

16 Paradigmatic Structure: Inflectional Paradigms and Morphological Classes

ANDREW CARSTAIRS-MCCARTHY

1 Introduction

For a language to have inflectional morphology, there must be some words (more precisely, some lexemes) which occur in a variety of forms (or word forms), with the choice between these forms being determined by the syntactic context. Each word form can be thought of as expressing the lexical content of the lexeme plus some morphosyntactic property or combination of morphosyntactic properties. For example, English *gave* can be thought of as expressing the lexeme GIVE plus the property Past, while *gives* expresses GIVE plus the properties Third Person, Singular and Non-Past.

There are widely differing views of how what we here call ‘morphosyntactic properties’ should be handled in syntax. Some linguists treat them on a par with lexical stems as occupants of terminal positions in syntactic structures (with or without a phonological shape); others treat them as features of lexical stems which are ‘spelled out’ outside the syntactic component. The linguists who have devoted most attention to paradigms and inflection classes in recent years have generally preferred the latter approach. In principle, however, the issues discussed here arise independently of the way in which inflection is handled syntactically, and the use of the term ‘morphosyntactic property’ in this article should not be taken as necessarily implying that entities such as ‘Past’, ‘Third Person’ or ‘Plural’ (or their phonological realizations) can never be syntactic constituents.

Each of the actually or potentially distinct word forms belonging to a lexeme is associated, then, with some morphosyntactic property or combination of properties. (The significance of ‘actually or potentially’ will be discussed in section 2.) The entire set of these properties or property combinations constitutes

the ‘paradigm’ for that lexeme, and each individual property or property combination within this set can be called a ‘cell’. For example, the paradigm for the English lexeme GIVE consists of five cells. Possible labels for these cells, in terms of morphosyntactic properties, are given here alongside the corresponding word forms:

(1) Past	<i>gave</i>
Third-Person Singular Non-Past	<i>gives</i>
Perfective <i>or</i> Passive	<i>given</i>
Progressive	<i>giving</i>
Basic (used in all other syntactic contexts)	<i>give</i>

In the form labelled ‘Perfective *or* Passive’, comparison with the corresponding forms in other lexemes such as *spoken, taken, eaten* allows us to distinguish a suffix *-en* which is added to the lexical stem. But not all English verbs take *-en* in the form corresponding to this cell; for example, the verb SING, as illustrated here, does not:

(2) Past	<i>sang</i>
Third-Person Singular Non-Past	<i>sings</i>
Perfective <i>or</i> Passive	<i>sung</i>
Progressive	<i>singing</i>
Basic (used in all other syntactic contexts)	<i>sing</i>

This difference between GIVE and SING can be expressed by saying that they conform to the same paradigm, but belong to different morphological classes or ‘inflection classes’. An inflection class is a set of lexemes which share a paradigm and whose word forms are alike in respect of the realization of the morphosyntactic properties in every cell.

Before we go further, some comments on terminology are needed. The sense of ‘inflection class’, as defined here, is well established. It applies to words of any word class; older writers, however, used the term ‘declension’ (class) for an inflection class of nouns or adjectives and ‘conjugation’ (class) for an inflection class of verbs. The sense of ‘paradigm’, as defined here, is also well established, although one also encounters it used in the sense of ‘set of inflectional realizations appropriate to a given inflection class’ (Carstairs 1987, Matthews 1991). In that sense, GIVE, SING and BAKE will have distinct ‘paradigms’. The term ‘cell’, as defined here, is not yet well established, but there is a clear need for some term with this denotation.

Many questions arise from these introductory definitions and illustrations:

- 1 Must all lexemes in a given word class share the same paradigm?
- 2 What morphosyntactic properties typically distinguish the cells in paradigms, and how do these properties interact? For example, do some properties presuppose or exclude others? And do properties of

one morphosyntactic category (e.g. Case or Tense) tend to be realized consistently in the same position relative both to the stem and to properties of other categories, whether in individual languages or cross-linguistically?

- 3 For any given lexeme, how is the form for any one cell affected by forms for other cells? Do some forms act as the bases from which others are formed? And are there any generalizations to be made about patterns of homonymy or syncretism between forms?
- 4 How much inflection class proliferation is possible within a language, and how is inflection class membership represented lexically?

These questions will be addressed in sections 2–5 respectively.

2 Paradigm consistency

It is conceivable that some members of a word class might be inflected for one set of morphosyntactic properties while other members of that word class in the same language are inflected for quite a different set of properties. In such a language, there would be two quite distinct paradigms applicable to lexemes of the same word class. But this kind of paradigm inconsistency never, or almost never, arises (Carstairs 1987: 10–11). Instead, all, or nearly all, members of the same word class in a given language are inflected for the same properties. Deviations from this pattern, when they occur, may be of two kinds: a lexeme may unexpectedly lack some forms, or it may possess unexpected 'extra' forms. As regards the first kind of deviation, gaps in the set of word forms associated with certain lexemes nearly always turn out not to be unexpected after all, in that they are attributable to clear-cut semantic or syntactic factors; for example, mass nouns may lack Plural forms, or intransitive verbs may lack Passive forms. In the few lexemes which lack expected word forms without any such motivation, what we find is that the same morphosyntactic properties apply as in the usual paradigm, but that some cells are simply left idiosyncratically unrealized. These are 'defective paradigms' (which in our terminology would more aptly be called 'defective inflection classes'). An example is the archaic English verb QUOTH 'say', used to indicate direct quotations, which has a form only for the Past cell (*quoth*). The second kind of deviation, involving 'extra' forms, will be discussed in section 4.2.

The strong tendency towards paradigm consistency is understandable from a syntactic point of view. If word classes were paradigmatically inconsistent, then there would be a risk that the syntax would need to make available different constructions with the same syntactic function for use with different lexemes. Paradigm consistency also makes it reasonable to allow the same word form to be associated with more than one cell. For example, most English verbs, unlike GIVE and SING, have only four, not five, forms (e.g. *bake*, *baked*,

bakes, baking for the lexeme *BAKE*), yet we typically ascribe to all English verbs the same paradigm, thus:

(3) Past	<i>baked</i>
Third-Person Singular Non-Past	<i>bakes</i>
Perfective <i>or</i> Passive	<i>baked</i>
Progressive	<i>baking</i>
Basic (used in all other syntactic contexts)	<i>bake</i>

The Past and Perfective cells are treated as morphosyntactically distinct for all verbs, because the associated word forms, though usually identical, are actually distinct in the small but frequently occurring class of irregular verbs such as *GIVE* and *SING* (see (1) and (2)).

The label ‘Perfective *or* Passive’ illustrates the fact that one word form may realize a disjunction rather than a combination of morphosyntactic properties. The term ‘morphome’ has been proposed for such disjunctions of properties with consistently shared realizations (Aronoff 1994). It appears that the morphological system of a language may treat certain morphosyntactic properties alike, even if they have no special relationship in the syntax.

3 Morphosyntactic categories and the internal structure of paradigms

In Indo-European languages with relatively elaborate inflectional morphology, nouns typically inflect for the morphosyntactic categories Number and Case, adjectives for Number, Case and Gender, and verbs for Person, Number, Tense, Mood (e.g. Indicative, Imperative, Subjunctive), sometimes Voice and sometimes Aspect (e.g. Imperfective versus Perfective). But how typical are Indo-European languages in this respect? What cross-linguistic generalizations, if any, can be made about the morphosyntactic categories which are relevant to particular word classes? Bybee (1985) studied categories expressed inflectionally in verbs in a sample of fifty languages from different language families and different cultural and geographic areas. She found there to be nine such categories: Valence (relating to the number and role of arguments of the verb), Voice, Aspect, Tense, Mood, Number, Person of subject, Person of object and Gender. Of these, the most frequent is Mood (expressed inflectionally in 68 per cent of the fifty languages), followed by Person (56 per cent), Number (54 per cent), Aspect (52 per cent) and Tense (48 per cent); least frequent are Gender (16 per cent) and Valence (6 per cent). She seeks to explain these figures in terms of the two notions of ‘relevance’ and ‘generality’. The commonest inflectional categories are those which are both highly relevant, in that their semantic content ‘directly affects or modifies’ the semantic content of the stem (Bybee 1985: 13), and highly general, in that their semantic content is applicable to all or almost all verbs. Gender and Valence are relatively rare as inflectional

categories, because they fail this criterion, but in opposite ways. Gender is said to be highly general but low in relevance; Valence is said to be high in relevance but low in generality, because valence-changing operations such as causative formation often alter the meanings of stems so unpredictably that the new form becomes lexicalized, and therefore no longer belongs to the same lexeme as its base.

A similar cross-linguistic survey has yet to be done for nominal and adjectival categories. There are in any case often difficulties in deciding objectively which of two categories, such as Tense and Aspect, has a greater effect on the semantic content of stems. But, independently of these semantic issues, Bybee points out generalizations concerning the order in which the nine categories are realized in relation to each other and to the verb stem. For example, when a language in Bybee's sample has both Aspect and Tense as inflectional categories and their order of realization can be determined (i.e. they are not realized cumulatively), Aspect is realized closer to the stem than Tense; similarly, Tense is nearly always realized closer to the stem than the less 'relevant' property Person (of subject). A related observation concerns stem allomorphy. It is not uncommon for languages to have special stem allomorphs or suppletive stem forms distributed according to Aspect or Tense, but it is very rare for stem allomorphy to be distributed according to Person (Rudes 1980). This tends to support the view that the property contrasts which partition a paradigm are not all equally important; rather, the distinction between cells which differ in Tense or Aspect is more fundamental, in some sense, than the distinction between cells which differ in Person – a less 'relevant' category in Bybee's hierarchy.

If, in general, all the morphosyntactic categories which are manifested in a paradigm were equally fundamental, one might expect to find a cell for every possible combination of the properties belonging to them. The English data in (1)–(3) show this to be wrong, however. There is a non-Past Third-Person Singular cell but no Past Third-Person Singular, and the Tense and Person distinctions do not apply to the Perfective or the Progressive forms. In other languages too it is common for a certain category to be excluded (or for property contrasts within it to be neutralized) when a particular property belonging to some other category is present. For example, adjectives in German and Russian inflect for Number, Case and Gender; however, there are no cells for distinct Gender forms in the Plural. It seems likely that the excluded category will generally be either low in 'relevance' (in Bybee's terms) or realized relatively far from the stem, while the property which imposes the exclusion will be high in relevance or realized relatively close to the stem, or both; but this has not yet been systematically investigated.

Bybee's observations also provoke the question as to whether all the properties of one category are always realized in the same position, or 'slot', relative to the stem. The answer is no; for example, in the Present Tense of Georgian verbs the Person and Number of the subject are sometimes cumulated as a prefix, sometimes cumulated as a suffix, sometimes realized separately, with Number as a suffix *-t*, and sometimes not realized overtly at all:

(4)		Singular	Plural
	Person: 1	v-xedav 'I see (him/her/it/them)'	v-xedav-t
	2	xedav	xedav-t
	3	xedav-s	xedav-en

Even so, it is more usual for all properties in a category to occupy a consistent position in all or most of the relevant word forms, as in the following Swahili example, where the prefixal positions labelled I, II and III are occupied by object markers, Tense markers and subject markers respectively:

(5)		III	II	I	Stem	
		a-	ta-	ni-	penda	'he will like me'
		a-	ta-	ku-	penda	'he will like you (sg.)'
		a-	li-	ni-	penda	'he liked me'
		a-	li-	ku-	penda	'he liked you (sg.)'
		wa-	ta-	ni-	penda	'they will like me'
		wa-	ta-	ku-	penda	'they will like you (sg.)'
		wa-	li-	ni-	penda	'they liked me'
		wa-	li-	ku-	penda	'they liked you (sg.)'

Detailed consideration of phenomena of this kind has led Stump (1992, 1993c) to re-emphasize the importance in morphological theory of 'position classes', meaning classes of affixes which are mutually exclusive and occupy the same position (or positions) relative to the stem (cf. also Stump, INFLECTION). (Zwicky 1990 mentions the related principle of 'slot competition'.)

In the example at (5), each of the positions is occupied by affixes realizing properties belonging to only one category (subject Person, object Person, or Tense). This is not always so, however; for example, in Georgian the prefix position immediately before the verb stem in the Present tense (see (4)) may be occupied by an object Person marker such as *g-* 'you (sg.)' as well as by the subject marker *v-* 'I'. When the two are in competition for this slot, it is *g-* which wins: *g-xedav* 'I see you (sg.)'. It remains to be seen how tightly the roles of slot competition and position classes can be constrained within morphological theory.

4 Word forms, syncretism and the internal structure of paradigms

Alongside the mutual relationship of the morphosyntactic properties which define the cells of a paradigm, one can consider the mutual relationship of the corresponding word forms. These word forms may simply be distinct, implying no special mutual relationship. On the other hand, one word form may appear

to be built on another, or one word form may be identical with another. These situations will be illustrated and discussed in sections 4.1 and 4.2.

One constraint on nominal paradigms has been proposed which relates simply to the number of phonologically distinct exponents (affixes, etc.) available for the properties to be expressed. Plank (1986: 46) suggests that 'the number of exponents potentially available for nominal inflexion in any language is limited to about 30'. It follows that in a language with cumulative exponence of Case and Number (as in Latin), with two Numbers (Singular and Plural) and with two inflection classes for nouns, there can be no more than about seven Cases; on the other hand, a language with agglutinative exponence of Case and Number (as in Hungarian or Turkish) can have twenty or more Cases. Observationally this appears correct, but it is not clear why the crucial limit should be around thirty rather than (say) ten or sixty.

4.1 Word forms as bases for other word forms

Word forms which appear to be built on other word forms are common in the nominal inflection of Daghestanian languages such as Archi (Kibrik 1991a and ARCHI (CAUCASIAN – DAGHESTANIAN)). In these languages, each noun typically has an 'oblique' stem, which is used in some but not all Case forms, and which differs between Singular and Plural. The oblique stem may be identical to a particular Case form, such as the Ergative, as in Lezgian (Haspelmath 1993), so that the question arises of whether or not we should analyse this actual Case form as the base on which others are built. Mel'čuk (1986), in contrast to Kibrik (1991), favours this approach, distinguishing 'secondary' Cases whose forms are built on other Case forms from 'primary' Cases whose forms are independent. It is not clear how far the primary–secondary dichotomy, so defined, correlates with any other partitions of the paradigm discussed in this section and section 3. A better-known example of the same kind involves the Future Active Participle forms of Latin verbs, which are sometimes said to be built on the Past Passive Participle forms, as illustrated in (6):

(6)	Present stem	Past Passive Participle stem	Future Active Participle stem
	am- 'love'	ama:t-	am-a:t-u:r-
	po:n- 'put'	posit-	posit-u:r-
	fer- 'carry'	la:t-	la:t-u:r-
	rump- 'break'	rupt-	rupt-u:r-

Matthews (1972) describes the formation of the Future Active Participle as 'parasitic' on the Past Passive Participle; on the other hand, Aronoff (1994) argues that both participles have equal status as derivatives from one stem which is itself a purely morphological entity, expressing no morphosyntactic properties.

Bybee (1985) approaches this issue from the point of view of markedness relations between properties within a category. She emphasizes the historical tendency for forms which realize relatively unmarked properties within a category to function as the base for new forms realizing relatively more marked properties. An illustration is the development of Preterite inflection in the modern Charente dialect of Provençal, as shown in (8), from the Old Provençal pattern in (7):

(7) Old Provençal:

	Singular	Plural
Person: 1	améi 'I loved'	amém
2	amést	amétz
3	amét	améren

(8) Charente dialect:

	Singular	Plural
Person: 1	cantí 'I sang'	cantétem
2	cantétei	cantétei
3	cantét	cantéten

Bybee attributes this to the reanalysis of the Third-Person Singular form. Having been originally a combination of a stem *cant-* and a Tense–Person–Number suffix *-et*, it was reanalysed as Preterite Tense form *cant-ét* with no overt realization of the morphosyntactically unmarked Person–Number combination Third-Person Singular. As such, it came to function as the base for the derivation of a new set of Preterite forms in the Charente dialect (except in the First-Person Singular). In Mel'čuk's terms, one could say that all but the First-Person Singular are secondary forms derived from the primary Third-Person Singular form. Morphosyntactic unmarkedness is one of the factors which, according to Bybee, contribute to the autonomy of certain forms of a lexeme – that is, to the likelihood that they have separate lexical representations. Kuryłowicz (1945–9) discusses similar phenomena, but with different terminology.

4.2 *Syncretism*

In principle, there could be both factors encouraging and factors inhibiting inflectional homonymy, or syncretism, both at the level of morphosyntactic content and at the level of morphological realization. This yields four possible types of factors. In practice, there is evidence that factors of all four types exist, though there are considerable differences between them as regards strength and theoretical status.

Universal homonymy, for all lexemes, in the realization of two cells in a paradigm is by definition impossible, for in any such purported situation there would be no ground for recognizing those two cells as distinct. Even

near-universal homonymy is grounds for suspicion. For example, one might recognize two distinct Locative Singular cells in Russian on the grounds of the distinct forms *lesu* and *lese* of LES 'forest' in *v lesu* 'in the forest' and *o lese* 'concerning the forest', and likewise in a few other nouns, even though for the majority of nouns the corresponding forms are identical. But synchronic descriptions of Russian have tended to say, rather, that for the majority of nouns there is only one Locative Case, thus implicitly attributing to Russian a departure from strict paradigm consistency on the part of the LES class (see section 2).

There are, however, clear-cut instances of inflectional homonymy which cannot be ascribed to paradigm inconsistency, and which are not attributable to phonological factors either. A standard illustration is the Dative–Ablative Plural syncretism in all Latin nouns, adjectives and determiners. Its universality in Latin confirms that it is a systematic feature of the morphological system, and the fact that two phonologically quite distinct affixes, *-is* and *-ibus*, realize Dative–Ablative Plural in different lexical contexts shows that this homonymy cannot just be a phonological accident. A second illustration is the Past and Perfective syncretism in regular English verbs, shown in (3). How should morphological theory accommodate such phenomena?

From a naïve common-sense point of view, all inflectional homonymy should impair communicative efficiency by increasing the chance of misunderstanding through ambiguity. In practice, the pragmatic, semantic and lexical context of any utterance nearly always prevents misunderstanding due to homonymy, whether lexical or inflectional. Even so, it has been suggested that there are some inflectional homonymies which are avoided because of the morpho-syntactic ambiguity which they create. Plank (1979, 1980) argues that the development in Vulgar Latin of Genitive and Nominative Singular homonymy in certain nouns was inhibited for prototypical possessors (humans) and for certain typical possesses (body parts, etc.) because of the need to maintain an overt distinction between possessor and possessee in constructions with a head noun as possessee and a dependent noun in the Genitive as possessor. The same need is held to explain certain otherwise puzzling patterns of acceptability in German; for example, *Benachteiligungen andersgläubiger Frauen* 'acts of discrimination against heterodox women' is acceptable because the suffix *-er* on *andersgläubiger* 'heterodox' unequivocally marks the phrase *andersgläubiger Frauen* as dependent; on the other hand, **Benachteiligungen Frauen* 'acts of discrimination against women' is not acceptable because there is no overt indication of the roles of the two nouns.

Jakobson (1936: 85–8) proposed for Russian nouns a variety of implicational statements relating to syncretism possibilities, with the conditions expressed in terms of morphosyntactic content; for example, if a noun has distinct forms for the Accusative and the Nominative, then either the Accusative and Genitive or the Dative and Locative must be homonymous. Such conditions can hardly be generalized to other languages, however. Bierwisch (1967) described certain German syncretisms in terms of the sharing of syntactic features between

cells. E. Williams (1981b, 1994a) also uses syntactic features to organize hierarchically the paradigm cells for any lexeme class into a branching tree structure, such that forms which are systematically homonymous for any lexeme are all and only those cells dominated by some node on the tree. One may question, however, whether this hierarchical organization is independent of the syncretisms which it is designed to explain; and in any case there are both syncretisms which might occur but do not, and (it has been suggested) some syncretisms which do occur despite the fact that the cells concerned do not form a 'constituent' within the tree structure (Joseph and Wallace 1984). So the search for morphosyntactic-feature configurations which strongly favour syncretism has had only modest success so far.

A characteristic of many syncretisms is that one can distinguish a morphosyntactic property (or properties) which provides the context in which the syncretism occurs. This does not apply to the Past–Perfective syncretism in regular English verbs, for which the conditioning factor is purely lexical (a matter of belonging to the regular inflection class). On the other hand, it does apply to the Latin Dative–Ablative Plural syncretism, where the conditioning factor is morphosyntactic (a matter of being Plural rather than Singular). For such syncretisms, one can investigate the possibility of generalizations concerning the realization of the contextual property in the word forms concerned. Is it, for example, always realized cumulatively with the properties whose expression is rendered homonymous by the syncretism? Carstairs (1987) suggests that the answer to this question is generally yes, and that there is, moreover, a quasi-functional motivation for it. For a lexeme with cumulative inflection, a syncretism reduces the amount of morphological material (affixes, etc.) to be distributed among the cells of the paradigm. On the other hand, for a lexeme with agglutinative inflection, a syncretism does not reduce the amount of this material, and may even increase it; and, in any case, the syncretism will complicate the distribution of this material among the cells, by comparison with a similar pattern without syncretism. Carstairs and Stemberger (1988) show also that syncretism with cumulation is easy to model in a connectionist framework, while syncretism without cumulation is difficult. There remains, however, the problem of the minority of morphosyntactically conditioned syncretisms where the contextual property is not cumulated with the homonymously realized properties. Carstairs (1987) suggests that they are subject to a generalization involving the mutual ranking of the properties in Bybee's relevance hierarchy; but even if this generalization is correct, the reason for it is not clear.

5 Inflection class organization

Two main questions have been asked about inflection class organization in recent years: (a) How is the assignment of a lexeme to its inflection class to be represented grammatically? (b) Are there any constraints on the number of

inflection classes which words of a given class in a given language can have? These questions, though seemingly independent, have turned out to have answers (at least provisional ones) which impinge on one another.

Many pedagogical grammars and dictionaries of languages such as Latin and Russian indicate inflection class membership by means of essentially arbitrary labels, such as 'Class I' or 'second declension'. These labels are a shorthand for full sets of inflected forms illustrated by one or more exemplary lexemes for each class. Labels of this kind have been criticized as arbitrary by both generative and non-generative grammatical theorists. In the generative tradition, Lieber (1981) suggested that the full inflectional behaviour of any lexeme might be predictable from its pattern of morphological alternation (essentially, its set of stem allomorphs). However, this strong claim cannot be sustained, and another generative linguist, James Harris (1991), reverts to traditional numerical labels to encode lexically the inflectional behaviour of Spanish substantives.

An alternative tradition has been to cite certain word forms, or 'principal parts', for each lexeme, from which its whole inflectional behaviour can be determined by reference to rules. This approach is adopted and refined theoretically by Wurzel (1984, 1987). Wurzel points out that simply to label distinct inflection classes as 'Class 1' etc. obscures three facts: (a) that inflection classes are typically not entirely distinct in their realizations for all cells, but rather resemble each other in some or most cells; (b) that inflection class membership is frequently influenced by extramorphological factors such as Gender, meaning (e.g. animateness) and phonological shape; (c) that within a set of inflection classes whose membership is not influenced by such extramorphological factors and which are therefore in a sense competing for the same pool of lexemes ('complementary classes', in Wurzel's terminology), there is typically one class which is unmarked in the sense that it is the class to which new words are assigned, the class to which words of the other classes are 'wrongly' assigned by infant learners, and the class into which members of the other classes drift in the course of language change. Wurzel seeks to account for these facts by a model of inflection class membership which incorporates 'paradigm structure conditions' (PSCs) and 'reference forms'.

A PSC is a statement to the effect that if a lexeme has characteristic X , then it must have realizations R_1, \dots, R_n for cells C_1, \dots, C_n . Such conditions are conceived as structuring the implicit knowledge that native speakers of inflectionally complex languages acquire. 'Characteristic X ' may be more or less elaborate, and 'cells C_1, \dots, C_n ' may cover greater or lesser portions of the paradigm. In the simplest case, characteristic X is membership of a given word class (e.g. verb), and cells C_1, \dots, C_n extend to the whole paradigm. A PSC of this kind is appropriate for a word class with no inflection class distinctions. But what if there are two or more complementary inflection classes within a word class (or an extramorphologically defined subset of a word class)? Here Wurzel invokes 'reference forms', which are akin to traditional 'principal parts'. Lexemes within the unmarked class have no lexical specification, and obey the

dominant PSC for the set of complementary classes – one which specifies realizations for all cells in the paradigm. On the other hand, lexemes within a marked class contain in their lexical specification a reference form for at least one cell in the paradigm, which overrides the dominant PSC in respect of that cell. This reference form may also function as part of ‘characteristic X’ in a more specific PSC, which, so far as it extends, overrides the dominant condition; for cells not covered by the more specific PSC, however, these lexemes conform to the dominant PSC, which thus constitutes the default PSC for the whole set of complementary classes. The possibility of PSCs being wholly or partly shared by different inflection classes is what accounts for inflection class resemblances in Wurzel’s framework, and the drift in membership to unmarked classes is seen as lexical simplification through the loss of reference forms, whereby the lexemes in question become wholly subject to the dominant PSC.

The illustration in (9) shows a hypothetical set of complementary classes for nouns, with Class A assumed to be unmarked and a–h representing distinct inflectional realizations.

(9)	Class A	Class B	Class C	Class D
Cell 1	a	a	f	f
Cell 2	b	e	e	e
Cell 3	c	c	h	h
Cell 4	d	d	d	g

In (10) we see how this pattern of classes would be represented in Wurzel’s framework:

- (10) PSCs: (a) Noun → {a/1, b/2, c/3, d/4}
 (b) g/4 → h/3 → {f/1, e/2}

Reference forms included in lexical specification:

- Class A: none
- Class B: e/2
- Class C: h/3
- Class D: g/4

The fact that PSC (a) is dominant follows from the fact that it is more general than PSC (b), being framed so as to apply to all nouns, not just those with a particular reference form; and the fact that Class A is unmarked is represented by the fact that it has no reference form in its lexical specification. PSC (b) incorporates two implications, and so serves for both Class C and Class D. For Class C words, however, only the second implication (h/3 → {f/1, e/2}) has effect, and the form for cell 4 (*viz.* d) is supplied by PSC (a). Wurzel has applied this sort of analysis to inflection class systems in a variety of mainly Indo-European languages, particularly German, Icelandic and Latin.

Carstairs (1987) asks a different question from Wurzel: What is the largest number of inflection classes which a given array of inflectional resources can be organized into? Looking at (10), one can see that the smallest conceivable number of inflection classes, given the number of distinct realizations available for each cell, is two; on the other hand, the largest conceivable number of inflection classes, given that each of the five cells has two realizations, is 2^5 or 32. Carstairs (1987) proposes a Paradigm Economy Principle to the effect that, subject to certain qualifications, the actual restriction on the number of distinct inflection classes for any word class in any language is very tight: it must be no more than the conceivable minimum. An actual language which had precisely two inflections available for each cell could therefore have no more than two inflection classes, not four, as posited in (9).

It is clear that Carstairs and Wurzel disagree, in that Wurzel's framework can handle without difficulty a pattern, such as (9), which Carstairs claims cannot occur. Carstairs-McCarthy (1991) has suggested that the sort of behaviour allowed by Wurzel's multi-stage PSCs, such as (b) in (10), can be found only in non-affixal inflection (such as stem allomorphy and stress alternation), not in affixal inflection. If correct, this suggests that the two types of inflection are subject to different constraints.

More recently, Carstairs-McCarthy (1994) has suggested that inflection classes are constrained by a principle related to the Principle of Contrast proposed by Clark (1987) for lexical acquisition: every two forms contrast in meaning. As Carstairs-McCarthy applies this to affixal inflection, it has the effect of requiring that each word form should either identify unambiguously the inflection class to which its lexeme belongs or else supply no positive information about inflection class membership at all, exhibiting the sole default, or multi-class, realization for that cell. If correct, this claim would rule out the inflection class organization in (9), but for a different reason from the Paradigm Economy Principle; for cells 1 and 3 each have two realizations neither of which either unambiguously identifies its inflection class or constitutes the sole default realization for that cell.

The outcome of the comparison of Wurzel's and Carstairs-McCarthy's approaches will depend on further detailed study of actual inflection class systems; it seems clear already, however, that inflection class organization is by no means a language-particular free-for-all without interest for the morphological theorist.