6

CASE MARKING

6.1 THE DISCRIMINATORY FUNCTION OF CASES

In this chapter, we are going to look at one way in which consideration of data from a wide range of languages has enabled us to gain important new insights into a general linguistic phenomenon, insights that would probably not have been obtained solely by the investigation of a single language, and certainly not from the detailed, abstract analysis of English. If one looks at the accounts given of the uses of cases in traditional, and many non-traditional, grammars, there is usually the assumption – in many instances, justified – that the use of a given morphological case will correlate highly either with a given semantic role, or with a given grammatical relation. Thus the locative case is said to be the case for expressing location, the ablative for expressing motion away from, and so on; the nominative is described as being the case for the subject, the accusative for the direct object (or, in frameworks that eschew the distinction between semantic and syntactic cases, nominative correlates with agent and accusative with patient). In addition to case marking systems based on semantic and/or syntactic criteria, recent linguistic research has also uncovered languages where pragmatic criteria are important in assigning case, as in Japanese and Tagalog, for instance.

In addition, however, to languages where some or all of the cases can be accounted for in this way, there remains a set of recalcitrant data, where on the basis of semantic roles or grammatical relations or pragmatic roles there remain some cases that do not correlate directly with any syntactic or semantic or pragmatic role, but rather seem to be used for a given role, but only in certain, limited circumstances. The aim of this chapter is to investigate some of these examples, in particular examples concerned with subjects and direct objects (or, more accurately, with S, A, and P). The reason why this discussion fits well into our general discussion of universals and

typology is that the kinds of non-correspondence that we shall be looking at are found to recur in a wide variety of languages from different genetic and areal groupings, i.e. we are dealing with a significant phenomenon from the viewpoint of language universals. Moreover, not only can we establish a general pattern of similar distribution across languages, we can actually go a long way towards finding an explanation for this cross-language similarity.

We shall begin our discussion by considering the nominative–accusative and ergative–absolutive case marking systems, already introduced in passing in chapter 5. If we take S, A, and P as our primitives, and assume for the moment that we are restricting ourselves to languages that treat each of these three relations homogeneously, i.e. do not have different cases for different types of S, etc., then it is clear that there are not just two logically possible kinds of case marking system, but five. The nominative–accusative system groups S and A (nominative) together against P (accusative). The ergative–absolutive system groups S and P (absolutive) together against A (ergative). Both of these systems are widespread across the languages of the world. The neutral system would have the same form for all three primitives, but since this is tantamount to lack of case marking for these relations, it is not directly relevant to our considerations: as a system, it is, of course, widespread in the languages of the world, but most languages with this system have other means, such as verb agreement or word order, to indicate which noun phrase is A and which is P in the transitive construction. The fourth possible type, tripartite, would have distinct cases for each of the three primitives. The fifth type would group A and P together as against S.

The tripartite system is found, but is very rare. In a number of languages, as we shall see in more detail below, it is found with a subset of the noun phrases in a language, namely where nominative–accusative and ergative–absolutive systems co-existing in a language intersect. But there is only one language for which it has been reported that this tripartite system exists for all noun phrases in the language, namely Wanggumara. Thus we can say with confidence that this system is very rare across the world's languages. The last type, with A/P–S alignment, seems to be equally rare: the only reliable attestations known to us are for certain classes of noun phrases in certain Iranian languages, where it represents an intermediate diachronic stage in the breakdown of an earlier ergative–absolutive case marking system in the direction of a nominative–accusative system. The question arises immediately why, of four logically possible case marking systems, two should account for almost all the languages of the world that have a case marking system that consistently distinguishes among S, A, and P. If we compare the noun phrase arguments of intransi-

tive and transitive constructions, as in (1)–(2) (irrespective of word order), then a possible motivation for this distribution emerges:

$$S \quad V_{intransitive} \tag{1}$$

$$A \quad P \quad V_{transitive} \tag{2}$$

In the intransitive construction, there is only a single argument, so there is no need, from a functional viewpoint, to mark this noun phrase in any way to distinguish it from other noun phrases. In the transitive construction, on the other hand, there are two noun phrases, and unless there is some other way (such as word order) of distinguishing between them, ambiguity will result unless case marking is used. Since it is never necessary, in this sense, to distinguish morphologically between S and A or S and P (they never cooccur in the same construction), the case used for S can be used for one of the two arguments of the transitive construction. The nominative–accusative system simply chooses to identify S with A, and have a separate marker for P; while the ergative–absolutive system chooses to treat S the same as P, with a separate marker for A. The tripartite system is unnecessarily explicit, since in addition to distinguishing A from P, it also distinguishes each of these from S, even though S never cooccurs with either of the other two. The A/P – S system is, from a functional viewpoint, singularly inefficient, failing to make the most useful distinction (between A and P), and making a useless distinction (between A and S, likewise between P and S). Whatever may be the value of functional explanations in general in linguistics and language universals in particular, here we do have a good example where the predictions of the functional approach appear to fit in very well with the observed distribution of case marking systems across the languages of the world.

In fact, the functional approach makes a further prediction that is borne out by actual distribution. In a case system where one of the two cases used for indicating these three primitives is formally less marked than the other, for instance where one of the forms is simply the stem of the noun in question whereas the other has some overt affix, it is nearly always the case that the formally unmarked item is used to indicate S, whence also A in the nominative–accusative system and P in the ergative–absolutive system. This is Greenberg's universal number 38: 'where there is a case system, the only case which ever has only zero allomorphs is the one which includes among its meanings that of the subject of the intransitive verb', although a very few counterexamples to this generalization have since been uncovered, all with a nominative case more marked than the accusative, e.g. such as Mojave Yuman languages where the nominative takes the suffix -*č* and the accusative has no suffix. If, however, we restrict ourselves

to the more general pattern, then we can see that in the nominative–accusative system, a special marker is added to P to distinguish it from A, which like S is unmarked. In the ergative–absolutive system, a special marker is added to A to distinguish it from P, which like S is unmarked. The functional explanation of these two case marking systems may also explain why there is so often a discrepancy between the case marking system and the syntactic orientation of the language in question, as discussed in chapter 5: the cases do not relate directly to grammatical relations, but rather directly to distinguishing between A and P.

We would emphasize one point before proceeding further with the functional model of case marking and its implications. We are not claiming that the sole function of case marking is discriminatory in this sense, since there is a whole host of instances where the function of a given case can be correlated with semantic parameters. What we are claiming is that there do exist many instances where this functional approach is necessary in order to guarantee a full understanding of the role of case marking.

6.2 NATURAL INFORMATION FLOW IN THE TRANSITIVE CONSTRUCTION

From section 6.1, it emerges that the discriminatory function of case marking will show itself most clearly in the transitive construction, where there is a need to distinguish between A and P, rather than in the intransitive construction, where S alone occurs. Where one finds different cases used for different occurrences of S in a language, the conditioning factor is usually semantic (to the extent that it is not lexically idiosyncratic): for instance in Tsova-Tush, as discussed in chapter 3 (sentences (1)–(2)), the distinction between the ergative and absolutive cases for intransitive subject is dependent on the degree of control exercised by the S over the situation described. There are also instances where differential case marking on A and/or P can be readily handled in semantic terms without appeal to functional factors. For instance, in Finnish the P stands in the partitive case if only partially affected by the action (e.g. if only some of an entity is affected), but in a non-partitive case if totally affected:

Hän otti rahaa (PARTITIVE). (3)
'He took some money.'

Hän otti rahan (ACCUSATIVE). (4)
'He took the money.'

In this section, however, we will be concerned with formal case distinctions that do not correlate this closely with a combination of semantic or syntactic factors, in particular trying to account for the following facts: a large number of languages have special cases for animate and/or definite Ps, distinct from the cases used for other Ps, and also not used elsewhere as markers of definiteness; conversely, many languages have a special case used only for As of low animacy, and not otherwise used as indicators of either A or low animacy.

Before proceeding to the data here, we will outline the explanation, following on from the discussion of the preceding section, that we will be appealing to, as this will make the citation of the individual pieces of data more comprehensible. In the transitive construction, there is an information flow that involves two entities, the A and the P. Although in principle either of A and P can be either animate or definite, it has been noted that in actual discourse there is a strong tendency for the information flow from A to P to correlate with an information flow from more to less animate and from more to less definite. In other words, the most natural kind of transitive construction is one where the A is high in animacy and definiteness, and the P is lower in animacy and definiteness; and any deviation from this pattern leads to a more marked construction. This has implications for a functional approach to case marking: the construction which is more marked in terms of the direction of information flow should also be more marked formally, i.e. we would expect languages to have some special device to indicate that the A is low in animacy or definiteness or that the P is high in animacy or definiteness. This is precisely what we will try to document in the remainder of this section.

In the immediately preceding discussion, we have introduced the two terms animacy and definiteness. We will return to definiteness in more detail later on in this chapter, but for the moment we can work with the general definition of definiteness as the presupposition that the referent of a definite noun phrase is identifiable by the hearer; in terms of English structure, a definite noun phrase will either be a pronoun, a proper name, or a common noun introduced by the definite article or a demonstrative or preposed possessor. Animacy is a much more complex phenomenon, to which we return in chapter 9. For the moment, suffice it to say that a noun phrase is higher in animacy if it is to the left on a continuum some of whose main points are: first/second persons pronouns > other human noun phrases > animal noun phrases > inanimate noun phrases.

If a given transitive construction has to be marked to show that it does not correspond to the normal direction of flow of information, then there are (at least) three ways in which this marking could be made. First, one could mark the construction as a whole, say by marking the verb, to indicate an unexpected constellation of A and P; we examine this possibility in section 6.2.1. Secondly, one of the noun phrases (or both of them) could be marked, say by having a special marker for unexpected As (those low in definiteness or animacy) and/or for unexpected Ps (those high in definiteness or animacy); such examples are discussed in section 6.2.2.

6.2.1 INVERSE FORMS

A number of languages have special verb forms to indicate whether the transitive action is initiated by an A higher in animacy than the P or lower in animacy than the P (with the third possibility, A and P equal in animacy, being treated arbitrarily as the one or the other). Perhaps the most famous instance of this in the linguistic literature is in the Algonquian languages, where one set of verb forms, the so-called direct forms, are used when the A is higher in animacy than the P, while the so-called inverse forms are used where the P is higher than the A. The actual animacy hierarchy of Algonquian languages takes the form: second person > first person > third person proximate > third person obviative. The distinction between two subtypes within third person, proximate and obviate, the former higher in animacy than the latter, guarantees that there will never, in fact, be a transitive construction where A and P are equal in animacy.

The examples below are from Fox, though the general principle holds for Algonquian languages as a whole. The suffix *-aa* in these examples indicates the direct form, while *-ek* indicates inverse form. The prefix *ne-* indicates first person: this illustrates another important property of the Algonquian verb forms, namely that the prefix invariably encodes the participant higher in animacy, irrespective of its grammatical role:

ne *-waapam-aa* *-wa.* (5)
1SINGULAR see DIRECT 3
'I see him.'

ne *-waapam-ek* *-wa.* (6)
1SINGULAR see INVERSE 3
'He sees me.'

6.2.2 DIFFERENTIAL MARKING OF A AND P

The most widespread indication of unnatural combinations of A and P across languages, however, is not by marking the verb, but rather by marking one or both of the noun phrase arguments. The following patterns in particular are found: (a) mark a P high in animacy, i.e. the accusative case is

restricted to Ps that are high in animacy; (b) mark a P high in definiteness, i.e. the accusative case is restricted to definite Ps; (c) mark an A that is low in animacy, i.e. the ergative case is restricted to noun phrases that are low in animacy. Somewhat embarrassing is the absence of clear attestations of the fourth expected type, namely marking of an indefinite A; languages seem rather to avoid this particular construction by outlawing or discouraging transitive sentences with an indefinite A, either recasting them as passives or by using a presentative construction (like English *there is/are ...*). In English, although the sentences *a bus has just run John over* and *a bird is drinking the milk* are surely grammatical, more natural ways of expressing these pieces of information would be *John has just been run over by a bus* and *there's a bird drinking the milk*. In most languages that use the three methods outlined above for indicating less natural combinations of A and P, the case marking of A and P is determined independently, i.e. any A below a certain degree of animacy is marked ergative, irrespective of the P; conversely, any P above a certain degree of definiteness or animacy is marked accusative, irrespective of the A. This contrasts with the inverse verb forms discussed in section 6.2.1, where it is usually the relation of A to P that is important. Finally, before proceeding to detailed exemplification, we should note that there are some languages where the occurrence of the special ergative or accusative marker is conditioned not by any specific rigid cut-off point on the animacy or definiteness hierarchies, but rather by a more general condition of the kind: use the special marker only if there is likelihood of confusion between A and P; the assessment of likelihood of confusion is left to the speaker in the particular context. Hua is an example of a language of this type.

For the relevance of animacy, particularly clear data are provided by Australian languages, almost all of which have split case marking determined by the animacy hierarchy. As would be expected from our discussion above, a special accusative case is often restricted to noun phrases towards the top of the animacy hierarchy: thus in Dyirbal it is found only with first and second person pronouns; in Arabana only with human noun phrases; and in Thargari only with animate noun phrases. Conversely, the special ergative case is found only towards the bottom of the hierarchy, though usually, in fact, in these languages extending quite high up the hierarchy: thus most Australian languages have a separate ergative case for all non-pronominal noun phrases (e.g. Dyirbal), sometimes extending further up the hierarchy into the pronouns. Since the determination of the case of A and P is independent, it sometimes happens that accusative and ergative case marking meet neatly in the middle of the hierarchy without any overlap or gap, but quite frequently there is overlap in the middle of the hierarchy, which means that some noun phrases have a tripartite case

marking system; and it sometimes happens that there is a gap in the middle of the hierarchy, some noun phrases having the neutral case marking system. Thus Ritharngu, for instance, has a nominative–accusative case marking system for pronouns; the tripartite system for humans and intelligent animals; and ergative–absolutive case marking for other nouns, i.e. for non-intelligent animals and inanimates. In some languages, the middle ground in the hierarchy may be shared by both the tripartite and neutral case marking systems, as was discussed in section 3.4 for the Saibai dialect of Kala Lagaw Ya, which thus combines within one language nominative–accusative, ergative–absolutive, tripartite, and neutral case marking.

One result of the split case marking pattern is that a single sentence, in addition to having a nominative A and an accusative P, or an ergative A and an absolutive P, can also have one of the patterns: ergative A and accusative P; nominative A and absolutive P. These possibilities were often effectively discounted in earlier work on ergativity, with its rigid distinction between nominative and ergative constructions. The following illustrations are from Dyirbal:

Balan dᵛugumbil *baŋgul yaṛaŋgu* *balgan.* (7)
woman-ABSOLUTIVE man-ERGATIVE hit
'The man hit the woman.'

ŋadʸa *ŋinuna* *balgan.* (8)
I-NOMINATIVE you-ACCUSATIVE hit
'I hit you.'

ŋayguna *baŋgul yaṛaŋgu* *balgan.* (9)
I-ACCUSATIVE man-ERGATIVE hit
'The man hit me.'

ŋadʸa *bayi yaṛa* *balgan.* (10)
I-NOMINATIVE man-ABSOLUTIVE hit
'I hit the man.'

Although the most spectacular evidence for the relevance of animacy in the A does seem to come from Australian languages, it is also found in other languages. For instance, in some North-East Caucasian languages (e.g. Lak), nouns have an ergative–absolutive case marking system, but personal pronouns have a neutral system. This is particularly interesting in that it goes against an otherwise largely valid generalization that pronouns tend to distinguish more categories than do nouns.

The restriction of accusative marking to nouns that are high in animacy is very widespread across the languages of the world, and we will limit ourselves to a few examples. Even English provides relevant data here, since it has a nominative–accusative distinction with (many) pronouns, e.g. *I – me*, whereas it does not have any comparable distinction for other noun phrases. A particularly clear set of instances is provided by the Slavonic languages, where animacy is one of the key parameters determining whether a noun phrase will have a separate accusative case or not. In Russian, for instance, masculine singular nouns of the declension Ia have a separate accusative case (with the ending *-a*) if animate, but not otherwise:

Ja videl mal'čik-a/begemot-a/dub/stol. (11)
'I saw the boy/hippopotamus/oak/table.'

In Russian, all animate nouns in the plural have a separate accusative case, while no inanimate nouns do. In Polish, only male human nouns have a special accusative case in the plural, instantiating a different cut-off point on the animacy hierarchy:

Widziałem chłopców/dziewczyny/psy/dęby/stoły. (12)
'I saw the boys/girls/dogs/oaks/tables.'

The forms of the last four nouns are identical with the nominative plural, whereas the nominative plural of 'boys' is *chłopcy*.

There are data from a wide range of languages for special marking of definite direct objects: again, a few examples will suffice. In Turkish, only definite direct objects take the special accusative case suffix *-ı* (or its vowel harmony variants), all other direct objects being in the same suffixless form as is used for subjects (A or S):

Hasan öküz-ü aldı. (13)
Hasan ox ACCUSATIVE bought
'Hasan bought the ox.'

Hasan bir öküz aldı. (14)
Hasan a ox bought
'Hasan bought an ox.'

(In Turkish, *Hasan öküz aldı* is also possible, although it leaves open now many oxen were bought, i.e. 'Hasan bought an ox or oxen'.) In Persian, the suffix *-rā* is used to indicate definite direct objects:

Hasan ketāb-rā did. (15)
Hasan book ACCUSATIVE saw
'Hasan saw the book.'

Hasan yek ketāb did. (16)
Hasan a book saw
'Hasan saw a book.'

(As in Turkish, Persian also allows *Hasan ketāb did* 'Hasan saw a book or books'.)

What is particularly interesting in this respect is that some languages, in determining whether or not a P is to take the special accusative form or not, use both parameters of animacy and definiteness. In Hindi, for instance, a human direct object will normally take the postposition *ko* whether or not it is definite; only occasionally, and with affective value, does one find indefinite human noun phrases without *ko* in P position. Non-human, especially inanimate, Ps, however, never take *ko* if they are indefinite, though they may, and usually do, take *ko* if they are definite:

Aurat bacce ko bulā rahī hai. (17)
woman child ACCUSATIVE calling PROGRESSIVE is
'The woman is calling the/a child.'

?Aurat baccā bulā rahī hai. (18)

(The oblique form *bacce*, of *baccā*, is automatic before a postposition.)

Un patrõ ko paṛhie. (19)
those letters ACCUSATIVE read-POLITE
'Please read those letters.'

Ye patr paṛhie. (20)
these letters read-POLITE
'Please read these letters.'

Patr likhie. (21)
letters write-POLITE
'Write letters please.'

Thus, in order to know whether to assign *ko* to a P in Hindi, one must weigh against one another its position on both animacy and definiteness hierarchies, and even then there is room in the middle for subjective judgement.

A somewhat similar situation is observed in Spanish, in connection with the use of *a* to mark certain direct objects. Normally, this preposition is only used for human Ps, but such Ps must moreover be high in definiteness: in particular, human Ps that are non-specific occur without the preposition:

El director busca el carro/al empleado/a un empleado/un empleado. (22)
'The manager is looking for the car/the clerk/a (certain) clerk/a clerk.'

In this example, the difference between *a un empleado* and *un empleado* in P position is that the former implies that there is some specific individual that the manager is seeking, whereas the second implies simply that he needs any clerk.

Although we have treated animacy and definiteness as if they were unproblematic categories in the brief preceding discussion, this is in fact far from the case. In chapter 9, we return to examining animacy in more detail, but to conclude the present chapter we will turn to some problems concerning definiteness. One problem when we compare categories across languages is that we should have some basis on which to identify the same category in different languages. Thus, if we say that definite direct objects go into the accusative case in both Turkish and Persian, then we should be able to justify using the same term definite in referring to both these languages, and also to English, where the category definiteness exists but does not condition case marking. Failure to ensure this cross-language comparability would mean that we are not doing language universals research, but are simply analysing each language as an independent unit – and, unlike those linguists who maintain that this is the only way to study languages, we would be doing so surreptitiously by pretending, through use of the same term, that our results are comparable across languages. We will show below that a problem of this kind seems to arise in connection with definiteness, but that a solution to this problem is in fact forthcoming, a solution which, moreover, actually strengthens the universal base of our discussion.

The problem is that certain Ps in Persian and Turkish stand in the accusative case even though they are clearly not definite. In Persian, for instance, if one wants to say 'give one of them to me', then although the noun phrase 'one of them' is clearly, by definition, indefinite, yet still Persian here requires the definite marker *-rā*:

Yeki az ānhā -rā be man bedehid. (23)
one of them ACCUSATIVE to me give

In sentences (14) and (16) we illustrated the absence of the accusative marker in Persian and Turkish with the indefinite article *yek* or *bir*. However, although the direct object introduced by the indefinite article is clearly indefinite, both languages allow the accusative suffix here, so that the full range of data is actually:

Hasan bir öküz aldı. (24)

Hasan bir öküz-ü aldı. (25)

Hasan yek ketāb did. (26)

Hasan yek ketāb-rā did. (27)

The existence of the second example in each language might seem to quash any possibility of identifying the concept called definite in these languages with that called definite in the discussion of English.

An indication of the route out of this dilemma is, however, indicated by our discussion of animacy. Animacy is clearly not a single dichotomy between animate and inanimate, but rather a continuum along which we can range entities according to their degree of animacy, so that for instance people are more animate than animals, and animals more animate than inanimate objects. In describing definiteness cross-linguistically, we can make use of a similar notion of continuum, i.e. a continuum of definiteness (or specificity). Definiteness in the highest degree means, as in English, that the speaker presupposes that the hearer can uniquely identify the entity being spoken of. In Persian example (23) we are clearly not dealing with definiteness in this extreme degree, rather what is at issue is that the referent of the noun phrase has been delimited by specifying a certain set, which can be identified (namely *ānhā* 'them'), and then indicating that the entity which is to be given, while not uniquely identifiable, must still be a member of this identifiable set. This can be described by the term definite superset, meaning that the identity of the entity is not determinable absolutely, but some headway can be made in identifying it because it must be a member of a delimited set.

Turkish example (25) and Persian example (27) represent a different realization of the notion degree of definiteness/specificity. Although both members of each pair of sentences in (24)–(27) are translated the same way into English, they are far from equivalent in the original languages. The versions with the accusative marking on the P noun phrases suggest that the reference of the noun phrase in question is important, relevant for the discourse as a whole. In other words, in a discourse that started with (25) or (27) we would expect the ox or the book to recur in the discourse. The versions without the accusative suffix, however, are quite neutral in this

respect, and could be used, for instance, in simply relating the various events that happened to Hasan, without any particular interest in the ox or the book. We can refer to this distinction as relevance of referent identification. The absence of the accusative suffix advises the hearer not to bother about identifying the referent, while presence of this suffix advises him that the referent of this noun phrase, though not yet determinable by the hearer, will be of relevance to the ensuing discourse. So all uses of the accusative case can be linked together in terms of a hierarchy of definiteness: at one extreme we have complete identifiability of the referent; further down the hierarchy we have partial identifiability (definite superset); and further down still we have indication that identification of the referent is relevant; at the bottom, identification of the referent is neither possible nor relevant. If we then compare accusative case marking in Persian and Turkish with definiteness (say, the occurrence of the definite article with common nouns) in English, then we see that the same parameter is involved throughout, only the cut-off points are different in the various languages.

6.3 SUMMARY

To conclude this chapter, we may note that case marking, which has so often been viewed as an area of language-specific idiosyncrasy, often lacking in generalization even internal to a single language, can be the subject of fruitful language universals, fruitful not only in the sense that they involve cross-language generalizations about case marking, but also because they point the way to more adequate analyses of other areas of language structure.

NOTES AND REFERENCES

The discussion of the five homogeneous systems for case marking of S, A, and P is taken from Comrie (1978b, 330–4). The Wanggumara data are discussed by Blake (1977, 11; 1986, 21–2). The A/P – S system, considered unattested by Comrie (1978b), is documented by Payne (1979, 443) for Roshani. For Mojave case marking, see Munro (1976, 18).

The presentation in section 6.2 stems from some of the ideas contained in Comrie (1978b, 384–8), as modified by the similar results obtained independently by Silverstein (1976). The explanation has been modified slightly in the direction of ideas presented in DeLancey (1981). In particular, as noted by DeLancey, and also by Hopper & Thompson (1980), it is misleading to claim that Ps are typically inanimate/indefinite, rather than

just less animate/definite than As. The structure of the verb in Fox is discussed by LeSourd (1976).

Much of the discussion in section 6.2.2 follows Comrie (1977b, 1978c, 1979b); many of the data are taken from these articles and sources cited there. The ergative in Hua is discussed by Haiman (1979, 59–61; 1980, 360–64). The Australian data are from Silverstein (1976), Heath (1976) (for Ritharngu), Blake (1977, 13–15), and Dixon (1972, 59–60). For the closing remarks on a continuum of definiteness, see further Comrie (1978a).