

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Lingua xxx (2005) xxx–xxx

Lingua

www.elsevier.com/locate/lingua

Gender meets the Usage-Based Model: Four principles of rule interaction in gender assignment

Tore Nessel*

Center for Advanced Study in Theoretical Linguistics, University of Tromsø, N-9037 Tromsø, Norway

Received 29 August 2003; received in revised form 15 April 2004; accepted 14 June 2004

Abstract

Exploring four principles of gender assignment from the perspective of Langacker's [Langacker, 1991, Langacker, 1999] Usage-Based Model, the present article has important implications both for theories of the way gender is assigned and for the Usage-Based Model itself. The model simultaneously facilitates the implementation of principles of rule ordering and "rule counting" and thus provides a unified account of these approaches, which have generally been held to be antagonistic. However, in addition to discussing the implementation and interaction of principles proposed by other students of gender assignment, the present study also introduces the *Core Semantic Override Principle*. While in general the proposed analysis lends support to the Usage-Based Model, the discussion of the Core Semantic Override Principle motivates certain amendments.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Gender assignment; Usage-Based Model; Cognitive linguistics

While the category of gender has been widely discussed among linguists in recent years, the relationship between theories of gender assignment and contemporary linguistic theories is an area that has received less attention in scholarly literature. In the present paper I relate the problem of rule interaction in gender assignment to the Usage-Based Model, a model proposed by Langacker (1991, 1999) within the broader

* Tel.: +47 77645633; fax: +47 77646340.

E-mail address: tore.nesset@hum.uit.no.

framework of cognitive linguistics. The purpose of the present study is not to analyze new data or to provide full-fledged analyses of particular gender systems. Instead, I shall explore the implementation in the Usage-Based Model of four principles, viz. Gender Tally (Steinmetz, 1986, see also Doleschal, 2000), Elsewhere Condition (Kiparsky, 1982), default hierarchies (Steinmetz, 1986, this volume and Rice, this volume) and what I shall refer to as the “Core Semantic Override Principle”. My contribution can be summarized as follows:

- Gender Tally falls out as an automatic consequence of the architecture of the Usage-Based Model.
- In the same way, ordering of rules by the Elsewhere Condition comes for free in the Usage-Based Model.
- The Usage-Based Model offers a unified account of “rule counting” and rule ordering since Gender Tally and the Elsewhere Condition derive from the same principle in the model.
- While it is technically possible to implement default hierarchies in the model, the approach is at variance with the model’s strong emphasis on low-level schemas.
- Corbett and Fraser’s (2000a) claim that semantic rules take precedence universally in gender assignment appears to be too strong. As an alternative, I advance the Core Semantic Override Principle.
- Since the Core Semantic Override Principle arguably does not follow from the architecture of the model, it is suggested that the model must be supplemented with substantial principles of rule ordering.

On a more general level, the application of the Usage-Based Model to the field of gender assignment seems fruitful both for the model and the field. The Usage-Based Model brings together principles that have hitherto been considered independent stipulations. This lends additional support to the principles as such, while at the same time providing support for the model. However, as a result of the application of the Usage-Based Model to gender assignment, the need for certain extensions of the model also become apparent.

After a brief introduction of the Usage-Based Model in Section 1, Section 2 discusses the relationship between the model and “rule counting”. I then address formal principles of rule ordering in terms of the Elsewhere Condition (Section 3) and default hierarchies (Section 4). Section 5 explores substantial principles of rule ordering, in particular the Core Semantic Override Principle, before the contribution of the paper is summarized in Section 6.

1. Cognitive linguistics and the Usage-Based Model

Cognitive linguistics is a family of broadly compatible theoretical approaches that share the fundamental assumption that analyses of language should reflect research in other cognitive activities. Well-known versions of cognitive linguistics are found in the work on metaphor by George Lakoff and his associates (cf. e.g. Lakoff and Johnson, 1981, 1999). In the following, however, we shall be concerned with Ronald Langacker’s version of cognitive linguistics, more specifically with what he calls the Usage-Based Model




Fig. 1. Categorization network.

(Langacker, 1991, 1999; Kumashiro, 2000; see also Barlow and Kemmer, 2000 and Bybee, 2001 for related frameworks with the same name). Fundamental in Langacker's work is what he refers to as the "content requirement":

- (1) Content requirement:
 The only structures permitted in the grammar of a language [...] are
 (1) phonological, semantic or symbolic structures that actually occur in
 linguistic expressions; (2) schemas for such structures; and (3) categorizing
 relationships involving the elements in (1) and (2). (Langacker, 1987: 53f)

The content requirement can best be explained by referring to the categorization network given in Fig. 1, which concerns some Norwegian words denoting vehicles. The boxes represent schemas. They are essentially Saussurean signs combining content and form or, as Langacker puts it, semantic and phonological information. In the figure, I have separated the semantic and phonological information by slashes. The schema for *personbil* 'private car' is activated whenever this word is uttered and thus represents a generalization over symbolic structures actually occurring in linguistic expressions. The schemas are connected in terms of two types of categorization relation. The solid arrows represent what Langacker calls instantiation. They connect compatible schemas of different degrees of specificity. The arrow points at the more specific schema. For instance the schema for a private car is more specific than that of car, since the former is an example of the latter, but not vice versa. The dashed arrows represent the second type of categorization relation, extension, which connects a prototype with peripheral members of the category. Arguably, a private car is a prototypical car, whereas police cars and racing cars are more peripheral types of car.¹

¹ For simplicity, the phonological information in the schemas in Fig. 1 is represented orthographically, since the details of the sound shape of the nouns in question do not bear on my argument. This practice is adopted throughout the paper. Notice that the schema for 'vehicle' includes no phonological information since there appear to be no generalizations to be made about the phonology of the words in question. Accordingly, the phonological pole in the schema is represented by suspension points.




Fig. 2. Assignment of gender to Russian *slédovatel'* 'investigator'.

Langacker's Usage-Based Model has been further developed by his student Fumiko Kumashiro in her doctoral dissertation (Kumashiro, 2000). She models the alternative strategies that speakers may employ when activating schemas in categorization networks. Instead of introducing the model on a purely theoretical level I would like to draw the reader's attention to a concrete example depicted in Fig. 2. Assume, for the sake of the argument, that a speaker wonders whether to assign masculine or feminine gender to the Russian noun *slédovatel'* 'investigator'. Accordingly, the model includes two candidates given at the bottom level of the figure. Throughout the paper I shall let full-drawn boxes represent masculine gender and dashed boxes feminine gender. In order to find out which candidate to select, the speaker evaluates them against the categorization network given in the rectangle that represents a fragment of Russian grammar.² As can be seen from the figure, I assume five schemas, which are related to each other and the candidates in terms of instantiation relations. Schema (a) is a representation of the phonological generalization that nouns in palatalized sonorants with non-final stem stress tend to be masculine in Russian. The capital sigma represents a stressed syllable, and the capital L represents any sonorant. The apostrophe following the L symbolizes palatalization. Although nouns with these phonological properties display a clear affinity towards masculine gender, there are exceptions to this generalization, e.g. the feminine *mébel'* 'furniture'. It is possible that speakers of

² The rectangle marked "Grammar" represents (a small subset of) the conventionalized units in the Russian language, as opposed to the candidates that are not conventionalized. Notice, that the notion of "grammar" should not be taken to suggest a clear-cut distinction between grammar and lexicon. In cognitive linguistics, grammar and lexicon are assumed to represent a continuum.

Russian are aware of feminine nouns of this type, and that they have included general schemas for them in their mental grammars. In order to represent this situation, I have included schema (d) in Fig. 2 as a feminine counterpart to the masculine schema (a). Schema (b) represents the semantic generalization that animate nouns tend to be masculine (unless they specifically denote females). Counterexamples to this generalization like the feminine *rys* ‘lynx’ are accounted for by means of schema (e) linking animates to the feminine gender. Schema (c) captures the morphological generalization that nouns ending in the agentive suffix *-tel’* are masculine. Since there are no feminine nouns containing this suffix in Russian, there is no dashed schema corresponding to (c) in the figure.

Further discussion of details pertaining to the schemas in Fig. 2 is not necessary as the figure only serves illustrative purposes and no specific claims about gender assignment in Russian are made at this point. Of importance in the present context, however, is the general question as to what determines the choice between competing candidates. Langacker (1991: 282, 1999: 105f.) suggests that two general factors are relevant, *inherent ease of activation* and *conceptual overlap*. The former is assumed to correlate with what Langacker calls entrenchment (Langacker, 1999: 105), which in turn is a function of type frequency. Other things being equal, then, a schema that covers a large class of items is more likely to be activated than a schema for a smaller class. Therefore, a candidate that is connected to the schema of a large class has an advantage in the competition with other schemas. In the figure, type frequency correlates with the thickness of the line in the relevant box. In other words, candidates associated with schemas with thick lines have better chances to survive in the competition than candidates connected with schemas with thinner lines. The rationale behind the second factor is that a schema in the grammar is more easily activated if it shows a high degree of overlap with the candidate. Thus, candidates that share many features with the potential categorizing structures in the grammar have better chances to survive in the competition than candidates sharing fewer features with the relevant schemas in the grammar. In Fig. 2 the full-drawn masculine candidate receives the best score on both inherent ease of activation and conceptual overlap as it is connected to three schemas that are all more entrenched than the competing two feminine schemas. The Usage-Based Model hence predicts that *sledovatel’* is masculine, a prediction that is borne out by the facts. For the convenience of the reader, I have supplied the winning candidate with a happy face, a practice that will be followed throughout the article.

The two factors, inherent ease of activation and conceptual overlap, sometimes point in different directions. A case in point is the assignment of gender to Norwegian *bonde* ‘peasant’, illustrated in Fig. 3. According to Trosterud (2001: 43), Norwegian nouns with umlaut in the plural are normally feminine. This generalization is captured by schema (c), where the symbol *Ö* represents umlaut. However, Trosterud points out that this is not a rule without exceptions, so on the assumption that the speakers are aware of the exceptions and have formed generalizations about them I have included the less entrenched schema (a) as a masculine counterpart to (c). (We shall return to the treatment of “exceptions” in the Usage-Based Model in Sections 2 and 4 below.) Another important rule in Trosterud’s (2001) analysis of Norwegian is the generalization that nouns ending in *-e* are feminine. This is captured by schema (d). However, if a noun in *-e* refers to a person, the noun is




Fig. 3. Assignment of gender to Norwegian *bonde* ‘peasant’.

masculine, as is shown by schema (b).³ Schema (c) is more entrenched than schema (a), and (d) is more entrenched than (b). The factor of inherent ease of activation therefore favors feminine gender. However, degree of conceptual overlap points in the opposite direction. While (a) and (c) involve the same information, schema (b) is more specific than (d), since (b) includes the additional piece of information that we are dealing with a person. Given that schemas (a) and (b) contribute more information, we have a higher degree of conceptual overlap between these schemas and the masculine candidate, than between (c) and (d) and the feminine candidate. The noun *bonde* is in actual fact assigned masculine gender. This suggests that the factor of conceptual overlap takes precedence in rule interaction when it conflicts with inherent ease of activation. Langacker’s (1991: 278ff.) analysis of the English past tense suggests that the picture may be somewhat more complex, and more research is clearly needed on the relative strength of conceptual overlap and inherent ease of activation. However, I am not aware of any counterexamples from gender assignment at least, so for the purposes of the present study, I shall regard the former as decisive whenever it conflicts with the latter.

2. The Usage-Based Model and “rule counting”

Enger (2002) characterizes the theory of gender assignment advocated in Steinmetz (1986, this volume) and Doleschal (1993, 2000) as a “rule counting approach” because it invokes a principle that Steinmetz (1986) calls Gender Tally:

³ The claim that personal nouns in *-e* are masculine is a simplification in that such nouns are invariably feminine if they denote *female* persons, e.g. *jente* ‘girl’, *dame* ‘lady’ and *kone* ‘wife’. The assignment of the correct gender to these nouns requires a schema that links female sex to feminine gender. I assume this schema takes precedence over schema (b) in Fig. 3 by the Core Semantic Override Principle to be discussed in Section 5 below.

(2) Gender Tally:

Count the number of times each gender is assigned and assign the noun the gender with the highest value. (Steinmetz, 1986: 193)

A principle to the same effect has been formulated in Doleschal (2000: 142). In order to see how this works, consider the assignment of gender to German *Gemüse* ‘vegetable’, for which Steinmetz (this volume) assumes the following rules to be relevant:

- (3) a. German nouns ending in *-e* are feminine.
- b. German nouns with the prefix *ge-* are neuter.
- c. Superordinate nouns in German are neuter.

The term “superordinate” in (3c) may require clarification. *Gemüse* is a superordinate noun in the sense that it is a cover term for the semantic field of vegetables (*Gurke* ‘cucumber’, *Bete* ‘beetroot’, etc.). According to Steinmetz’ rule (3c), such superordinate terms are neuter in German. Since in the case of *Gemüse* two rules point towards neuter (i.e. 3b–c) and one towards feminine gender (i.e. 3a), Gender Tally predicts neuter gender, a prediction that is borne out by the facts.

Before we relate the Gender Tally principle to the Usage-Based Model, a remark on the notion of “counting” is in order. It is a widely held belief that linguistic rules obey a locality principle, i.e. that their application is restricted to small and bounded domains. As observed by Hayes (1995: 307), locality is often interpreted in terms of restrictions on counting in grammar. In phonology, for instance, no language appears to posit a rule stressing, say, the fifth syllable from the end of the word. Likewise, no syntactic rule specifies the movement of a constituent past, say, four constituents. In this sense, languages do not count. Since Steinmetz’ theory has been referred to as a “rule counting approach” (Enger, 2002), it is legitimate to ask whether Gender Tally is at variance with the ban on counting in grammar. I think, however, that the question can be answered in the negative. Rice (this volume) demonstrates that Gender Tally can be represented in terms of violation marks in OT tableaux. Asking whether Gender Tally involves counting is therefore tantamount to asking whether the evaluation of candidates in OT tableaux presupposes counting. This problem is addressed in Prince and Smolensky (1993: 27) who point out that the number of violations associated with a candidate is of no relevance for the evaluation. All the evaluation does is to compare the set of violation marks for the relevant candidates in order to identify the candidate with the smallest set—the optimal candidate. This comparison can be carried out without involving counting, as explicated by Prince and Smolensky (1993: 76). Let us assume a competition between two candidates, each associated with a set of violation marks. In order to identify the smaller set, we start by establishing whether one of the sets is empty. In the case that one of them is, we have a winner. If not, one violation mark is removed from each set, before the comparison is carried out again. The procedure is repeated until an empty set is identified. Three conclusions can be drawn from this discussion. First, Gender Tally does not hinge on counting in the sense that reference to numbers like 4 or 7 is not presupposed. Secondly, Gender Tally is not more controversial than OT. Thirdly, despite its intuitive appeal, the term “rule counting” is in fact misleading as a characterization of the theories advanced by

Steinmetz (1986) and Doleschal (1993, 2000). Therefore, “rule counting” is included in quotation marks throughout this study. For further discussion of this issue, see Rice (this volume).

Since Gender Tally has not received much attention in the scholarly literature, it may be worthwhile pointing out that the principle receives support from connectionist processing (cf. e.g. McClelland and Elman, 1986). When a target (in our case a noun) activates certain units in a network (assignment rules in our case), one of the factors facilitating the selection of a certain unit is the amount of conceptual overlap. The higher the degree of overlap, the greater are the chances that a certain unit is selected. This is analogous to Gender Tally. When a majority of rules competes with a minority, the majority represents the higher degree of conceptual overlap. In other words, Gender Tally assigns gender on the basis of conceptual overlap. While this does not indicate that one has to believe in connectionism in order to adopt Gender Tally, the parallelism is nevertheless interesting.

In view of the parallelism between Gender Tally and connectionism, it should not come as a surprise that Steinmetz’ principle can be accounted for in the Usage-Based Model, which is compatible with connectionism (cf. Langacker, 1999: 105). Consider Fig. 4, which models the assignment of gender to German *Gemüse* in that each schema corresponds to one of Steinmetz’ rules in (3). Schema (a) represents the generalization that nouns in *-e* are feminine, while schemas (b) and (c) capture the affinity of the *ge-* prefix and superordinates to the neuter gender. Notice that for convenience boxes with dashed lines (---) stand for feminine gender, while a “dash-dot” pattern (-.-) represents neuter gender. Since the feminine candidate instantiates one schema and the neuter candidate two, the latter involves the highest degree of conceptual overlap. In Section 1, the assumption was developed that inherent ease of activation only comes into play when degree of conceptual overlap does not facilitate the selection of a unique candidate. Given this assumption, we do not need to take inherent ease of activation into account in Fig. 4. Degree of conceptual overlap points towards neuter, and therefore the Usage-Based Model correctly predicts the assignment of neuter to German *Gemüse*.




Fig. 4. Assignment of gender to German *Gemüse* ‘vegetable’.

Fig. 4 shows that degree of conceptual overlap has the same effect as Gender Tally, in that the gender supported by the majority of the schemas is selected. Notice that while Gender Tally is specifically designed to deal with gender, degree of conceptual overlap is a more general principle. It is pivotal in the Usage-Based Model, which seeks to model the strategies speakers may employ when activating schemas in categorization networks in general. Although the Usage-Based Model is primarily a *linguistic* model, it is based on principles that are assumed to be valid for processing and storage of information in cognition in general (Langacker, 1987). In other words, Gender Tally follows from very general linguistic and cognitive principles. Once the Usage-Based Model has been adopted, Gender Tally comes for free, in that Steinmetz' principle for gender assignment falls out as an automatic consequence of the very architecture of the model. This has two important corollaries. First, by Occam's razor Gender Tally need not be stipulated as a part of our theory of gender, since it follows from more general principles. On the other hand – and this is the second corollary – the fact that Gender Tally is in harmony with general cognitive and linguistic principles lends additional support to a “rule counting” approach to gender assignment.

Despite the close parallelism between the analyses of gender assignment in the Usage-Based Model and Steinmetz' “rule counting” approach, there are differences. As the reader will recall from Section 1, in the Usage-Based Model degree of conceptual overlap interacts with another factor, viz. inherent ease of activation. There is nothing, however, in Steinmetz' approach that corresponds to inherent ease of activation. A second look at the assignment of gender to German *Gemüse* is illustrative. While the rules in (3) represent clear tendencies, there are exceptions to all of them. Although most words in *-e* are indeed feminine, it is not too hard to find examples of neuters in *-e*, e.g. *Interesse* ‘interest’, *Auge* ‘eye’—as well as *Gemüse* itself. While the *ge-* prefix in general indicates neuter gender, nouns like *Gewalt* ‘violence’ and *Gemeinde* ‘community’ are feminine. Last but not least, feminine nouns like *Waffe* ‘weapon’ are exceptions from the rule that superordinate terms are neuter. Fig. 4 models a speaker's mental grammar without any general schemas for the exceptions. Whether this is because s/he is not aware of the exceptions at all, or has just failed to form any generalizations about them, is of no importance here. Let us, however, consider a situation where the speaker *has* formed general schemas for the exceptions. This yields the six schemas given in Fig. 5, where the entrenched schemas (a), (e) and (f) correspond to the rules in (3) and the schemas in Fig. 4, while the less entrenched (b), (c) and (d) represent the exceptions to these rules. Since in Fig. 5 three schemas point towards feminine and three towards neuter gender, the two candidates tie with regard to degree of conceptual overlap. This factor is therefore not sufficient to select a candidate alone, and inherent ease of activation comes into play. Since feminine gender gains support from only one highly entrenched schema, while neuter is favored by two such schemas, neuter is assigned. Thus, the Usage-Based Model correctly predicts that *Gemüse* is neuter in German.

Whether Figs. 4 and 6 offers the best analysis of the relevant part of the grammar of German is an empirical question that may be tested, e.g. in terms of psycholinguistic experiments. Indeed, it is possible that Fig. 4 best reflects the mental grammar of some speakers, while that of others may be better represented as Fig. 5. To settle this issue, however, is well beyond the scope of the present paper. However, the discussion of *Gemüse*




Fig. 5. Assignment of gender to German *Gemüse* 'vegetable' (second version).

illustrates not only the similarities, but also the differences between the Usage-Based Model and Steinmetz' "rule counting" approach. While the rules are treated in essentially the same way, the exceptions are handled differently. As we have seen, in the Usage-Based Model inherent ease of activation is pivotal in the account of exceptions, a factor that does not have a direct counterpart in Steinmetz' "rule counting" approach. The investigation of the empirical implications of this difference, however, must be left for future research.




Fig. 6. Assignment of gender to Ukrainian *poni* 'pony'.

3. The Usage-Based Model and rule ordering: the Elsewhere Condition

Having explored the relationship between the Usage-Based Model and “rule counting” in some detail, we now turn to the model’s relationship to rule *ordering*. We shall see that ordering by means of Kiparsky’s (1982) Elsewhere Condition is relevant for gender assignment, and that it falls out as an automatic consequence of the architecture of the Usage-Based Model. Interestingly, the Elsewhere Condition follows from the same factor that was seen to imply “rule counting”, viz. degree of conceptual overlap. Thus, the Usage-Based Model provides a unified account of “rule counting” and rule ordering, two approaches that have been considered antagonistic in the literature (cf. Enger, 2002).

Rule ordering – the phenomenon that a rule takes precedence over a conflicting rule – is well known from many areas of linguistics. The Elsewhere Condition regulates the order of application of rules of different degrees of specificity. If rule A refers to a proper subset of the nouns referred to by rule B, A takes precedence over B.⁴ (This takes place no matter whether A belongs to a majority of rules favoring a certain outcome.) While this principle was not devised in order to deal with gender in particular, it is clearly relevant for gender assignment also. Consider, as an example, the case of so-called indeclinable nouns in Ukrainian, i.e. nouns taking a zero ending throughout their inflectional paradigm. In Ukrainian, nouns of this type tend to belong to the neuter gender, e.g. *sari* ‘sari’ and *frykase* ‘fricassée’. However, indeclinable nouns denoting animates are masculine, e.g. *flamingo* ‘flamingo’ and *poni* ‘pony’ (Pugh and Press, 1999: 56f.). The following two rules capture these generalizations:⁵

- (4) a. Indeclinable nouns are neuter.
- b. Indeclinable nouns denoting animates are masculine.

Since indeclinable nouns denoting animates constitute a proper subset of indeclinable nouns, rule (4b) takes precedence over (4a) by the Elsewhere Condition.

The question now arises as to whether and how the Elsewhere Condition can be accounted for within the Usage-Based Model. Consider Fig. 6 where schema (a) corresponds to rule (4a) and schema (b) to rule (4b). Notice that an instantiation arrow leads from (a) to (b). This is because the two schemas are compatible, but (b) of higher specificity than (a). Now, since (b) contains more specific information than (a) the degree of conceptual overlap is greater between (b) and the masculine candidate than between (a) and the neuter candidate. Given that degree of conceptual overlap is the primary factor

⁴ The generalization that specific information takes precedence is also known as “Proper Inclusion Precedence” (Koutsoudas et al., 1974), “Wilensky’s law” (Lakoff, 1987) and “Panini’s Theorem” (Prince and Smolensky, 1993).

⁵ Notice in passing that there are some systematic exceptions to these rules. For instance, according to Pugh and Press (1999: 57) indeclinable nouns denoting languages tend to be feminine, e.g. *bengali* ‘bengali’ and *urdu* ‘Urdu’. Indeclinable common nouns like *madam* ‘madame’ and *ledi* ‘lady’, as well as a few indeclinable female first names like *Esfir* are feminine contra rule (4b) because Ukrainian has a general rule assigning feminine gender to nouns denoting females. In order to avoid unnecessary complications, the cases mentioned in this footnote will not be discussed in the following.

facilitating the selection of a candidate, the Usage-Based Model correctly predicts *poni* ‘pony’ to be masculine.

The implications of the *poni* example in Fig. 6 go well beyond the study of Ukrainian. As pointed out by Langacker (1999: 106), the parameter of conceptual overlap implies that lower-level schemas containing information of a high degree of specificity have a built-in advantage in the Usage-Based Model. The more specific a schema is, the more features it contains that can potentially be shared with a candidate. In other words, degree of specificity correlates with conceptual overlap, and hence facilitates the selection of a particular candidate. This is important, because it implies that the Elsewhere Condition does not need to be stated separately in the Usage-Based Model. It follows from the architecture of the model. The fact that specific information takes precedence is an epiphenomenon that follows from more fundamental cognitive properties of processing and storage.

It is interesting to notice that it is the parameter of degree of conceptual overlap that is responsible for the incorporation of the Elsewhere Condition in the Usage-Based Model. As we saw in the previous section, this parameter also enabled us to account for “rule counting” (gender tally) effects. Since both the Elsewhere Condition and “rule counting” follow from the same fundamental principle in the model, viz. degree of conceptual overlap, the Usage-Based Model provides a unified account of these phenomena and thus brings together “rule counting” and rule ordering. This is an important result, since in the (albeit limited) literature on “rule counting” (Steinmetz, 1986; Doleschal, 1993, 2000; Enger, 2002) its relationship to the Elsewhere Condition has not been acknowledged.

4. The Usage-Based Model and rule ordering: default hierarchies

In the previous section we saw how specific schemas take precedence over more general schemas by the Elsewhere Condition. In such situations, the more general schema acts as a local default, in the sense that it is the “unmarked” gender for a certain group of nouns (e.g. for indeclinable nouns in Ukrainian, for which neuter was shown to be the default). The local default gender is selected whenever no other rules apply to the nouns of the relevant type. In the following, we shall explore the notion of “global default gender”, i.e. the default gender of the language as a whole. The relevance of global defaults seems fairly uncontroversial in gender assignment (cf. e.g. Fraser and Corbett, 1997). However, Steinmetz (1986) has proposed an extension of this idea in that he assumes that each language has several global default rules that are organized in a hierarchy. As this idea has not received much attention in the literature on gender assignment, it seems fair to say that more data, i.e. more analyses of entire gender systems, is needed in order to establish whether default hierarchies have a role to play in gender assignment (see, however, Rice, this volume). Such evidence will not be presented here; in the following I shall limit myself to discussing whether and how default hierarchies are compatible with the Usage-Based Model. While we shall see that it seems technically possible to represent default hierarchies in the model, the proposal is somewhat problematic in the context of cognitive linguistics. Furthermore, the empirical predictions of Steinmetz’ framework and its adaptation to the Usage-Based Model are slightly different.

While Gender Tally presented in (2) in Section 2 is the most important principle in Steinmetz' approach, this concept does not always enable us to assign gender to a noun. A case in point is German *Frucht* 'fruit', for which Rice (this volume) assumes the following assignment rules to be relevant:⁶

- (5) a. German nouns ending in /uxt/ are feminine.
- b. German superordinate nouns are neuter.

Both rules apply to *Frucht*, which not only has the right phonological shape, but also is a superordinate term subsuming the names of the various fruits in German. Since one rule points towards the feminine, and one towards the neuter gender, Gender Tally does not facilitate the selection of either gender. (Obviously, masculine is not assigned given that no rules favor this gender.) In order to account for cases like this, Steinmetz invokes the following default hierarchy:

- (6) Default hierarchy for Modern German:
 Masculine $\gg$ Feminine $\gg$ Neuter ($\gg$ = "outranks")

The prediction of this hypothesis is that whenever we face a tie in German between the masculine and another gender, masculine gender is assigned. In cases of ties between the feminine and the neuter, feminine gender is assigned. This latter situation applies to *Frucht*. Since Gender Tally fails to select a gender, and feminine is ranked above neuter in the default hierarchy, Steinmetz' approach predicts assignment of feminine gender to *Frucht*. This is confirmed by the data—*Frucht* is indeed feminine in German.

A global default rule for gender assignment is essentially a rule stating that a certain gender is selected whenever no specific information points towards another gender. In the context of the Usage-Based Model we might represent this by means of a maximally general schema, i.e. a schema containing no information except the reference to the relevant gender. Examples are the "empty" schemas (a), (c) and (e) in Fig. 7 which assign masculine, feminine and neuter gender, respectively, without invoking any specific information favoring these genders.⁷

The question now arises as to whether and how the hierarchical organization of global defaults can be implemented in the Usage-Based Model. Despite their central position in Steinmetz' approach, he provides no detailed theoretical discussions of the principles on which default hierarchies are based. A major criterion, however, appears to be type frequency. At least, Steinmetz (this volume) advances statistics on the relative frequency of the genders in order to motivate the ranking of default genders. The gender with the highest

⁶ Köpcke (1982: 97f.) states this rule somewhat differently in that he assigns feminine gender to nouns ending in a non-sibilant fricative followed by /t/. As pointed out by an anonymous referee, there are several non-feminine nouns in /axt/, /ext/ and /ixt/, so I prefer the more restrictive version of the rule referring to nouns in /uxt/ only.

⁷ To refer to the schemas as "empty" is a simplification as they are not completely devoid of any information. For instance, in order to capture the generalization that, e.g. verbs are not specified for gender, the schemas must include information about parts of speech. However, since information of this sort is outside the focus of the present study, it has been omitted from the schemas throughout the paper. The default schemas therefore appear empty.




Fig. 7. Assignment of gender to German *Frucht* 'fruit'.

type frequency, i.e. the gender with most members, receives the highest ranking. Since the central notion of entrenchment reflects frequency, it is possible to represent default hierarchies in the Usage-Based Model as empty schemas with varying degrees of entrenchment. Thus for modern German I suggest representing the hierarchy in (6) by assuming the masculine-assigning schema (a) to be more entrenched than the feminine (c), which in turn outdoes the neuter (e) with regard to entrenchment.⁸

Let us now consider Fig. 7 in more detail. We have three candidates, one for each gender. In order to select a winner, we must first consider the degree of conceptual overlap between each candidate and the schemas in the grammar. The masculine candidate cannot win because it only instantiates the empty schema (a), while the feminine and neuter candidates instantiate two schemas each. In addition to the empty default schemas, the feminine candidate is connected to schema (b), which represents rule (5a), and the neuter candidate to schema (d), which corresponds to rule (5b). It seems reasonable to assume the simple schemas (b) and (d) to involve the same level of specificity. Since schemas (c) and (e) also contain the same amount of information, namely zero, we seem to be facing a tie between the feminine and neuter genders in terms of degree of conceptual overlap. Given that degree of conceptual overlap turns out to be indecisive, we turn to the second factor, i.e. inherent ease of activation. In the figure I have considered schemas (b) and (d) to be on approximately the same level with regard to entrenchment since both schemas cover restricted and not very large classes of nouns. Therefore, the entrenchment of schemas (b) and (d) does not enable us to pick a winner. However, given that schema (c) is more entrenched than (e) because there are more feminines than neuters, the feminine candidate

⁸ According to Köpcke (1982: 45), who provides relevant statistics for monosyllabic nouns, masculine is the gender with most members in German. While there are more neuter than feminine monosyllabic nouns in Köpcke's corpus, it seems reasonable to assume that feminine is more frequent than neuter in the language as a whole. As Köpcke points out, the German feminine comprises nouns with extremely productive suffixes like *-e*, *-ung* and *-heit*, which are only found in polysyllabic nouns.

is preferable in terms of inherent ease of activation. The Usage-Based Model therefore correctly predicts feminine gender for *Frucht*.

While specific details concerning the German example could be disputed, Fig. 7 illustrates that it is technically possible to represent default hierarchies in the Usage-Based Model if the hierarchical ordering is interpreted in terms of differences in entrenchment. However, the proposed solution is open to critique. First, adopting default rules for each gender in a language and assuming a hierarchical ordering of them implies supplementing the theory of gender assignment with extensive machinery. In other words, the proposal may be attacked for its lack of economy. As pointed out by Langacker (1991: 262f., 1999: 91f.), however, arguments based on economy generally carry limited weight in the Usage-Based Model. The reason is the model's emphasis on psychological reality. If empirical investigation should suggest that speakers prefer complex grammars with massive redundancy, the best linguistic description is one that reflects this state of affairs accurately—even if it is not very economical. In other words, in the Usage-Based Model economy is considered an empirical, not a methodological issue.

From the perspective of the Usage-Based Model a more serious challenge concerns the importance of maximally general, high-level schemas in the implementation of default hierarchies sketched above. As suggested by the name, the Usage-Based Model attaches primary importance to low-level schemas covering small classes of items actually occurring in language use. Very general, high-level schemas for large classes enjoy a more peripheral status in the model. However, in the account of default hierarchies outlined above, the maximally general “empty schemas” are in fact crucial for the assignment of gender. While Fig. 7 illustrates that an analysis in terms of default hierarchies can be made to work, it certainly goes against a fundamental property of the model, viz. its strong focus on low-level schemas.

It is also worth noticing that there are differences between the Usage-Based Model and the approach to gender assignment advocated by Steinmetz. First, in Fig. 7 I made the assumption that the competing schemas (b) and (d) involve approximately the same level of entrenchment. However, this is a simplification, at least if we ignore compounds; Mater (1970) lists only eight simplex words in /uxt/, a number that is clearly smaller than the number of superordinate terms in the language. Now, if the difference in entrenchment is big, it may have an impact on gender assignment. Assume, for the sake of the argument, a situation where schema (d) displays such a high degree of entrenchment that the two schemas (d) and (e) will receive a higher score on inherent ease of activation than (b) and (c), and hence facilitate the selection of the neuter candidate instead of the feminine. A detailed discussion of the relative strength of schemas (b) and (d) in German is beyond the scope of the present study. However, the example illustrates an important difference between the Usage-Based Model and Steinmetz' approach. In the former, the entrenchment of other schemas may interfere, while such interference is not expected in the latter since entrenchment has no formal status in this model.

Another difference between the two models concerns schemas for exceptions. In Fig. 7 I have tacitly assumed that the mental grammar of the speakers do not involve any schemas over exceptions to rules (5a) and (5b) above. Consider, however, a situation where the speakers are aware of exceptions like *Waffe* ‘weapon’ to rule (5b) (but not of any exceptions to (5a) which seems not to encounter any). This situation is depicted in Fig. 8. Here the




Fig. 8. Assignment of gender to German *Frucht* 'fruit' (second version).

feminine candidate outdoes the neuter candidate with regard to degree of conceptual overlap, since the former instantiates three schemas, and the latter two. Therefore, we are in a position to select a candidate without invoking inherent ease of activation, and hence without reference to the default hierarchy. We shall not discuss whether Fig. 8 offers a more realistic model than Fig. 7 of the mental grammars of native speakers. The purpose of Fig. 8 is only to illustrate the differences between the Usage-Based Model and Steinmetz' approach as far as exceptions are concerned.

While this section has shown that it is possible to represent default hierarchies in the Usage-Based Model, we have seen that the approach is at variance with the model's strong emphasis on low-level schemas. Furthermore, it has been demonstrated that the Usage-Based Model is slightly different from Steinmetz' framework. Investigation of the empirical implications of these differences represents an interesting topic for future research, but cannot be attempted in the present study.

5. The Usage-Based Model and rule ordering: the Core Semantic Override Principle

So far we have seen that the Usage-Based Model enables us to implement both "rule counting" (Gender Tally) and rule ordering (Elsewhere Condition and—with certain qualifications—default hierarchies). The ordering principles considered hitherto concern the logical relationship between rules (or schemas). For this reason I shall refer to them as *formal* principles of rule ordering. In the following I shall suggest that formal principles

must be supplemented by what I shall call *substantial* ordering principles, i.e. principles referring to particular types of information, not to the relationships between rules or schemas. The substantial principle I shall explore in the present section will be referred to as the “Core Semantic Override Principle”.

The problem to be addressed in this section can be illustrated by means of an example from Arapesh, a Toricelli language spoken on the north coast of Papua New Guinea.⁹ As argued by Corbett and Fraser (2000a), this language displays a strong correlation between gender and inflection class. In most cases, membership in one of the 13 genders in this language can be predicted on the basis of a noun’s declension. However, nouns denoting male or female beings constitute systematic exceptions to this general pattern. A case in point is *nakor* ‘husband’s father’ discussed by Dobrin (1999). Since this noun ends in /r/, the prediction is that it belongs to declension 18 and therefore is assigned to gender X. However, on the other hand, there is a rule in the language that assigns gender VII to nouns denoting male beings. This gender resembles the masculine genders in Indo-European languages in that it includes all animals and human beings of male biological sex. In actual fact, *nakor* belongs to the “masculine” gender VII. Thus, the conflict between the two assignment rules is resolved in favor of the gender assigned by the semantic rule referring to biological sex. The question is why.

Dahl (2000) observes that examples of this type are not isolated. In fact, there is solid typological evidence in favor of a privileged position of sex-based rules in gender assignment. According to Dahl (2000: 101f.), sex is the “major criterion” for gender assignment of gender in languages with more than one gender for animates. While Dahl’s term “major criterion” may seem opaque, it is clear from his discussion that it implies that sex-based gender assignment tends to take precedence. Notice that the provision “tends to” does not indicate that we are dealing with a mere statistical generalization. Rather, the set of cases where sex-based rules are overridden is limited and well defined. Dahl (2000: 103) isolates the following:¹⁰

- (7)
 - a. Special morphological rules may take precedence for augmentative and diminutive derivations.
 - b. Special semantic rules may take precedence for nouns denoting young or small animates.
 - c. Special semantic rules may take precedence for certain kinds of animals.
 - d. The “wrong” gender may be used in order to obtain special rhetorical effects (“downgrading” and “upgrading”).

German diminutives in *-chen* and *-lein* are well known examples of (7a). As an illustration of special treatment for nouns denoting young or small animates in (7b), Dahl (2000: 103) mentions the assignment of neuter gender to unmarried women in certain

⁹ The gender system of Arapesh is discussed in Aronoff (1994), Fraser and Corbett (1997), Corbett and Fraser (2000a) and Dobrin (1997, 1999). Data are from Fortune (1942/1977), but Dobrin has also carried out fieldwork on Papua New Guinea.

¹⁰ Dahl also mentions arbitrary exceptions, but I have not included that in the list in (7) since we are interested in the systematic properties of gender systems.

Polish dialects (see also Corbett, 1991: 100). As for (7c), in the Australian language Ngangikurrunggurr nouns denoting animals hunted for meat are relegated to a special gender (Corbett, 1991: 140). Finally, the special effects obtained by the use of *s/he* about inanimate objects and *it* about humans in American English serve to illustrate downgrading and upgrading in (7d) (Dahl, 2000: 105). A detailed discussion of cases of these types is beyond the scope of the present study. Suffice it to say that Dahl's typological evidence strongly suggests that sex-based rules take precedence universally in gender assignment, with the exception of the four well-defined cases in (7). For explicitness, I suggest formulating the following principle:

- (8) The Core Semantic Override Principle:
Rules referring to biological sex take precedence in gender assignment.

I refer to (8) as the “Core Semantic Override Principle” because biological sex is of fundamental importance for the category of gender.

Could the principle in (8) have been stated more inclusively so as to embrace all semantic rules, not only those involving biological sex? In his seminal monograph, Corbett (1991: 68f.) claims that “[i]f there are conflicting factors at work, semantic factors usually take precedence”. In a more recent work, Corbett and Fraser (2000a) adopt an even stronger position:

- (9) As is universally the case, the formal gender assignment rules [...] are dominated by the semantic gender assignment rules.
(Corbett and Fraser, 2000a: 321)

While this seems correct for languages like Arapesh with a strong correlation between declension and gender, languages where this correlation is less strong appear to be problematic. Examples include Germanic languages like German (Köpcke and Zubin, 1984, 1995 and references therein) and Norwegian (Trosterud, 2001; Enger, 2002). Since in languages of this type the morphological rules cover a smaller portion of the vocabulary, researchers have postulated numerous semantic rules, not all of which refer to biological sex. For instance, Trosterud (2001) assumes 28 semantic rules for Norwegian, although some of them cover relatively small classes of nouns. It seems fair to say that at present the interaction of semantic and other rules in complex systems of this type is not well understood. Trosterud (2001) explicitly avoids making strong claims about rule interaction on the grounds that the rules themselves are not sufficiently well understood. In a similar vein, Köpcke and Zubin (1984: 42) find it premature to draw conclusions on the question as to whether semantic rules always take precedence in gender assignment.

As an example of a somewhat problematic case for Corbett and Fraser's position, let us consider the interaction of the following three assignment rules for German (after Steinmetz, 1986: 190):

- (10) a. Superordinate nouns are neuter (e.g. *das Möbel* ‘furniture’).
b. Nouns in *-e* are feminine (e.g. *die Hose* ‘trousers’).
c. Nouns in /uxt/ are feminine (e.g. *die Bucht* ‘bay’).




Fig. 9. Assignment of gender to Arapesh *nakor* 'husband's father'.

From Corbett and Fraser's position we would expect the semantic rule (10a) to override the morphological (10b) and the phonological (10c). However, despite this superordinate nouns like *die Waffe* 'weapon' and *die Pflanze* 'plant' are feminine in accordance with (10b) and *die Frucht* 'fruit' in accordance with (10c). Clearly, we cannot draw strong conclusions on the basis of such examples without being sure that the rules in (10) are actually correctly stated. Furthermore, we cannot know whether the rules are correct before we have considered the German gender system in its entirety. Nevertheless, the German examples suggest that Corbett and Fraser's proposal may be too strong, and that it may be wise to adopt a somewhat more cautious position. Until a fuller understanding of the rule interaction in languages like German is arrived at, I suggest adopting the Core Semantic Override Principle in (8).

The question now arises as to whether and how the Core Semantic Override Principle can be implemented in the Usage-Based Model. Let us return to the example of *nakor* 'husband's father' in Arapesh. As a first approximation to the problem, consider Fig. 9 where schema (a) represents the generalization that nouns denoting males belong to gender VII, whereas schema (b) captures the generalization that declension 18 nouns ending in /r/ are in gender X.¹¹ The "masculine" gender VII is represented in terms of solid lines in the figure, while gender X is given as a "dash-dot" pattern. Here, there is one schema pointing towards each gender, so the factor of conceptual overlap does not enable us to select a winning candidate. As for entrenchment (inherent ease of activation), the situation is uncertain, since the written sources do not indicate the relative size of the relevant classes of nouns. However, semantic assignment consistently takes precedence in Arapesh, and

¹¹ For descriptive convenience, in the Fig. 1 represent this generalization not by invoking declension 18, but rather the final /r/ characteristic of this declension class. This enables us to avoid an irrelevant discussion of the representation of inflection class in the Usage-Based Model.

Lise Dobrin (p.c.) informs me that it appears to do so regardless of the size of the classes of nouns in question. Thus, Arapesh displays examples where the right gender cannot be assigned on the basis of conceptual overlap and entrenchment, and so the Usage-Based Model is in jeopardy. The inability to choose a winner is represented by the absence of a happy face under both candidates in the figure.

Is it possible to solve the problem posed by the Core Semantic Override Principle without introducing extra machinery in the Usage-Based Model? It may be worthwhile to consider an attempt in terms of default hierarchies, although I shall conclude that it is not viable. Recall from Section 4 that according to Steinmetz' (1986) theory, the gender higher up in the default hierarchy is selected in cases of ties. According to Fraser and Corbett (1997, 2000a), the global default in Arapesh is gender VIII, so this gender must be on the top of the hierarchy. In order to make the analysis work, we must furthermore assume gender VII to outrank gender X. This yields the following partial, hypothetical default hierarchy for Arapesh:

- (11) Partial Default Hierarchy for Arapesh:
VIII $\gg$... VII $\gg$... X

An analysis based on this hierarchy is sketched in Fig. 10. Each candidate is instantiated by two schemas. While it seems impossible to select a winning candidate on the basis of degree of conceptual overlap, however, the factor of inherent ease of activation enables us to choose gender VII. In Section 4 I implemented default hierarchies in terms of frequency/entrenchment. Since VII outranks X in the hierarchy postulated in (11), I have represented the empty schema (a) as more entrenched than schema (c). Therefore, gender VII is favored




Fig. 10. Assignment of gender to Arapesh *nakor* 'husband's father' (second version).

with regard to inherent ease of activation, and the candidate of gender VII is chosen as the winner.

Do analyses in terms of default hierarchies offer an adequate solution to the problem of why sex-based rules universally take precedence in gender assignment? The analysis of Arapesh in Fig. 10 hinges on the assumption that gender VII outranks X in the default hierarchy. Is there any independent evidence in favor of this assumption? While Steinmetz (1986, *this volume*) does not give a detailed discussion of the criteria for establishing default hierarchies, in Section 4 I suggested an interpretation in terms of type frequency in the context of the Usage-Based Model. A gender with more members receives a higher ranking than a gender with fewer members. As mentioned, however, the written sources do not indicate the relative size of the genders in Arapesh, so at present there is not necessary support for the hierarchy postulated in (11) above. Without independent evidence to back up the hierarchy, the analysis only states that gender VII takes precedence over gender X—which is exactly what one wants to explain. Clearly, more data is needed in order to make the analysis more than an ad hoc solution.

On a more general level, the example from Arapesh shows that if it is possible for all cases where sex-based rules take precedence to establish adequate default hierarchies and come up with independent evidence in favor of them, the “Core Semantic Override Principle” might turn out to be an epiphenomenon. In the context of the Usage-Based Model discussed in the present study the prediction would be that genders encompassing nouns denoting males or females universally have more members than competing genders. However, even if this prediction turned out to be true for all gender systems of the world’s languages, I find it somewhat hard to believe that the universally consistent override of sex-based rules in gender assignment conflicts is a mere coincidence conditioned by the relative sizes of genders. In any case, the onus of proof is on those who would want to argue this. Hence, at our present level of knowledge about gender assignment, it appears premature to conclude that default hierarchies offer a general solution to the problem of why sex-based rules take precedence in gender assignment (see, however, Rice, 2002 for a somewhat different point of view).

Two typological arguments also go against a solution in terms of default hierarchies. First, as pointed out by an anonymous referee, the idea that it cannot be a coincidence that sex-based rules take precedence receives further support from the animacy hierarchy (cf. Comrie, 1989: 185ff.). Nouns for biological males and females refer to human beings and higher animals, and thus occupy positions high up in the animacy hierarchy. It is well known that the animacy hierarchy bears on a wide range of linguistic phenomena, in that words in different parts of the hierarchy behave differently. Rules for biological males and females taking precedence in gender assignment is another example of such behavior. Since this behavior is in harmony with such an important typological principle as the animacy hierarchy, it would seem problematic to maintain that the behavior is a mere coincidence created by the relative sizes of genders.

A second typological argument against a solution based on default hierarchies concerns the status of semantic rules in gender systems. It is not difficult to conceive of a language that is identical to Arapesh except that gender X is higher than gender VII in the default hierarchy. In such a quasi-Arapesh, *nakor* ‘husband’s father’ would belong to gender X rather than the “masculine” VII by the logic of the default-based approach.

A default-based analysis would even be compatible with a language where semantic rules were never decisive for the assignment of gender, since there is no principle ensuring that semantic rules take precedence. This is highly problematic because such languages are not attested:

- (12) “In a sense all gender systems are semantic in that there is always a semantic core to the assignment system”
(Corbett, 1991: 8 based on Aksenov 1984: 17f.).

Thus, a pure default-based approach yields dubious typological predictions, in that it is compatible with unattested gender systems without a semantic core.

The upshot of this discussion is threefold. First, it is not clear whether a solution in terms of default hierarchies would work in all cases. Secondly, even if it would it is doubtful whether it would have the same explanatory power as the “Core Semantic Override Principle”. Thirdly, an approach based on default hierarchies appears to have unwanted typological implications. Given this, it seems that the “Core Semantic Override Principle” is an independent principle that does not follow automatically from any of the mechanisms described in Sections 2–4. While those mechanisms are “formal” in the sense that they involve the logical relationships between schemas, the “Core Semantic Override Principle” is a *substantial* ordering principle. It says that schemas take precedence if they involve a particular kind of information, viz. reference to biological sex. As a descriptive convention I suggest representing schemas of this type as shaded boxes. On the assumption that the “Core Semantic Override Principle” is adopted as part of the Usage-Based Model, candidates that instantiate shaded schemas take precedence. Thus, in Fig. 11 the




Fig. 11. Assignment of gender to Arapesh *nakor* 'husband's father' (third version).

“masculine” gender VII is selected because the relevant candidate is an instantiation of the shaded schema (a) in the grammar.

The account proposed in Fig. 11 raises two additional questions. First, we must ask whether the “Core Semantic Override Principle” pertains to gender, or whether it is valid for language in general. In other words, is the principle something speakers have to learn about the category of grammatical gender, or is it part of their general knowledge about language? As the distinction between male and female sex is of fundamental importance in life, it would intuitively not be surprising if it proved relevant beyond the category of gender. However, there are at least some parts of language where sex-based rules do not consistently take precedence. A case in point is inflection classes, as pointed out by Corbett and Fraser (2000b). In Russian, for instance, the “Core Semantic Override Principle” correctly predicts that nouns like *djadja* ‘uncle’ belong to the masculine gender. However, *djadja* pertains to the same inflection class as feminine nouns in *-a*. It is not inflected like most masculine nouns even though it denotes a male person. Evidently, therefore, biological sex is not decisive for the assignment of inflection class.

Despite the differences, it might be possible to advance a unified account of the assignment of both gender and inflection class. A potential strategy would be to assume the “Core Semantic Override Principle” to be a special case of a more general principle according to which prototypes take precedence. Now, this principle is vacuous unless the prototype of a category can be established on a principled basis cross-linguistically. For gender, the biological sexes stand out as good candidates for such universal prototypes. If a principled way of establishing prototypes for inflection classes in terms of formal features can be established, the precedence of formal features in the assignment of inflection class would follow from the same principle that governs gender assignment. Needless to say, this is not a trivial task, but in view of the “Core Semantic Override Principle” discussed in the present paper it represents an important challenge for future research.

The second question actualized by the analysis of Arapesh in Fig. 11 pertains to the status of the “Core Semantic Override Principle” in cognition. Does the principle gain support or motivation from what we know about cognition? It has often been argued in cognitive linguistics that cognition is “embodied”, i.e. that conceptualization is grounded in experience with our bodies and how we use them in interacting with our surroundings (Johnson, 1987; and Lakoff and Johnson, 1999). Since biological sex is of primary importance for our bodies, the privileged position of sex-based rules comes as no surprise from the perspective of a theory of embodied cognition. To pinpoint the exact nature of the relationship between this theory and the “Core Semantic Override Principle” is an interesting implication for future research.

6. Conclusion

From the perspective of the Usage-Based Model, I have considered four principles of relevance for a theory of gender assignment: Gender Tally, Elsewhere Condition, default hierarchies and what I have called the Core Semantic Override Principle. The contribution of the paper can be summed up in four points. First, it has been argued that Gender Tally (“rule counting”) falls out as an automatic consequence of the architecture of the

Usage-Based Model and hence need not be stipulated. This shows that the model is capable of accounting for “rule counting” in gender assignment, and at the same time it lends additional support to the concept of “rule counting”, since it follows from more general principles.

Secondly, we have seen that the Elsewhere Condition is relevant for gender assignment, but that it need not be stipulated as it follows automatically from the architecture of the Usage-Based Model. Interestingly, the Elsewhere Condition has been shown to follow from the same principle of the Usage-Based Model as Gender Tally. In this way, the Usage-Based Model provides a unified account of both “rule counting” and rule ordering, thus suggesting that these approaches are in fact closely related.

As a third point, it has been shown that it is technically possible to implement default hierarchies in the Usage-Based Model. However, representing hierarchies in terms of empty schemas of varying degrees of entrenchment appears to be at variance with the model’s strong emphasis on low-level schemas.

Last but not least, it has been suggested that Corbett and Fraser’s claim that semantic rules take precedence universally in gender assignment may be too strong. As an alternative, I have advanced the Core Semantic Override Principle, according to which only semantic rules invoking biological sex take precedence. This principle arguably does not follow from the architecture of the Usage-Based Model, thus suggesting that formal principles of rule ordering must be supplemented with substantial principles like the Core Semantic Override Principle.

The analysis pursued in the present study yields several implications for future research. In Section 2, we saw that despite the close parallelism to Steinmetz’ “rule counting” approach, the Usage-Based Model handles exceptions differently. While it is possible to represent default hierarchies in the Usage-Based Model (cf. Section 4), the analysis I have explored differs from Steinmetz’ approach. In the Usage-Based Model the effect of default hierarchies can be neutralized by highly entrenched schemas representing non-default patterns. The empirical consequences of these differences stand out as interesting questions for future investigation. Other questions include the exact scope of the Core Semantic Override Principle and its relationship to the theory of embodiment and prototypes. Whatever the final answers to these questions turn out to be, the present study has shown that the study of gender assignment may lead to clarification and revision of important aspects of the Usage-Based Model—and that the model provides a fruitful theoretical platform for investigating gender assignment.

References

- Aronoff, M., 1994. *Morphology by Itself*. MIT Press, Cambridge, MA/London/England.
- Barlow, M., Kemmer, S. (Eds.), 2000. *Usage based Models of Language*. CSLI Publications, Stanford, CA.
- Bybee, J., 2001. *Phonology and Language Use*. Cambridge University Press, Cambridge.
- Comrie, B., 1989. *Language Universals and Linguistic Typology*, second ed. Blackwell, Oxford/UK/Cambridge/USA.
- Corbett, G.G., 1991. *Gender*. Cambridge University Press, Cambridge.
- Corbett, G.G., Fraser, N.M., 2000a. Gender assignment: a typology and a model. In: Senft, G. (Ed.), *Systems of Nominal Classification*. Cambridge University Press, Cambridge, pp. 293–325.

- Corbett, G.G., Fraser, N.M., 2000b. Default genders. In: Unterbeck, B., Rissanen, M. (Eds.), *Gender in Grammar and Cognition*. Mouton de Gruyter, Berlin/New York, pp. 55–97.
- Dahl, Ö., 2000. Animacy and the notion of semantic gender. In: Unterbeck, B., Rissanen, M. (Eds.), *Gender in Grammar and Cognition*. Mouton de Gruyter, Berlin/New York, pp. 99–115.
- Dobrin, L.M., 1997. The morphosyntactic reality of phonological form. In: Booij, G., van Marle, J. (Eds.), *Yearbook of Morphology 1997*. pp. 59–81.
- Dobrin, L.M., 1999. Phonological form, morphological class, and syntactic gender: the noun class systems of Papua New Guinea Arapeshan. Doctoral dissertation, University of Chicago.
- Doleschal, U., 1993. Genus als grammatische und textlinguistische kategorie. Eine kognitiv-funktionalistische Untersuchung des Russischen. Doctoral dissertation, University of Vienna.
- Doleschal, U., 2000. Gender assignment revisited. In: Unterbeck, B., Rissanen, M. (Eds.), *Gender in Grammar and Cognition*. Mouton de Gruyter, Berlin/New York, pp. 117–165.
- Enger, H.-O., 2002. Stundom er ein sigar berre ein sigar. *Maal og Minne* 2, 135–151.
- Fortune, R., 1942. *Arapesh*. J.J. Augustin Publisher, New York (Reprinted 1977 by AMS Press, New York).
- Fraser, N.M., Corbett, G.G., 1997. Defaults in Arapesh. *Lingua* 103, 25–57.
- Johnson, M., 1987. *The Body in the Mind*. University of Chicago Press, Chicago.
- Kiparsky, P., 1982. *Explanation in Phonology*. Foris, Dordrecht.
- Köpcke, K.-M., 1982. *Untersuchungen zum Genusystem der deutschen Gegenwartssprache*. Max Niemeyer Verlag, Tübingen.
- Köpcke, K.-M., Zubin, D., 1984. Sechs Prinzipien für die Genuszuweisung im Deutschen: Ein Beitrag zur natürlichen Klassifikation. *Linguistische Berichte* 93, 26–51.
- Köpcke, K.-M., Zubin, D., 1995. Prinzipien für die Genuszuweisung im Deutschen. In: Lang, E., Zifonum, G. (Eds.), *Deutsch-Typologisch*. Walter de Gruyter, Berlin/New York, pp. 473–491.
- Koutsoudas, A., Sanders, G., Noll, C., 1974. The application of phonological rules. *Language* 50 (1), 1–28.
- Kumashiro, F., 2000. *Phonotactic interactions: a non-reductionist approach to phonology*. Doctoral dissertation, University of California, San Diego.
- Lakoff, G., 1987. *Women, Fire and Dangerous Things*. University of Chicago Press, Chicago.
- Lakoff, G., Johnson, M., 1981. *Metaphors We Live By*. University of Chicago Press, Chicago.
- Lakoff, G., Johnson, M., 1999. *Philosophy in the Flesh*. Basic Books, New York.
- Langacker, R.W., 1987. *Foundations of Cognitive Grammar*, vol. 1. Stanford University Press, Stanford, CA.
- Langacker, R.W., 1991. *Concept, Image, and Symbol*. Mouton de Gruyter, Berlin/New York.
- Langacker, R.W., 1999. *Grammar and Conceptualization*. Berlin/New York, Mouton de Gruyter.
- Mater, E., 1970. *Rückläufiges Wörterbuch der deutschen Gegenwartssprache*. Verlag Enzyklopädie, Leipzig.
- McClelland, J.L., Elman, J.L., 1986. The TRACE model of Speech Perception. *Cognitive Psychology* 18, 1–86.
- Prince, A., Smolensky, P., 1993. *Optimality theory. constraint interaction in generative grammar*. Report no. RuCCS-TR-2, Rutgers University, New Brunswick.
- Pugh, S.M., Press, I., 1999. *Ukrainian. A Comprehensive Grammar*. Routledge, London and New York.
- Rice, C., 2002. Optimizing gender. Paper presented at the conference “The Grammar of Gender”, Oslo, November 28–29, 2002.
- Steinmetz, D., 1986. Two principles and some rules for gender in German: inanimate nouns. *Word* 37, 189–217.
- Steinmetz, D., 2002. Gender shifts in Germanic and slavic. Paper presented at the conference “The Grammar of Gender”, Oslo, November 28–29, 2002.
- Trosterud, T., 2001. Genustilordning i norsk er regelstyrt. *Norsk Lingvistisk Tidsskrift* 19, 29–57.