Chapter 1

Introduction

The problem of bases in inflectional paradigms is neither new nor forgotten in linguistic theory. Consider, for example, the full set of forms which a noun could take in Sanskrit (Whitney 1924), a language with a moderately large set of distinct case forms:[1]

(1) Full paradigm for Skt. *pā́d* 'foot'

	Sg.	Dual	Pl.
Nom.	*pā́t*	*pā́dāu*	*pā́das*
Acc.	*pā́dam*	*pā́dāu*	*padás*
Instr.	*padā́*	*padbhyā́m*	*padbhís*
Dat.	*padé*	*padbhyā́m*	*padbhyás*
Abl.	*padás*	*padbhyā́m*	*padbhyás*
Gen.	*padás*	*padós*	*padā́m*
Loc.	*padí*	*padós*	*patsú*
Voc.	*pā́t*	*pā́dāu*	*pā́das*

It is commonly observed, starting at least as far back as Paul (1920, chap. 5), that for languages like Sanskrit, it would be impractical to memorize every form of every word, since there are so many forms and so many words. Even more dramatic examples include Hungarian, which has 924 possible forms for each noun (Tihany 1996), Archi, which is claimed to have up to 1,502,839 forms for each verb, if deverbal and commentative forms are included (Kibrik 1998), and Shona, which may have up to 16,000,000,000,000 verbal forms (Odden 1981).

Fortunately, it is intuitively clear that individual inflected forms are not unrelated, isolated words; rather, if we compare them (and consider a variety of other nouns), we can see that they stand in definite relations to one another. For example, we might observe that for this particular Sanskrit noun, the genitive and locative duals are identical, and the locative plural can be obtained by concatenating the nominative singular with a suffix *-su*, shortening the [a:], and shifting the accent to the second syllable; or we might even observe that all of the forms share a common root (*pā́d-*), with a certain set of suffixes used to mark the cases. Statements of this sort, which tell the speaker how to create forms based on other forms, can greatly reduce the

[1] I will use the following abbreviations here: *sg.* = singular, *pl.* = plural, *nom.* = nominative, *acc.* = accusative, *instr.* = instrumental, *dat.* = dative, *abl.* = ablative, *gen.* = genitive, *loc.* = locative, *voc.* = vocative; for features, *cons.* = consonantal, *cont.* = continuant, *cor.* = coronal, *lab.* = labial, *dors.* = dorsal

amount of information that must be learned. Rather than memorizing the entire paradigm of every word, the speaker must simply memorize a single form for each word—either a privileged surface form, or an abstract underlying form. In addition, the speaker must learn a grammar of morphological and phonological rules that derive the remainder of the paradigm from the memorized base form. The base therefore serves as the *input* to the grammar; it is what the grammar operates on to produce the remaining forms of the paradigm. With a grammar in place, the speaker is ideally able to produce any form of any word without having memorized it—or, at least, of most words, with a small residue of exceptions that must be memorized as such. The goal of this thesis is to explore how bases can be identified algorithmically, and to look for evidence that human language learners employ a similar approach.

Bases are more than just a computational convenience. Memorization is not only impractical in highly inflected languages, but it is also inadequate: speakers need to be able to produce and comprehend forms that they have never encountered before, and have thus had no chance to memorize. There are many sources of evidence showing that speakers are able to construct some forms in the paradigm based on information from other forms. First, there is the simple and easily observed fact that speakers can utter — often without hesitation — forms of words that they almost certainly have never heard before. Spanish speakers, for example, have no trouble producing the first person singular indicative forms of 'to decaffeinate' (*descafeíno*) or 'to Italianize' (*italianizo*), even though few, if any, have ever encountered or had occasion to use these forms before. Thus, even if we were to accept the idea that speakers memorize every inflected form that they have encountered (the *full listing hypothesis*, Butterworth 1983), we would still need some explanation for how speakers can produce novel forms based on other, previously listed forms.

In addition, there is abundant historical evidence that speakers construct relations between parts of the paradigm, since forms are often rebuilt on the basis of other forms within the paradigm. Consider, for example, the change in Yiddish verb paradigms shown in (2), in which the 2nd and 3rd singular forms of 'to dig' and the plural forms of 'to know' have been rebuilt to match the 1sg form. (The notation $*A \Rightarrow B$ indicates that the expected form A has been replaced by the analogical form B.)

(2) Leveling in Yiddish verbal paradigms

a. *grɔbn* 'to dig'

	sg.	pl.
1st	*grɔb*	*grɔbn*
2nd	**grebst* ⇒ *grɔbst*	*grɔbt*
3rd	**grebt* ⇒ *grɔbt*	*grɔbn*

b. *visn* 'to know'

	sg.	pl.
1st	*veys*	**visn* ⇒ *veysn*
2nd	*veyst*	**vist* ⇒ *veyst*
3rd	*veys(t)*	**visn* ⇒ *veysn*

Changes of this sort, in which alternations are eliminated by replacing some members of the paradigm, are known as *paradigm leveling* or *analogical leveling*. Such changes poses a well-known problem in historical linguistics, and there are two basic approaches to explaining

them. On the one hand, they occur frequently, and are natural and unsurprising. For this reason, it is often assumed that the drive towards nonalternating paradigms is simply a primitive of language, sometimes referred to as "Humboldt's Universal" (one form for one meaning). The analysis that this implies is that paradigm leveling is an output-output (OO) effect between related surface forms. Even in a language with paradigmatic alternations (such as *grɔb* ∼ *grebst*), there is some force that compels speakers to consider the possibility of uttering non-alternating forms (e.g., *grɔbst*), and furthermore, when speakers are faced with this possibility, they find the innovative forms appealing because of their resemblance to other forms within the paradigm. The idea that paradigms like to be uniform was never formalized in rule-based generative phonology (though on the need for it, see Kenstowicz and Kisseberth 1977, p. 74, as well as Hock 1991, p. 260). It has been formalized in recent years in Optimality Theory (OT) (Prince and Smolensky 1993), however, as UNIFORM EXPONENCE, LEVEL, or PARADIGM UNIFORMITY (Kenstowicz 1997b; Kenstowicz 1997a; Steriade 2000; Kager 2000; Raffelsiefen 2000; Kenstowicz 2002). Under these approaches, the force that suggests non-alternating forms is GEN, and the force that prefers them is the set of Uniform Exponence or Paradigm Uniformity constraints. Thus, OO constraints provide us with a formalism to describe the way that surface forms might influence one another to favor paradigm leveling.

A blanket preference for nonalternating paradigms can only go so far in explaining paradigm levelings, however. It is often noted that a paradigm uniformity preference can tell us that an alternation is likely to be leveled, but it cannot necessarily tell us when, or in which direction. For example, why was the desire for uniform paradigms stronger in Yiddish than in other German dialects, most of which have retained alternating paradigms (cf. Modern German *grabe, gräbst, gräbt*)? Why was it the 1sg form that was extended, and not some other form, such as the 3sg, yielding paradigms like **greb*, **grebst*, **grebt*? In this and other cases, the challenge is to explain why the change went in this direction, and not in other, logically possible directions, such as in (3):

(3) Other, logically possible (but unattested) changes in Yiddish

a. *grɔbn* 'to dig'

	sg.	pl.
1st	*grɔb* ⇏ **greb*	*grɔben* ⇏ **greben*
2nd	*grebst*	*grɔbt* ⇏ **grebt*
3rd	*grebt*	*grɔben* ⇏ **greben*

b. *visn* 'to know'

	sg.	pl.
1st	*veys* ⇏ **vis*	*visn*
2nd	*veyst* ⇏ **vist*	*vist*
3rd	*veys(t)* ⇏ **vis(t)*	*visn*

Many proposals over the years have attempted to explain the direction of analogical change. The usual approach, pioneered by Kuryłowicz (1947) and Mańczak (1958) and continued by Bybee (1985) and others, has been to focus on tendencies, or groups of factors that may compete in making one form the *base*, or *pivot* of a change. Some of the factors that seem to play a role include: (1) presence or absence of suffixes, such that leveling is often to forms with no affixes or with shorter affixes (Mańczak 1958; Hayes 1995, Bybee 1985, pp. 50-52), (2) token frequency,

with leveling often extending the members of the paradigm with the highest token frequency (Mańczak 1980, pp. 284-285), and (3) some sense of morphosyntactic markedness, with leveling often extending the "unmarked" member of the paradigm (Jakobson 1939; Greenberg 1966; Bybee and Brewer, 1980; Tiersma 1982).[2]

Often these factors all coincide to favor a single form, so in language after language, it is the 3sg present form of verbs, or the nominative singular of nouns, that is extended. In Polish, for example, the paradigms of many diminutives have been rebuilt on the basis on the nominative singular. In the language in general, there is a regular [ɔ] ∼ [u] alternation before underlyingly voiced obstruents, with /ɔ/ occurring as [ɔ] in open syllables (especially those followed by a lax vowel), and [u] in closed syllables (Gussmann 1980, chap. 4; Kraska-Szlenk 1995, pp. 108-114; Kenstowicz 1997), shown in (4) for the words *dół* [duw] 'ditch' and *krowa* [krova] 'cow':[3]

(4) [ɔ] ∼ [u] alternations in Polish nouns

a. [duw] 'ditch' (masc.)

	sg.	*pl.*
nom.	[d**u**w]	[dɔwɨ]
gen.	[dɔwu]	[dɔwuf]
dat.	[dɔwovi]	[dɔwom]
acc.	[d**u**w]	[dɔwɨ]
instr.	[dɔwem]	[dɔwami]
loc.	[dɔle]	[dɔwax]

b. [krɔva] 'cow' (fem.)

	sg.	*pl.*
nom.	[krɔva]	[krɔvɨ]
gen.	[krɔvɨ]	[kr**u**f]
dat.	[krɔvje]	[krɔvom]
acc.	[krɔvẽ]	[krɔvɨ]
instr.	[krɔvã]	[krɔvami]
loc.	[krɔvje]	[krɔvax]

The Polish diminutive suffix *-(e)k* (masc.)/*-(e)ka* (fem.) contains an initial deletable vowel (a so-called "yer") that disappears when the following syllable contains a full vowel. The result is that sometimes this suffix is vowel-inital, and sometimes it is consonant-initial. We would expect, therefore, that it should condition [ɔ] ∼ [u] alternations in the final syllable of the noun, just as in (4); the expected forms for *dołek* [dɔwek] 'little ditch' and *krówka* [krufka] 'little cow' are shown in (5):

[2]The definition of "unmarked" is often problematic, though various authors have attempted to find a non-circular basis for deciding that one part of the paradigm is "less marked" than another.

[3]This alternation was originally a lengthening of [ɔ] to [ɔː] conditioned by a following coda voiced obstruent; this lengthening was subsequently made opaque by raising/tensing of [ɔː] to [u], and the conditioning environment of open and closed syllables has been made opaque in some cases by the loss of a vowel (the "yer") in some forms of some suffixes, including the diminutive suffix *-ek* (Gussman 1980, p. 30). Furthermore, loans and paradigm leveling have introduced many exceptions to the raising alternation seen in (4), and it has been argued that this alternation is no longer synchronically productive (Buckley 2001, Sanders 2001).

(5) Expected paradigms of Polish diminutives

a. [dɔwek] 'little ditch' (masc.)

	sg.	*pl.*
nom.	[dɔwek]	[d**u**wki]
gen.	[d**u**wka]	[d**u**wkuf]
dat.	[d**u**wkovi]	[d**u**wkom]
acc.	[dɔwek]	[d**u**wki]
instr.	[d**u**wkjem]	[d**u**wkami]
loc.	[d**u**wku]	[d**u**wkax]

b. [krufka] 'little cow' (fem.)

	sg.	*pl.*
nom.	[kr**u**fka]	[kr**u**fki]
gen.	[kr**u**fki]	[krɔvek]
dat.	[kr**u**ftse]	[kr**u**fkom]
acc.	[kr**u**fkẽ]	[kr**u**fki]
instr.	[kr**u**fkã]	[kr**u**fkami]
loc.	[kr**u**ftse]	[kr**u**fkax]

In fact, this alternation has been leveled in many diminutives; in all cases, it is the form that is expected in the nominative that has been extended to the remainder of the paradigm (6):

(6) Actual paradigms of Polish diminutives

a. [dɔwek] 'little ditch' (masc.)

	sg.	*pl.*
nom.	[dɔwek]	[dɔwki]
gen.	[dɔwka]	[dɔwkuf]
dat.	[dɔwkovi]	[dɔwkom]
acc.	[dɔwek]	[dɔwki]
instr.	[dɔwkjem]	[dɔwkami]
loc.	[dɔwku]	[dɔwkax]

b. [krufka] 'little cow' (fem.)

	sg.	*pl.*
nom.	[kr**u**fka]	[kr**u**fki]
gen.	[kr**u**fki]	[kr**u**vek]
dat.	[kr**u**ftse]	[kr**u**fkom]
acc.	[kr**u**fkẽ]	[kr**u**fki]
instr.	[kr**u**fkã]	[kr**u**fkami]
loc.	[kr**u**ftse]	[kr**u**fkax]

The fact that we get an [ɔ] in [dɔwki] but an [u] in [kr**u**fki] is not conditioned by any phonological difference in the suffixes of these forms; rather, the oblique and plural forms seem to be influenced by the (phonologically expected) difference in the nominative singular forms. In Polish, as in many other languages, the nominative singular is also the unsuffixed (or, at least, the least suffixed) form, the morphosyntactically unmarked form, and possibly also the most frequent form, so its privileged status in the paradigm is unsurprising.

However, it is not always true that paradigm leveling extends the nominative singular. A famous counterexample, discussed by Hock (1991, pp. 179-180), Kenstowicz (1997b), and many others, is a change that occurred in the history of Latin. In pre-classical Latin, rhotacism of /s/ to [r] intervocalically created [s] ∼ [r] alternations within noun paradigms, as in (7a). This alternation was leveled out in the late pre-classical period, extending the [r] of the oblique and plural forms to the nominative singular.

(7) Change of Latin *honōs* to *honor*

a. Pre-leveling

	sg.	*pl.*
nom.	[hono:**s**]	[hono:**r**e:s]
gen.	[hono:**r**is]	[hono:**r**um]
dat.	[hono:**r**i:]	[hono:**r**ibus]
acc.	[hono:**r**em]	[hono:**r**e:s]
abl.	[hono:**r**e]	[hono:**r**ibus]

b. Post-leveling

	sg.	*pl.*
nom.	[hono**r**]	[hono:**r**e:s]
gen.	[hono:**r**is]	[hono:**r**um]
dat.	[hono:**r**i:]	[hono:**r**ibus]
acc.	[hono:**r**em]	[hono:**r**e:s]
abl.	[hono:**r**e]	[hono:**r**ibus]

The comparison of Yiddish, Polish and Latin in Table 1.1 illustrates the basic conundrum. In Yiddish and Polish, it was the unsuffixed member of the paradigm (the 1sg and nominative singular, respectively) that was extended, while in Latin, it was a suffixed form. In Polish, it was the universally unmarked member of the paradigm (the nominative singular) that was extended, while in Yiddish and Latin, marked forms (the 1sg and a non-nominative form) were extended. In Latin, it was perhaps the member of the paradigm with the highest token frequency that was extended (see section 4.4.2), but in Yiddish, the form that was extended was most likely not the most frequent member of the paradigm, or even the most frequent alternant.[4] In Yiddish and Latin, the form that was extended was the form that occurred in the majority of slots in the paradigm, but in Polish masculines, it was the minority form. The conclusion that is generally drawn from such facts is that no single factor guarantees that a particular form will be extended in paradigm leveling.

This is a problem for theories that try to explain basehood using "static" factors like frequency or markedness. Every language has differences in the frequency of forms, differences in the "degree of suffixation" of forms, differences in markedness, and so on, but it appears that speakers weight these factors differently in deciding which form should get extended in leveling. Thus, proposals that use such factors as an explanation allow us to derive typological predictions, but not to make predictions about a given language at a given time, because we do

[4] I do not have token frequency counts for the various members of the Yiddish verb paradigm. However, Bybee (1985, p. 71) gives some equivalent counts for Spanish, showing that the 3sg is almost twice as frequent as the 1sg. Furthermore, between 65% and 67% of the tokens are either a 2sg or 3sg, meaning that the umlaut alternant would have been the most frequent alternant, but it was not extended.

Table 1.1: Comparison of factors encouraging the extension of a form

Language	Form	Unsuffixed	Unmarked	Highest Freq. Form	Highest Freq. Alternant	Majority Form
Yiddish	1sg	*yes*	*no*	*no*	*no*	*yes*
Polish	nom.sg.	*yes*	*yes*	??	*yes*	*no*
Latin	oblique	*no*	*no*	*no* (?)	*yes*	*yes*

not know which factors will win in that particular case.[5] As Bybee and Brewer (1980, p. 215) state:

> A hypothesis formulated in such a way makes predictions of statistical tendencies in diachronic change, language acquisition and psycholinguistic experimentation. It cannot, nor is it intended to, generate a unique grammar for a body of linguistic data.

In this thesis, I will take a different approach, in the spirit of Paul (1920) and Kiparsky (1965), that focuses on the role of language learners in historical change. I will treat paradigm leveling not as an output-output effect, but rather as an effect of the way that speakers use their grammar to project unknown forms — that is, as an input-output (IO) effect. In particular, I will pursue the hypothesis that learners impose structure on paradigms, as part of an effort to construct phonological and morphological grammars that generate unknown forms as accurately or as confidently as possible. The way that they do this, I will claim, is by seeking a base form within the paradigm that is "maximally informative" — that is, that suffers the least serious phonological and morphological neutralizations — and then deriving the remaining forms in the paradigm from the base form. Under this approach, we can use the direction of the grammar (base form → derived forms) to predict the direction of possible analogical change.

Before beginning with the task of identifying bases in paradigms, it is useful to recognize from the outset that different models of morphology operate on radically different types of inputs. In some, termed "Item and Process" (IP) models by Hockett (1958), morphological rules create words from other words — for example, "add *-su* to the nominative to create the locative plural." In such a model, morphological and phonological rules are assumed to operate on a free-standing surface form from somewhere within the paradigm; examples of such models include Aronoff's word-based model (Aronoff 1976), Anderson's Extended Word and Paradigm (EWP) model (Anderson 1992), Bochner's Lexical Relatedness model (Bochner 1993), and Ford and Singh's whole word morphology (Ford and Singh, 1996; Neuvel, to appear; Neuvel and Fulop, to appear). For models that operate on words, a base selection procedure must be able to choose a surface form that will serve as the base. For example, in the case of Sanskrit, we might choose the accusative, and then formulate a set of rules that changes the suffix, moves the accent, and performs various phonological adjustments to derive the remainder of the paradigm.

[5] In fact, this lack of predictiveness is considered appropriate by many, since leveling is classified as an analogical change, and as such is held not to be rule-governed. I will take the opposite approach here, of trying to pursue a hypothesis that makes strong and falsifiable predictions about possible changes (Gvozdanović 1985).

In other models, termed "Item and Arrangement" (IA) by Hockett, morphological rules combine sub-parts of words (stems and affixes) to create surface forms — for example, "the locative plural morpheme is *-su*, and it occurs after the nominal root." In this type of model, the grammar is assumed to operate on a set of underlying forms, which are combined and readjusted to yield surface forms. This is in fact the type of model usually assumed by phonologists. The underlying forms of morphologically complex words are typically represented in phonological analyses as something like /*root*+*affixes*/, where the /*root*/ and /*affix*/ are independent entities. Phonologists do not, on the whole, devote much attention to the question of how the morphemes came to be in that particular configuration, except to suppose that there is a separate morphological module that takes them out of the lexicon and arranges them somehow. A few recent formalizations of how this is actually done using an IA approach include Distributed Morphology (Halle and Marantz 1994), DATR (Evans and Gazdar 1996), and Lieber's syntactic approach to morphology (Lieber 1992). For models that operate on (possibly bound) stems, a base selection procedure must be able to discover the underlying forms of word roots, and whatever morphological information is necessary to determine which affixes they should combine with. For Sanskrit, this would involve learning that the underlying form for 'foot' is something like /*pā́d*/ (in spite of surface variations), and that it is in the class of nouns that take *-am* in the accusative, *-as* in the plural, etc.

On the face of it, it would seem that the different requirements of word-based and stem-based approaches would demand fundamentally different base identification procedures: for an IP approach, we need to isolate a surface form in the paradigm that will act as the base, while for an IA approach, we need to be able to compare the surface forms to arrive at a (possibly abstract) underlying form for the stem of each word. What I will argue in the course of this thesis, however, is that although we could imagine very different strategies for selecting whole-word bases vs. stems, they are both compatible with a range of possible strategies, and furthermore, in many cases the empirical evidence drives us to parallel conclusions for both.

1.1 The problem of bases in word-based morphology

As noted above, selecting a base has the potential to greatly reduce the amount of information that must be learned; rather than having to memorize 30 forms for each word, in most cases learners can simply memorize one form, and use a set of rules to derive the rest of the paradigm. Unfortunately, the rules listed above were only two of the many possible rules that could be formulated to relate parts of the Sanskrit noun paradigm. In addition to relating the genitive dual to the locative dual, and the nominative singular to the locative plural, we could equally well have observed that the instrumental singular is formed by taking the locative dual and replacing the *-os* suffix with *-ā*, the dative dual is equal to the ablative plural minus the final *-s* and adding a final *-am*, and that the nominative singular is the accusative plural minus the *-as* suffix, possibly shifting the accent to the first syllable, and changing the [c] to a [k]. If we were to say simply that a grammar contains rules relating forms in the paradigm to one another and leave it at that (as, for example, Bochner 1993, Barr 1994, and Neuvel and Singh 2001 do), we would be left with an enormous number of pairwise relations to include in the grammar. Intuitively, the answer to this problem is that speakers probably do *not* learn rules for every pairwise relation in the paradigm, but rather for only a subset of the possible relations. The

question, then, is which relations are part of the grammar, and which relations are not. Stated differently, what are the possible structures of the paradigm?

A common way to restrict the set of possible paradigm structures in word-based morphology is to place restrictions on bases. Take, for example, the four-member paradigm *walk*, *walks*, *walked*, and *walking*. In a completely unrestricted system, we could say that *walked* is derived from *walk* by the rule $[\,\mathrm{X}\,]_{\mathrm{pres.}} \rightarrow [\,\mathrm{X}\ \textit{ed}\,]_{\mathrm{past}}$, from *walks* by the rule $[\,\mathrm{X}\ \textit{s}\,]_{\mathrm{3sg.pres.}} \rightarrow [\,\mathrm{X}\ \textit{ed}\,]_{\mathrm{past}}$, or from *walking* by the rule $[\,\mathrm{X}\ \textit{ing}\,]_{\mathrm{pres.participle}} \rightarrow [\,\mathrm{X}\ \textit{ed}\,]_{\mathrm{past}}$. Suppose, however, that we restrict grammars such that they may only operate on certain forms within the paradigm. These privileged forms, which serve as the input to the grammar, we will call *bases*. One possible restriction would be to say that each slot in the paradigm must be derived from at most one unique base, but different slots may be derived using different bases . Now there can only be one way to derive a given form, and we would have to pick just one of the rules deriving *walked*; intuitively, we would probably want to derive *walked* from *walk* by the rule $[\,\mathrm{X}\,]_{\mathrm{pres.}} \rightarrow [\,\mathrm{X}\ \textit{ed}\,]_{\mathrm{past}}$.

Requiring that each slot in the paradigm have at most one base greatly reduces the number of pairwise relations that are included in the grammar. If we assume a particular version of this restriction in which one form in the paradigm is an underived base form and all other forms are derived by a single relation, then we reduce the number of pairwise relations in the grammar to $n - 1$. The resulting set of possible paradigm structures is much more constrained, but still perhaps larger than we might like. When we consider a hypothetical paradigm with four members (*A*, *B*, *C*, *D*), even if we assume that A is always a base and that there is only one way to derive each form, there are still sixteen possible paradigm structures, as in (8). In particular, there is one structure using A to derive all three of the other forms (A_3), there are six structures using it to derive just two of the other forms (A_2), and there are nine structures using it to derive just one of the other forms (A_1).

(8) Possible paradigms using just one mapping per form, *A* as base

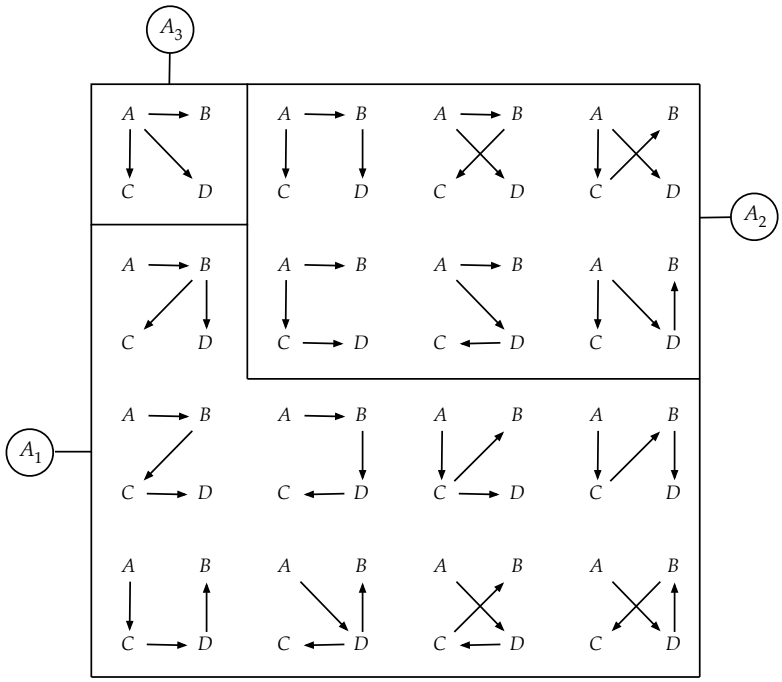

The paradigm structures in (8) are more manageable than a completely unrestricted system would be, and in chapter 6 I will present some evidence that structures like these (in particular, the A_2 structures) may sometimes be needed. However, for now we may note that this restriction still leaves us with quite a large number of hypotheses to explore. In practice, many analysts seem to assume an even more restrictive hypothesis, which is that, in the usual case, the entire inflectional paradigm is derived from a single base form. If we adopt this assumption, then for a paradigm of n forms, we need to consider only n possible candidates for base status. Furthermore, once we have selected a base form, we will have only one possible paradigm structure, consisting of statements about only $n-1$ pairwise relations. For example, in the case of Sanskrit, if we choose the nominative singular as the base, then we are left with a maximum of 23 statements about how to relate the other forms to the nominative. Furthermore, a learner operating under this restriction must consider only n possible paradigm structures of n –1 relations each. This is shown schematically in (9).

(9) The single base hypothesis: only n possible paradigms

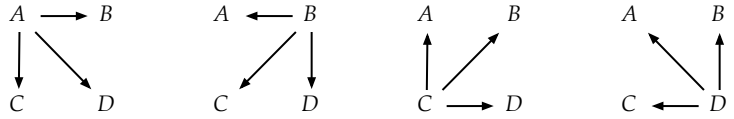

The single base hypothesis is appealing not only from a learning point of view, but also because of its restrictiveness as a linguistic analysis. Once we have limited the possible paradigm structures to those in (9), all that remains is to construct a grammar that derives each of the non-basic forms in the paradigm.

For purposes of concreteness, I will make the following assumptions about the mechanisms by which speakers can produce forms: first, forms may be produced by retrieving them from the lexicon ready-made, already bearing whatever features are required for the syntactic context. I will refer to this mode of production as retrieval of listed forms, or resorting to memorized word-specific knowledge. For obvious reasons, this option is available only if the speaker has memorized the relevant form ahead of time. Second, forms may be produced by synthesizing them with the rules of the grammar. This option is available only if the speaker has one or more rules to derive the desired part of the paradigm, and if the speaker knows the base form (the input) for the word in question. For example, if you need form B of word w (w_B), you need two things: a rule to derive form B from another form (e.g., $A \rightarrow B$), and you need to know word w in form A, so you can apply the rule to derive $w_A \rightarrow w_B$. I will refer to this mode of production as synthesis, or derivation by the grammar.

In some linguistic theories, limitations are placed on what forms may or may not be stored in the lexicon, requiring a theory of how speakers decide which forms to memorize and which to exclude. In this thesis, I will make what I take to be the rather simpler assumption that speakers are potentially able to memorize any form that they have had sufficient exposure to, whether it is a base form or an inflected form. Indeed, there are several situations in which this capability is necessary. First, when learners are just starting out and have not yet constructed any grammatical rules, they will need to memorize whatever forms they hear, both as a production mechanism until rules for synthesis are in place, and also to use as input data for morphological learning.

Second, even after speakers have constructed a grammar of rules, there may be some forms that the grammar cannot productively derive the correct output for (such as lexical exceptions). In such cases, the correct form can be produced only by using some form of word-specific knowledge—either by retrieving it as a listed exception, or by some sort of lexically-specified rule. We may assume, following Aronoff (1976), that forms produced using word-specific knowledge take precedence over or *block* productively synthesized forms. For inflected forms that the grammar *can* derive productively, memorizing the inflected form is not necessary, but it also not harmful. In the discussion that follows, I will assume that speakers may memorize regular, grammatically derivable forms at least some of the time — following, for example, Baayen, Dijkstra and Schreuder (1997) and Gordon and Alegre (1999) — but this is not crucial.

The hypothesis to be tested here, then, is that although there are no a priori restrictions on what forms can be memorized, there *are* a priori restrictions on the structure of grammar. In particular, the single base hypothesis means that for one form in the paradigm (the base), there are no rules that can be used to synthesize it, and memorization is the only option. Other forms in the paradigm may be memorized or may be synthesized, but synthesis must be done via operations on the base form. Since we are assuming here a word-based model of morphology, the base is a fully formed surface member of the paradigm, and for this reason, I will call this the *single surface base* hypothesis.

The single surface base hypothesis makes strong predictions about the types of errors that a speaker may make. There is only one way to produce base forms (retrieving them from the lexicon), so if lexical access fails for some reason, the speaker will have no way to synthesize a base form. In other words, base forms will be produced correctly, or not at all. In contrast, non-base forms may either be retrieved from memory as listed inflected forms, or they may be synthesized, using the base form together with the relevant rules. In the case of regular, grammatically derivable forms, either method will yield the same result. For exceptional non-base forms, however, only the stored inflected form will yield the right result; the grammatically synthesized form will be an overregularization (Marcus, Pinker, Ullman, Hollander, Rosen, and Xu 1992). This is shown in Figure 1.1.

This model predicts several asymmetries: first, it is possible to produce incorrect non-base forms, but there is no way to derive an incorrect base form. Second, it is possible to overregularize non-base forms by uttering the grammatically expected form instead of a lexically listed exceptional form, but there is no way to "overirregularize" by creating new listed exceptions without any positive evidence for them. For any given language, if you know which form is the base form and what the grammar for deriving the rest of the paradigm looks like, there is only one class of items that should be open to change: exceptional forms in non-basic parts of the paradigm. The goal of this thesis is to develop an algorithm to determine what the base form is and what the grammar looks like.

For the first several chapters of this thesis, I will consider relatively small, "local" paradigms involving just one tense or a handful of noun cases, and I will these small examples to explore the procedure by which one might identify a single, privileged base form. It is important to keep in mind, however, that more complicated structures may also be necessary, especially when we consider larger paradigms or multiple tenses, moods, etc. Language descriptions frequently refer to multiple stems or bases for a single lexical item—for example, Latin nouns are listed in the dictionary in two forms (nominative and genitive), and Latin verbs are listed in four principal parts. Descriptions involving multiple stems, or multiple listed root allomorphs all

a. One way to produce basic forms:

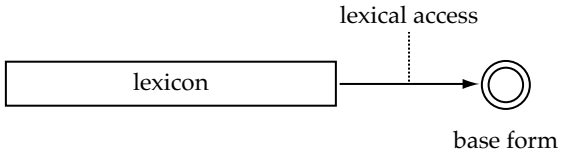

b. Two ways to produce non-basic, regular forms, with identical outcomes:

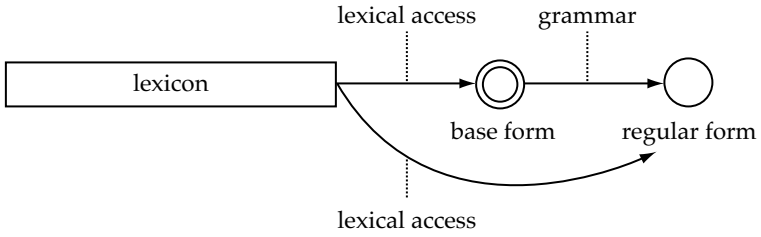

c. Two ways to produce non-basic, exceptional forms, with different outcomes

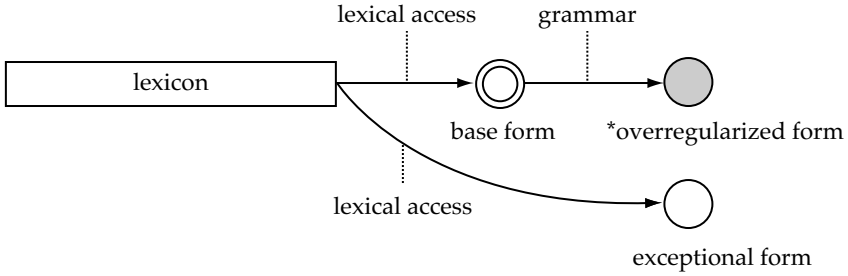

Figure 1.1: Routes for deriving different types of words

seem to require more than just a single base form within the paradigm. In chapter 6, I will discuss a case from Spanish that appears to require multiple, local bases, and I will propose an extension of the base identification algorithm that could be used to identify when a more complicated paradigm structure might be helpful.

1.2 The problem of bases in stem-based morphology

The base identification algorithm proposed in chapter 3 is couched in a word-based model of morphology, and much of the discussion here will focus on the task of selecting a whole word base. It is important to remember, however, that many models of morphology, including those usually assumed (implicitly or explicitly) by phonologists, do not convert surface forms to other surface forms. Rather, they combine independent morphemes (typically stems and affixes) to construct the surface forms of the paradigm. Under such an approach, the input to the morphology is a set of underlying forms, which may or may not match any of the surface forms. What would base identification consist of in a stem-based model?

Within generative phonology, the most serious attempt to develop an all-purpose model of input discovery came in the late 1970s, as part of a quest to constrain the abstractness of underlying representations (URs). As in word-based morphology, there are many possible theories of what should be allowed to serve as bases (or URs) in a stem-based model, and many possible restrictions that could be placed on them. Kenstowicz and Kisseberth (1977, chap. 1) review a series of intuitively reasonable and appealing hypotheses about how to restrict underlying forms, including:

(1) requiring that URs match unsuffixed forms

(2) allowing URs to come from suffixed forms, but requiring that they come from the same place in the paradigm for all words

(3) allowing URs to come from anywhere in the paradigm, but requiring that the form chosen as the UR always occur in the most slots in the paradigm

(4) allowing the UR to come from different parts of the paradigm for different words, but requiring each UR to match a surface form somewhere in the paradigm

(5) allowing URs to combine information from different parts of the paradigm, but requiring that each segment in the UR must surface as such somewhere in the paradigm.

In each case, Kenstowicz and Kisseberth present well motivated arguments showing that if we were to adopt the restriction, we would be unable to construct phonological rules to predict certain alternations. The end conclusion is that we must allow URs to contain abstract structure that is found nowhere in the paradigm, but is needed to derive alternations.

Unfortunately, this conclusion also makes the job of the UR discovery mechanism much more difficult. It is perhaps telling that the field has yet to produce a general purpose model of UR selection that can compare all the forms in the paradigm, find the necessary segments from each one, and posit abstract segments where needed. This is not to say that it could not be done. However, in the course of this thesis, I will present some arguments in favor of adopting one of the more restrictive hypotheses above. In particular, an algorithm that selects a single surface

base form within the paradigm is essentially the same as restriction (2) above. In chapters 2 and 4, I will show that in some cases (Yiddish and Latin), this hypothesis makes more specific predictions than the more conventional, unrestricted use of URs, while in chapter 5 I will show that in other cases (such as Lakhota) it makes completely different predictions; moreover, in all cases, these predictions appear to be correct.

1.3 The problem of bases in correspondence theory

I have been using the term "base" here to refer to the entity from which complex forms are derived — that is, as the input to the morphological grammar. It is worth noting, however, that the motivation behind an algorithm to identify bases is not just morphological; in fact, many phonological theories also have a vested interest in identifying bases. Within the framework of OT, various output-based versions of Correspondence Theory (McCarthy and Prince 1995, Steriade 2000, Benua 1997, (Kenstowicz 1997a, 1997b), McCarthy 1998) have made explicit reference to bases. This usually takes the form of faithfulness constraints which demand that derived forms preserve properties of the base in their paradigm. In the strongest versions of output-output correspondence, which seek to eliminate URs altogether and rely completely on surface relations (e.g., Burzio 1996, Cole and Hualde 1998), it is especially vital that bases be correctly identified. The problem of identifying bases has largely been ignored or deferred in the OT literature. In the case of reduplication, it is only a minor problem, since there are only two entities involved (the base and the reduplicant), we generally know that one was used to create the other, and to a certain extent, it doesn't matter which one is which.

In the larger context of transderivational output-output constraints, there seem to be two opposing camps in the literature. The first argues that output-output constraints are inherently symmetrical, and that any form may potentially influence any other form, in order to achieve a globally more harmonic paradigm (However, it should be noted that change towards a more "basic" form may be enforced in a roundabout way by the existence of a UR, and IO-Faithfulness constraints demanding surface forms to match the UR.) An alternative approach, advocated by Benua (1997), Kenstowicz (1998), and others, is to treat output-output correspondence as an asymmetrical relation, where one form is given priveleged *base* status, and allowed to influence other forms of the paradigm. (See Kiparsky, in prep., for a discussion of the pros and cons of these two approaches.)

If we rely on asymmetrical output-output faithfulness constraints, then we need a procedure for determining which form should be considered the base. This is not a trivial problem; Noyer (1998) and Buckley (1999) discuss several cases in which derived forms are apparently faithful to things other than their smallest or most immediate consitutents. Note also that the base identification problem is difficult to solve using an OT-internal learning strategy, because we would need to evaluate three things that are changing simultaneously: the set of hypothesized URs, the hypothesized base, and the ranking of OO faithfulness with respect to markedness and IO faithfulness. Finding the right base for the purposes of evaluating OO correspondence constraints will not be the focus of my discussion here, but it is hoped that the same considerations may hold in both tasks, and that a procedure like the one outlined here could provide an independent means of identifying such bases.

1.4 Plan of the thesis

The puzzle that must be solved, then, is as follows: bases appear to play a role in several areas of linguistics, including in historical changes like paradigm leveling, in psycholinguistic models of lexical organization, in evaluating output-output faithfulness between surface forms. However, the form that serves as the base seems to vary somewhat from language to language. Therefore, learners must be equipped to learn somehow what forms to use as the bases in their language. The factors which would lead learners to use different bases in different languages, and a procedure for exploring hypotheses about basehood as part of the acquisition process, remain largely unexplored and unformalized; I am pursuing here the strong hypothesis that the choice of base is always determined by a single, universal principle. The crucial observation is that previous attempts have met with only partial success because they have focused solely on inherent properties of the base forms themselves — their frequency, their morphosyntactic markedness, etc. The hypothesis which I will explore in this thesis is that bases are identified in the process of learning the relations between forms:

- Learners begin by exploring all relations that are available to them (i.e., all relations between forms that they have actually encountered)
- The goal of morphological acquisition is to find the relations which make the morphological projection problem "easier," in a way which can (and will) be quantified. In other words, they are looking for what would make the best bases.
- Once a global decision has been made about the best all-purpose base, learners concentrate on relations from that form to the rest of the paradigm

I will start with a schematic example in chapter 2, showing for one language (the older, pre-leveling stage of Yiddish, shown in (2)) how one might go about comparing the informativeness of different parts of the paradigm and selecting the form that is globally most informative. It turns out that this form is the 1sg for Yiddish verb paradigms, and furthermore, Yiddish has subsequently undergone widespread paradigm leveling to precisely this form. This example is meant to show conceptually how one could derive predictions about paradigm leveling using such a restrictive model of base identification.

In order to make testable predictions about bases in different languages, it is useful to have a computational model of morphological acquisition. Therefore, in chapter 3, I will propose a formal system for modeling the acquisition of basehood computationally, building on the system for learning morphological rules developed by Albright and Hayes (1999a). I will show how this system can be used to identify bases in several small artificial languages. Then, in chapter 4, I will discuss its application to the more difficult and realistic problem of the Latin *honor* analogy, a change which violates numerous typological generalizations. I will show that the model is able to select an oblique form as the base for Latin noun paradigms, and that the resulting grammar predicts leveling for exactly the right set of words, in the right direction. The Latin example will serve as a demonstration of how the model is able to select a form other than the nominative singular as the base for this change, and how it can predict paradigm leveling without relying on a formal notion of paradigm uniformity or uniform exponence.

In chapter 5, I will turn to a rather different type of paradigmatic change: the creation of new paradigm types. Using data from Lakhota, I will show how the single surface base hypothesis

makes predictions not only about paradigm leveling, but also about other types of analogical change as well. The changes that have taken place in Lakhota are especially interesting because they are completely unexpected under a less restrictive model of UR or stem discovery, for reasons that will be discussed.

Finally, in chapter 6, I will return to some comparisons with other traditional explanations of paradigm leveling. I will contrast the proposed model with explanations that rely on factors like markedness, token frequency, and frequency of occurrence within the paradigm, showing that the current model makes stronger predictions, which appear to be correct. I will also discuss a class of cases, in which it appears that analogical changes have been based on *less* informative members of the paradigm; this includes Korean (Hayes 1995; Kenstowicz 1997b), Maori (Hale 1973), and others. I will suggest a way in which the proposed model could be extended to these cases as well, by making use of the idea that not all forms in the paradigm are actually available in equal numbers to learners. I will show how this extension of the system could be used to handle many of the typological tendencies observed by Kuryłowicz, Mańczak, Bybee, Hock, and others. Finally, I will discuss some possible ways in which the single base hypothesis could be relaxed to allow local bases, in order to handle larger inflectional paradigms with multiple tenses, moods, and aspects.