally harmonic, but the contexts are skewed such that V 's category occurs more frequently, then $RR(V)$ will be lower than if the counts for V are the same but the categories are distributed symmetrically. In turn, this latter $RR(V)$ will be lower than with the same counts for V , but contexts skewed such that the category to which V does not belong occurs more often. This property can be seen by comparing $RR(V)$ values across a row in Table 1. As such, $RR(V)$ takes into account not only how harmonic V is based on its raw counts, but also how harmonic we would expect it to be given the overall distribution of the categories; the higher the $RR(V)$ value, the more harmonic V is, compared to what we would expect given the category distribution. Any asymmetry in the category distribution will translate accordingly into an asymmetry in $RR(V)$ values.

It is worth noting that within a category, vowels can be compared simply by their own ratio (i.e. $\#V$ in context $C_1 / \#V$ in context C_2 , for each V in category C_1), but this is not the case across categories. For example, in the toy example in Table 1, say that there is one vowel, V_a , that is generally harmonic, and a second vowel, V_b , that is evenly distributed between the categories. If V_a and V_b both belong to category C_1 , then the ratio between their $RR(V)$ values will be constant regardless of the distribution of their categories; in Table 1, this ratio is 5, as can be seen by comparing values within a column. The reason for this consistency is that the ‘correction factor’ for category frequency, namely the ratio of category distributions, is the same in both cases if both vowels are in the same category (see (4-5)). In contrast, if V_a belongs to category C_1 and V_b belongs to category C_2 , then it is crucial to use $RR(V)$, rather than a pure ratio, as the ‘correction factor’ is different. This fact becomes important when the distribution is skewed towards one of the categories. For example, if the numbers are as in Table 1, V_a is generally harmonic and in category C_1 , V_b is evenly distributed and in category C_2 , and the distribution is skewed towards category C_1 , then $RR(V_a)$ should be 1.25 and $RR(V_b)$ should be 4. These numbers indicate that V_b is more harmonic than V_a , even though V_b is evenly distributed across categories, because the categories themselves are not evenly distributed. Not correcting for the distribution of categories in such an example would make it seem like V_a is the more harmonic vowel, which is false in this case. Thus, $RR(V)$ allows for comparison across

categories, whereas a simple ratio would not.

5 Case studies

For $RR(V)$ to be a valid and useful measure, it needs to correspond to other notions of how much a vowel participates in harmony. While O/E does not provide a single, unified measure of how harmonic a vowel is, it does give relevant information. Specifically, using O/E values, we would consider a vowel highly harmonic if it is highly under-represented in disharmonic contexts (O/E close to 0) and over-represented in harmonic contexts (O/E greater than 1). Vowels with such patterning should have the highest $RR(V)$ values. In contrast, we would consider a vowel not particularly harmonic if its representation in harmonic and disharmonic contexts is either approximately what we would expect from random distribution (O/E close to 1), or if it is under-represented in harmonic contexts and over-represented in disharmonic ones. Vowels of this type should have the lowest $RR(V)$ values.

While the toy example in Table 1 already shows that $RR(V)$ does correspond to the notion of how harmonic a vowel is, relative to category frequency, it is ideal to illustrate it with actual language data. Thus, to test this measure, I examine corpora of three languages with front/back harmony: Chuvash, Tatar, and Mari (Luutonen et al. 2007; Luutonen et al. 2016). All of these languages are spoken in a similar area in Russian; the former two are Turkic languages, while Mari is Uralic. Using the software Phonological CorpusTools (Hall et al. 2015), counts of all pairs of vowels adjacent on a vowel tier were obtained for all of the corpora. These were used to compute O/E values for all vowel pairs, as well as $RR(V)$ for each vowel. This work is exploratory, and does not make any concrete claims about harmony in any of these languages. Notably, the corpora are dictionary lists that include derived forms and are written in orthography, not transcription. The entirety of the corpora were used in the computational analysis to get a picture of broad harmonic behaviour throughout the lexicon. However, there may be complications due to the inclusion of derived forms and loans, as well as in the orthography-to-phonology mapping, that could skew or distort the counts, and further work should examine these issues.

As mentioned, the intuitive notion of how much a vowel participates in harmony can be captured by looking at O/E values; essentially, if a vowel occurs with very low O/E values with vowels of the opposite harmonic class, then it is highly harmonic, and vice versa. To check whether $RR(V)$ is a valid measure of participation, we must compare it to what we observe through O/E values across a vowel's co-occurrences with all other vowels. Vowels that are highly harmonic given their O/E values should have high $RR(V)$, and less harmonic vowels should have lower $RR(V)$.

If $RR(V)$ is a valid measure, then we can next examine the kinds of patterns that emerge in participation. One hypothesis is that relative $RR(V)$ values should correspond to categorical cross-linguistic participation trends. In particular, while [i] is not consistently described as neutral in the languages examined here, it is often categorically neutral in front/back harmony in related languages, like Finnish and Hungarian, in the Uralic family with Mari, and Uyghur in the Turkic family with Chuvash and Tatar (e.g. Ringen & Heinämäki. 1999 on Finnish; Siptár & Törkenczy 2000 on Hungarian; Lindblad 1990 on Uyghur).⁴ Moreover, as described in Section 2, [i] is considered a good neutral vowel in front/back harmony. We might therefore predict that the $RR(V)$ value of [i] should be low (i.e. among the lowest of the values) across these languages, and across both environments (V1 and V2 positions). Specifically, we might expect $RR(i)$ to be lower than the $RR(V)$ value of front vowels that are lower and/or rounded, as these are known factors for reducing the likelihood of a vowel being neutral to front/back harmony (e.g. Benus and Gafos 2007).

5.1 Chuvash Table 2 gives the O/E values for all combinations of Chuvash vowels. Values above 1 indicate over-representation, while values below 1 indicate under-representation; the closer a value is to 0,

⁴ Whether [i] is neutral in the languages examined here is not entirely clear. The descriptions of vowel harmony in Chuvash and Tatar are limited, but do not mention neutrality; moreover, Chuvash [i] has a back counterpart in some contexts, and there is a diphthong often taken to be a back counterpart of [i] in Tatar. (See e.g. Krueger 1961 and Clark 2015 on Chuvash; Berta 2015 on Tatar.) These facts suggest that [i] is likely not categorically neutral in Chuvash and Tatar. The case of Mari is less clear; Kangasmaa-Minn (1998) says that [i] and [e] are neutral vowels in Mari harmony, while Vaysman (2009) does not describe any neutrality in the language.

the closer the co-occurrence is to categorically absent. The darker lines in the table separate the back and front categories, while the grey boxes indicate instances where harmonic co-occurrences are under-represented or disharmonic ones are over-represented. It is particularly worth noting that the only O/E values of disharmonic co-occurrences (i.e. the upper right and lower left quadrants) that approach or exceed 1 involve [i], and that 5 of the 8 disharmonic co-occurrences involving [i] have high O/E values, above 0.8. These numbers are suggestive of [i] co-occurring regularly with back vowels.

V1 ↓ V2 →	a	ǔ	u	u	e	ǝ	i	y
a	1.30	1.14	1.36	1.54	0.32	0.08	1.49	0.06
ǔ	1.41	2.03	0.94	0.84	0.02	0.11	0.27	0.09
u	1.26	2.22	0.34	0.89	0.09	0.02	0.35	0.10
u	1.23	1.62	1.56	1.20	0.31	0.14	0.92	0.15
e	0.48	0.15	0.52	0.61	2.26	1.64	1.55	3.98
ǝ	0.04	0.04	0.62	0.26	2.70	4.43	0.24	3.18
i	0.88	0.24	0.86	0.91	1.79	1.49	1.21	0.45
y	0.03	0.02	0.00	0.17	3.03	4.24	0.15	1.95

Table 2: Chuvash O/E values

Table 3 gives the RR(V) values for all Chuvash values, in both V1 and V2 context, while Figure 1 shows the same information in graphical representation. Noteworthy is the fact that [a,u,i] have very low RR(V) in both contexts, while in general the more marked vowels, like [y], have higher RR(V). The one exception is [u] in V2 context; however, the distribution of this vowel is restricted in non-initial position in Chuvash (Krueger 1961), suggesting that its behaviour here may be more reflective of a very small number of examples. Besides [u], the vowels [a,u,i] have the lowest RR(V) in both contexts, and in particular, [i] has lower RR(V) than front vowels that are lower and/or rounded.

V	RR(V) (V1 position)	RR(V) (V2 position)
a	2.09	2.77
ǔ	13.32	10.64
u	10.31	1.94
u	3.06	2.09
e	5.13	10.26
ǝ	40.94	25.28
i	2.30	1.03
y	66.85	28.81

Table 3: Chuvash RR(V) values

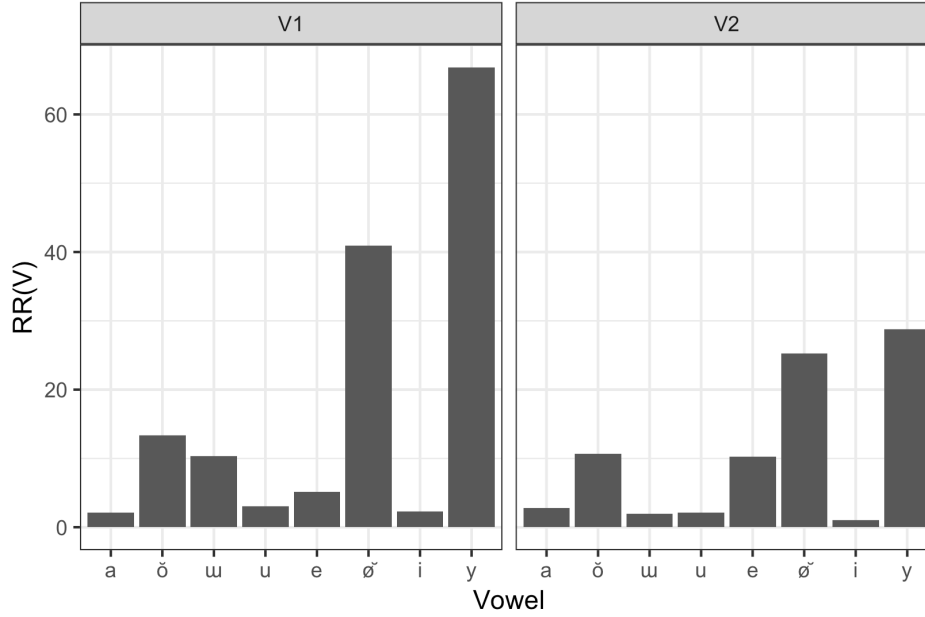


Figure 1: Chuvash RR(V)s

5.2 Tatar Table 4 shows the O/E values for combinations of Tatar vowels. As in Chuvash, Tatar [i] occurs regularly with back vowels, and often not as regularly as expected with front vowels. In Tatar, [o] demonstrates a similar property, occurring fairly regularly with front vowels but under-represented in many back vowel contexts. For disharmonic co-occurrences (upper-right and lower-left quadrants), the co-occurrences that approach or exceed 1 are [i] with [a, o] in either order, and [e] with [o] in either order. However, [i] and [e] differ in terms of their co-occurrences with front vowels, in that the O/E of co-occurrences of [e] with front vowels consistently exceeds 1, while 5 of the 9 cases of [i] with front vowels have O/E less than 1. In terms of [o], 5 of the 7 co-occurrences of [o] with back vowels have O/E below 1, and all back-back co-occurrences with O/E less than 1 involve [o] as one of the two vowels.

V1 ↓ V2 →	a	o	ɤ	u	æ	ø	e	y	i
a	1.24	0.90	1.48	2.01	0.21	0.49	0.31	0.03	1.29
o	1.32	2.60	0.94	0.42	0.04	0.46	0.82	0.03	1.80
ɤ	1.27	0.18	2.50	2.27	0.09	1.22	0.14	0.04	0.08
u	1.94	0.44	1.76	0.61	0.08	1.23	0.32	0.03	0.84
æ	0.23	0.03	0.08	0.07	2.68	0.86	1.80	4.75	0.63
ø	0.17	0.08	0.03	0.04	3.40	3.98	2.97	0.71	0.30
e	0.61	1.34	0.27	0.21	1.52	2.17	1.88	2.12	1.09
y	0.12	0.08	0.06	0.06	4.23	1.48	2.49	0.59	0.31
i	1.04	2.05	0.24	0.27	1.50	0.39	1.43	0.57	1.23

Table 4: Tatar O/E values

Table 5 and Figure 2 illustrate RR(V) values for Tatar vowels. Notably, [o] and [i] show the lowest values in both contexts, and [a, u, e] are also quite low. As in Chuvash, marked vowels like [y] and [ɤ] show higher RR(V) values, though [ø] in V2 context has one of the lowest values. Worth noting again is that [i] has a lower RR(V) value than front vowels that are lower and/or rounded.

V	RR(V) (V1 position)	RR(V) (V2 position)
a	2.91	2.35
o	1.60	0.97
ɤ	15.89	9.04
u	4.33	9.93
æ	15.52	14.81
ø	22.91	1.89
e	3.15	5.09
y	25.94	74.81
i	1.65	0.84

Table 5: Tatar RR(V) values




Figure 2: Tatar RR(V)s

5.3 Mari Table 6 illustrates O/E values for combinations of Mari vowels. Again, [i] co-occurs regularly with back vowels; the same is true of [e], with most of its O/E values with back vowels nearly at 1. In fact, O/E values for [e] are consistently close to 1 across back vowels, while [i] has lower O/E values with [ɤ]. However, as in Tatar, the situation in front-front co-occurrences is different for [i] versus [e], in that [i] shows a far greater tendency than [e] to not co-occur with other front vowels (other than itself and [e]). Interestingly, in Mari, [æ] and [ø] in V1 position co-occur with each other at much higher than expected rates, while rarely occurring with any other vowels of either category. In V2 position, these vowels again occur most often with only each other. It is notable that in at least some dialects of Mari, [æ] is described as occurring only as a result of harmony (Vaysman 2009).

V1 ↓ V2 →	a	o	u	ɤ ⁵	æ	ø	y	e	i	ə
a	1.32	0.71	1.17	1.12	0.06	0.16	0.93	0.85	1.26	0.05
o	0.96	2.05	0.96	1.09	0.04	0.20	0.64	1.05	1.43	0.02
u	1.01	1.50	1.36	1.38	0.04	0.37	0.77	0.84	0.99	0.04
ɤ	1.41	0.60	0.63	1.37	0.01	1.20	0.88	0.92	0.26	0.02
æ	0.06	0.10	0.24	0.08	8.46	0.07	0.84	0.26	0.27	6.14
ø	0.72	0.27	0.60	1.16	1.64	13.90	1.72	1.01	0.25	1.81
y	0.68	0.37	0.66	1.44	1.15	9.32	4.37	1.11	0.13	1.52
e	0.71	1.09	1.19	0.62	1.49	0.39	1.34	1.60	1.78	1.29
i	0.67	1.84	1.93	0.46	1.55	0.44	0.58	1.40	1.69	1.22
ə	0.02	0.11	0.18	0.02	6.50	0.43	0.60	0.60	0.12	8.63

Table 6: Mari O/E values

Table 7 and Figure 3 show RR(V) values for Mari vowels. What is most noteworthy in Mari is that two vowels, [æ] and [ə], have substantially higher RR(V) values than the remaining vowels, for both V1 and V2 contexts. Among the remaining vowels, the differences are so small, though [i] again has one of the lowest RR(V) values in both contexts.

V	RR(V) (V1 position)	RR(V) (V2 position)
a	1.86	2.36
o	1.55	1.14
u	2.12	0.94
ɤ	2.88	2.23
æ	37.75	79.87
ø	1.74	3.82
y	1.39	1.59
e	2.05	1.30
i	1.78	1.16
ə	74.11	90.27

Table 7: Mari RR(V) values

⁵ This vowel is often transcribed as [ɛ̃] in Uralic literature (e.g. Kangasmaa-Minn 1998); I have written it as [ɤ] to emphasize its status as a back mid unrounded vowel, but it should be noted that it is reduced.




Figure 3: Mari RR(V)s

6 Discussion

The results of all three languages show a close correspondence between RR(V) and the notion of how much each vowel is participating in harmony, based on its O/E values. For example, in Chuvash, the O/E values of [y] in V1 position before a back vowel are very low, ranging from 0 to 0.17; correspondingly, RR(y) for V1 is very high, at 66.85. In contrast, the O/E values for [i] in V1 position before a back vowel in Chuvash are very high, ranging from 0.24 to 0.91, with most of the cases close to expected representation; correspondingly, RR(i) is very low in the V1 condition, at 2.3. The vowel [u] in Chuvash is intermediate between [i] and [y], with O/E values of 0.02 to 0.35 when it occurs in V1 position before a front vowel, and RR(u) in V1 condition intermediate at 10.31. A similar correspondence holds across other vowels and other languages, suggesting that RR(V) is indeed a good measure of how much a vowel participates in harmony, where participation is viewed in terms of rarely co-occurring with vowels of the opposite harmonic class.

The concept of RR(V) allows us to look for patterns in gradient participation. Noticeable across all three languages is the fact that RR(i) is quite low, meaning it has a low degree of participation in the harmony system. In particular, in both Chuvash and Tatar, [i] has the lowest RR(V) of any front vowel in both positions, and one of the lowest overall. This result is particularly interesting because, across front/back harmony systems, [i] is a common transparent vowel, and this pattern has often been connected to the lack of a back, high, unrounded counterpart to [i] in the inventory in languages with transparent [i], like Finnish and Hungarian (e.g. Vago 1973). As described above, [i] is not consistently described as neutral to harmony in any of the languages examined here, and in Chuvash it even has a back counterpart, yet RR(i) shows that it nonetheless participates less in the harmony system, in a gradient way. Specifically, [i] in both Chuvash and Tatar shows less overall participation in harmony than lower and/or rounded front vowels, which are less frequently transparent and neutral in harmony systems. This result suggests that the cross-linguistic behaviour of [i] as a neutral vowel could extend beyond categorical properties and be due to reasons beyond harmonic pairing. As such, it suggests a need to re-examine more carefully the factors involved in harmony participation.

Beyond [i], it is notable that cardinal vowels in general tend to have lower RR(V), and are therefore less harmonic. For example, [a] and [u] also have low RR(V) values in Chuvash, and [a], [o], and [e] have low values in Tatar. One possible explanation for this result is that it is due to non-harmonic loans from Russian, which has neither harmony nor front rounded or back unrounded vowels. Indeed, Chuvash, Tatar, and Mari are all spoken in Russia and could be expected to have many recent, disharmonic Russian loans.

Removing loans from the corpus would test this hypothesis. Alternatively, this result could illustrate a more general property about how vowels participate in harmony systems, with more marked vowels being more required to harmonize. Indeed, across all three languages, particularly in Chuvash and Tatar, the more marked vowels have consistently higher RR(V) values, except in cases where their distribution is limited (e.g. V2 position of [u] in Chuvash). This observation is consistent with existing ideas in the literature on categorical harmony, in which harmony has been argued to be a way to enhance the perceptibility of marked features (e.g. Kaun 1995).

The results from Mari are more complex and puzzling than those of Chuvash and Tatar; two vowels have very high RR(V) values, and the others are approximately equivalent. The vowels with high RR(V) were [æ] and [ə]; the former is perhaps expected, since it is more marked, but the latter raises questions about the nature of the front reduced vowel in Mari. Specifically, there is a general relationship in Chuvash and Tatar between markedness and high RR(V) values, yet the schwa symbol typically denotes an unmarked vowel, so that relationship does not seem to hold in this case in Mari. Moreover, the other Mari vowels all behave essentially the same way, regardless of whether or not they are marked or exist in Russian, and the reasons for this result are unclear. Whether it says something about the Mari phonological system, the morphology (e.g. complex words with non-harmonizing vowels tending to avoid [æ] and [ə]), or loanwords from Russian remains to be investigated.

7 Conclusion

In summary, this paper has explored a method of calculating segment-specific participation in vowel harmony, using relative risk. I have shown how this computation captures the notion of participation by examining relative risk in comparison to O/E values across corpora of three different front/back harmony systems: Chuvash, Tatar, and Mari. I then examined what the results imply about gradient participation. Specifically, I showed that [i] has a low degree of participation across these harmony systems, which corresponds to the fact that [i] is often neutral cross-linguistically, but is interesting because this neutrality is often considered categorical in other languages, in a way not consistently described for the languages considered here. The relative risk measure provides a way of looking at participation in harmony in a more nuanced way, and it allows for cross-linguistic comparison; this approach can therefore provide significant new insights into harmony systems.

References

- Agresti, Alan. 2013. *Categorical data analysis*. John Wiley & Sons.
- Andersen, Torben. 1999. Vowel harmony and vowel alternation in Mayak (Western Nilotic). *Studies in African Linguistics* 28(1), 1-30.
- Anderson, Lloyd B. 1980. Using asymmetrical and gradient data in the study of vowel harmony. In Robert Vago (ed.), *Issues in vowel harmony*, 271–340. Amsterdam: John Benjamins.
- Archangeli, Diana & Douglas Pulleyblank. 1994. *Grounded Phonology*. Cambridge, Massachusetts: MIT Press.
- Arsenault, Paul. 2012. Retroflex Consonant Harmony in South Asia. PhD dissertation, University of Toronto.
- Baković, Eric. 2000. *Harmony, Dominance and Control*. PhD dissertation, Rutgers University.
- Beddor, Patrice, Rena A. Krakow, & Stephanie Lindemann. 2001. Patterns of perceptual compensation and their phonological consequences. In Elizabeth Hume & Keith Johnson, *The role of perceptual phenomena in phonology*, 55-78.
- Benus, Stefan. 2005. *Dynamics and transparency in vowel harmony*. PhD dissertation, New York University.
- Benus, Stefan & Adamantios Gafos. 2007. Articulatory characteristics of Hungarian transparent vowels. *Journal of Phonetics* 35:271-300.
- Berta, Árpád. 2015. Tatar and Bashkir. *The Turkic Languages*, ed. Lars Johanson & Éva Á. Csató. New York : Routledge.
- Brown, Jason Camy. 2008. *Theoretical aspects of Gitksan phonology*. PhD dissertation, University of British Columbia.
- Clark, Larry. 2015. Chuvash. *The Turkic Languages*, ed. Lars Johanson & Éva Á. Csató. New York : Routledge.
- Coetzee, Andries W., & Joe Pater. 2008. Weighted constraints and gradient restrictions on place co-occurrence in Muna and Arabic. *Natural Language & Linguistic Theory*, 26(2), 289-337.

- Finley, Sara. 2015. Learning nonadjacent dependencies in phonology: Transparent vowels in vowel harmony. *Language* 91(1), 48-72.
- Frisch, Stefan A., Janet B. Pierrehumbert & Michael B. Broe. 2004. Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22: 179–228.
- Graff, Peter, & T. Florian Jaeger. 2009. Locality and feature specificity in OCP effects: Evidence from Aymara, Dutch, and Javanese. In *Proceedings from the annual meeting of the Chicago linguistic society* (Vol. 45, No. 1, pp. 127-141). Chicago Linguistic Society.
- Hall, Kathleen C., Blake Allen, Michael Fry, Scott Mackie & Michael McAuliffe. 2015. Phonological CorpusTools, Version 1.1.0. [Computer program]. Available from <https://sourceforge.net/projects/phonologicalcorpustools/>.
- Harrison, David, Emily Thomforde & Michael O’Keefe. 2004. The vowel harmony calculator. Available from http://www.swarthmore.edu/SocSci/harmony/public_html/.
- Hayes, Bruce & Zsuzsa Cziráky Londe 2006. Stochastic phonological knowledge: The case of Hungarian vowel harmony, *Phonology* 23, 59–104.
- Kangasmaa-Minn, Eeva. 1998. Mari. *The Uralic languages*, ed. Daniel Abondolo. New York: Routledge.
- Kaun, Abigail R. 1995. *The Typology of rounding harmony: an optimality theoretic approach*. PhD dissertation, UCLA.
- Kiparsky, Paul & Karl Pajusalu. 2003. Toward a typology of disharmony. *The Linguistic Review* 20, 217-241.
- Krueger, John R. 1961. *Chuvash manual*. Indiana University.
- Lindblad, Vern M. 1990. *Neutralization in Uyghur*. Master’s thesis, University of Washington.
- Luu-tonen, Jorma, et al. (ed.). 2007. *Electronic Word Lists: Mari, Mordvin and Udmurt. With SFOu WordListTool 1.3*. Lexica Societatis Fenno-Ugricae XXXI:1. ISBN 978-952-5150-98-8. Helsinki.
- Luu-tonen, Jorma, et al. (ed.). 2016. *Electronic Word Lists: Komi, Chuvash and Tatar. With SFOu WordListTool 1.4*. Lexica Societatis Fenno-Ugricae XXXI:2. ISBN 978-952-5667-79-0. Helsinki.
- Martin, Andrew T. 2007. *The evolving lexicon*. PhD dissertation, University of California, Los Angeles.
- Ozburn, Avery. In press. A target-oriented approach to neutrality in vowel harmony: Evidence from Hungarian. To appear in *Glossa*.
- Pierrehumbert, Janet B. 1993. Dissimilarity in the Arabic verbal roots. *Proceedings of NELS* 23: 367–381.
- Ringen, Catherine & Orvokki Heinämäki. 1999. Variation in Finnish vowel harmony: An OT account. *Natural Language & Linguistic Theory* 17(2), 303-337.
- Rose, Sharon, & Lisa King. 2007. Speech error elicitation and co-occurrence restrictions in two Ethiopian Semitic languages. *Language and Speech* 50(4): 451-504.
- Siptár, Péter & Miklós Törkenczy. 2000. *The phonology of Hungarian*. Oxford: Oxford University Press.
- Tupper, Paul, and John Alderete. In prep. Risk ratios for linguists.
- Vago, Robert. 1973. Abstract vowel harmony systems in Uralic and Altaic languages. *Language* 49, 579-605.
- Vaysman, Olga. 2009. *Segmental alternations and metrical theory*. PhD dissertation, MIT.
- Wilson, Colin & Marieke Obdeyn. 2009. Simplifying subsidiary theory: statistical evidence from Arabic, Muna, Shona, and Wargamay. Ms, Johns Hopkins University.