Phonology **23** (2006) 59–104. © 2006 Cambridge University Press
doi:10.1017/S0952675706000765 Printed in the United Kingdom

*Stochastic phonological knowledge: the case of Hungarian vowel harmony**

Bruce Hayes
Zsuzsa Cziráky Londe
University of California, Los Angeles

In Hungarian, stems ending in a back vowel plus one or more neutral vowels show unusual behaviour: for such stems, the otherwise general process of vowel harmony is lexically idiosyncratic. Particular stems can take front suffixes, take back suffixes or vacillate. Yet at a statistical level, the patterning among these stems is lawful: in the aggregate, they obey principles that relate the propensity to take back or front harmony to the height of the rightmost vowel and to the number of neutral vowels. We argue that this patterned statistical variation in the Hungarian lexicon is internalised by native speakers. Our evidence is that they replicate the pattern when they are asked to apply harmony to novel stems in a 'wug' test (Berko 1958). Our test results match quantitative data about the Hungarian lexicon, gathered with an automated Web search. We model the speakers' knowledge and intuitions with a grammar based on the dual listing/generation model of Zuraw (2000), then show how the constraint rankings of this grammar can be learned by algorithm.

1 Introduction: irregularity in phonology

Linguists sometimes have the luxury of working with systematic, exceptionless data. More often, we encounter data which cannot be reduced to a single general pattern. In phonology, the variation is usually lexical: a subset of stems fails to adhere to the most frequent data pattern. This article addresses the question of what the language-learning child does when she confronts such cases. We focus on the appearance of front- *vs.* back-vowel suffix allomorphs in Hungarian vowel harmony, as

* We would like to thank Stephen Anderson, Arto Anttila, Andrew Garrett, Matthew Gordon, Gunnar Hansson, Sharon Inkelas, Patricia Keating, Paul Kiparsky, Joe Pater, Janet Pierrehumbert, Catherine Ringen, Péter Siptár, Donca Steriade, Robert Vago, Colin Wilson and Kie Zuraw for helpful advice. We also received valuable input from three reviewers and the associate editor. As is usual, they are not to be held responsible for defects. We would also like to thank our many Hungarian language consultants for sharing their native speaker intuitions.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

for example in [fɔl-nɔ̲k] *falnak* 'wall-DAT' *vs.* [kɛrt-nɛ̲k] *kertnek* 'garden-DAT'.

Most previous work concerning irregularity in phonology has adopted an approach in which the majority pattern is characterised as regular, with some mechanism chosen to deal with the residual cases. For instance, in Halle & Mohanan's (1985) account of English past tenses, irregular forms are lexically marked to undergo minor rules. Thus, for instance, the irregular past tense form *clung* is lexically listed as /klɪŋ/ with a special diacritic mark, which causes it to undergo Backing Umlaut, a rule which has the effect of shifting the stem vowel from [ɪ] to [ʌ] in the past tense form.

Another possibility (Pinker & Prince 1988, Pinker 1999) is to use grammar to derive only the regular forms, and simply list all the irregulars in the lexicon; thus for Pinker & Prince, *clung* is underlyingly just /klʌŋ/. At first blush such an approach seems inadequate, as without amplification it cannot account for the fact that irregulars usually occur in patterns, such as *cling–clung*, *fling–flung*, *sling–slung*. There is evidence (Bybee & Moder 1983, Prasada & Pinker 1993, Albright & Hayes 2003) that such patterns are partly productive, so a pure-listing account fails to capture the native speaker's knowledge. Therefore, Pinker & Prince (1988) amplified their proposal with the idea that the memorised lexical entries for irregulars are embedded in a kind of associative network. This network would be able to generate novel irregulars by some sort of analogy, thus accounting for whatever productivity their patterns may have.

This approach is unsatisfactory to the extent that irregulars can be shown to be derived on the basis of principles of phonological theory, which presumably would be included only in the grammar, and not in the analogical network. Albright & Hayes (2003), comparing machine implementations of grammar and analogy, argue that English irregular past tenses must indeed be derived by a grammar and not by analogy; the cases thought by Pinker & Prince to be analogical are in fact better derived by grammatical principles acting on a small scale.

The Hungarian data reviewed below arguably form an especially clear case: they are the result of completely ordinary mechanisms of vowel harmony, so it is quite difficult to justify relegating the minority patterns to a totally different mechanism. Indeed, the minority patterns sometimes compete with the majority on a near-equal basis, which makes it especially arbitrary to claim that the majority results from grammar and the minority from analogy.

We argue here instead for a unified approach to irregularity, as developed by Zuraw (2000). In Zuraw's theory, all of the competing patterns are expressed in a single grammar, along with a characterisation of their relative strength. The grammar is implemented in stochastic Optimality Theory (Boersma 1997, Boersma & Hayes 2001). Individual inflected forms, which usually show invariant behaviour, are lexically listed where necessary (just as in Pinker & Prince's theory), and their invariance is guaranteed by high-ranking faithfulness constraints. When a speaker must

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

provide a novel inflected form (for instance, because she has never heard the stem in the relevant inflectional category), the stochastically ranked constraints of the grammar provide a *range* of options, each with a probability of being output. This probability is determined in the course of language learning, and approximately reflects the frequencies of the competing patterns as they appear in the lexicon. The basic prediction of the model is that the lexical frequencies, insofar as they reflect relevant phonological properties of stems, should give rise to a grammar that generates outputs at frequencies approximating the lexical frequencies.

This is a testable prediction, since we can get speakers to use their grammars to generate novel forms by asking them to inflect stems they have never heard before – this is the classical 'wug' test paradigm (Berko 1958). Zuraw's own work demonstrates a fairly good match between the frequency patterns of Nasal Substitution in the Tagalog lexicon and the intuitions of her wug-test subjects, and she develops a stochastic grammar that links the two. Similar work in other frameworks has likewise found a match between statistical patterns in the lexicon and gradient speaker intuition: Eddington (1996, 1998, 2004), Coleman & Pierrehumbert (1997), Berkley (2000), Bailey & Hahn (2001), Frisch & Zawaydeh (2001), Albright (2002), Albright & Hayes (2003), Ernestus & Baayen (2003) and Pierrehumbert (in press).

Here, we present a study similar to Zuraw's, focusing on Hungarian vowel harmony. As with earlier studies, we find that native speakers are good frequency matchers; in a wug test, the forms they volunteer closely match the statistical pattern of the Hungarian lexicon.

With this result in hand, we proceed to analysis, with three goals in mind. First, we develop a grammar within Zuraw's framework that accurately describes the wug-test intuitions of our consultants. Second, we extend the analysis to an area not addressed by Zuraw, namely the question of *impossible* harmony patterns; i.e. the characterisation of patterns that do not exist and (as we will claim) could not exist. Finally, we turn to the question of learnability, suggesting a way in which algorithms proposed in earlier work can be used to learn the variable Hungarian harmony pattern.

2 Hungarian vowel harmony

Hungarian vowel harmony has been the focus of great deal of research, which we have freely mined for generalisations and analytic insights. A non-exhaustive list of earlier work includes Esztergár (1971), Vago (1974, 1976, 1980a), Ringen (1975), Kontra & Ringen (1986), Hare (1990), Kornai (1991), Ringen & Vago (1995, 1998), Dienes (1997) and Siptár & Törkenczy (2000).

The vowels of Standard Hungarian are given below in (1), both in IPA and in orthography.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(1)

	front unrounded		front rounded		back	
high	[i] *i*	[iː] *í*	[y] *ü*	[yː] *ű*	[u] *u*	[uː] *ú*
mid		[eː] *é*	[ø] *ö*	[øː] *ő*	[o] *o*	[oː] *ó*
low	[ɛ] *e*				[ɔ] *a*	[aː] *á*

The vowel transcribed as [ɛ] is classified as low, since it is paired with the low vowel [ɔ] in the harmony system.

In what follows, it will be useful to express the vowel sequences of words in formulas, and for this purpose we adopt the following abbreviations: N (mnemonic for 'neutral') will designate the front unrounded vowels, F the front rounded vowels and B the back vowels. For example, in this notation, the word [ɔlbeːrløː] *albérlő* 'lodger' is BNF.

Hungarian is a richly inflected language, with dozens of suffixes. We will deal here only with the large class of suffixes that show a two-way alternation in backness; thus we will be ignoring the harmonically invariant suffixes (Vago 1980a: 15–18, Siptár & Törkenczy 2000: 65–66), as well as suffixes that show a three-way alternation based on backness and rounding (Vago 1980a: 18–19, Siptár & Törkenczy 2000: 72–74). Since the two-way suffixes generally behave alike, it suffices for present purposes to discuss just one of them, namely the dative, which appears as [-nɔk] *-nak* or [-nɛk] *-nek*, according to the principles of vowel harmony.

Vowel harmony depends on the vowels that appear near the end of the stem. For instance, if the last vowel of a stem is back, then no matter what vowels come earlier, the suffix must also be back, as shown in the examples of (2):

(2) BB [ɔblo̲ːk-nɔk] *ablaknak* 'window-DAT'
NB [biːro̲ː-nɔk] *bírónak* 'judge-DAT'
FB [glyko̲ːz-nɔk] *glükóznak* 'glucose-DAT'

Likewise, if the last vowel of a stem is a front rounded vowel, then the suffix vowel must be front:

(3) F [y̲ʃt-nɛk] *üstnek* 'cauldron-DAT'
BF [ʃofø̲ːr-nɛk] *sofőrnek* 'chauffeur-DAT'

Again, it does not matter what vowels occur earlier in the stem.

Most stems whose vowels are all front unrounded (N) take front suffixes:

(4) N [kɛ̲rt-nɛk] *kertnek* 'garden-DAT'
N [tsi̲ːm-nɛk] *címnek* 'address-DAT'
NN [rɛpɛ̲s-nɛk] *repesznek* 'splinter-DAT'

However, there are a few dozen exceptional all-N stems that take back suffixes, even though they contain no back vowels:

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(5)	N	[hiːd-nɔk]	*hídnak*	'bridge-DAT'
	N	[ʃiːp-nɔk]	*sípnak*	'whistle-DAT'
	NN	[dɛreːk-nɔk]	*deréknak*	'waist-DAT'

Following earlier usage, we will refer to these as *híd* stems, after the word for 'bridge' in (5). All of the *híd* stems but two are monosyllabic, and of these, most contain the vowel /iː/.

The remaining cases are those in which a harmonic vowel (F or B) precedes a string of one or more neutral vowels at the end of the stem. Of these, the stems of the form … FN and … FNN all take front harmony:

(6)	FN	[fyːsɛr-nɛk]	*fűszernek*	'spice-DAT'
	FNN	[øːriːzɛt-nɛk]	*őrizetnek*	'custody-DAT'

The most complex examples, which are the focus of this article, are stems of the type … BN, … BNN, etc. Here, we find extensive lexical idiosyncrasy (Vago 1980a: 14, 22, Siptár & Törkenczy 2000: 70–72): individual stems can require back suffixes or front suffixes, or allow both in free variation. Thus, for instance, Siptár & Törkenczy cite [hɔvɛr] *haver* 'pal' as a stem that takes only back suffixes ([hɔvɛr-nɔk]), [hotɛl] *hotel* 'hotel' as a stem that can take either front or back [hotɛl-nɔk] ~ [hotɛl-nɛk]) and [koːdɛks] *kódex* 'codex' as a stem that only takes front suffixes ([koːdɛks-nɛk]). A triplet from our own data, in this case with /… Beː/, is the following:

(7)	BN	[pɔlleːr-nɔk]	*pallérnak*	'foreman-DAT'
	BN	[ɔrzeːn-nɔk, ɔrzeːn-nɛk]	*arzénnak, arzénnek*	'arsenic-DAT'
	BBN	[mutɔgeːn-nɛk]	*mutagénnek*	'mutagen-DAT'

A particular speaker of Hungarian must therefore be assumed to memorise the vowel-harmony behaviour of individual BN(N) stems.

3 The statistical patterning

While it is not predictable in general whether a BN or BNN stem will take front or back harmony, there are clear tendencies present. If one knows what vowels such a stem contains, it is possible to guess, with far better than chance frequency, what kind of harmony it will take. The crucial generalisations have been studied by Vago (1974), Anderson (1980), Kontra & Ringen (1986), Farkas & Beddor (1987), Siptár & Törkenczy (2000), Benus (2005) and other scholars.

The first generalisation is what we will call the HEIGHT EFFECT, based on the height of the rightmost vowel in … BN: the phonologically low vowel [ɛ] occurs with front suffixes more often (that is, in proportionately more stems) than the mid vowel [eː], which occurs with front suffixes more often than the high vowels [i] and [iː]. Second, there is a COUNT EFFECT: BNN stems take front suffixes more often than BN stems do.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Since the sources cited above do not agree on the precise nature of the height and count effects, we have sought to collect as many data as possible on these quantitative patterns. We have followed two methods: elicitation from native speakers of a large number of stems, and a machine-based search of the World Wide Web. Since the latter has turned up more data, we will cover it first.

3.1 A search-engine study of the Hungarian lexicon

The basic method of collecting quantitative patterns for phonology by using a Web search engine was pioneered by Zuraw (2000). The idea is that where forms occur in free variation, we can measure their relative frequencies by counting the hits returned for each.

To see how this works, consider the forms from (7) above. A query for these forms using the Google search engine (12 May 2004) yielded the hit counts shown in (8).

(8)

		hits	
[mutɔgeːn-nɔk]	*mutagénnak*	0	0%
[mutɔgeːn-nɛk]	*mutagénnek*	32	100%
[ɔrzeːn-nɔk]	*arzénnak*	32	73%
[ɔrzeːn-nɛk]	*arzénnek*	12	27%
[pɔlleːr-nɔk]	*pallérnak*	15	100%
[pɔlleːr-nɛk]	*pallérnek*	0	0%

These hit counts agree with findings obtained by casual elicitation from native speakers; namely that [mutɔgeːn] takes front suffixes, [pɔlleːr] takes back suffixes and [ɔrzeːn] can take either. In other words, native speaker intuition matches native speaker behaviour, i.e. the behaviour of Hungarian speakers who happen to be using the dative form of these stems when composing a Web page.

In our study, we searched not just a few representative stems, but a long list, taken from a Hungarian electronic lexicon (Müller 2002). We did this with a computer program that queried the search engine automatically.[1] The forms that were fed to the program were constructed by adding *-nak* and *-nek* to each stem and applying the rule of Low Vowel Lengthening (Vago 1980a: 3–4, Siptár & Törkenczy 2000: 56–58) where appropriate. We kept data for words in which the search yielded ten or more total hits ([-nɔk] and [-nɛk] summed). The total number of stem types in our database was 10,974, and the total number of word tokens ([-nɔk] and [-nɛk] together) was about 14 million.

[1] Our software, called 'Query Google,' was programmed by Timothy Ma of UCLA. It is implemented as a publicly accessible Web applet, available (March 2006) at http://www.linguistics.ucla.edu/people/hayes/querygoogle/.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Since obtaining phonological data from the Web is a fairly new technique, we mention here a few precautions. First, it should be remembered that a search engine does not count actual tokens of the target form, but only the number of Web pages that contain it. We believe that where the focus of interest is *relative* frequency (here, of front *vs.* back endings), this factor will impose only minimal distortion, particularly where the stems under investigation are not especially common.

Second, at least some of the data retrieved in a Web search will be nonsensical in some way. For instance, speakers who lack Hungarian keyboards occasionally leave off umlauts or acute accents, which can distort the results for generalisations based on backness or length. Compound words always take the harmony of the second member, which can create errors if they are mistakenly counted as monomorphemic. Borrowed words are occasionally spelled as in the source language but take harmony according to their Hungarian pronunciation; thus, for example, *Birmingham* is pronounced ['børmiŋgɛm] and thus takes front harmony, despite its orthographic *a*.

To ward off trouble from these sources, we did some hand checking. We went through every instance of BN and BNN stems in the corpus, eliminating the illegitimate examples. We also checked all baffling instances such as B stems taking front harmony or F or FN stems taking back, and found that (with just a tiny residue of completely mysterious forms), the exceptions could be accounted for on the basis of the above categories.

As always in such studies, we must consider whether to count TOKENS (e.g. '365,822 occurrences of BN stems in the corpus take [-nɔk]') or TYPES (e.g. '603 BN stems in the corpus take [-nɔk]'). The literature suggests that when the extension of morphological patterns in the lexicon is at stake, it is type frequency that is primarily relevant; for discussion see Bybee (1995, 2001), Pierrehumbert (2001), Albright (2002) and Albright & Hayes (2003). Our findings are in harmony with these earlier claims, as we found that types provide a better fit to native speaker judgment (see § 4.2). We will therefore report only type frequencies here.

In counting type frequencies we assigned vacillators to the front and back categories according to the percentage breakdown of each type; thus for a vacillator that took [-nɔk] 20% of the time and [-nɛk] 80% of the time, we would add 0·2 to the total of [-nɔk] stems in its category and 0·8 to the total of [-nɛk] stems.

We report our data with a 'backness index': for any particular phonological category (such as BN), the backness index is the proportion of stems in that category that took [-nɔk], counting vacillators as just noted. The backness index for a category takes the value 1 when every stem always takes [-nɔk] and 0 when every stem always takes [-nɛk].

Lastly, to help relate our findings to previous work, we also include a sorting into 'back', 'front' and 'vacillator' forms, where 'back' is arbitrarily defined as taking back suffixes at least 97% of the time, 'front' as taking front suffixes at least 97% of the time and 'vacillator' any other form.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

stem type	back	vacillator	front	total stems	backness index
B	6251	39	0	6290	0·999
BN	603	78	83	764	0·831
Bi	458	17	0	475	0·989
Bi:	52	0	1	53	0·980
Be:	93	18	9	120	0·845
Bɛ	0	43	73	116	0·104
BNN	6	21	44	71	0·206
BNi	1	12	17	30	0·223
BNi:	1	7	0	8	0·358
BNe:	4	2	6	12	0·421
BNɛ	0	0	21	21	0
N	14	23	259	296	0·078
NN	0	4	939	933	0·002
F	0	0	698	698	0

Table I
Findings of the Google survey.

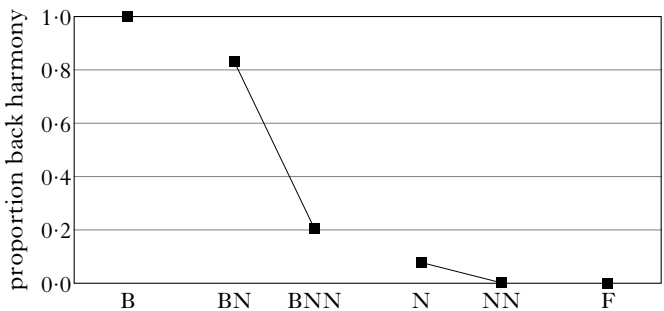


Figure 1
Google data: basic stem types.

The Google survey strongly confirmed the generalisations stated in § 3 above. Figure 1 provides a coarse classification of our data, ignoring vowel height for the moment. As can be seen, stems ending in F always take front harmony (that is, the backness index for the 698 stems examined was zero). Stems ending in B virtually always take back harmony (0·999; we assume that the exceptions were typographical errors). Stems with all neutral vowels (N, NN) are occasionally *híd* stems when monosyllabic (0·078) and only rarely when disyllabic (0·002). Crucially, comparing the

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

BN *vs.* BNN stems, we find a strong confirmation for the count effect in the much lower backness index for BNN (0·206) *vs.* BN (0·831).[2]

Turning to the height effect, we first sort the relevant forms (all BN and BNN) by the height of their last vowel: Bi, Bi:, BNi and BNi: all have high rightmost vowels and form the High category, Be: and BNe: all have mid rightmost vowels and form the Mid category, and Bɛ and BNɛ all have low rightmost vowels and form the Low category.

stem type	back	vacillator	front	total stems	backness index
high	512	36	18	566	0·938
mid	97	20	15	132	0·806
low	0	43	94	137	0·089

Table II
Google data: height effect.

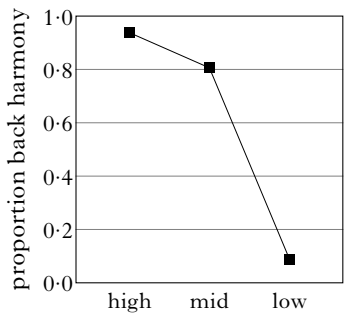


Figure 2
Google data: height effect.

Both the High–Mid and the Mid–Low differences are highly significant; see note 2.

We must also consider how the two effects interact: do both BN and BNN forms have a height effect? The data here are equivocal, as shown in Fig. 3.[3] The height effect is clearly evident in the BN forms, which are numerically preponderant. However, the numbers for BNi and BNe: are surprisingly reversed with respect to Bi *vs.* Be:. This fact will be relevant below when we consider the preferences of Hungarian speakers for novel forms.

[2] We submitted the data to chi-square tests, dividing the vacillators between front and back in the same way described above. For BN *vs.* BNN, $\chi^2 = 146{\cdot}856$, $p < 0{\cdot}001$; for NN *vs.* N, $\chi^2 = 65{\cdot}248$, $p < 0{\cdot}001$. For the height effect described in the next paragraph: High *vs.* Mid, $\chi^2 = 23{\cdot}489$, $p < 0{\cdot}001$; Mid *vs.* Low, $\chi^2 = 140{\cdot}205$, $p < 0{\cdot}001$.

[3] Looking ahead to comparison with our experimental data, we omit forms with /i:/, of which there are very few.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

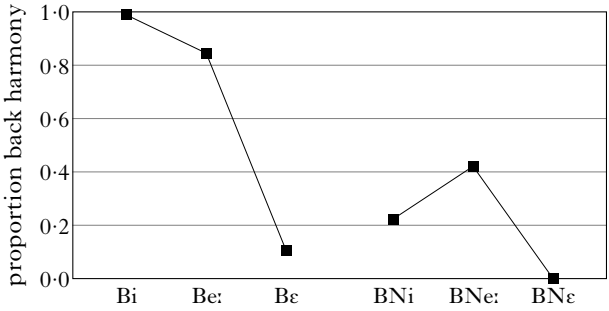


Figure 3
Google data: height effect in BN *vs*. BNN.

Lastly, it appears that when both the height and count effects are maximally present, the lexicon is variation-free: as Siptár & Törkenczy (2000: 71) point out, all BNɛ stems take front suffixes.

3.2 Verification with native speaker consultants

As a check on the search-engine method, we selected from our lexical list all instances of BN, BNN and N, plus representative cases of NN, for a total of 1130 stems, and asked two adult native speakers of Hungarian from Budapest to indicate for each whether they preferred to use [-nɔk], preferred [-nɛk] or could use either. The speakers did not rate exactly the same stems as the Google survey, because a few of the Google words were unfamiliar to them, and the speakers rated a number of words that failed to reach our threshold of ten hits on the Google survey. For the words examined in both studies, the agreement seems quite good, as Fig. 4 shows.

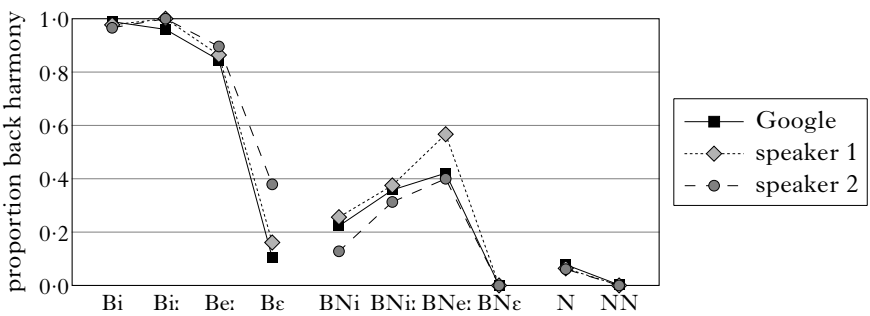


Figure 4
Two native speakers compared with Google data.

To quantify this agreement, we created backness indices for the native speaker data by assigning a value of 0·5 to each form judged to be a vacillator, 1 to every back form and 0 to every front form. For the 768 forms shared between Speaker 1 and the Google data, the correlation of backness

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

indices was r = 0·951. For the 767 forms shared between Speaker 2 and the Google data, the correlation was r = 0·937. We conclude that the search-engine method, which assesses naturalistic language use, gives results quite similar to the metalinguistic judgments of native speaker consultants.

4 The productivity of the pattern: a wug test

Are the height and count effects mere statistical patterns of the Hungarian lexicon, or are they actually internalised by Hungarian speakers and extended productively? The usual test for answering this question is the 'wug' test (nonce probe task), pioneered by Berko (1958), in which productivity is assessed by asking speakers to inflect novel stems. In the wug test we conducted, we gave speakers new, made-up stems in the nominative case (that is, with no suffix), and set up the experiment to elicit these stems with the dative suffix, either [-nɔk] or [-nɛk], as the consultant chose. Our experiment extends and complements work by Kontra & Ringen (1986) and Gósy (1989), who tested loanwords.

4.1 Procedure

We chose our wug stems on the basis of several criteria. First, we included both BN and BNN stems, with at least one stem of each type ending in each of the vowels [i eː e]. In order to sample the rest of the stem inventory, we included stems ending in F and B as well as monosyllabic and disyllabic neutral-vowel stems. We made two such sets of 15 wug stems; any particular consultant saw just one of the two sets, chosen at random. The two sets were as shown in (9).

(9)

	set 1	*set 2*
Bi	[moɲil, ʧaːdik]	[kaːnit, pozin]
	monyil, csádik	*kánit, pozin*
Beː	[haːdeːl, koleːn]	[vaːneːl, vuseːk]
	hádél, kolén	*vánél, vuszék*
Bɛ	[oːrɛl, bontɛl, kaːzɛn]	[rɔɲɛl, uɲɛg, ʧuːltɛk]
	órel, bontel, kázen	*ranyel, unyeg, csúltek*
BNi	[poribit]	[lolivit]
	poribit	*lolivit*
BNeː	[laːɲiteːg]	[aːɲiveːl]
	lányitég	*ányivél*
BNɛ	[faːɲɛdɛg, lutɛkɛr]	[aːlɛndɛl, moːlɛtɛr]
	fányedeg, luteker	*álendel, móleter*
N	[hiːɲ]	[ɲiːʃ]
	híny	*nyís*
NN	[zɛfeːt]	[pɛtleːr]
	zefét	*petlér*
F	[ɟylyt]	[høʃøg]
	gyülüt	*hösög*
B	[sɔndɔt]	[bortog]
	szandat	*bortog*

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

In constructing these wug stems, we attempted to make them sound as phonologically ordinary in Hungarian as possible. This was done by extracting from our electronic dictionary the most common initial, medial and final consonants and consonant clusters, and incorporating them into the wug stems. To minimise the effects attributed to a strong direct resemblance to any particular existing stem (see Bailey & Hahn 2001, Frisch & Zawaydeh 2001), we tried to avoid any stems that might invoke such a resemblance. We also attempted to avoid stems that would be likely to be interpreted as compounds.[4] This was done by generating large numbers of candidates for each type and checking them according to the native intuition of the second author and several other native speakers.

We administered the wug test as a written questionnaire.[5] The wug words were given in paragraph frames, meant to give the participants practice in using them before constructing their dative forms. The first sentence of the paragraph provided the nominative form of the wug stem, the second required the consultant to repeat the nominative by filling in a blank and the third provided a grammatical context requiring the consultant to use the dative case. Frames and instructions were composed with the goal of encouraging the subjects to treat the stems as long-forgotten but authentic words of Hungarian, rather than as recent loans. Here is an example of one of our frames translated into English; italics indicate expected responses:

(10) *Sample wug-test frame*

hádél

Women in the Middle Ages used hádél to wash clothing. Back then, __*hádél*__ grew abundantly in the fields. It is very hard to find nowadays, but it is said that __*hádélnak* or *hádélnek*__ had a wonderful fragrance.

We used multiple versions of each test, to make sure that no wug stem was consistently affiliated with a particular frame, and we also changed the order of the wug stems and frame sentences at random in the various versions.

Two experimenters, of whom one was the second author, distributed copies of the test forms in two Hungarian cities where the standard variety is spoken: Budapest (161 consultants) and Tiszafüred (10 consultants). The experimenters gave out forms to people whom they knew, mostly

[4] This turned out to be harder than we thought: Péter Siptár has pointed out to us the resemblances [vuseːk] *vuszék* ~ [seːk] *szék* 'chair', [aːɲiveːl] *ányivél* ~ [veːl] *vél* 'think, opine', [rɔɲɛl] *ranyel* ~ [ɲɛl] *nyel* 'swallow' and [moɲil] *monyil* ~ [ɲiːl] *nyíl* 'arrow'. Future work in this area should probably check for such resemblances by machine as well as by eye.

[5] As it turned out, written presentation provided an important additional control, since there is now some evidence that there are tiny phonetic differences between Hungarian stems that have the same basic vowels but take different harmony; see Benus (2005) and Benus & Gafos (in press). Hungarian orthography provided the subjects with an unbiased characterisation of the gross vowel phonemes present, without the possible interference of microphonetic distinctions.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

young adults. Potential participants were told that we were conducting a 'survey concerning one of the peculiarities of the Hungarian language', and those who chose to participate filled out their forms voluntarily, without pay. It appears that our questionnaire was interesting to the participants, as participation and completion rates were very high.

The forms began by asking the speakers about their linguistic background, and in our analysis we excluded data from any consultants who said that they were not native speakers of Hungarian, or that they had not spoken Hungarian in childhood.

4.2 Results

Since wug stems of the same type (e.g. [oːrɛl], [bontɛl] are both /Bɛ/) received quite similar scores, we pooled their scores into single categories, obtaining the results shown in Fig. 5 and Table III. For comparison, Fig. 5 also includes the Google data from the previous section.

stem type	back	front	invalid response[6]	backness index
B	165	2	4	0·988
Bi	322	16	4	0·953
Beː	127	211	4	0·376
Bɛ	35	456	22	0·071
BNi	48	119	4	0·287
BNeː	12	153	6	0·073
BNɛ	5	330	7	0·015
N	10	148	13	0·063
NN	0	171	0	0·000
F	1	164	6	0·006

Table III
Source data for Fig. 5.

Examining the individual cases, we can see that the experiment yielded sensible results for areas where the suffix choice is obvious: virtually all instances of F stems took front harmony, and virtually all instances of B stems took back. Thus our consultants appear to have understood the task and performed reliably.

In stem types for which the Hungarian lexicon includes both front and back cases, the aggregate behaviour of the consultants tended to

[6] The invalid responses were mostly blanks and forms given with no suffix.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

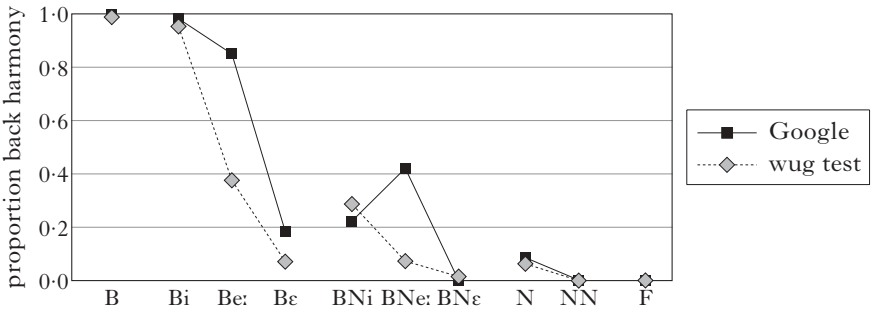


Figure 5
Wug-test data compared with Google data.

statistically match the proportions found in the lexicon. For instance, about 7·8 % of the monosyllabic N stems in the Google data are *híd* stems, taking back harmony. In the wug experiment, 6·3 % of our consultants interpreted [hiːɲ] and [ɲiːʃ] as if they were *híd* stems, attaching [-nɔk]. By way of comparison, NN stems are very seldom of the *híd* type (0·2 % in the Google survey), and none of our consultants produced back harmony for either of our NN stems [zɛfeːt] and [pɛtleːr].

The height and count effects found in the Google survey for BN and BNN stems also emerged in the subject responses for the wug experiment. As can be seen in Fig. 5, the lower the final stem vowel, the more front responses we obtained; and we obtained more front responses for BNN than for BN.

We verified the height and count effects statistically with a repeated-measures analysis of variance (ANOVA). There were two factors: Height (three levels: High, Mid and Low) and Count (two levels: BN and BNN). The analysis showed significant main effects for both Height ($F(1{\cdot}964, 320{\cdot}203) = 431{\cdot}446$, $p < 0{\cdot}0001$)[7] and Count ($F(1, 163) = 370{\cdot}862$, $p < 0{\cdot}0001$).

The interaction of Height and Count was also significant ($F(1{\cdot}762, 287{\cdot}229) = 113{\cdot}554$, $p < 0{\cdot}0001$). This arises because there is a larger height effect for BN than for BNN (or to put it differently, there is a larger count effect for higher vowels). In § 5.6, we will see that this interaction can be naturally modelled with constraint ranking in a stochastic OT framework.

To check the height and count effects in fine detail, we performed two-tailed paired t-tests on all logically adjacent categories in the data: {Bi/Beː, Beː/Bɛ, BNi/BNeː, BNeː/BNɛ}, along with {Bi/BNi, Beː/BNeː, Bɛ/BNɛ}. All comparisons were statistically significant ($p = 0{\cdot}003$ for BNeː/BNɛ; $p < 0{\cdot}0001$ for all others); thus the height effect comprises both a high/mid effect and a mid/low effect, in both BN and BNN stems, and there is also a count effect at all three heights.

[7] Where applicable, we employed the Huynh-Feldt correction for sphericity, which reduces the degrees of freedom.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Lastly, we assessed the overall degree of agreement between the Google survey and the wug-test results. To do this, we took each of the 30 stems tested (see (9)), and paired it with the backness value obtained in the Google survey for its general category (such as Bi, BNi, etc., as in Fig. 5). The correlation found was $r = 0{\cdot}896$, indicating fairly close agreement. If the proportions in the Google data are calculated from token rather than type frequencies, this correlation emerges as somewhat lower ($r = 0{\cdot}820$); cf. discussion above in § 3.1.[8]

4.3 Smoothing

Some of the discrepancies between the Google survey and the wug-test data seem of potential importance. As noted above, in the wug data there is an across-the-board height effect even in the BNN forms: BNi stems took more back suffixes than BNe: stems, which in turn took more back suffixes than BNɛ stems. The data from the Google survey contradicted this pattern, with more back responses for BNe: than for BNi.

In our view, it is the Google data that most likely are aberrant. At this level of phonological detail, there are only a few relevant stems in the lexicon. The backness value of 0·421 found for BNe: is based on just 4 back stems, 2 vacillators and 6 front stems, for a total of 12.

What is interesting is that if the aberrant figure of 0·421 does represent the Hungarian lexicon as a whole, the aberrance is evidently not registered by native speakers, whose wug-test values for BNN indicate a straightforward height effect, with higher vowels taking more back suffixes. We conjecture that the speakers have in some sense SMOOTHED the data. Rather than reflecting every small idiosyncrasy in the Hungarian lexicon, they formulate more general patterns based on natural phonological dimensions, namely the height effect and the count effect.

The surprising 1·5 % of cases where our consultants volunteered [-nɔk] for BNɛ stems, contradicting the unanimous lexical pattern, are plausibly also an instance of smoothing. The fact that BNN stems in general can take [-nɔk] and B…ɛ stems in general can take [-nɔk] may have led our consultants to arrive at the marginal possibility that BNɛ stems, which intersect these two categories, can take [-nɔk].

4.4 Caveats

Before continuing with a formal analysis of our data, we discuss a possible confound and a puzzle.

First, it has been suggested to us that our wug-test results merely reflected a mixture of differing idiolects. For example, the 0·376 backness

[8] A reviewer made the intriguing suggestion that we match the BNN wug-test forms against only those Google forms that share exactly the same NN sequence; thus, for example, using only Bii forms for predicting the wug-test responses for [poribit]. We find that this lowers the overall correlation, to 0·883. See also § 8 below.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

value for Be: stems could have resulted from 37·6% of the speakers having an idiolect that always assigns back endings to Be:, and 62·4% having an idiolect that always assigns front endings. We checked this hypothesis by testing cases where the very same consultant rated two different stems with the same pattern. For instance, [ha:de:l] and [kole:n], both Be:, appeared on the same questionnaires, and thus were encountered in the same session (though not consecutively) by the same speakers. In a series of chi-square tests, we found that consultants who gave back responses for [ha:de:l] were no more likely to give back responses for [kole:n] than consultants who gave front responses for [ha:de:l]. We obtained similar results for all other pairs where enough data were available for testing. We conclude that idiolect variation played at most a minor role in our results. The level of variation is not between individuals, but within the individual: when confronted with a novel wug stem, each speaker behaved stochastically, in a way that matched the frequencies of the lexicon.

There is one respect not yet discussed in which the wug-test results diverged from the lexicon: overall, in comparison to the Google survey, the wug-test subjects preferred front suffixes; averaging across the ten categories given in Fig. 5, the wug-test values are 0·091 more front. The frontness preference is particularly strong in the Be: forms. We conjecture that this is because the Be: forms have medial backness indices, so that the consultants' judgments are not stabilised by floor or ceiling effects.[9]

Concerning why there should be an overall frontness preference, we offer the following conjecture. For reasons we do not understand, there is a weak connection in the Hungarian lexicon between stem frequency and frontness: the rarer the stem, the more likely it is to take front suffixes. Thus, we find that within just the BN stems, the correlation of the frequency of the bare stem (which we also measured in the Google survey) with the proportion of back responses is $r = 0{\cdot}141$.[10] Wug stems are, by definition, the rarest of stems (frequency zero), and this may have contributed to their front-preferring behaviour. In principle, this factor could be entered into the model described below, but we will not attempt to do this here.

5 A theoretical model of variation in Hungarian vowel harmony

We turn to the task of developing an explicit analysis of our experimental findings, drawing on various notions from current phonological theory. To achieve descriptive adequacy, our model must accomplish three tasks.

[9] Indeed, because the endpoints of the scale are anchored, and the medial values show a frontness preference in the wug test, we find that the Google data correlate better with the square root of the wug-test values (which lower the unanchored medial scores) than with the raw scores; $r = 0{\cdot}923$ *vs.* 0·896.

[10] This connection might explain Gósy's (1989) finding that younger children tend to give backer responses in wug testing for the relevant word classes. Younger children would be less likely to be familiar with rarer words.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

First, it should accommodate STEM-SPECIFIC BEHAVIOUR, permitting speakers to list (in some form yet to be addressed) whether a particular BN or BNN stem takes front suffixes or back suffixes, or is a vacillator (cf. (7) above).

Second, an adequate model should characterise the native speaker's EXPECTATIONS about what suffixes a novel stem will take – in particular, it should be able to account for our wug-test data.

Third, a model should characterise the LIMITS OF STEM-SPECIFIC BEHAVIOUR. While it is true that BN and BNN stems can have their own specific behaviour, B or F stems cannot; their harmony pattern is completely predictable. For example, our wug test included the B stems [bortog] and [sɔndɔt], and these virtually always took back harmony; likewise for the F stems [ɟylyt] and [høʃøg], which took front harmony. We claim that forms like *[bortog-nɛk] or *[ɟylyt-nɔk] are simply unacceptable in Hungarian, and that a phonological analysis should capture this fact.

5.1 Theory

We assume Optimality Theory (Prince & Smolensky 1993), in which the outcomes of phonological derivations depended on the ranking of conflicting constraints. Constraint conflict arises here when a stem contains both front and back vowels. A variety of constraints require that the suffix vowels match the stem vowels in backness, and when a stem contains both front and back vowels, these constraints will conflict.

We use a STOCHASTIC variant of Optimality Theory (Boersma 1997, Hayes & MacEachern 1998, Boersma & Hayes 2001), in which the ranking of constraints is probabilistic: every constraint pair (A, B) is associated with a probability (0–1) specifying how likely it is that A will dominate B on any given speaking occasion. The reason for using stochastic OT is that it permits precise predictions about the relative proportions of forms produced in free variation.

Lastly, we assume the DUAL LISTING/GENERATION MODEL of Zuraw (2000). In this model, grammars may contain sets of markedness constraints that are stochastically ranked with respect to each other, but subordinated to faithfulness constraints. This means that existing forms, which are covered by a particular lexical entry and protected by faithfulness, surface without variation; whereas newly inflected forms, where faithfulness constraints are inapplicable, are derived stochastically, according to the pattern of the subordinated markedness constraints.

The following sections cover underlying forms, constraints and rankings, and an assessment of the model's performance in describing our data.

5.2 Underlying forms

The underlying phonological form of Hungarian suffixes is somewhat difficult to establish, since they normally appear harmonised to a preceding vowel. The occurrence of case suffixes as independent stems, in

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

constructions like [nɛk-ɛm] *nekem* 'me-DAT', could in principle justify an underlying backness value for these suffixes (Vago 1973). However, we will see later on that there is reason to let both allomorphs of a harmonising suffix serve as underlying forms, and so we will assume here that the use of the case suffixes as stems represents an arbitrary lexical choice and does not determine a unique underlying form. For discussion, see Reiss (2003). Below, in cases where it is not crucial to assert an underlying backness value, we will simply use capital letters to designate the general category of the suffix vowel; thus /A/ is a vowel that alternates between [ɔ] and [ɛ], so the underlying form of [-nɔk] ~ [-nɛk] will be shown as /-nAk/.

5.3 Markedness constraints governing harmony

We assume that BN and BNN stems vary in their harmony because they have two triggers: one which is strong but non-local (B), and another which is weak but local (the rightmost N, which is closest to the suffix). Suffix variation results from stochastic ranking of the conflicting constraints that require suffixes to agree in backness with these triggers. The constraints are assumed to be members of the AGREE family (Lombardi 1999, Kiparsky & Pajusalu 2003), relativised to distance.

In our usage, a vowel-harmony constraint is LOCAL if it assigns violations to vowel sequences that are separated only by a (possibly null) consonant string. A constraint is DISTAL if it assigns violations without regard to intervening material.[11] In formalising the constraints, we assume that phonology makes available a vowel projection (Vergnaud & Halle 1979) or tier (Archangeli & Pulleyblank 1987, Clements & Hume 1995) which expresses just the vowels of the string, so that consonants can be ignored in the structural description of constraints. Thus the constraints we will call LOCAL[B] and DISTAL[B] are stated below:

(11) a. LOCAL[B]
*[+back][−back]
Assess a violation when the closest vowel following a [+back] vowel is [−back].

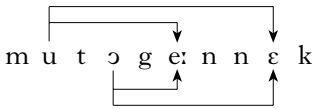
b. DISTAL[B]
*[+back]X[−back]
Assess a violation when a [+back] vowel is followed somewhere in the word by a [−back] vowel.

For example, in the candidate form in (12), DISTAL[B] incurs four violations, as shown by the arrows.

[11] For earlier analyses that also assume non-local mechanisms, see Kiparsky & Pajusalu (2003), which uses constraints, and Esztergár (1971: 29) and Vago (1976: 252), which use rules.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(12) *mutagénnek*



In contrast, LOCAL[B] assesses violations only for vowel pairs separated by (at most) a consonant string, e.g. just one violation for [mutɔgeːn-nɛk]:

(13) m u t ɔ g eː n n ɛ k

We assume additional agreement constraints defined on particular vowels or natural classes of vowels, formalised analogously to (11). In the present analysis the following constraints will be employed:

(14) a. LOCAL[F], where F = [−back, +round]
b. DISTAL[F]
c. LOCAL[i]
d. LOCAL[eː]
e. LOCAL[ɛ]

In principle, LOCAL[iː] should also be included, but we will ignore it here since we have no wug-test data for this vowel (in the Google data, it matches /i/ fairly closely).

We also need a constraint to enforce the more frequent appearance of front suffixes in BNN stems. Walker (2001) has noted that it is possible for harmony processes to have 'double triggers'; i.e. harmony occurs only when two in a row of the relevant triggering class are present. Walker analyses the phenomenon in depth; for present purposes we will just stipulate a constraint LOCAL[NN], violated when NN is followed by B.

(15) LOCAL[NN]
*[−back][−back][+back]

In the suffixed forms of BN stems, DISTAL[B] conflicts with one or more of the three constraints LOCAL[i], LOCAL[eː] and LOCAL[ɛ]; in BNN stems, it additionally conflicts with LOCAL[NN].[12]

[12] An alternative to LOCAL[NN] is to fragment DISTAL[B], splitting it into constraints requiring agreement with B two syllables away (governing BN forms) *vs.* three (governing BNN). We have explored this kind of analysis and find it works about as well as the one presented in the text. For reasons of space we will not present both analyses here.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

5.4 Faithfulness constraints

Following Ringen & Vago (1998), we assume two faithfulness constraints governing backness; one is limited to root vowels, while the other is simply the general IDENT constraint for this feature:

(16) a. IDENT-IO[back]$_{root}$
Assess a violation if a vowel belonging to a morphological root differs in surface representation from its underlying correspondent in its value for the feature [back].

b. IDENT-IO[back]
Assess a violation if a vowel differs in surface representation from its underlying correspondent in its value for the feature [back].

As we will see, IDENT-IO[back]$_{root}$ must be ranked higher than IDENT-IO[back], reflecting the greater immutability of root vowels relative to suffix vowels in Hungarian. The pattern of greater faithfulness in roots is often observed cross-linguistically; see for example McCarthy & Prince (1995), Beckman (1997) and Casali (1997).

5.5 Constraint rankings: strict

We can now consider how the constraints given above can be ranked to characterise the data. We begin with some straightforward rankings that are non-stochastic (in the theory assumed, they are associated with the probability value 1).

To begin, IDENT-IO[back]$_{root}$ must be ranked strictly over LOCAL[B] and DISTAL[B]. This is because Hungarian stems, unlike suffixes, are not in general required to respect harmony: there are many BN and NB stems, and also a number of borrowings like [glykoːz] *glükóz* 'glucose', with FB, and [ʃoføːr] *sofőr* 'chauffeur', with BF. The need for a strict ranking is shown below in tableau (17): /fɔrmɛr/ *farmer* 'blue jeans' survives intact, despite its violations of the two harmony constraints.[13]

(17)

/fɔrmɛr/	IDENT-IO[bk]$_{rt}$	LOCAL[B]	DISTAL[B]
☞ a. fɔrmɛr		*	*
b. fɔrmɔr	*!		

There are also strict rankings among the Markedness constraints. Thus, LOCAL[B] must strictly dominate DISTAL[F], because in all stems in

[13] There is more to the problem than this, in that B combines with N freely in stems, but far less often with F; the BF and FB stems are all borrowings and are sometimes felt to be foreign. While we will not try to integrate this fact into the analysis, we believe the necessary apparatus is at hand. Kiparsky & Pajusalu (2003) propose constraints that penalise BF/FB but not BN/NB; and Itô & Mester (1995), among others, have suggested rankings particular to vocabulary strata such as foreign words.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

which the last vowel is B and some other vowel is F, the suffix must surface as back – the LOCAL[B] vowel wins out as trigger over the distal F vowel. This can be seen in the following tableau for [glyko:z-nɔk] 'glucose-DAT'.

(18)

/glykɔ:z-nAk/	IDENT-IO[bk]$_{rt}$	LOCAL[B]	DISTAL[F]
☞ a. glyko:z-nɔk[14]			**
b. glyko:z-nɛk		*!	*
c. glykø:z-nɛk	*!		
d. gluko:z-nɔk	*!		

For the same reason, LOCAL[F] must strictly dominate DISTAL[B], for example to obtain [ʃofø:r-nɛk] 'chauffeur-DAT', not *[ʃofø:r-nɔk].

(19)

/ʃofø:r-nAk/	IDENT-IO[bk]$_{rt}$	LOCAL[F]	DISTAL[B]
☞ a. ʃofø:r-nɛk			**
b. ʃofø:r-nɔk		*!	*
c. ʃofo:r-nɔk	*!		
d. ʃøfø:r-nɛk	*!		

5.6 Constraint rankings: stochastic

In the model of stochastic OT adopted here, the probabilistic rankings of the constraints are expressed by assigning them values along a numerical scale of 'ranking strength'; from this scale the relative ranking probabilities can be deduced with a standard mathematical formula, given in Boersma (1997: 45). In the discussion that follows, we will present the pairwise probabilities, since these are more readily interpretable.[15]

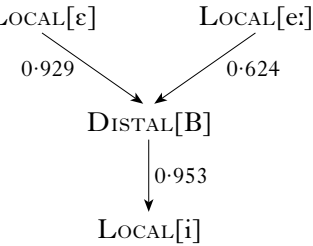
Three crucial probabilistic rankings in the analysis are those of DISTAL[B] against its competitors among the weak front-harmony triggers, namely LOCAL[ɛ], LOCAL[e:] and LOCAL[i]. We propose that these ranking probabilities should be as shown in (20):

[14] The rounding of the low back short vowel [ɔ] is non-contrastive (Vago 1980a: 3), and is straightforwardly derived by ranking a ban on short low back unrounded vowels over the faithfulness constraint IDENT[round]. We omit this detail from our tableaux.

[15] One set of ranking values that yields the probabilities proposed in this section is: LOCAL[ɛ] = 104·154, LOCAL [NN] = 101·802, LOCAL[e:] = 100·894, DISTAL[B] = 100·000, LOCAL[i] = 95·263. Strict rankings, such as those from the previous section, can be implemented with any difference in ranking value (e.g. 20) that translates into something very close to 1.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(20)



The 0·624 probability proposed for the ranking of Local[eː] over Distal[B] means that given an input like /haːdeːl-nAk/ ([haːdeːl] is the wug stem appearing in (10)), there is a probability of 0·624 that the grammar will output [haːdeːl-nɛk]. This is shown in tableaux (21) and (22).

(21)

/haːdeːl-nAk/	Ident-IO[bk]$_{rt}$	Local [B]	Local [ɛ]	Local [eː]	Distal [B]	Local [i]
0·624 ☞ a. haːdeːl-nɛk		*			**	
b. haːdeːl-nɔk		*		*!	*	
c. haːdaːl-nɔk	*!					

In (21), the probability of 0·624 that Local[eː] will dominate Distal[B] implies the same probability that [haːdeːl-nɛk] will be output by the grammar. The opposite ranking, generating [haːdeːl-nɔk] with a probability of 0·376, is given in (22):

(22)

/haːdeːl-nAk/	Ident-IO[bk]$_{rt}$	Local [B]	Local [ɛ]	Distal [B]	Local [eː]	Local [i]
a. haːdeːl-nɛk		*		**!		
0·376 ☞ b. haːdeːl-nɔk		*		*	*	
c. haːdaːl-nɔk	*!					

Thus, over a large number of trials, we would expect [haːdeːl-nɛk] to be the winner in about 62·4% of the trials, and [haːdeːl-nɔk] to be the winner in about 37·6%. In fact, in our wug test, with stems of the Beː type, this was the percentage obtained from the participants as a whole; our hypothesised ranking values were set up with the express purpose of mimicking this frequency.[16] The remaining values in (20) can similarly be used to derive the correct wug-test percentages for Bi and Bɛ stems: the

[16] This practice seems to have resulted in some misunderstanding, so we attempt to clarify. In the research methodology of generative grammar, the analyst seeks to devise a set of rules or other principles that generates the pattern of the language. Our purpose, at this stage, is simply to develop an analysis of this sort, taking our wug test to be our data for what the Hungarian pattern is. No essential difference arises from the fact that the system generates quantitative outcomes.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

probability of 0·953 that DISTAL[B] ⫸ LOCAL[i] predicts 95·3% back suffixes for Bi stems, and the probability of 0·929 that LOCAL[ɛ] ⫸ DISTAL[B] predicts 92·9% front suffixes for Bɛ stems.

The results so far merely indicate that the constraint set is sufficiently rich to discriminate between the Bi, Beː and Bɛ categories – we have set three ranking probabilities, and have derived three relative frequencies. More interesting is the task of extending the analysis to the BNN stems, specifically BNi, BNeː and BNɛ. In the wug-test data, these stems exhibited both the height effect and the count effect; the two effects are additive in the sense that the stems that are most likely to take front endings are the BNN stems with low final vowels. We propose that this can be modelled simply by assigning an appropriate probability to the ranking LOCAL[NN] ⫸ DISTAL[B]: this will shift the percentages of back suffixes downward in BNN stems relative to analogous BN stems, because two constraints rather than one are working against LOCAL[B].

We have calculated that a probability of 0·738 for LOCAL[NN] ⫸ DISTAL[B] best fits the wug-test data. The crucial part of the grammar is thus as in (23):

(23)

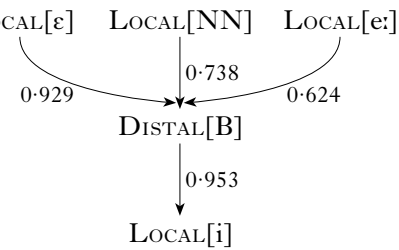


Tableau (24) illustrates how the constraints interact in the case of a representative BNeː stem.

(24)

/aːɲiveːl-nAk/	LOCAL [ɛ]	LOCAL [NN]	LOCAL [eː]	DISTAL [B]	LOCAL [i]
0·836 ☞ a. aːɲiveːl-nɛk				⁎⁎⁎(!)	
0·164 ☞ b. aːɲiveːl-nɔk		⁎(!)	⁎(!)	⁎⁎	

0·738, 0·624

In effect, (24) abbreviates six tableaux, one for each of the six possible rankings of LOCAL[NN], LOCAL[eː] and DISTAL[B]. From a series of calculations, we have determined that, assuming the probabilities shown in (24), a back suffix outcome for a BNeː stem will win in 83·6% of trials and a front suffix outcome in 16·4%.[17] We performed similar calculations for the remaining cases, BNi and BNɛ, where three constraints interact.

[17] We use *OTSoft* 2.1 (Hayes *et al.* 2003) to perform all calculations for stochastic OT reported here. For the present task, the software takes a large number of samples,

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Summarising, our hand-ranked model generates the frequencies given in (25), along with the wug-test frequencies we were attempting to mimic.

(25)

	wug test	*model*		*wug test*	*model*
Bi	0·953	0·953	BNi	0·287	0·257
Beː	0·376	0·376	BNeː	0·073	0·164
Bɛ	0·071	0·071	BNɛ	0·015	0·045

The grammar does not achieve an exact quantitative match for the BNN forms, but it is not far off, and moreover it captures the correct qualitative generalisations: there is a count effect, a height effect and also an interaction – the size of the height effect is reduced in BNN forms, and the size of the count effect is reduced for lower vowels. This illustrates the ability of the stochastic OT framework to capture such quantitative interactions.

The grammar also accomplishes the two instances of 'smoothing' we noted in § 4.3. Although in the Hungarian lexicon, BNeː unexpectedly takes [-nɔk] more often than BNi, our model in fact gives BNeː-nɔk a lower frequency than BNi-nɔk – just as our wug-test subjects did. Moreover, the wug testees unexpectedly volunteered a small number of BNɛ-nɔk forms, despite their absence in real Hungarian. Our model likewise generates a small number of these forms. In both cases, the cause of the smoothing is the same: the constraints responsible for the height effect are ranked on the basis of all of the data, not just the BNN forms. The statistical patterns found in the (more numerous) BN cases are carried over to some extent to BNN.

5.7 The basis of the height effect

It can be observed that in (23), the constraints Local[ɛ], Local[eː] and Local[i] are stochastically ranked so as to make lower vowels 'stronger triggers' for front harmony; that is, better able to compete with Local[B]. We suggest that this ranking is not accidental.

Kaun (1995, 2004), in a study of the typology of rounding harmony, proposes that the differences in the strength of vowels as harmony triggers depend on the phonetic salience with which they manifest the harmonic feature: harmony is triggered preferentially by perceptually inferior vowels, i.e. the ones that lack an extreme phonetic realisation of their category. In the case of rounding harmony, these are the low rounded vowels, which (relative to their high counterparts) are phonetically less rounded and acoustically less distinct from unrounded vowels. In her

using Gaussian distributions centred about the ranking values given in note 15, and calculates the probability by summing over the samples. We performed ten trials, each with 100,000 samples per form. Although this calculation method is not exact, the random fluctuations are quite small: the greatest standard deviation across the ten trials for the proportion of back suffixes derived was never greater than 0·0014 for any input form.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

survey, Kaun found that low rounded vowels often trigger harmony in contexts where high rounded vowels do not. Her functional explanation for this tendency is that the low rounded vowels, which most need help in identification, are more likely to obtain this help by spreading their rounding feature across the word.[18]

Pursuing the same approach for backness harmony, we note that of the front vowels of Hungarian discussed here, it is the lower front vowels that have the lowest second formant frequencies, and thus are perceptually inferior relative to the higher front vowels. Following Kaun's approach, we expect the strength of the front triggers to be determined by their height, with [ɛ] the best trigger, [eː] the second best and [i] the worst. In grammatical terms, this is manifested in an *a priori* ranking preference:

(26) LOCAL[ɛ] ≫ LOCAL[eː] ≫ LOCAL[i]

As we have just seen, this ranking, in a looser stochastic form, is what is needed for the analysis of the Hungarian data.[19]

For suggestions that the height effect may not be unique to Hungarian (and thus deserves a general explanation), see Esztergár (1971) and Anderson (1980); for a different phonetic account, see Benus (2005) and Benus & Gafos (2005).

5.8 Treatment of existing stems: the role of faithfulness

Recall that the Hungarian variation is primarily stem-by-stem variation and not token-by-token variation; only the vacillators actually permit the two outcomes generated by the grammar thus far, while most stems impose an invariant suffix choice. This information must be encoded in the lexicon: part of the task of learning Hungarian is to memorise the harmonic behaviour (front, back or vacillating) of the stems that fall into the unpredictable categories (cf. (7) above).

There are various forms of representation that could be used by speakers to memorise whether a stem takes front or back suffixes. These include diacritics, floating backness autosegments (Goldsmith 1979, Lieber 1987, Wolf, in press) and (following Zuraw 2000) simply the full lexical listing of the inflected forms. These possibilities are shown for the stems in (7) in (27).

[18] Kaun also demonstrates a tendency for harmony to apply only to vowels of matched heights. Readers have asked us: might the reason that [-nɛk/-nɔk] harmonises preferentially after low vowels be that its own vowel is low? The answer is no: as a Google check confirms, lower vowels are stronger harmony triggers in Hungarian even when the suffix vowel is high.

[19] Constraint families based on phonetic scales like (26) have been implemented in various ways. The approach in (26) affiliates one constraint with each member of the scale, and ranks the constraints so as to match the scale. Another approach (Prince 1997, de Lacy 2004) implements the scale by stating the constraints with cut-off points, e.g. 'Agree in backness if the trigger is mid *or lower*'. We have implemented our analysis under both approaches and achieved equally good matches to the data. For brevity we only report the first approach here.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(27) a. *Diacritic*

i. *takes front suffixes*	/mutɔgeːn/ [−back harmony]
ii. *vacillator*	/ɔrzeːn/ [∅back harmony]
iii. *takes back suffixes*	/pɔlleːr/ [+back harmony]

b. *Floating autosegment*

i. *takes front suffixes*	/ m u t ɔ g eː n / +b +b −b −b
ii. *vacillator*	/ ɔ r z eː n / / ɔ r z eː n / +b −b −b +b −b +b
iii. *takes back suffixes*	/ p ɔ l l eː r / +b −b +b

c. *Full lexical listing*

i. *takes front suffixes*	/mutɔgeːn-nɛk/
ii. *vacillator*	/ɔrzeːn-nɛk/, /ɔrzeːn-nɔk/
iii. *takes back suffixes*	/pɔlleːr-nɔk/

With each of these options, the faithfulness constraints of the grammar must be stated to enforce the particular form(s) listed in the lexicon. This could be done by requiring a proper match between diacritic specification and suffix allomorph (27a), by requiring surface realisation of the floating autosegment (27b) or simply by requiring the maintenance of underlying suffix vowel backness (27c).

Here, we adopt the full-listing proposal (27c), and give some tentative evidence in its favour below. The vowels in listed suffix allomorphs are protected by the general faithfulness constraint IDENT-IO[back], stated in (16).

Here is an example of how this works. The stem [ɔtseːl] *acél* 'steel' falls into the Beː class, which in the grammar developed so far permits variation in suffix choice. In fact, it is a lexical property of [ɔtseːl] that it takes only back suffixes. Thus, there is a listed entry /ɔtseːl-nɔk/ that emerges as the winner, as shown in tableau (28):

(28)

/ɔtseːl-nɔk/	IDENT-IO[bk]$_{rt}$	LOCAL[B]	IDENT-IO[bk]	LOCAL[eː] (0·624)	DISTAL[B]
☞ a. ɔtseːl-nɔk		*		*	*
b. ɔtseːl-nɛk		*	*!		**
c. ɔtsaːl-nɔk	*!		*		

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

The crucial ranking is IDENT-IO[back] ≫ LOCAL[eː], which cancels the possibility that LOCAL[eː] could force the outcome *[ɔtseːl-nɛk]. Were it not for IDENT-IO[back], this candidate would win 62·4 % of the time. The candidate with stem-internal harmony, *[ɔtsaːl-nɔk], is ruled out by undominated IDENT-IO[back]$_{root}$.[20]

The outcome for [ɔtseːl] should be compared to the phonologically similar wug stem [haːdeːl]. Wug stems lack lexical entries for their suffixed forms, because the subjects heard only the unsuffixed stems. Thus, they cannot specify whether they take [-nɔk] or [-nɛk]. Because of this, neither [haːdeːl-nɛk] nor [haːdeːl-nɔk] violates IDENT-IO[back], and this constraint therefore would not affect the outcome in such cases.

(29)

/haːdeːl-nAk/	IDENT-IO[bk]$_{rt}$	LOCAL[B]	IDENT-IO[bk]	LOCAL[eː]	DISTAL[B]
0·624 ☞ a. haːdeːl-nɛk		*			**
0·376 ☞ b. haːdeːl-nɔk		*		*!	*
c. haːdaːl-nɔk	*!		*		

0·624

It can be seen that this grammar respects a memorised suffix choice when there is one, but performs stochastically (i.e. like a Hungarian speaker) when given a wug stem. To cover the full range of cases, the rankings needed are as in (30):

(30) IDENT-IO[back] ≫
{DISTAL[B], LOCAL[i], LOCAL[eː], LOCAL[ɛ], LOCAL[NN]}

That is to say, the bloc of stochastically ranked AGREE constraints in (23) is generally subordinated to IDENT-IO[back].[21,22]

[20] Zuraw's theory further assumes a constraint USELISTED, which requires a listed entry to be employed, thus blocking the possibility of a winning candidate *[ɔtseːl-nɛk], created afresh by the morphology. In the present analysis, USELISTED may be assumed to be undominated.

[21] Hungarian speakers often feel that 'wrong' choices among BN and BNN words (say [ɔtseːl-nɛk]) are not crashingly bad, and that they might accept them if heard from other speakers. This suggests that the main ranking in (30) might actually not be completely strict; and therefore lets through the 'wrong' choice as a weak (improbable) alternative. We lack the data that would be needed to establish this stochastic ranking precisely, and will for the remainder of the article assume strict ranking for purposes of exposition.

[22] A reviewer notes that full lexical listing cannot be a complete theory of exceptions in phonology: we need to cover phonology that is exceptionally triggered, or not triggered, or not undergone, by particular affixes. For discussion, see Pater (in press). In our view, even the morpheme-specific constraints of the type Pater proposes would often have to be ranked stochastically, to cover cases like Tagalog Nasal Substitution (Zuraw 2000: § 2.2.2) or Spanish diphthongisation (Eddington 1996, 1998, 2004).

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

5.9 The representation of vacillators

For vacillating stems, we assume that there are two rival underlying representations, as for example (27.ii) (/ɔrzeːn-nɔk, ɔrzeːn-nɛk/). This is in principle no different from cases like English *envelope*, where /ˈɛnvəˌloʊp/ and /ˈɑnvəˌloʊp/ often compete even within single idiolects. It is likely that the rival underlying forms are represented in a way that assigns them a quantitative 'strength', which is reflected by their frequencies in actual usage. The observed Google frequencies of such doublets are seldom actually 50–50 (as the simplest double-listing approach would predict), but vary over the full range of values.

This provides an argument for Zuraw's pure lexical-listing theory for idiosyncratic forms (27c). Unlike the diacritic and floating-autosegment theories, Zuraw's account implies the possibility there could be stems that favour a *different mix* of front and back allomorphs for (say) the dative than for some other suffix. Some supporting cases are given in Kontra & Ringen (1986: 10); a particularly dramatic case is the common NN stem [feːrfi] *férfi* 'man', which is a vacillator in the dative ([-nɔk]/[-nɛk]), but allows only back [-ɔk] *-ak* in the plural. In Zuraw's theory, such differences would follow from particular suffixed lexical entries for individual inflected forms. The overall tendency for a stem to take the same backness of suffixes throughout its paradigm would best be attributed to output-to-output correspondence constraints (Benua 1997) governing suffix backness, though we will not attempt to flesh out this proposal here.

5.10 Keeping lexical entries in check

The constraint IDENT-IO[back] permits individual stems to force particular suffix choices, even in the face of the phonological agreement constraints. Yet such suffix preferences should not be allowed unchecked, because the resulting grammar would overgenerate. In Hungarian, there are absolutely no stems of the following types:

(31) *Impossible forms*
- a. F stems that take back suffixes (e.g. *[ɟylyt-nɔk])
- b. B stems that take front suffixes (e.g. *[bortog-nɛk])
- c. F(N)* stems that take back suffixes (e.g. *[ɟylit-nɔk])

Native speakers vigorously reject such forms, and (other than the occasional random error) do not volunteer them on a wug test.

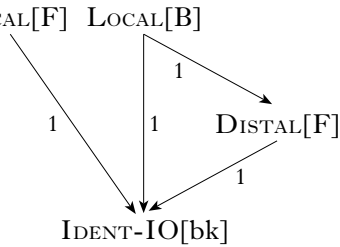
The normal approach for excluding impossible forms in Optimality Theory follows the principle of the Rich Base (Prince & Smolensky 1993): we show that *if* such an item occurred as a lexical entry, then the grammar would derive from it an unfaithful well-formed output. For example, if there were a lexical entry like /ɟylyt-nɔk/, the grammar would output

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

[ɟylyt-nɛk] instead. This approach to constraining idiosyncrasy has been applied earlier in the theories of exceptions developed by Kager (in press) and Pater (in press).[23]

In the grammar under discussion, the crucial restraint can be achieved by ranking the strongest vowel-harmony constraints above IDENT-IO[back], with a probability of 1:

(32)



Under this ranking, for the hypothetical lexical entry /ɟylyt-nɔk/ the winning candidate would be well-formed [ɟylyt-nɛk], in which the underlying /ɔ/ of the suffix surfaces as front:

(33)

/ɟylyt-nɔk/	LOCAL[F]	IDENT-IO[bk]
☞ a. ɟylyt-nɛk		*
b. ɟylyt-nɔk	*!	

Similar impossible forms are ruled out analogously: because LOCAL[B] strictly dominates IDENT-IO[back], there could be no forms like *[bortog-nɛk] even if the lexicon 'asked for' them ([bortog-nɔk] wins); and because DISTAL[F] strictly dominates IDENT-IO[back], there could be no *[ɟylit-nɔk] ([ɟylit-nɛk] wins). However, for all markedness constraints ranked below IDENT-IO[back] (see (30)), an invariant listed form violating that constraint *can* assert itself in the output.

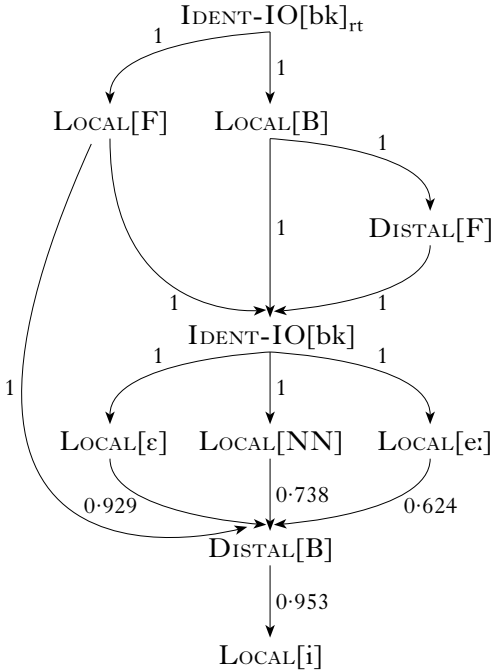
5.11 Summary and assessment

The complete set of constraints and rankings in our analysis is summarised in the following Hasse diagram:

[23] Like Zuraw, Kager uses full lexical listing to encode exceptions, but lists the allomorphs of morphemes, rather than full words. Kager's theory does not aspire to account for statistical patterning of exceptions, and thus could not be used to address our wug-test data.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(34)



We claim that the analysis meets the goals set out at the beginning of this section. For novel BN and BNN stems, it generates outputs in proportions that fairly closely match those produced by native speakers, as a result of stochastic rankings among the lowest-ranked constraints (§ 5.6). Further, it permits lexical entries of stems of these types to specify particular suffix choices, as a result of IDENT-IO[back] (§ 5.8). Lexical entries are not permitted to specify impossible forms, due to the ranking of LOCAL[F], LOCAL[B] and DISTAL[F] above IDENT-IO[back] (§ 5.10). Lastly, because $\text{IDENT-IO[back]}_{\text{root}}$ is at the top of the grammar, harmony cannot alter stems (§ 5.5). The model achieves a fairly close match to the wug-test data, diverging slightly in the BNN forms, as shown in Fig. 6.

To characterise the match quantitatively, we took each of the 26 wug stems for which the model makes a prediction[24] and paired the wug-test score of the stem with the prediction made by the model for stems in its class. The two sets of values thus obtained are highly correlated; $r = 0{\cdot}991$. This is not surprising, given the fairly rich constraint set used. More important, we claim that the constraints themselves are not arbitrary, but follow general principles of phonological theory – that is, they are all either markedness constraints of the AGREE family or faithfulness constraints for backness.

[24] The model includes no constraints for N or NN stems; see § 5.12 below.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

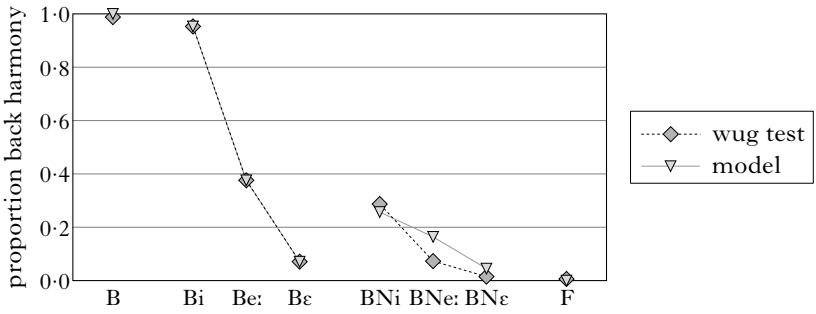


Figure 6
Match-up of model to wug-test data.

5.12 Further issues

A number of issues surrounding the analysis remain for future research.

First, the use of constraints like DISTAL[F] and DISTAL[B] glosses over a problem in the analysis of non-local harmony in the AGREE framework: as stated, they fail to distinguish forms like BFN from FBN, which incur the same violations. Such forms are rare in Hungarian, but it seems fairly clear that BFN stems take front harmony and FBN stems behave like the corresponding BN stems. In other words, the principle 'closest trigger wins' is not predicted by our analysis in the general case.

In the older autosegmental approach, the 'closest trigger wins' principle was the automatic consequence of the ban on crossed association lines (Clements 1977). However, in light of recent research (Hansson 2001, Frisch *et al*. 2004, Rose & Walker 2004) it appears that the 'closest trigger wins' principle is not exceptionless; and it is entirely incompatible with the view taken here, in which variation in Hungarian is attributed to conflicting harmony triggers. What is needed, we think, is a more articulated theory of non-local AGREE constraints in which the principle becomes negotiable, subject to constraint ranking.

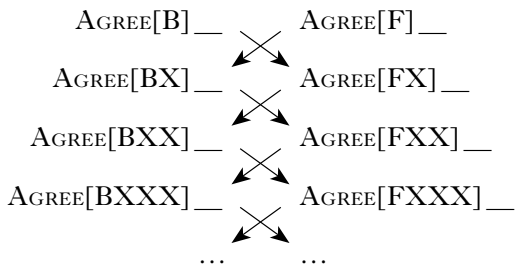
One possibility for such a theory is to suppose that a constraint like DISTAL[B] is actually an intrinsically ranked infinite schema, along these lines:

(35) AGREE[B]__ ≫
AGREE[BX]__ ≫
AGREE[BXX]__ ≫
AGREE[BXXX]__ ≫
…

What needs to be worked out is a constrained system for ranking the members of this schema which would appropriately implement the general idea of 'closest trigger wins'. For instance, we expect the default ranking of the DISTAL[B] schema and the DISTAL[F] schema to be as in (36):

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(36)



Since developing such a system explicitly would lead us rather far from present concerns, we will not attempt it here.

We also have not dealt with N or NN stems. In our overall approach, the occasional volunteering of back suffixes for N stems in our wug test (as in [hiːŋ-nɔk]; see § 4.2) cannot be the result of IDENT-IO[back], since wug stems are assumed to lack lexical entries. Rather, they must reflect constraints, perhaps language-particular, that require dissimilation (cf. Ringen 1980: 139). Most likely, the principal dissimilation constraint singles out monosyllables in /iː/, since most of the *híd* stems of Hungarian take this form. Given the relative rarity of *híd* stems, the dissimilation constraints must be stochastically ranked rather low (lower than the constraints of the LOCAL[N] family), so that wug forms attaching [-nɔk] to neutral stems are output only occasionally.

6 Learning the grammar

Thus far, we have followed the classical procedure of generative linguistics, inventing a formal hypothesis intended to describe the native speaker's tacit knowledge. But the goal of explanatory adequacy (Chomsky 1965, Chomsky & Halle 1965) implies that linguistic theory should not just provide an accurate and principled account of language-particular grammars, but offer a mechanism for how a child exposed to learning data could arrive at the correct grammar. Recent work on learnability in Optimality Theory (Boersma 1997, Boersma & Hayes 2001, Hayes 2004, Prince & Tesar 2004) makes possible a sketch of how the Hungarian vowel-harmony system as outlined here could be learned. The discussion will be confined to the learning of the rankings shown in (34); we assume that the constraint inventory itself is either innate (Tesar & Smolensky 2000) or else is accessible to the child through some form of inductive learning (Hayes 1999).

6.1 Factoring the learning task

The learning task at hand can be divided into three parts.

(i) The child must discover WHAT IS POSSIBLE; that is, she learns to distinguish harmonically possible words from harmonically impossible ones.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

For instance, she must ultimately come to know that forms such as *[ɟylyt-nɔk], *[bortog-nɛk], *[ɟylit-nɔk] (all from (31)), *[glykoːz-nɛk] (18) and *[ʃoføːr-nɔk] (19) are all impossible in Hungarian, and must also learn that forms like [haːdeːl-nɔk] and [haːdeːl-nɛk] are both, in principle, possible Hungarian forms.

(ii) Another aspect of the learning task involves the LEXICON: the child must learn which particular stems take which kinds of suffixes, internalising lexical entries along the lines of § 5.8. We take no stance on how lexical learning takes place, since nothing in our model bears on this question.[25]

(iii) Lastly, it is evident from our experimental results that Hungarian-learning children ultimately develop a STATISTICAL MODEL of the lexicon, hence the ability to project novel forms stochastically in proportions matching their lexical frequencies. In our model, this is done by assigning a stochastic ranking to the constraints at the bottom of the grammar in (34).

We will ignore here the (non-trivial) task of learning thousands of lexical entries and focus instead on tasks (i) and (iii). In principle, they could be accomplished by a single algorithm, but here we find it necessary to use two.[26] The scheme invoked here is to use one algorithm to learn what is possible, by establishing a set of ranked constraint strata, plus a second algorithm to fine-tune these strata with statistical information. The two algorithms could in principle run simultaneously; what is crucial is that the statistical fine-tuning must respect the overall stratal structure.

6.2 Learning what is legal

The task of learning what is possible in Hungarian vowel harmony confronts a classical conundrum: no negative evidence is available. The child is never informed that words like *[ɟylyt-nɔk], *[bortog-nɛk], etc., are ill-formed, but rather comes to know it through some combination of Universal Grammar and data processing capable of detecting the systematic gaps in the learning data.

For Optimality Theory, this kind of problem has been addressed with two proposed constraint ranking algorithms: Low Faithfulness Constraint

[25] Plainly, lexical entries must be learned and maintained for unpredictable cases like [ɔtseːl-nɔk] *acélnak* (cf. (28)). Further, at the stage before harmony is learned, we assume that the child also memorises even completely predictable cases like [ɔblɔk-nɔk] *ablaknak* 'window-DAT', in order to have a data set to learn from. Once the child knows that B stems always take back suffixes, it is safe for her to delete such entries from her lexicon, since their form is predictable. However, our guess is that children do not carry out this deletion: the experimental literature (see Baayen *et al.* 2002 and work cited there) includes ample evidence that speakers memorise large numbers of high-frequency regular forms. These plausibly represent the regulars that were memorised while the grammar was being learned.

[26] Zuraw (2000), which inspired our study, uses only one algorithm to rank constraints, but this is because her work did not extend to the task of excluding logically possible but non-occurring patterns of alternation.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Demotion (Hayes 2004) and Biased Constraint Demotion (Prince & Tesar 2004). These algorithms were invented to cover purely phonotactic distributions, but nothing precludes using them on a 'morphotactic' problem such as the distribution of affix allomorphs. In both, the fundamental goal is to rank faithfulness constraints as low as possible, since it is high-ranking faithfulness that leads to overgeneration. In our simulations we tried both algorithms.

We fed the algorithms data that specified the kinds of forms that exist in Hungarian, but without any frequency information. The data were schematic forms like 'B-nɔk', intended to express whole classes of real Hungarian stems that share the same constraint violations. The learning data were as follows:

a. B-nɔk[27]	c. Bi-nɔk Bi-nɛk	f. BNi-nɔk BNi-nɛk	i. FB-nɔk	l. BF FB FNB BNF
b. F-nɛk	d. Beː-nɔk Beː-nɛk	g. BNeː-nɔk BNeː-nɛk	j. BF-nɛk	
	e. Bɛ-nɔk Bɛ-nɛk	h. BNɛ-nɛk	k. Fi-nɛk, Feː-nɛk, Fɛ-nɛk	

Table IV
Forms used in first stage of learning.

The presence of both Bi-nɔk and Bi-nɛk in the learning data (see cell c) meant that a ranking permitting both types had to be found; this turns out to be IDENT-IO[back] ≫ {DISTAL[B], LOCAL[i]}, which can be seen in (34) above. The fact that B-nɛk is *not* in the learning data (cell a) means that a ranking must be found that excludes it; this turns out to be LOCAL [B] ≫ IDENT-IO [back]. Cells b and i–k similarly support the rankings in (34) that result in these gaps. The fact that FB is in the learning data (cell l) means that a ranking must be discovered that permits it; this turns out to be IDENT-IO[back]$_{root}$ ≫ LOCAL F; and similarly for the other forms in the same cell.

The form BNe-nɔk is not in the learning data (see cell h); as noted above, such forms do not occur in Hungarian. However, if we are to account for our wug-test data, where forms like BNe-nɔk actually were volunteered by native speakers, the learned grammar must be able to generate it, even though it is not part of the learning data.

[27] B and F are meant to subsume, among other forms, NB and NF. We included no NB or NF forms in the simulation, because our constraint set includes no DISTAL[N], which would be the only constraint that would 'care' about the non-final N. We have verified that in a full simulation DISTAL[N] would be ranked, as expected, in the bottom stratum (the ranking is needed to derive back harmony in disyllabic *híd* words).

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Both of our algorithms, Low Faithfulness Constraint Demotion and Biased Constraint Demotion, require a set of losing candidates for learning (the rankings are learned by comparing these losers with winning candidates). We obtained our losing candidates using the method given in Tesar & Smolensky (2000) and Prince & Tesar (2004): they are simply the wrong guesses made by preliminary versions of the grammar.

We will not review the specific courses followed by the algorithms, but simply give the results they obtained. The two algorithms learned identical, correct grammars,[28] in the form of the strictly ranked constraint strata given below:

(37) IDENT-IO[bk]$_{rt}$ ≫
{LOCAL[B], LOCAL[F]} ≫
DISTAL[F] ≫
IDENT-IO[bk] ≫
{DISTAL[B], LOCAL[i], LOCAL[eː], LOCAL[ɛ], LOCAL[NN]}

From (37), all the non-stochastic rankings (probability = 1) of Hasse diagram (34) can be deduced. As already shown, these rankings guarantee that none of the impossible forms mentioned above can be generated, and all of the possible forms can.[29]

6.3 Learning statistical distributions

The other part of learning consisted of fine-tuning this overall ranking so as to match the frequencies of the lexicon. For this purpose, we used the Gradual Learning Algorithm (Boersma 1997, Boersma & Hayes 2001), operating under the constraint that it had to respect all of the pairwise rankings defined by the strata in (37). Under this regimen, most stem types cannot influence ranking, so we (harmlessly) restricted the learning set to the stem types that matter, namely Bi, Beː, Bɛ, BNi, BNeː and BNɛ. As a means of approximating the experience of real Hungarian learners, we chose as frequencies the type frequencies found in the Google survey reported in § 3.1, dividing the frequency share of each vacillator in proportion to the token counts. Thus the learning data assumed were as follows:

[28] A detail: Biased Constraint Demotion was run in the modified version devised by Hayes (2004), which adds a provision favouring specific over general faithfulness constraints. Without this modification, the algorithm wrongly puts IDENT-IO[back] at the top of the grammar. Preference for general faithfulness constraints is a problem with the original version of Biased Constraint Demotion, discussed in Hayes (2004: 192–194) and Prince & Tesar (2004: 288).

[29] For the reason just given, we consider BNɛ-nɔk to be a possible form.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(38)

	harmony pattern	*frequency*		*harmony pattern*	*frequency*
Bi	Bi-nɔk	469·8	BNi	BNi-nɔk	6·7
	Bi-nɛk	5·2		BNi-nɛk	23·3
Beː	Beː-nɔk	101·4	BNeː	BNeː-nɔk	5·1
	Beː-nɛk	18·6		BNeː-nɛk	6·9
Bɛ	Bɛ-nɔk	12·1	BNɛ	BNɛ-nɔk	0·0
	Bɛ-nɛk	103·9		BNɛ-nɛk	21·0

We further assumed that this part of learning is uninfluenced by faithfulness, in particular faithfulness to listed suffixed forms like (27c) /pɔlleːr-nɔk/. This assumption was needed to keep faithfulness from determining the outcome, which would have kept the Gradual Learning Algorithm from learning the aggregate statistical pattern of the lexicon. We therefore excluded the faithfulness constraints from this phase of the ranking; they are all ranked correctly in any event in the non-stochastic phase just described.[30]

We ran the Gradual Learning Algorithm for ten trials on these learning data in the way just described. All trials yielded similar outcomes; we report the least accurate one here.[31]

For the stochastic lower region of the grammar, which is what is at issue here, the algorithm learned the stochastic rankings given below. We give pairwise ranking probabilities for the four crucial cases, where DISTAL[B], which favours back suffixes, conflicts with some other constraint favouring front suffixes:

(39)

ranking	*probability*
a. LOCAL[ɛ] ≫ DISTAL[B]	0·903
b. LOCAL[NN] ≫ DISTAL[B]	0·739
c. DISTAL[B] ≫ LOCAL[eː]	0·871
d. DISTAL[B] ≫ LOCAL[i]	0·988

[30] A more nuanced approach would suppose that the irrelevance of IDENT-IO[back] arises not from simply turning faithfulness off, but rather from the fact that it takes time for lexical entries like /pɔlleːr-nɔk/ to get established – it would take multiple hearings for a learner to become confident that [pɔlleːr] is not a vacillator. During this period, when the learner hears [pɔlleːr-nɔk], she can only interpret it as a representative Beː stem. In this capacity, it would play a role in incrementally reranking the markedness constraints to favour back suffixes for such stems.

[31] The learning parameters were: 250,000 trials at each of the plasticity values 1, 0·1, 0·01 and 0·001; noise set at 2·0 for all trials; results tested for 100,000 trials. Strict rankings were enforced by maintaining a minimum distance of 20 along the ranking scale. The ranking values output by the Gradual Learning Algorithm were LOCAL[ɛ] = 105·176, LOCAL[NN] = 103·313, DISTAL[B] = 101·430, LOCAL[eː] = 98·315, LOCAL[i] = 95·079.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

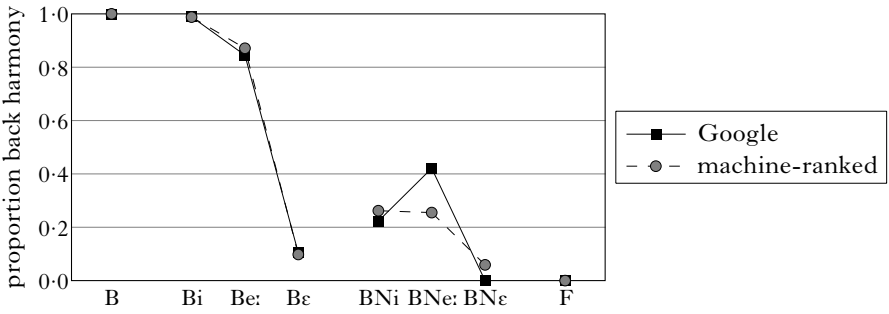


Figure 7
Match-up of machine-ranked model to Google data.

6.4 Evaluating the simulation

The learned grammar can be evaluated in two ways. First, we can ask the purely mechanical question of whether the algorithm was able to rank the constraints in a way that mimicked the frequencies (that is, the Google frequencies) in the learning data. From a more scientific viewpoint, we can ask if it mimicked a real Hungarian speaker: ideally, the learning system should be trained with real data, then behave like a native speaker when it is wug-tested – including any divergences from the pattern in the learning data.

The results for the criterion of corpus-mimicry were reasonably good (correlation for the eight values given: $r = 0{\cdot}988$) and are given in Fig. 7 and Table V.

stem type	wug test	hand ranking	machine ranking	Google frequencies
B	0·988	1·000	1·000	0·997
Bi	0·953	0·953	0·988	0·989
Be:	0·376	0·376	0·871	0·845
Bɛ	0·071	0·071	0·098	0·104
BNi	0·287	0·257	0·262	0·223
BNe:	0·073	0·164	0·255	0·421
BNɛ	0·015	0·045	0·059	0·000
F	0·006	0·000	0·000	0·000

Table V
Source data for Figs 7 and 8.

It is clear that the B and F forms were unproblematic, and that the algorithm was driven primarily by the frequencies of the BN forms, which were closely mimicked. The frequencies of BNN forms were not so

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

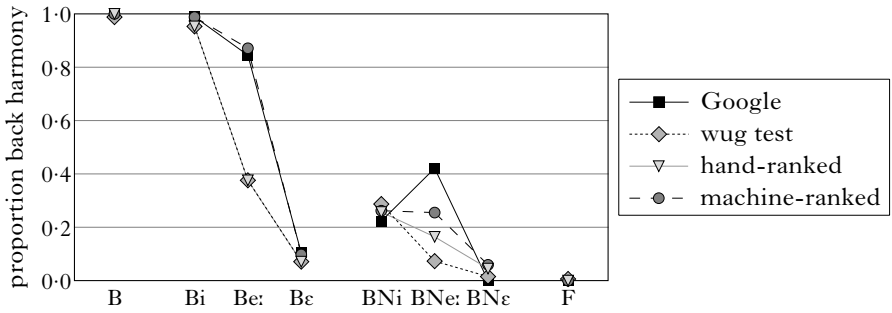


Figure 8
Match-up of hand-ranked and machine-ranked models to wug-test and Google data.

accurately reflected; however, we would claim that this is all to the good: like native speakers, the algorithm *smoothed* these forms (§§ 4.3, 5.6). The peak in BNe: forms is diminished in the learned grammar, and the fraction of BNɛ-nɔk forms is raised from zero to a modest level. These smoothings are the result of the algorithm carrying over patterns from the statistically preponderant BN forms.

Next, we compare the results of the learning model to the wug-test data. The match-up achieved by our model is in Fig. 8. For comparison we include as well the predictions made by the handcrafted grammar of (34).

The following observations seem pertinent:

(i) The machine-learned grammar did somewhat worse than the handcrafted grammar at matching the wug-test data (correlations: $r = 0{\cdot}908$ *vs.* $r = 0{\cdot}990$). This is to be expected, since the handcrafted grammar was deliberately made to match the wug-test data, whereas the machine-learned grammar was trained on the Google data.[32]

(ii) The greatest source of discrepancy between the machine-learned model and the wug-test data was in the Be: forms: for these, a best-fit ranking would have assigned a probability of 0·624 to the ranking LOCAL[e:] ≫ DISTAL[B], whereas the machine-learned grammar had a probability of just 0·129. This discrepancy arises primarily from the difference between the Google data and the wug-test data, which we discussed above in § 4.4. The mismatch is aggravated further by the presence of the aberrant BNe: forms, discussed in § 4.3.[33]

[32] Note that running a learning simulation in which the learning data are wug-test intuitions would lack scientific legitimacy; it would presuppose that language learners could directly access the intuitions of those around them rather than just their speech output.

[33] A reviewer asks if a simple 'bias' constraint, favouring front suffixes, might help here. We have tried this and it does not help; the predicted values in fact are very close to those given by the Gradual Learning Algorithm simulation described in the main text.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

(iii) As noted, the machine-learned grammar smoothed the frequencies of BNN forms. As a result, the correlation of the wug-test data (30 forms) with the predictions of the machine-learned grammar ($r = 0{\cdot}908$) was actually slightly higher than the correlation of the wug-test data with the Google data from which the machine grammar was learned (among the forms covered by the grammar, this is $r = 0{\cdot}896$). In other words, the machine-learned grammar's failures in mimicking its learning data actually made it a marginally better mimic of the wug-test data, a desirable outcome given our goal of modelling real language learners.

(iv) However, the machine-learned grammar evidently did not smooth as much as it should have. The handcrafted grammar, which was set up to match the wug-test data rather than the Google data, achieves a better degree of smoothing for BNe: (a lower value) than the machine-learned grammar.

Overall, we are encouraged by the degree of match between the learning model and the wug-test data, and by the ability of the model to smooth in a qualitatively appropriate way.

6.5 Why two algorithms?

In conclusion, we confess a sense of dissatisfaction that our simulation needs two algorithms to work; one to delimit the range of possible forms, the other to fine-tune the distribution. Our reviewers asked in particular why the Gradual Learning Algorithm alone would not suffice. The answer is that, in the absence of negative evidence, the algorithm has no way of ensuring restrictiveness (non-overgeneration). To be sure, it is possible to give it a crude bias for restrictiveness by starting out the faithfulness constraints ranked much lower than the markedness constraints. However, there is no reason why a head start – no matter how large – would guarantee a correct final grammar. We find, in fact, that if we use the Gradual Learning Algorithm alone, then in all but the most oversimplified, stripped-down simulations, IDENT-IO[back] inevitably rises too high in the grammar.

We gave our two-algorithm simulation to show that at least one form of automated learning can handle the Hungarian facts. However, we judge that our study also shows the need for more theoretical work, particularly to find an algorithm that is both stochastic and restrictive in the absence of negative evidence.

7 Other models

We emphasise that our model represents only one way to use the data of the Hungarian lexicon to make predictions about the harmony behaviour of novel stems. Our particular interest in this model is that it is grounded in a well-developed approach to phonological analysis, Optimality Theory.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

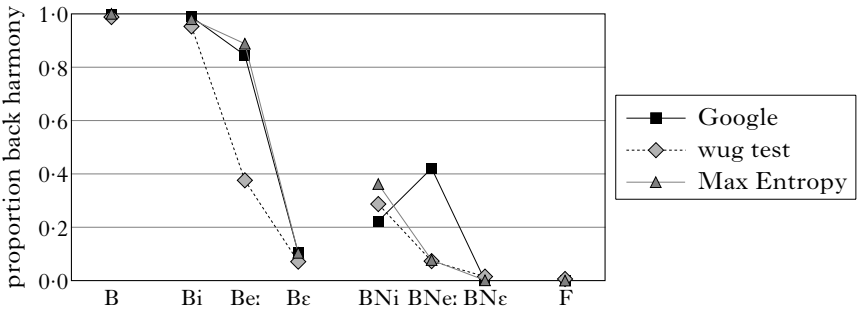


Figure 9
Match-up of Maximum Entropy model to wug-test and Google data.

This is a reasonable research strategy, because OT has been successfully used in the description and analysis of a great deal of other phonological data. There are, however, other important contenders, which certainly should not be ruled out at this stage of research. We will mention a few here.

Fairly closely related to OT is the Maximum Entropy model described by Goldwater & Johnson (2003). Instead of the stochastic OT ranking we used here, this model assigns penalty weights to constraints. To evaluate a candidate, its violations for each constraint are counted, and the count is multiplied by the weight of the constraint. The score assigned to a candidate is the sum of these values for all of the constraints. There is also a learning algorithm for Maximum Entropy models, described in Goldwater & Johnson's work, which finds the set of weights that best fits the data. Applying this model[34] to the Stage II learning data given above, we found that the correlation of the Maximum Entropy model with the wug-test data was r = 0·904. This is marginally better than the 0·902 value for our learned stochastic OT model, but not as good as the 0·991 for the hand-ranked OT model. More strikingly, the Maximum Entropy model proved completely successful in predicting the 'smoothing' effect for BNe: forms described above. The predictions of the Maximum Entropy model are given in Fig. 9.

Both constraint-based models mentioned so far use constraints provided *a priori* under the theory. An alternative is for the model to induce the constraints from the learning data, as in the non-OT models of Ellison (1994) and Albright & Hayes (2002, 2003, in press). The ability to incorporate learned constraints into the model employed here may be crucial in extending it to highly irregular paradigmatic patterns, such as those found in English past tenses, where the constraints would likely have to be partly language-specific.

[34] We have used an implementation devised by Colin Wilson of UCLA, to whom many thanks. Multiple runs yielded identical outcomes.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

The above models use constraints that are violated discretely; thus, for example, for any given vowel pair, LOCAL[B] is either violated or obeyed. An alternative would be to unify LOCAL[B] and DISTAL[B] into a single, gradiently violated constraint, as Frisch *et al*. (2004) do for the similarity-based phonotactics of Arabic roots. We find this prospect intriguing, but at present we do not see how to integrate gradient constraints into a general phonological theory; nor have we yet seen algorithms that could be used for learning and properly weighting such constraints. We view the exploration and theoretical integration of gradient constraints as a task for further research.

Still further afield are models that do not use constraints at all, but rather attempt to predict novel forms on the basis of analogy with existing forms. Important models in this area include the Tilburg Memory-Based Learner (TiMBL; Daelemans *et al*. 2004), Analogical Modelling of Language (AML; Skousen 1989, 2002), and connectionist models, such as the one implemented in Ernestus & Baayen (2003). We judge that the greatest challenge for such models is to derive regular forms with complete reliability; for discussion see Albright & Hayes (2003: 149–152). However, we have not determined whether the regular forms of Hungarian, such as B or F stems, would cause trouble for these models.

It is clear that future research should involve comparisons between rival models. For Hungarian, this comparison will be more effective once we have been able to increase the number of wug forms that have been submitted to native speakers for judgment. For now, we have posted our corpus data and experimental findings on the Internet (http://www.linguistics.ucla.edu/people/hayes/hungarianvh/), and hope that modellers of all persuasions will obtain them and use them to test their proposals.

8 Conclusions

Our main empirical result, from the wug test, is that Hungarian speakers know not just the legal patterns of harmony, but also the frequency of these patterns, and they actively use this knowledge in guessing the harmonic behaviour of novel stems. This is linguistic knowledge that most previous models of phonology and morphology (see § 1) have not captured.

It is not implausible to suppose that this knowledge is in fact useful to speakers and thus worth acquiring. It permits them to guess more accurately when they must produce an inflected form for a stem they have never heard with a suffix – probably a common experience for children. It may also serve them in speech perception, by providing rational top-down biases for recognising suffixes uttered by other speakers.

On the theoretical side, we have suggested that the way Hungarian speakers internalise the frequencies is not through some kind of raw data table, but instead in their grammars. We have found that by ranking some of the constraints stochastically in a 'subterranean' grammar of the kind

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

proposed by Zuraw (2000), we can model the native speaker's intuition fairly accurately. The constraints that are needed are ordinary constraints of Optimality Theory; all that is different is the possibility of stochastic ranking. The fact that stochastic OT also allows rankings that are essentially non-stochastic (probability vanishingly close to 1) means that our model can also rule out impossible forms.

Lastly, we have made an initial attack on the problem of learning such grammars. Our conjecture is that the full set of rankings can be learned through a combination of algorithms, one of which learns the basic range of possibilities, while the other fine-tunes the grammar to match lexical frequencies.

A theme of our work has been to show that tools now exist to permit study of phonological systems in greater detail than would otherwise be possible. Our web-corpus study, experimentation and stochastic theoretical modelling show that the native speaker of Hungarian possesses a richer knowledge of the harmony system than can be adequately described under the older rules-and-exceptions approach.

This said, we believe that a great deal of further progress is needed. In particular, while we have demonstrated the considerable detail with which the native speaker of Hungarian learns the harmony pattern of the lexicon, our study was not designed to find out any upper limit on what is learned: are there statistically reliable patterns in the lexicon that cannot be detected by the human phonological capacity?[35] A positive answer to this question would be very informative concerning the nature of that capacity. We think the methods laid out here might serve to address this question, notably by expanding greatly (perhaps by use of the Web) the scope of wug testing. We also anticipate that further progress will follow from improvements in the underlying phonological theory and in the theory of phonological learning.

REFERENCES

Albright, Adam (2002). Islands of reliability for regular morphology: evidence from Italian. *Lg* **78**. 684–709.

Albright, Adam & Bruce Hayes (2002). Modeling English past tense intuitions with minimal generalization. In Mike Maxwell (ed.) *Proceedings of the ACL-02 Workshop on Morphological and Phonological Learning*. Philadelphia: Association for Computational Linguistics. 58–69.

Albright, Adam & Bruce Hayes (2003). Rules *vs.* analogy in English past tenses: a computational/experimental study. *Cognition* **90**. 119–161.

Albright, Adam & Bruce Hayes (in press). Modeling productivity with the Gradual Learning Algorithm: the problem of accidentally exceptionless generalizations. In Gisbert Fanselow, Caroline Féry, Matthias Schlesewsky & Ralf Vogel (eds.) *Gradience in grammar : generative perspectives*. Oxford: Oxford University Press.

[35] In the Google data, we find small effects from the height of the penultimate vowel of BNN, and also (as in Finnish; Ringen & Heinämäki 1999) of the main stressed (initial) vowel of the word. Given the few forms that we wug-tested, we lack the data to determine whether these effects are actually internalised by native speakers of Hungarian.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Anderson, Lloyd B. (1980). Using asymmetrical and gradient data in the study of vowel harmony. In Vago (1980b). 271–340.
Archangeli, Diana & Douglas Pulleyblank (1987). Maximal and minimal rules: effects of tier scansion. *NELS* **17**. 16–35.
Baayen, Harald, Robert Schreuder, Nivja de Jong & Andrea Krott (2002). Dutch inflection: the rules that prove the exception. In Sieb Nooteboom, Fred Weerman & Frank Wijnen (eds.) *Storage and computation in the language faculty*. Dordrecht: Kluwer. 61–92.
Bailey, Todd M. & Ulrike Hahn (2001). Determinants of wordlikeness: phonotactics or lexical neighborhoods? *Journal of Memory and Language* **44**. 568–591.
Bateman, Leah, Michael O'Keefe, Ehren Reilly & Adam Werle (eds.) (in press). *University of Massachusetts Occasional Papers in Linguistics* 32: *Papers in Optimality Theory III*. Amherst: GLSA.
Beckman, Jill N. (1997). Positional faithfulness, positional neutralisation and Shona vowel harmony. *Phonology* **14**. 1–46.
Benua, Laura (1997). *Transderivational identity : phonological relations between words*. PhD dissertation, University of Massachusetts, Amherst.
Benus, Stefan (2005). *Dynamics and transparency in vowel harmony*. PhD dissertation, New York University.
Benus, Stefan & Adamantios I. Gafos (2005). Qualitative and quantitative aspects of vowel harmony: a dynamics model. In Bruno G. Bara, Lawrence Barsalou & Monica Bucciarelli (eds.) *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. CD-ROM. Mahwah, N.J.: Erlbaum. 226–231.
Berkley, Deborah (2000). *Gradient OCP effects*. PhD dissertation, Northwestern University.
Berko, Jean (1958). The child's learning of English morphology. *Word* **14**. 150–177.
Boersma, Paul (1997). How we learn variation, optionality, and probability. *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam* **21**. 43–58.
Boersma, Paul & Bruce Hayes (2001). Empirical tests of the Gradual Learning Algorithm. *LI* **32**. 45–86.
Bybee, Joan (1995). Regular morphology and the lexicon. *Language and Cognitive Processes* **10**. 425–455.
Bybee, Joan (2001). *Phonology and language use*. Cambridge: Cambridge University Press.
Bybee, Joan & Carol Moder (1983). Morphological classes as natural categories. *Lg* **59**. 251–270.
Casali, Roderic F. (1997). Vowel elision in hiatus contexts: which vowel goes? *Lg* **73**. 493–533.
Chomsky, Noam (1965). *Aspects of the theory of syntax*. Cambridge, Mass.: MIT Press.
Chomsky, Noam & Morris Halle (1965). Some controversial questions in phonological theory. *JL* **1**. 97–138.
Clements, G. N. (1977). Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. *NELS* **7**. 49–64.
Clements, G. N. & Elizabeth V. Hume (1995). The internal organization of speech sounds. In Goldsmith (1995). 245–306.
Coleman, John & Janet Pierrehumbert (1997). Stochastic phonological grammars and acceptability. In John Coleman (ed.) *Proceedings of the 3rd Meeting of the ACL Special Interest Group in Computational Phonology* (*SIGPHON 3*). Somerset, NJ: Association for Computational Linguistics. 49–56.
Daelemans, Walter, Jakub Zavrel, Ko van der Sloot & Antal van den Bosch (2004). *TiMBL : Tilburg Memory-Based Learner* (*version 5.1*). *Reference guide*. Tilburg: ILK.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

de Lacy, Paul (2004). Markedness conflation in Optimality Theory. *Phonology* **21**. 145–199.
Dienes, Péter (1997). Hungarian neutral vowels. *The Odd Yearbook 1997*. 151–180.
Eddington, David (1996). Diphthongization in Spanish derivational morphology: an empirical investigation. *Hispanic Linguistics* **8**. 1–35.
Eddington, David (1998). Spanish diphthongization as a non-derivational phenomenon. *Rivista di Linguistica* **10**. 335–354.
Eddington, David (2004). *Spanish phonology and morphology : experimental and quantitative perspectives*. Amsterdam & Philadelphia: Benjamins.
Ellison, T. Mark (1994). The iterative learning of phonological constraints. Ms, University of Edinburgh. Available (March 2006) at ftp://ftp.cogsci.ed.ac.uk/pub/phonology/papers/ellison-iterative.ps.Z.
Ernestus, Mirjam R. & Harald Baayen (2003). Predicting the unpredictable: interpreting neutralized segments in Dutch. *Lg* **79**. 5–38.
Esztergár, Marianne (1971). *A generative phonology of nouns and vowel harmony in Hungarian*. PhD dissertation, University of California, San Diego.
Farkas, Donka F. & Patrice S. Beddor (1987). Privative and equipollent backness in Hungarian. *CLS* **23:2**. 91–105.
Frisch, Stefan A., Janet B. Pierrehumbert & Michael B. Broe (2004). Similarity avoidance and the OCP. *NLLT* **22**. 179–228.
Frisch, Stefan A. & Bushra Adnan Zawaydeh (2001). The psychological reality of OCP-place in Arabic. *Lg* **77**. 91–106.
Goldsmith, John A. (1979). *Autosegmental phonology*. New York: Garland.
Goldsmith, John A. (ed.) (1995). *The handbook of phonological theory*. Cambridge, Mass. & Oxford: Blackwell.
Goldwater, Sharon & Mark Johnson (2003). Learning OT constraint rankings using a Maximum Entropy model. In Jennifer Spenador, Anders Eriksson & Östen Dahl (eds.) *Proceedings of the Stockholm Workshop on Variation within Optimality Theory*. Stockholm: Stockholm University. 111–120.
Gósy, Mária (1989). Vowel harmony: interrelations of speech production, speech perception, and the phonological rules. *Acta Linguistica Hungarica* **39**. 93–118.
Halle, Morris & K. P. Mohanan (1985). Segmental phonology of Modern English. *LI* **16**. 57–116.
Hansson, Gunnar (2001). *Theoretical and typological issues in consonant harmony*. PhD dissertation, University of California, Berkeley.
Hare, Mary (1990). The role of similarity in Hungarian vowel harmony: a connectionist account. *Connection Science* **2**. 123–150.
Hayes, Bruce (1999). Phonetically driven phonology: the role of Optimality Theory and inductive grounding. In Michael Darnell, Edith Moravcsik, Frederick Newmeyer, Michael Noonan & Kathleen Wheatley (eds.) *Functionalism and formalism in linguistics*. Vol. 1: *General papers*. Amsterdam & Philadelphia: Benjamins. 243–285.
Hayes, Bruce (2004). Phonological acquisition in Optimality Theory: the early stages. In Kager *et al*. (2004). 158–203.
Hayes, Bruce & Margaret MacEachern (1998). Quatrain form in English folk verse. *Lg* **74**. 473–507.
Hayes, Bruce, Bruce Tesar & Kie Zuraw (2003). *OTSoft 2.1*. Software package, http://www.linguistics.ucla.edu/people/hayes/otsoft/.
Itô, Junko & R. Armin Mester (1995). Japanese phonology: constraint domains and structure preservation. In Goldsmith (1995). 817–838.
Kager, René (in press). Lexical irregularity and the typology of contrast. In Kristin Hanson & Sharon Inkelas (eds.) *The structure of words*. Cambridge, Mass.: MIT Press.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Kager, René, Joe Pater & Wim Zonneveld (eds.) (2004). *Constraints in phonological acquisition*. Cambridge: Cambridge University Press.
Kaun, Abigail R. (1995). *The typology of rounding harmony : an optimality theoretic approach*. PhD dissertation, UCLA.
Kaun, Abigail R. (2004). The typology of rounding harmony. In Bruce Hayes, Robert Kirchner & Donca Steriade (eds.) *Phonetically based phonology*. Cambridge: Cambridge University Press. 87–116.
Kiparsky, Paul & Karl Pajusalu (2003). Towards a typology of disharmony. *The Linguistic Review* **20**. 217–241.
Kontra, Miklós & Catherine O. Ringen (1986). Hungarian vowel harmony: the evidence from loanwords. *Ural-Altaische Jahrbücher* **58**. 1–14.
Kornai, András (1991). Hungarian vowel harmony. In István Kenesei (ed.) *Approaches to Hungarian*. Vol. 3: *Structures and arguments*. Szeged: JATE. 183–240.
Lieber, Rochelle (1987). *An integrated theory of autosegmental processes*. Albany: SUNY Press.
Lombardi, Linda (1999). Positional faithfulness and voicing assimilation in Optimality Theory. *NLLT* **17**. 267–302.
McCarthy, John J. & Alan S. Prince (1995). Faithfulness and reduplicative identity. In Jill N. Beckman, Laura Walsh Dickey & Suzanne Urbanczyk (eds.) *Papers in Optimality Theory*. Amherst: GLSA. 249–384.
Müller, Keve (2002). Hungarian word list. Available (March 2006) at http://sourceforge.net/projects/wordlist-hu/.
Pater, Joe (in press). The locus of exceptionality: morpheme-specific phonology as constraint indexation. In Bateman *et al*. (in press).
Pierrehumbert, Janet B. (2001). Stochastic phonology. *Glot International* **5**. 195–207.
Pierrehumbert, Janet B. (in press). An unnatural process. In Louis Goldstein, Douglas Whalen & Catherine T. Best (eds.) *Laboratory Phonology 8*. Berlin & New York: Mouton de Gruyter.
Pinker, Steven (1999). *Words and rules : the ingredients of language*. New York: Harper Collins.
Pinker, Steven & Alan Prince (1988). On language and connectionism: analysis of a parallel distributed processing model of language acquisition. *Cognition* **28**. 73–193.
Prasada, Sandeep & Steven Pinker (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes* **8**. 1–56.
Prince, Alan (1997). Stringency and anti-Paninian hierarchies. Paper presented at LSA Institute, Cornell University. Available (March 2006) at http://ling.rutgers.edu/gamma/talks/insthdt2.pdf.
Prince, Alan & Paul Smolensky (1993). *Optimality Theory : constraint interaction in generative grammar*. Ms, Rutgers University & University of Colorado, Boulder. Published 2004, Malden, Mass. & Oxford: Blackwell.
Prince, Alan & Bruce Tesar (2004). Learning phonotactic distributions. In Kager *et al*. (2004). 245–291.
Reiss, Charles (2003). Deriving the feature-filling/feature-changing contrast: an application to Hungarian vowel harmony. *LI* **34**. 199–224.
Ringen, Catherine O. (1975). *Vowel harmony : theoretical implications*. PhD dissertation, Indiana University. Published 1988, New York: Garland.
Ringen, Catherine O. (1980). A concrete analysis of Hungarian vowel harmony. In Vago (1980b). 135–154.
Ringen, Catherine O. & Orvokki Heinämäki (1999). Variation in Finnish vowel harmony: an OT account. *NLLT* **17**. 303–337.
Ringen, Catherine O. & Robert M. Vago (1995). A constraint based analysis of Hungarian vowel harmony. In István Kenesei (ed.) *Approaches to Hungarian*. Vol. 5: *Levels and structures*. Szeged: JATE. 309–319.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press

Ringen, Catherine O. & Robert M. Vago (1998). Hungarian vowel harmony in Optimality Theory. *Phonology* **15**. 393–416.
Rose, Sharon & Rachel Walker (2004). A typology of consonant agreement as correspondence. *Lg* **80**. 475–531.
Siptár, Péter & Miklós Törkenczy (2000). *The phonology of Hungarian*. Oxford: Oxford University Press.
Skousen, Royal (1989). *Analogical modeling of language*. Dordrecht: Kluwer.
Skousen, Royal (2002). *Analogical modeling : an exemplar-based approach to language*. Amsterdam & Philadelphia: Benjamins.
Tesar, Bruce & Paul Smolensky (2000). *Learnability in Optimality Theory*. Cambridge, Mass.: MIT Press.
Vago, Robert M. (1973). Abstract vowel harmony systems in Uralic and Altaic languages. *Lg* **49**. 579–605.
Vago, Robert M. (1974). *Hungarian generative phonology*. PhD dissertation, Harvard University.
Vago, Robert M. (1976). Theoretical implications of Hungarian vowel harmony. *LI* **7**. 243–263.
Vago, Robert M. (1980a). *The sound pattern of Hungarian*. Washington: Georgetown University Press.
Vago, Robert M. (ed.) (1980b). *Issues in vowel harmony*. Amsterdam: Benjamins.
Vergnaud, Jean-Roger & Morris Halle (1979). Metrical structures in phonology. Ms, MIT.
Walker, Rachel (2001). Round licensing, harmony, and bisyllabic triggers in Altaic. *NLLT* **19**. 827–878.
Wolf, Matt (in press). For an autosegmental theory of mutation. In Bateman *et al.* (in press).
Zuraw, Kie Ross (2000). *Patterned exceptions in phonology*. PhD dissertation, UCLA.

https://doi.org/10.1017/S0952675706000765 Published online by Cambridge University Press