Chapter 5

Extension of *ablaut* in Lakhota

In the preceding chapters, I have outlined a model of base identification that compares various slots in the paradigm, assessing their effectiveness in projecting the remainder of the paradigm. I showed that in the case of both Yiddish (chapter 2) and Latin (chapter 4), one form in the paradigm preserved more contrasts than any other form, and moreover, in subsequent paradigm levelings, contrasts that were maintained in these forms were preserved, while contrasts that were neutralized in these forms were lost.

Most of the neutralizations discussed so far have been *asymmetrical*; that is, the contrast between two segments is better preserved in one form than in another, making mappings in one direction obviously easier than mappings in the other direction. This is illustrated in Figure 5.1. There are many cases that do not fit this pattern, however; often, a neutralization is *symmetrical*, in the sense that neither surface form shows the full range of possibilities. One common type of symmetrical neutralization occurs when a language has three surface patterns involving two phonemes: non-alternating [A], non-alternating [B], and alternating [A] ∼ [B]. In this case, neither direction is obviously better than the other, since each of the surface forms has one ambiguous phoneme (Figure 5.2).

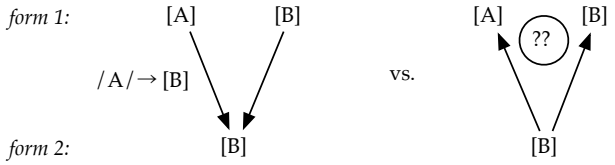

Figure 5.1: An asymmetrical neutralization

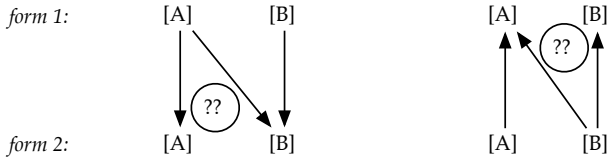

Figure 5.2: A symmetrical neutralization

Lakhota is an example of a language with this type of symmetrical neutralization. There are two types of verbs in Lakhota. The first has invariant final vowels, of any quality; examples of such verbs, with a variety of final vowels, are shown in (59), in the 3sg and 3pl.[1]

(59) Invariant final vowels

3sg	3pl	gloss
*gleshk**a***	*gleshk**a**-pi*	'be spotted'
*low**ã***	*low**ã**-pi*	'sing'
*washt**e***	*washt**e**-pi*	'be good'
*man**i***	*man**i**-pi*	'walk'
*naj**ĩ***	*naj**ĩ**-pi*	'stand'
*man**u***	*man**u**-pi*	'steal'
*nax'**ũ***	*nax'**ũ**-pi*	'hear'
*th**o***	*th**o**-pi*	'be blue'

A second type of verb has a variable final vowel, which surfaces as *-e* in unsuffixed forms (such as the 3sg), and *-a* in the suffixed forms (such as the 3pl) (60). This alternation is known in the Siouanist literature as *ablaut*.

(60) Variant final vowels (*ablaut*)

3sg	3pl	gloss
*chep**e***	*chep**a**-pi* (**chep**e**-pi*)	'be fat'
*kagh**e***	*kagh**a**-pi* (**kagh**e**-pi*)	'do, make'
*khat**e***	*khat**a**-pi*	'be hot'
*naphop**e***	*naphop**a**-pi*	'pop'
*yatk**e***	*yatk**ã**-pi*	'drink'

The puzzle, therefore, is how to distinguish the three-way contrast between invariant [a], invariant [e], and variant ablaut [e]∼[a]:

(61) Three surface patterns

*gleshk**a***	*gleshk**a**-pi*	(invariant [a])
*washt**e***	*washt**e**-pi*	(invariant [e])
*chep**e***	*chep**a**-pi*	(alternating [e]∼[a])

Lakhota ablaut is an "everywhere ambiguous" or symmetrical neutralization. On the face of it, such patterns pose a challenge for the hypothesis that the base must match a single surface form, and must come from the same part of the paradigm for all lexical items. The fact that the verb *chepe∼a* has [e] in some forms and [a] in others cannot be recovered from any single surface form; it is only by comparing two forms that the learner can come to the conclusion that a particular verb exhibits the *ablaut* alternation.

A traditional approach to this type of problem is to encode the difference in the URs of the words, by positing some sort of three-way underlying distinction. Under the standard approach to UR discovery (Chomsky and Halle 1968; Kenstowicz and Kisseberth 1977), it is assumed that all surface forms of a word are derived from a single UR, and wherever possible, surface contrasts should be derived in a lawful way from underlying distinctions. In other words, if it

[1] All Lakhota examples are given in a practical orthography, to be described in section 5.1.1, p. 78.

is possible to distinguish surface distinction ([A] vs. [B]) by using an underlying difference (/A/ vs. /B/), this is preferable to using some other technique, such as marking all B's with a diacritic (/A/ vs. /A/$_{[+A\rightarrow B]}$ or using one UR (/A/) and listing all Bs as lexical exceptions.

In addition to this bias for using distinct URs for contrasting surface patterns wherever possible, it is also generally assumed that for any given word, learners can compare various parts of the paradigm, observe whatever alternations occur, extract all of the unpredictable information, and set up URs that maintain all of the observed contrasts. To take a trivial and uncontroversial example, consider the pattern of alternations caused by final devoicing in German:

(62) German final devoicing

	[ra**t**]	[ra**t**-e]	'advice-nom./dat.'
vs.	[ra**t**]	[ra**d**-e]	'wheel-nom./dat.'

In this case, learners would be able to compare the word for 'advice' and the word for 'wheel', observe that the voicing of the root-final segment is unpredictable in the dative form, and encode this unpredictable voicing specification as part of the UR: /rat/ vs. /rad/. Given these URs, it is also straightforward to formulate a rule of final devoicing that neutralizes the underlying contrast in the nominate form.

A slightly more complicated example comes from Turkish. Like German, Turkish has a general process of final devoicing; however, Inkelas (1994) claims that in addition to words with non-alternating [t] and alternating [t]∼[d] (like German), there are also words with non-alternating [d], such as [etyd]:[2]

(63) Three-way contrast in Turkish final devoicing

	[sanat**t**]	[sanat**t**-ɨ]	'art-nom./acc.'
vs.	[kanat**t**]	[kana**d**-ɨ]	'wing-nom./acc.'
vs.	[ety**d**]	[ety**d**-y]	'etude-nom./acc.'

Turkish final devoicing is therefore a symmetrical neutralization: there are three surface patterns ([t], [d], and [t]∼[d]), so a simple two-way underlying contrast (/t/ vs. /d/) is inadequate. One solution that has often been adopted in the literature is to use underspecification to create an underlying phonemic difference between alternating and non-alternating segments (Inkelas 1994; Inkelas, Orgun, and Zoll 1997; Krämer 2000). For example, following the Prague School practice of including in underlying forms only those specifications are common to all surface forms, we might say that the non-alternating [t] and [d] of Turkish are underlyingly /t/ and /d/, whereas alternating [t] ∼ [d] is an archiphoneme (/D/), with no underlying voicing specification (Trubetzkoy 1962; Anderson 1985, pp. 107-113). The [±voice] specification of underlying /D/ would then be filled on the surface by rules or by markedness constraints, such as no final voiced obstruents (*[+voi,-son]/___]$_\sigma$) and no intervocalic voiceless obstruents (*[-voi,-son]/V___V).

The underspecification/archiphonemic analysis is consistent with the basic tenets of generative phonology, but it is important to remember that an alternative solution is also available in such cases. In particular, if we relax the requirement that all forms in the language must

[2]Not all Turkish speakers seem to agree on whether the nominative singular of 'etude' should be pronounced [etyd] or [etyt]; it is possible that the pattern described by Inkelas represents an especially formal or educated speech style, in which French words are pronounced as faithfully as possible, even if this means violating final devoicing.

be rule-governed, we may use just two underlying phonemes (/t/, /d/), and list some forms as exceptions. In this scenario, we could set up non-alternating [t] as underlying /t/, alternating [t] ∼ [d] as underlying /d/ (with a rule of final devoicing), and non-alternating [d] as underlying /d/, marked in some fashion as an exception to the final devoicing rule. It is this line of analysis that the single surface base restriction forces us to.

The outline of the rest of this chapter is as follows: first, I will show that Lakhota is like Turkish, in that it has a three-way contrast ([a], [e], [e]∼[a]), but only two surface phonemes. Thus, it is a good candidate for a underspecification analysis – in fact, a better candidate than Turkish, because underspecification can account for not only final vowel alternations, but other processes in the language as well. I will then show that the underspecification analysis is nonetheless inadequate for Lakhota, and there is data that it cannot account for. In particular, historical changes show that many invariant [a]'s have switched to variant [e]∼[a], but other logically possible changes have not occurred (invariant [e] $\not\Rightarrow$ [e]∼[a], and [e]∼[a] $\not\Rightarrow$ invariant [a] or [e]). The result is new forms that are inconsistent with any UR in the old system, for reasons that will be explained in section 5.2. This is unexpected under an approach in which learners can compare various parts of the word to posit a UR that can neatly derive all of the surface forms. However, I will show that it follows straightforwardly if we assume that the single surface hypothesis holds not only for bases in output-output effects (paradigm leveling, word-based morphology), but also for the underlying forms that are the inputs to phonology. The end result is that the single surface base restriction appears to be relevant not only for models that consider the relations between surface forms, but also for models that use potentially abstract URs of stems.

I begin with a brief overview of the Lakhota segment inventory, and the processes involved.

5.1 Background on Lakhota

Lakhota is a Siouan language, spoken by roughly 6,000 speakers today in the Dakotas and surrounding areas (Grimes 2000). I draw my Lakhota data from the following sources, differentiating them where necessary: Boas and Deloria's grammar (1941), Buechel's Lakhota dictionary (1970), a verb list compiled in field work by Munro (1989), and notes from my own field work from 1999-2001 with Mary Rose Iron Teeth, a native speaker from the Pine Ridge Reservation in South Dakota.

5.1.1 Phoneme inventory and phonotactics

The Lakhota phoneme inventory is given in Tables 5.1 and 5.2; the practical orthography that I will be using here is given in italicized letters, and the IPA (where different) is given in brackets.

A phonotactic fact about Lakhota that will be relevant for this discussion is that there is a relatively large set of permissable CC onsets (including sequences like [kt], [xt], [mn], and so on), but codas are generally not allowed, especially in word-final position.

5.1.2 Final vowel alternations ("ablaut")

As described above, some Lakhota verbs have final vowel alternations between [e] and [a], in a process known as umlaut; the basic problem is to differentiate the following three types of

Table 5.1: Lakhota consonant inventory

unaspirated	*p*	*t*	*c* [tʃ]	*k*
aspirated	*ph* [pʰ]	*th* [tʰ]	*ch* [tʃʰ]	*kh* [kʰ]
ejective	*p'*	*t'*, *s'*	*c'* [tʃ'], *sh'* [ʃ']	*k'*, *x'*
fricatives		*s*, *z*	*sh* [ʃ], *j* [ʒ]	*x*, *gh* [ɣ]
nasals	*m*	*n*		*ng* [ŋ]
liquid		*l*		
glides			*y* [j]	*w*

Table 5.2: Lakhota vowel inventory

Oral		Nasal	
i	*u*	*ĩ*	*ũ*
e	*o*		
	([ɔ]*)		
([æ]*)	*a*		*ã*

*[æ] and [ɔ] are derived from /aya/, /awa/

words:

(64) Three surface patterns for final [a], [e]

*gleshk**a***	*gleshk**a**-pi*	'be spotted'	(invariant [a])
*washt**e***	*washt**e**-pi*	'be good'	(invariant [e])
*chep**e***	*chep**a**-pi*	'be fat'	(alternating [e]∼[a] = "ablaut")

Pursuing an underspecification approach along the lines of Inkelas (1994), we would start by inferring that alternating verbs like *chep*{*e*∼*a*} must end in something other than [e] or [a]. Using the strategy of creating an archiphoneme with just the shared feature specifications, this would lead us to conclude that such verbs end in an abstract segment that I will write as /A/:[3]

[3]An alternative analysis, suggested by Kim (2002), is that alternating *e*∼*a* is not underspecified, but rather *overspecified*, including not only [a]-features, but also a floating dorsal feature which combines with a floating coronal (front) feature to yield [e] before certain suffixes. This suggestion, which is in line with Lieber's autosegmental approach to morphologically-conditioned mutations (Lieber 1987; Lieber 1992), is problematic in various respects. First, using Kim's feature system, we might expect the combination of [a] with coronal and dorsal features to produce [æ] rather than [e], particularly since the language already has a surface [æ] that results from coalescence of /aye/ and /aya/. Second, the representation with floating features is supposed to unify the ablaut alternation with another coalescence process, of /ai/ to [e]. However, /ai/ to [e] coalescence is not a productive process in the language—surface [ai] sequences can easily be created by combining, for example, the valence-adding prefix *a-* with the instrumental/locative prefix *i-*. In addition, there is another process, not discussed here, in which alternating *e*∼*a* raises to [i] before certain morphemes (such as the future marker *-(n)kte* and the conjunctive clitic *-na*), so we would need to find some other floating feature to attach to these morphemes, and also provide a mechanism to delete the place features of /a/ so that it can raise to [i] in this context. Finally, the floating feature representation cannot explain why ablaut verbs also behave differently in reduplication (section 5.1.3).

(65) Feature specifications of /a/, /e/, /A/

/a/	/e/	/A/
+syllabic	+syllabic	+syllabic
-high	-high	-high
+low	-low	
+back	-back	

Underspecification can be used to derive surface alternations quite naturally in OT; underspecified segments have less to be faithful to, so general principles of markedness (that are needed in the grammar anyway) can play a greater role in determining their surface realization without incurring faithfulness violations. In this case, the crucial markedness constraint is "no word-final [a]", which is admittedly rather language-particular. However, the pattern falls out easily with the following rankings: first, all of the IDENT constraints for vowel features are ranked at the top of the grammar, forcing the surface form to preserve whatever feature values have been specified underlyingly (66a). Second, the general markedness constraint banning [e] is ranked above the constraint banning [a], and the language-particular constraint banning word-final [a] is ranked above both of these, forcing underspecified vowels to be realized as [e] word-finally, and [a] elsewhere (66b).

(66) Ident(V) $\gg$ *[a]/___# $\gg$ *[e] $\gg$ *[a]

a. Ident(V): violated when an underlyingly specified vowel is changed (/a/→[e], /e/→[a])

/gleʃka/ 'spotted'		Ident(V)	*[a]/___#	*[e]	*[a]
☞	a. [gleʃka]	√	*	*	*
	b. [gleʃke]	*!	√	**	√

/waʃte/ 'good'		Ident(V)	*[a]/___#	*[e]	*[a]
	a. [waʃta]	*!	*	√	**
☞	b. [waʃte]	√	√	*	*

b. Ident(V) satisfied by both /A/→[a] and /A/→[e]; realization falls to markedness constraints

/tʃ^{h}epA/ 'do, make'		Ident(V)	*[a]/___#	*[e]	*[a]
	a. [tʃ^{h}epa]	√	*!	√	**
☞	b. [tʃ^{h}epe]	√	√	*	*

/tʃ^{h}epA-pi/ 'do, make'		Ident(V)	*[a]/___#	*[e]	*[a]
☞	a. [tʃ^{h}epa-pi]	√	√	√	**
	b. [tʃ^{h}epe-pi]	√	√	*!	*

This analysis captures the pattern of final vowel ablaut, but requires positing an abstract, underspecified archiphoneme. Under a traditional approach to UR discovery, the mere existence of a three-way contrast is sufficient evidence for learners to infer that they need an abstract segment (either an archiphoneme, or a fully specified segment that never surfaces as such). Ideally, however, we might like some external evidence confirming this analysis, such as an indication that final ablaut vowels behave differently from nonalternating [e] and [a] in other respects as well. In fact, there is such evidence, in the form of differences in reduplication patterns.

5.1.3 Reduplication

Verbs can reduplicate in Lakhota, with a variety of meanings. In many cases, reduplication marks plurality, especially with stative verbs (67a). In other cases, it marks intensivity/iterativity/durativity (67b), while in other cases, the meaning is not so clear (67c).

(67) Meaning of reduplication in Lakhota

a. Plurality (mainly statives)

sha-sha	'red-pl.'
washte-shte	'good-pl.'

b. Intensive/iterative/durative

yushna-shna	'sprinkle'	(cf: *yushna* 'drop')
naphã-phã	'trample'	(cf: *naphã* 'stomp')
lowã-wã hiyaye	'went along singing'	(cf: *lowã* 'sing')

c. Meaning not so clear

gleshka-shka	'checkered/plaid'	(cf: *gleshka* 'spotted')

The basic pattern of reduplication is to copy the final syllable, as seen in (68):

(68) Reduplication of the final syllable

3sg	redup.	gloss
gleshka	*gleshka-shka*	'be spotted'
washte	*washte-shte*	'be good'
lowã	*lowã-wã*	'sing'
naxcha	*naxcha-xcha*	'blossom'
shakpe	*shakpe-kpe*	'be six in number'
yamni	*yamni-mni*	'be three in number'
zaptã	*zaptã-ptã*	'be five in number'
shakowĩ	*shakowĩ-wĩ*	'be seven in number'
wikcemna	*wikcemna-mna*	'be ten in number'

However, verbs with final *e* ∼ *a* alternations generally copy the "maximal penult",[4] sometimes with accompanying segmental changes (devoice fricatives, change /t/ → [l], etc.), as seen in (69):

(69) "Non-final reduplication"

3sg	redup.	gloss
chepe	*chep-chepe* (**chepe-pe*)	'be fat'
kaghe	*kax-kaghe* (**kaghe-ghe*)	'do, make'
khate	*khal-khate* (**khate-te*)	'be hot'
naphope	*na-pho-phope* (**naphope-pe*)	'pop'

The traditional analysis of this difference (Boas and Deloria 1941; Shaw 1980) is that final alternating (ablaut) vowels are completely absent underlyingly: /tʃʰep/, /kaɣ/, etc. Under this

[4]The 3sg form *chepe* is syllabified *che.pe*, so reduplicating just the penult should yield *che-chepe*. In this and many other cases, reduplication ignores syllabification of the base form, and copies as much as it can fit into a syllable.

analysis, the URs of these words have codas, which are prohibited on the surface (section 5.1.1). The illegal codas are then fixed by a process of epenthesis, which inserts an [e] word-finally ([tʃʰep**e**]), and an [a] word-internally ([tʃʰep**a**-pi]).[5] This allows us to say that reduplication is always final, and precedes epenthesis:[6]

(70) Rule ordering: reduplication precedes epenthesis (after Shaw 1980)

a. Simple forms

UR	/waʃte/	/gleʃka/	/tʃʰep/
REDUPLICATION	—	—	—
EPENTHESIS	—	—	tʃʰep**e**
SR	[waʃte]	[gleʃka]	[tʃʰepe]

b. Reduplicated forms

UR	/waʃte-RED/	/gleʃka-RED/	/tʃʰep-RED/
REDUPLICATION	waʃte-**ʃte**	gleʃka-**ʃka**	tʃʰep-**tʃʰep**
EPENTHESIS	—	—	tʃʰep-tʃʰep**e**
SR	[waʃte-ʃte]	[gleʃka-ʃka]	[tʃʰep-tʃʰepe]

Treating ablaut vowels as epenthetic is a more radical version of the underspecification analysis sketched above. The claim is that not only the distribution of [a] and [e] but the very occurrence of the vowel is predictable based on surface markedness considerations. The analysis of ablaut alternations would be much the same as in (66b) above, with the addition of a high-ranking *CODA constraint, and constraints ruling out the insertion of vowels other than [e] and [a].[7]

The epenthesis analysis has some obvious advantages. First, it captures the co-occurrence of two properties of words like *chepe*: they have final vowel alternations, and they have non-final reduplication. Furthermore, all words are completely rule-governed. If a speaker knows that there is an epenthesis process (resulting in [e] word-finally and [a] before a morpheme boundary), a final reduplication process (rendered opaque by epenthesis), and two types of URs (those with final consonants and those with final vowels), then it is possible to use the grammar to derive all of the surface forms correctly (70).

Let us now consider the various possible sources of acquisition-related error under this analysis, as we did for Yiddish in section 2.5. Suppose that a learner is faced with a new word, whose forms are not completely known. For example, suppose she hears a new 3sg form *pughe*

[5]This analysis recapitulates the history of verbs with ablaut alternations. It appears that Siouan did originally have consonant-final and vowel-final verbs, but at some point two post-verbal clitics (*-a* and *-e*) were reanalyzed as part of the verb stem, or as epenthetic vowels inserted to fix word-final codas: *chep-e* ⇒ *chepe* (Rood 1983).

[6]It is not easy to recast this analysis of the reduplication facts into OT. Intuitively, we want to penalize copying an epenthetic vowel, but Base-Reduplicant (BR) correspondence constraints do not know which base segments have incurred IO faithfulness violations (such as a DEP violation). The only other possibility is to use Input-Reduplicant (IR) correspondence to penalize having an epenthetic vowel in the reduplicant; however, this would require ranking DEP-IR above DEP-IO, which leads to undesirable typological consequences (McCarthy and Prince 1995, pp. 114-117). I will not pursue this problem here, since I will ultimately be arguing that the "ablaut vowel as epenthesis" analysis is wrong in any case.

[7]This could be accomplished either by faithfulness, with DEP-IO(i,u,o,æ,ĩ,ũ,æ̃), or else by markedness, with *[i], *[u], *[o], etc. The former approach looks more promising, since it seems questionable to claim that [i] is a more marked vowel than [e], which would be required in the ranking *[i] ≫ *[e].

'he snorted' – what might she conclude? One possibility is that she may assume that the [e] is underlying, setting up a UR /puɣe/ and predicting a plural form *pughe-pi* and a reduplicated form *pughe-ghe*. Another possibility is that she may assume that the [e] is not underlying (/puɣ/), and predict a plural form *pugha-pi* and a reduplicated form *pux-puɣe*. Conversely, suppose that the learner has heard a new 3pl form *puza-pi* 'they are dry.' In this case, she may either assume that the [a] is underlying (predicting 3sg *puza*, reduplicated *puza-za*), or she may assume that the [a] is epenthetic (predicting 3sg *puze*, reduplicated *pus-puze*).

As in chapter 2, it is difficult to make exact predictions about which errors we expect under a traditional model without an explicit theory of how learners reason about URs with incomplete information. A reasonable default assumption would be that learners do not posit underlying underspecification unless they have heard evidence that the word actually alternates. This is the principle behind the Prague School's use of archiphonemic underspecification, and it is also the principle behind Lexicon Optimization in OT (Prince and Smolensky 1993). In the present case, that would mean that learners with incomplete information would always set up a fully specified vowel, but that vowel should sometimes be /a/ and sometimes /e/, depending which form had been learned. A more subtle assumption is that speakers know the predominant patterns of their lexicon, and if the dominant pattern is alternation, then they are able to set up underspecified URs without actually hearing the alternation. This has been proposed by Inkelas (1996) as Alternant Optimization, and by Harrison and Kaun (2000) as Pattern-Responsive Lexicon Optimization. In the present case, if ablaut is the dominant pattern, then Pattern-Responsive Lexicon Optimization might lead learners to assume that partially-known words are underspecified in such a way that produces ablaut alternations and non-final reduplication.

Crucially, all of these theories share a common prediction: no matter what principles the learner uses to set up a UR using incomplete information, the result should resemble a valid existing paradigm. In particular, if she assumes that the final vowel of a word is underlying, then it should be invariant, and the word should have final reduplication. If, on the other hand, she assumes that the final vowel is underspecified, it should exhibit the ablaut alternation, and have non-final reduplication. In the next section, I will show that this prediction is wrong. As it turns out, two new "inconsistent" paradigm types have been created in Lakhota, both of which are incompatible with the analysis laid out thus far. After presenting the data, I will show that although these new paradigm types are unexpected under any version of the traditional analysis, they are in fact predicted by a single surface base approach.

5.2 Innovative paradigms in Lakhota

The verb types that I have discussed thus far are those that have a straightforward historical origin. In addition to the two paradigm types discussed so far (invariant vowels with final reduplication, and ablaut alternations with non-final reduplication), there have also arisen two innovative paradigm types in Lakhota. The first are paradigms with variant final vowels (ablaut),

but with final reduplication, as in (71):[8]

(71) Innovative paradigm type 1: ablaut plus final reduplication

3sg	*3pl*	*redup*		*gloss*
hãske	*hãska-pi*	*hãska-ska*	(**hã-hãske*[9])	'be tall'
hĩshme	*hĩshma-pi*	*hĩshma-shma*	(**hĩ-hĩshme*)	'be fuzzy'
ixat'e	*ixat'a-pi*	*ixat'at'a*	(**i-xa-xat'a*)	'laugh'
hoxpe	*hoxpa-pi*	*hoxpa-xpa*	(**hox-hoxpe*)	'cough'
naxme	*naxma-pi*	*naxma-xma*	(**nax-naxme*)	'hide'
kaxpe	*kaxpa-pi*	*kaxpa-xpa*	(**kax-kaxpe*)	'knock down'
katke	*katka-pi*	*katka-tka*	(**kal-katke*)	'choke'

The second innovation is a paradigm type with invariant final vowels, but non-final reduplication, as in (72):

(72) Paradigms with invariant final V, but non-final reduplication

3sg	*3pl*	*redup*		*gloss*
thokca	*thokca-pi*	*thok-thokca*	(**thokca-kca*)	'be different'
topa	*topa-pi*	*top-topa*	(**topa-pa*)	'be four in number'
ota	*ota-pi*	*ol-ota*	(**ota-ta*)	'be many'

It appears, then, that there have been been two changes, leading to the creation of two new paradigm types:

- The *-a*/*-e* alternation has been extended to some verbs that used to have invariant *-a* (**hinshma* ⇒ *hinshme* 'fuzzy-3sg')
- Nonfinal reduplication has been extended to some verbs that should have had final reduplication

These innovations are significant for two reasons. The first reason is that words belonging to the new paradigms are incompatible with any UR in the old system. The contradiction is

[8] There are several sources of evidence that these patterns are in fact innovative, and that the *-a*/*-e* alternation has been extended to forms which originally did not have it. First, there are verbs whose only vowel is an ablaut vowel (e.g., *t'e*∼*a* 'die'), and if ablaut vowels originated as reanalyzed clitics (fn. 5), then we would be forced to infer that these verbs were originally just a single consonant (*t'*). It seems more plausible to say that these verbs were originally CV (*t'a*), and that the ablaut alternation has been extended to them analogically – especially since there are sometimes words that appear to be etymologically related and have invariant *-a*, such as *t'at'a* 'listless, lazy.' In addition to this, some forms listed with *-a* in Boas and Deloria (1941) are now more common with *-e*/*-a* (e.g., *naxma* 'fled-3sg' ⇒ *naxme*). Finally, I have observed a fair amount of synchronic uncertainty or variation in whether a final *-a* should alternate or not, including even the use of both *-a* and *-e*/*-a* on the same verb in the same session. It should be noted, however, that some "impossible" forms also seem to be rather old – for example, *yatkan*/*e* is found in all sources and shared with other dialects, but appears to be a relatively local innovation in this branch of Siouan (Shaw 1980; Rood 1983). While it is interesting that this pattern is spreading to more and more verbs over time, what I am really concerned with here is what mechanism allowed the very first inconsistent paradigms to be created.

[9] It should be noted that in all cases, the non-occurring reduplications are phonotactically legal – so although on first glance, we might think of trying to explain the nonoccurrence of forms like *hã-hãska* as avoidance of sequences like [hãhã], perhaps due to its intervocalic [h], in fact such sequences are permitted in other words, like *hũke-shni* ∼ *hũhũka-pi-shni* 'weak'.

illustrated in Fig. 5.3; the fact that this word has a final vowel alternation would lead us to conclude that the final vowel is not specified underlyingly, while the fact that the final syllable reduplicates would lead us to conclude that the final vowel is present underlyingly.

The second reason is that the changes leading to new paradigm types have been *asymmetrical*; they have affected only words with original *-a* throughout the paradigm, and not *-e*. Thus, there are plenty of words like *hanske* which have switched from invariant *a* to alternating *e* ∼ *a* (73a), but no words that have switched from invariant *e* to alternating *e* ∼ *a* (73b).

(73) Changes have been asymmetrical

a. Attested:
naxma ⇒ *naxm**e*** *naxma-pi* *naxma-xma*

b. Not attested:
washte *washte-pi* ⇏ **washt**a**-pi* *washte-shte*

Thus, the Lakhota change poses two mysteries: first, how were new, "internally inconsistent" paradigms created? (This was not predicted in any of the incomplete learning scenarios discussed in the previous section.) Second, why were only [a]-final verbs affected? In the next section, I will show that both of these mysteries can be explained under a model that limits learners to choosing URs that match a particular surface form (the single surface base hypothesis). Under this restriction, learners are not always able to set up a UR that preserves all surface contrasts; in fact, in the case of symmetrical neutralizations, neither form alone can predict the paradigm of a word. As with Yiddish and Latin, I will compare the various forms in the Lakhota verb paradigm and see whether there is a form that, while not preserving *all* contrasts, at least preserves more contrasts than any other form. It will emerge that once both phonological and morphological neutralizations are taken into consideration, there is such a form (a second person form). Moreover, when we consider the grammar that would be needed to derive the remainder of the paradigm from the second person, it predicts two types of overregularization: extending ablaut and non-final reduplication to [a]-final verbs.

5.3 Restricting UR discovery to a single surface form

Let us now go back to the beginning, this time operating under the single surface base restriction. Recall that the basic analytical problem in (61) (repeated below), is that there are three

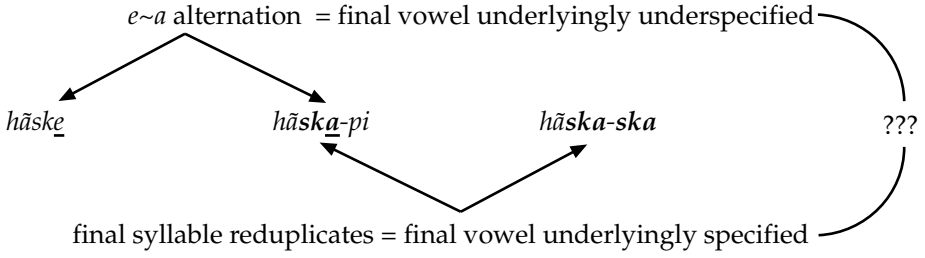

Figure 5.3: Innovative forms are incompatible with any UR

surface patterns, but only two phonemes involved. The challenge, therefore, is to come up with an underlying form for *chep*{*e*∼*a*}.

(61) Three surface patterns

*gleshk**a***	*gleshk**a**-pi*	(invariant [a])
*washt**e***	*washt**e**-pi*	(invariant [e])
*chep**e***	*chep**a**-pi*	(alternating [e]∼[a])

Under the single surface base restriction, we are now limited to choosing either /tʃʰepe/ or /tʃʰepa/. This leaves us with a number of possible analyses. We could, for example, choose /tʃʰepa/ with underlying /a/, and then posit a final raising rule (or its OT equivalent), as in (74):

(74) FINAL RAISING (ablaut): /a/ → [e] / ___#

This analysis would correctly derive [tʃʰepe] and [kaɣe] from their underlying forms /tʃʰepa/ and /kaɣa/, but it would fail for [gleʃka], incorrectly predicting the raised form *[gleʃke]. Thus, under this analysis, we would have to list [gleʃka] as an exception, which would block grammatically expected form [gleʃke]. Conversely, we could assume that there is no default final raising rule, and then list words with raising as exceptions ([tʃʰepe], [kaɣe]), or make Final Raising a lexically restricted rule, and mark /tʃʰepa/ and /kaɣa/ with [+Final Raising] diacritics. This is not an exhaustive list of all of the possible analyses, but it should be clear that no matter which UR we pick (/tʃʰepe/ or /tʃʰepa/), there will be some exceptions. The reason is that now we have only two URs available (/a/, /e/) to represent three surface patterns ([a], [e], [a]∼[e]). Some unpredictable information is going to have to be stored somewhere else, and that somewhere is the exception handling mechanism.

Once we recognize that exceptions are unavoidable, we can at least try to mitigate the problem by finding the set of URs and rules that requires the fewest listed exceptions. In order to do this, we will want to base the UR on the part of the paradigm that is "most informative" — that is, that has the fewest neutralizations, affecting the fewest lexical items. In order to assess this for Lakhota, we need to consider the neutralizations that might affect Lakhota verbs, and how many verbs are affected by each.

5.3.1 What is the most informative part of the Lakhota paradigm?

In order to evaluate the seriousness of various neutralizations in Lakhota, I selected a database of "simple" verbs. I began with the list of verbs compiled by Munro (1989), and then removed all entries that were morphologically complex according to one of the following criteria. First, I removed all "compound" entries, consisting of a combination of a verb plus verb, noun plus verb, preposition plus verb, and so on (e.g., *akan ishtima/e* 'sleep on', from *ishtima/e* 'sleep'). Next, I removed all entries derived by the valence-changing prefix *a-* (adds one argument), the causative suffixes *-ye* and *-khiye*, the reflexive marker *-c'i-*, and the possessive object marker *ki-* ('X one's own ___'). For example, the verb *akipsica/e* 'to jump over one's own' is derived by prefixing the valence-changing *a-* and possessive object *ki-* to *psica/e* 'jump'). I left in verbs containing derivational prefixes that are identifiable, but not predictable, such as *pa-* 'using hands', *na-* 'using feet', *ya-* 'using the mouth', etc. These prefixes are analogous to English *trans-* or *dis-*, in that they are easily segmented out as prefixes, but they are not productive,

Table 5.3: Number of words in each ablaut category (out of 545 total in database)

Pattern	*Count*
invariant *-a*	83
ablaut *-e/-a*	199
invariant *-e*	65
nasal ablaut *-e/-ã*	9
invariant *-ã*	41

they do not occur with all roots, and some verb roots that occur with them are bound roots. When these criteria were applied, a database of 545 simple verbs remained.

There are various sources of systematic unpredictability in Lakhota verbs. These include phonological unpredictability, such as whether or not a verb has ablaut alternations, and also morphological unpredictability, such as the location of person agreement. In addition to these wide-scale, systematic sources, there are also other sporadic irregularities that affect just a few verbs, and will not be discussed here.

Phonological unpredictability: ablaut

One major unpredictable property of a Lakhota verb is whether or not it has the ablaut alternation that has been the focus of discussion up until this point. In the examples thus far, I have limited the data to the three major patterns: invariant *a*, invariant *e*, and ablaut *e* ∼ *a*. There are, however, also a handful of words that display what I will call a *nasal ablaut* alternation between *e* and *ã*. Thus, the full range of possible surface patterns is as in (75).

(75) Ablaut alternations:

3sg (unsuffixed)	3pl (suffixed)	gloss	category
*gleshk**a***	*gleshk**a**-pi*	'spotted'	invariant *a*
*chep**e***	*chep**a**-pi*	'fat'	ablaut
*washt**e***	*washt**e**-pi*	'good'	invariant *a*
*yatk**e***	*yatk**ã**-pi*	'drink'	nasal ablaut
*yat**ã***	*yat**ã**-pi*	'light (a cigarette)'	invariant *ã*

Comparing the words in (75), we can see that both the unsuffixed 3sg and the suffixed 3pl suffer from neutralizations. The 3sg form neutralizes 3 types of words: invariant *-e*, ablaut *-e/-a*, and nasal ablaut *-e/-ã* all have *-e* in this form. Turning to the 3pl form, we see that 2 pairs of word types are neutralized: invariant *-a* and ablaut *-e/-a* are both *-a* in this form, and invariant *-ã* and ablaut *-e/-ã* are both *-ã*. Neither form is obviously better than the other in allowing us to predict which surface pattern a word should take; thus, as with Latin, we must compare the seriousness of the neutralizations by considering how many lexical items are affected by each.

The numbers of words instantiating each of the patterns in (75) are given in Table 5.3. As can be seen from the table, the (non-nasal) ablaut pattern is well represented, with almost 40% of verbs participating in it. There are also a fair number of invariant *a* and *e* verbs, with relatively fewer invariant *ã* verbs, and just a handful of nasal ablaut verbs.

Given these counts, let us now consider how informative the 3sg and 3pl forms are in practice in predicting the remaining of the paradigm. If we use the singular (unsuffixed) form,

we will have the following URs for the words in (75): /gleʃka/, /tʃ^{h}epe/, /waʃte/, /jatke/, and /jatã/. The problem here is the three forms with underlying /e/, which belong to three different surface classes. The majority of words with *e* in the 3sg are ablaut verbs with *a* in the plural (e.g., [tʃ^{h}epa-pi]), so if our goal is construct a grammar that can cover a majority of forms, we need to posit some sort of derived environment non-final lowering rule: /e/ → [a] /___+C. With this rule in place, plurals with [a] like [tʃ^{h}epa-pi] will be accounted for, and we just need to list non-lowerers like [waʃte-pi] (65 of them) and nasalizers like [jatkã-pi] (9 of them) as exceptions. Thus, choosing the 3sg as the UR would require 65 + 9 = 74 exceptions.

If, on the other hand, we were to use the plural (suffixed) form as the UR, we would have the following: /gleʃka/, /tʃ^{h}epa/, /waʃte/, /jatkã/, and /jatã/. In this case, there would be two problems: the two verbs with underlying /a/, and the two with underlying /ã/. Among those with underlying /a/, the majority have [e] like [tʃ^{h}epe] in singular, so we would need to posit a final raising rule (/a/ → [e] /___#). This would correctly derive /tʃ^{h}epa/ → [tʃ^{h}epe], but it would incorrectly predict *[gleʃke] for [gleʃka]. Therefore, we would need to list non-raisers like [gleʃka] as exceptions (83 exceptions). Among the underlying /ã/ words, the majority are invariant like [jatã], so we would not want to extend the final raising rule to cover nasalized vowels as well; rather, we would just list the nasal ablaut verbs like [jatke] as exceptions (9 exceptions).[10] Thus, choosing the 3pl as the UR would require 83 + 9 = 92 exceptions.

What we see from this comparison is that the unsuffixed (3sg) form is slightly better in predicting the final vowel of the suffixed (3pl) form than vice versa, requiring 18 fewer exceptions for this set of verbs (= 92 - 74). This advantage is rather small, however, and choosing the 3sg form as the UR relies on a rather questionable phonological rule (non-final lowering of /e/→[a] only before a suffix) in order to make ablaut verbs rule-governed. What I conclude from this section, therefore, is that the ablaut neutralization really is quite symmetrical, and any advantage that one form may have over the other will have to come from whatever other contrasts they may preserve.

Morphological unpredictability: person agreement

Another important unpredictable property of Lakhota verbs is the position of person agreement. Lakhota verbs fall into two classes, based largely (but not entirely) on whether they are active or stative. The subject markers for these two classes of verbs are given in Table 5.4. Note that *-pi* is a plural suffix for animate subjects, and therefore shows up in all of the plural cells; *-he*/*-ho* is a second person suffix, used in questions and second person declarative sentences (*-he* by female speakers, *-ho* by male speakers). Therefore, the the 2sg, 1pl, and 2pl forms usually occur with a suffix, as does the 3pl if it has an animate subject.

Membership of a verb in the active or stative class is more or less predictable given the meaning of the word; the position of the person agreement within the verb, on the other hand, is not. Subject markers in Lakhota may occur either as prefixes or as infixes, depending on the

[10] Shaw (1980) also treats nasal ablaut verbs as exceptions, marking them diacritically to take the /a/ → [e] ablaut rule even though they do not strictly provide the input for this rule, which is [a].

[11] When the 1sg marker *wa* occurs before a *y*, there is a morphophonological process that turns the *wa-y* sequence into *bl* – e.g., *wa-yatke* → *blatke* 'I drink'.

[12] When the 2sg marker *ya* occurs before a *y*, the *ya-y* sequence becomes *l* – e.g., *ya-yatke* → *latke* 'you drink'.

Table 5.4: Lakhota subject markers

a. Active (Munro's Type I)

	sg	*pl*
1st	wa[11]	un(k) ... pi
2nd	ya[12] ... he/ho	ya ... pi he/ho
3rd	∅	∅ ... pi

b. Stative (Munro's Type II)

	sg	*pl*
1st	ma	un(k) ... pi
2nd	ni ... he/ho	ni ... pi he/ho
3rd	∅	∅ ... pi

verb. Although I am not aware of any actual minimal pairs that differ only in the placement of person agreement, the verbs for 'to be lost' and 'to walk' in (76) are very similar phonologically, but get their subject markers in different positions.

(76) Variable position of subject markers

a. Sometimes prefixed

'be lost'	*sg*	*pl*
1st	**wa-nuni**	**un-nuni-pi**
2nd	**ya-nuni he**[13]	**ya-nuni-pi he**
3rd	nuni	nuni-**pi**

b. Sometimes infixed

'walk'	*sg*	*pl*
1st	ma-**wa-ni**	ma-**un-ni-pi**
2nd	ma-**ya-ni he**	ma-**ya-ni-pi he**
3rd	mani	mani-**pi**

The unpredictable location of person agreement is complicated even further by the fact that a small number of verbs take infixed person agreement in general, but prefixed agreement in the 1pl; for example, the verb *ahi* 'to bring someone somewhere':

(77) Mismatched location of person agreement:

'bring someone somewhere'	*sg*	*pl*
1st	a-**wa-hi**	**unk-ahi-pi (*a-un-hi-pi)**
2nd	a-**ya-hi he**	a-**ya-hi-pi he**
3rd	ahi	ahi-**pi**

In addition, there is occasionally free variation in the position of agreement for a single verb (e.g., *un-nawizi-pi* ∼ *na-un-wizi-pi* 'we are jealous', 3sg *nawizi*). Finally, there are a few words that take agreement in two locations simultaneously in the 1sg, 2sg, and 2pl (but not the 1pl). These complications affect relatively few forms, however, and including them would not influence the choice of base. Therefore, I will omit them from this discussion.

What does the variable position of person agreement mean for base or UR selection? The number of verbs with prefixing or infixing person agreement are summarized in Table 5.5; as it turns out, there are significant numbers of both prefixing and infixing verbs, so this is a

[13] For simplicity, I will list second person forms with just *-he*, as they would be said by a female speakers. The male version simply substitutes *-ho* instead.

Table 5.5: Number of words in each agreement pattern

position	*count*
prefixed	347
infixed	183
infixed, 1pl prefixed	12

serious neutralization. If we were to choose a third person form as the base, we would lose all information about where subject marking should go. We could then assume that agreement is prefixing by default, but this would force us to list 195 exceptions for the verbs in which it is infixed.

Fortunately, forms other than the third person reveal the position of person agreement more clearly, to varying degrees. The 1sg form unambiguously reveals the location of agreement and would allow us to project all other forms, except in two cases. The first is when the 1sg marker happens to be identical with the beginning of the verb root, as in *wawachi* 'I dance'. In these cases, it is impossible to tell whether the subject marking is the first *wa* (***wa***-*wachi*) or the second *wa* (*wa*-***wa***-*chi*). This ambiguity, which I will call the *wawa* problem, is more pervasive than one might imagine; it affects 24 verbs in the database of 545 "basic" verbs. Furthermore, although I am unable to quantify it, the *wawa* problem probably affects many more verbs than this in practice, because *wa-* is a productive prefix used to mark indefinite objects. The other case for which the 1sg form may be misleading is for the 12 "mismatch" verbs (77). For these verbs, the 1sg form would lead one to believe the person agreement should be infixed in the 1pl, but in fact it is exceptionally prefixed in this form. The upshot is that the 1sg form is much more informative about the position of person agreement than a third person form, but it is not perfect.

In the 1pl, there is an ambiguity analogous to the *wawa* problem, which occurs when the 1pl marker *un(k)* is added to a verb that already begins with *un(k)*, such as *unk-unpa-pi* 'we smoke' (the "*unkun(k)*" problem). This problem affects only six verbs in the database, which is probably an accurate estimate, because unlike *wa-*, there are no prefixes homophonous to *unk-* in the language. However, there are two other problems with the 1pl as a potential base form. The first is the set of "mismatch" verbs discussed above; these are prefixed in the 1pl, but infixed in the remainder of the paradigm. The second problem is that the 1pl subject marker is identical for the active verbs (5.4a) and the stative verbs (5.4b). As previously discussed, this is not a serious problem in most cases, because it is usually possible to predict which class a verb belongs to based on its semantics. Nonetheless, there will still be a residue of verbs that require memorization, and listing the 1pl form would not help in these cases. This number is small, and I will leave it unquantified, since quantifying it would require a specific semantic analysis of the distinction between these two series of verbs, and a word-by-word count of which verbs fit the analysis and which do not.

Finally, let us consider the second person forms. In theory, one would expect these forms to suffer from a *yaya* problem, exactly analogous to the *wawa* and *unkun(k)* problems. However, there is a morphophonological process turning /*ya-y*/ into [l] (see fn. 12), so prefixing *ya* to a *ya*-initial root does not yield an ambiguous *yaya* sequence. Therefore, the only case in which second person forms are ambiguous with respect to the position of person agreement is for the 12 mismatch verbs, which have a different location for marking in the 1pl. This makes the

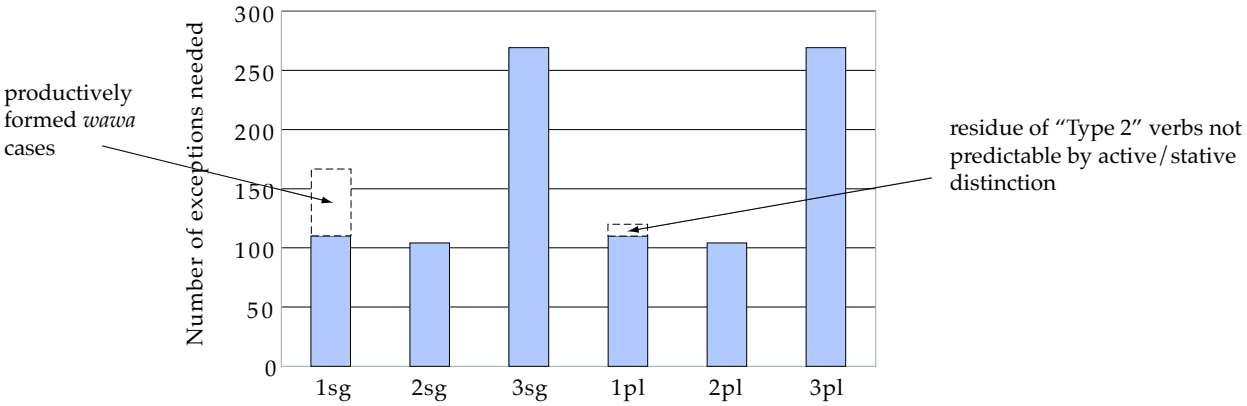

Figure 5.4: Comparing exceptions needed for each possible source of URs

second person forms the most informative, by a small margin, for purposes of predicting the location of subject marking.

Summary of unpredictability

When we compare the problems of predicting ablaut alternations and predicting the position of person agreement, we see that different forms have different advantages. The unsuffixed forms have a small advantage for maintaining ablaut contrasts (in particular, the contrast between invariant *a* and ablaut *e*∼*a*). However, this small advantage is far outweighed by the need to choose a base that reveals the position of person agreement. Unpredictable infixation favors choosing a first or second person form; moreover, the *wawa* and *unkun(k)* problems make the 1sg and 1pl forms problematic, while accidental facts about the language mean that there is no equivalent *yaya* problem affecting the second person. Therefore, this leads us to select a second person form as the all-around most informative part of the paradigm. This comparison is summarized in Figure 5.4.

Typologically, second person forms do not seem to serve as bases as often as third or first person forms; in fact, Bybee and Brewer (1980) hypothesize that second person forms might never serve as bases. However, I am aware of several other cases in which a second person form is claimed to be the base form. One is a change that occurred in the history of Eastern Scandinavian, discussed by Kuryłowicz (1947), in which a regular sound change made some verbs homophonous in the 2sg and 3sg, and this homophony was subsequently extended to all verbs in the language by replacing the 3sg forms with 2sg forms. Another type of evidence that second person forms can sometimes be bases is the fact that grammars sometimes describe the verbal inflection of a language by starting with a statement like "the root of the verb is the 2sg imperative"; this is the case in Tamil (Saravanan 2000), among others. Thus, positing that a second person form is the base in Lakhota does not seem completely anomalous.

The purpose of this section has been to show that a second person form is the most "informative" member of the Lakhota verb paradigm, and would thus be selected as the base, or UR, by a model that operates under the single surface base restriction. What remains to be shown, then, is that this makes the right prediction for the subsequent paradigmatic innovations dis-

cussed in section 5.2.

5.3.2 Consequences of using a 2nd person form as the UR

Suppose that you are a Lakhota learner, seeking the form in the paradigm with the most information about phonological and morphological properties of words. For verbs, this turns out to be the form found in the second person. We are now in a position to construct a grammar to derive the rest of the paradigm.

First, we must consider what the bases will be on this analysis. The second person forms are suffixed with the clitic *-he*, so alternating verbs have [a] (or [ã]) in this form. Therefore, the bases of alternating words will have /a/ or /ã/:

(78) Bases for Lakhota, under the single surface base restriction:

alternants	base	gloss
gleshka	ni-gleʃka-he	'be spotted'
chepe ∼ *chepa-*	ni-tʃʰepa-he	'be fat'
washte	ni-waʃte-he	'be good'
yatke ∼ *yatkã-*	latkã	'drink'
yatã	latã	'light (a cigarette)'

Note that this could also be translated into a model that seeks to discover underlying forms of verb roots by factoring out the person marking (removing *ni*, changing *l* to *y*, and removing the *he* suffix):

(79) URs under the single surface base restriction:

alternants	UR	gloss
gleshka	/gleʃka/	'be spotted'
chepe ∼ *chepa-*	/tʃʰepa/	'be fat'
washte	/waʃte/	'be good'
yatke ∼ *yatkã-*	/yatkã/	'drink'
yatã	/yatã/	'light (a cigarette)'

If the bases or URs of ablaut verbs like 'to be fat' have an underlying /a/, then we will also need a raising rule (or its OT equivalent) to derive the unsuffixed forms (/tʃʰepa/ → [tʃʰepe]), as in (74) above. This rule does not apply to /e/-final words like *washte*, and it correctly derives *chepe* from an underlying /a/. Words with invariant *a*, on the other hand, will need to be listed as exceptions to raising, to prevent incorrect unsuffixed forms like **gleshke*.

In addition to final raising, we will need two separate reduplication rules, since the difference between final and non-final reduplication (*chep-chepe* vs. *gleshka-shka*) can no longer be analyzed as a difference in their underlying forms, which both end in /a/ (/tʃʰep**a**/, /gleʃk**a**/). It appears that there are simply two competing reduplication processes: one favoring final

reduplication, and one favoring non-final reduplication.[14] Among /a/-final words, non-final reduplication is predominant, since there are more verbs like *chep-chepe* than like *gleshka-shka*; therefore, the grammar should be set up so that non-final reduplication applies by default in this environment, and forms like *gleshka-shka* must be listed as exceptions. In other environments, final reduplication prevails, and is the default rule.

What are the predicted errors if a word is not fully known? Suppose that a speaker has heard a 2sg form *ya-hoxpa he* 'you are coughing/are you coughing?', and has learned it as a base form, or has set up the UR /hoxpa/ for this verb. Since there is no evidence on the basis of this form alone that the verb is an exception to the Final Raising (ablaut) rule, the speaker will incorrectly apply raising to this verb, deriving the (etymologically) incorrect 3sg form **hoxpe*. Suppose, on the other hand, that the speaker has heard only a 3sg form of a verb, such as *kaze* 'he scoops'. In this case, the base form is not available, so the speaker simply memorizes this surface form and sets up no base or UR for the verb. (In the next section, I will discuss at greater length the idea of inferring nothing from non-basic forms.) Without a base, there is no way to derive an incorrect "undoing" of final raising, to predict incorrect suffixed **kaza-pi* (3pl) or **ya-kaza he* (2sg). There is no way to extend the [e]∼[a] alternation to invariant /e/ verbs (**washt**a**-pi*), since they have [e] in the base form, and [e] in the base always corresponds to [e] in the rest of the paradigm. Thus, there is an asymmetry: the only predicted error is on /a/-final verbs, by failing to learn that they are exceptions to final raising, and regularizing them to have final *e*∼*a* alternations. This is in fact the first innovation, shown in (71) on p. 84.

There is a similar asymmetry in the predicted reduplication errors. Suppose that a speaker has heard a ambiguous verb only in the 2sg, such as the (hypothetical) 2sg form *ya-t'apha he*. In this case, she would set up a base, or infer a UR /t'ap^{h}a/. The default reduplication pattern for /a/-final verbs is non-final reduplication, so in the absence of evidence that this verb takes final reduplication, she will apply the default (*t'ap-t'apha*). Suppose, on the other hand, that the speaker has heard an ambiguous verb only in the 3sg, such as the hypothetical 3sg form *sophe*. In this case, no base form has been learned, meaning there is no way to derive any reduplicated form (*sop-sophe* or *sophe-phe*). There is no way to apply incorrect final reduplication to ablaut verbs, since verbs with /a/ take penultimate reduplication, and ablaut verbs have /a/ in the base form. Furthermore, there is no way to derive incorrect non-final reduplication for invariant /e/ verbs, because they have /e/ in the base form, and final reduplication is the default for verbs that do not end in /a/. As a result, the only predicted reduplication error is for invariant /a/-final verbs, by failing to learn that they are exceptions to non-final reduplication, and incorrectly regularizing them to have non-final reduplication. This is the second innovation, shown in (72) on p. 84.

We see, then, that restricting bases or URs to a single surface form predicts only two types of errors, and both are attested in the new paradigm types in section 5.2. Furthermore, this ap-

[14] Nelson (to appear) points out that word-medial reduplication is a problem for OT because it does not satisfy either ANCHOR-L or ANCHOR-R. She goes on to argue that non-final reduplication patterns in cases like Lakhota are actually to be analyzed as stressed-syllable reduplication. This works for a majority of the Lakhota data, since most verb roots are di- or tri-syllabic, and stress is generally peninitial unless the second vowel is an ablaut vowel, meaning that the non-final syllable is usually (but not always) the stressed one for ablaut verbs. However, this analysis does not work completely; there are a number of verbs with nonfinal stress but final reduplication—e.g., [ˈhãske] ∼ [ˈhãska-ska] 'be tall', [wikˈdʒɛmna] ∼ [wikˈdʒɛmna-mna] 'be ten in number', etc. I do not have an alternative OT analysis of non-final reduplication at this time, but trust that it could be formulated somehow, perhaps using Nelson's insights about stress, or perhaps in some other fashion.

proach also explains the "de-coupling" of final vowel alternations from non-final reduplication, which were once predictably linked. In particular, final vowel raising and non-final reduplication are treated as the result of separate rules, rather than being derived from a common fact about underlying representations (underspecified final vowels). Since these are separate rules, they may each have their own lists of exceptions, and the fact that words like *gleshka* are an exception to both is purely an accident from the point of view of this analysis. If learners have evidence that a word is exceptional with respect to only one process, they may still regularize it with respect to the other. The result is "inconsistent paradigms", such as those that have arisen in Lakhota.

5.4 Inferring nothing from non-basic forms

A strong and perhaps uncomfortable assumption that was needed in the previous section was that if a speaker happens to have heard only non-basic forms of a word, she will memorize them as surface forms, but she will not infer a base form that can be used to derive other forms. As a consequence, there may be times when a speaker in some sense knows the word, but is unable to produce new forms of it. This assumption is potentially quite controversial—is there any way around it?

Consider a weaker version of the current theory, in which learners establish a base form by comparing the effectiveness of different forms in projecting the paradigm, but in which they retain the subgrammars needed to do mappings in *all* directions. Under this theory, learners prefer to derive forms using a base form as the input, since it is more reliable, but in the absence of such a form, they are able in a pinch to use a non-basic form as the input. This theory has some intuitive appeal, but it makes incorrect predictions about possible errors. In particular, it predicts that if Lakhota speakers happened to know only a 1sg or 3sg form, as must occur not infrequently, and that form ended in an *-e*, they would be able to reason backwards to infer that the suffixed form should end in *-a*, predicting errors like **washta-pi* instead of *washte-pi*. Similarly, in the case of Yiddish, if a speaker had heard an umlaut verb in only the 2sg or 3sg, she would have been able to project backwards to a 1sg with **e*, producing unattested errors like 1sg **fer* instead of *fɔr*.

Certainly, it would be difficult to argue that there is no such thing as backformation. However, asymmetries like these may show that it is not part of the ordinary, automatic workings of the synchronic morphological system.[15] The assumption that speakers infer nothing from

[15]Kiparsky (1982, pp. 21-22) makes the same claim for derivational morphology, following Marchand (1969). Given the fact that back-formations like *air-condition* do arise, Kiparsky and Marchand are forced to admit that back-formation exists, but only as a diachronic process. Kiparsky claims that synchronically, *air-condition* is the product of a N+V compounding process, which arose through reanalysis of N+N compounds ([*air* + [*condition*+*er*]]) as N+V+*er* compounds ([[*air* + *condition*] + *-er*]). This analysis is not totally satisfying, however, without a theory of possibly reanalyses; what allowed speakers (or learners) to reanalyze this form based on an unattested constituent? Crucially, whatever mechanism allows this reanalysis must *not* allow the reanalysis of [waʃte] as [waʃta] with final raising.

non-basic forms is needed here in order to explain the data.[16]

5.5 Other examples of inconsistent paradigms

Some readers may be wondering to what extent the changes discussed here are a result of the fact that Lakhota is an endangered language. The implied hope is that perhaps inconsistent paradigms arise only when the learning data is reduced or imperfect. Certainly, languages that are endangered experience far more radical and rapid changes than languages in which learners have access to a large sample of fluent monolingual speakers (see, e.g., Richards 2001 for a discussion of this in Lardil). Nevertheless, I believe that such factors merely facilitated the later stages of the Lakhota change, and that inconsistent paradigms can arise even in more stable environments.

For one thing, it appears that the changes discussed in this chapter probably began well before Lakhota was endangered. For example, the inconsistent paradigm of the verb *yatkã* 'drink' (*yatke*, *yatkã-pi*, *yatkã-tkã*) could have arisen only as an analogical extension of ablaut,[17] but it occurs in several related dialects that diverged before Lakhota was an endangered language (Rood 1983).

Furthermore, there seem to be examples of mixed behavior words in languages spoken more widely in monolingual environments. Tranel (1996) discusses one such case in French, in which a handful of indeclinable words behave like feminine forms in isolation, with their final consonants pronounced, but like masculine forms before a consonant-initial word, with the consonant deleted, as in (80). (See also L'Huiller 1999, p. 597.)

(80) Mixed behavior in French *huit* 'eight'

context	*petit* 'small' (masc.)	*huit* 'eight'	*petite* 'small' (fem.)
/ ___#V	*peti*[t]	*hui*[t]	*peti*[t]
/ ___#	*peti*[∅]	*hui*[t]	*peti*[t]
/ ___#C	*peti*[∅]	*hui*[∅]	*peti*[t]

Kenstowicz and Kisseberth (1977, p. 121) discuss a similar example from Chi-Mwi:ni, in which one exceptional verb behaves in some forms like it ends in a final /g/, and in others, like a final /k/. In chapter 6, I will discuss another possible example from Korean, which it appears that some nouns are more likely to appear with [t^h] in some forms, and with [s] in others.

I do not have an analysis of how such inconsistencies arose in French or Chi-Mwi:ni, nor do I have an estimate of how common such mixed-behavior or inconsistent words are in the world's languages. For present purposes, however, it suffices to note that the Lakhota case is not completely isolated, not does it appear to be a result of its current endangered status.

[16] A possible modification that would still explain the data would be to assume that whenever a speaker learns a new word in a non-basic form and does not know the base, she works her way backwards through the grammar to generate a set of possible base forms that could have yielded that derived form. If the set of possible bases has just one member, she infers it, otherwise she waits. Such a theory would allow speakers to set up underlying or base forms more rapidly, but strikes me as a rather perplexing strategy: why are speakers generally willing to guess about derived forms in the face of potential ambiguity, but not about base forms?

[17] The fact that the final syllable reduplicates and also the fact that it is nasalized indicate that it is etymologically an "underlying" vowel; if it had always been epenthetic, it would not be nasalized.

5.6 Local summary

In this chapter, I have shown that Lakhota presents an example of a three-way contrast ([e], [a], [e]∼[a]) that can be neatly described using archiphonemes or underspecification. This analysis is also supported by other facts in the language, since it can explain the co-occurrence of final vowel alternations and non-final reduplication. However, under this analysis, learners should always posit URs that produce "valid" paradigms, with the same set of properties as existing paradigms. This prediction is disproved by subsequent historical changes in Lakhota, which have resulted in the creation of two new paradigm types, inconsistent with any UR in the old system. These changes are puzzling not only because they have created novel paradigm types, but also because they have been asymmetrical: they have affected only verbs originally ending in invariant /a/. In section 5.3, I showed that by restricting learners to choosing a UR that matches a single surface form, and using the strategy of selecting the most informative part of the paradigm as the UR, we predict exactly these two errors and no others.

A consequence of this restriction is that learners are unable to capture certain generalizations about their language, such as the fact that ablaut and final reduplication are predictably linked with one another, since this cannot be deduced on the basis of any single form in the paradigm. It does allow them to capture other generalizations that the underspecification analysis does not allow, however, such as the fact that verbs that end in *-a* in suffixed forms tend to have *-e* in unsuffixed forms, and also tend to have penultimate reduplication. The historical evidence shows that these are in fact the generalizations that have been extended over time.

The Lakhota example complements the cases discussed in the previous chapters. In both Yiddish and Latin, alternations were leveled on the basis of other forms within the paradigm. In Lakhota, on the other hand, an alternation was extended, introducing new *e*∼*a* alternations into paradigms that did not originally have them. The difference between these cases, I have argued, is simply a difference in which pattern was predominant in the lexicon prior to the change. When the majority of words do not alternate, the best grammar to describe the language will not include a productive rule deriving the alternation, and exceptional alternating forms are open to replacement by overregularization to nonalternating forms. By contrast, when the majority of words alternate, it is more efficient to set up a rule producing alternations by default; in this case, we expect regularization to extend the alternation, rather than eliminating it.

This is a simple intuition, but as discussed at the end of the previous chapter, it is one that is not captured by the analysis of paradigm leveling as a universal preference for nonalternation. It is similar to the idea behind Harrison and Kaun's Pattern-Responsive Lexicon Optimization (discussed on p. 83), in that it allows speakers to assume that a new word has an alternation even if they have not actually heard it. Under Harrison and Kaun's proposal, however, there is still no reason to expect an asymmetry in the case of Lakhota. Alternating *-e*∼*-a* verbs outnumber both invariant *-e* and invariant *-a* verbs. Therefore, even if we assume that speakers may set up underspecified URs in response to the dominant patterns of the language, why would they do this only for *-a* verbs and not for *-e* verbs? What is missing is a theory of how "patterns" are defined; in the present case, the pattern is not merely a paradigm type with *-e* in some forms and *-a* in others, but rather a directional implication stating that if a word has *-a* in some forms, it will have *-e* in others. The single surface base restriction provides us with a theory of which patterns will be available to the speaker, namely, those involved in the mapping from the base

form to the remainder of the paradigm.

If this analysis of Lakhota is correct, then the requirement that URs must obey the single surface base restriction has widespread implications for phonological analysis. In particular, it calls into question an assumption that dates back at least to Bloomfield and Trubetzkoy, that speakers may respond to patterns of alternation by setting up lexical representations with abstract phonemes that are unlike any surface realization. I will consider some of issues raised by this proposal in the next chapter.