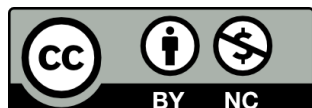Independent Pronoun Semantics: The pragmato-semanto-syntactic processing of pronominal reference

J.A. (Jac) Brown

Independent Pronoun Semantics:

The pragmato-semanto-syntactic processing of pronominal reference[1]

J.A. (Jac) Brown
Master Linguistics, specialization General Linguistics

Abstract: A comparative literature review is carried out on the topic of pronominal reference, looking specifically at pronoun gender agreement. Theoretical (Audring, 2013; Josefsson, 2006) and empirical accounts (Hammer et al., 2005; Dong et al., 2015) are compared. It is argued that these converge on the finding that pronouns, rather than being mere syntactic devices used for agreement, are selected on the basis of their own semantic load (referred to here as independent pronoun semantics); moreover, that independent pronoun semantics plays an important role in establishing pronominal reference. These findings are placed in a broader context, making a case for the psychological reality of Gricean pragmatics. Finally, a connection is made with Everett's model of language-culture interactions, that of cultural constraints on grammar (cf. Everett, 2005).

Keywords: personal pronoun, agreement, grammatical gender, language and culture

1.   Introduction

Agreement has traditionally been described as a syntactic phenomenon. In this view, agreement is a syntactic relation between two parts of a sentence, the *controller* and the *target*, where the target matches the controller in one or more syntactic features, such as number or gender (cf. Corbett & Fedden, 2016, p. 499). That is to say, the controller is that which gets matched with, and the target is that which matches. Consider the following example:

(1) Puell-a   pulchr-a   vir-um   crud-um   ama-t.

girl-NOM.   pretty-NOM.   man-ACC.   crude-ACC.   love-3SG.
SG.FEM   SG.FEM   SG.MASC   SG.MASC   PRES

"The pretty girl loves the crude man. "

In this Latin sentence, the adjective *pulchra* matches the noun *puella* in the features number and gender; the adjective *crudum* matches the noun *vir* in number and gender; and the verb *amat* matches *puella* in number[2]. Thus, we say that *puella* and *vir* are controllers, that *pulchra*, *crudum* and *amat* are targets, and that they agree with each other in these features.

    Recently, however, the notion of *semantic agreement* has been intro-

---

duced (cf. Audring, 2013; Josefsson, 2006, who speaks of *semantic gender*), where the idea is that the target matches the controller in one or more *semantic* features. In this paper, I will review this notion of semantic agreement from the point of view of pronominal reference, in an attempt to answer the research question "Is semantic or syntactic information primary when it comes to establishing pronominal reference?" Note that when I say "primary", I do not mean primary in a chronological sense, but rather "of primary importance". On the traditional, syntactic view of agreement mentioned above, syntactic information is of primary importance for establishing pronominal reference. However, as we will see, there are reasons to believe that semantic information is at least as important, if not more.

The structure of this paper is as follows. We will discuss recent literature on the topic. First we will consider theoretical accounts of gender agreement (Corbett & Fedden, 2016; Audring, 2013; Josefsson, 2006), with special attention to pronoun gender, and some potential problems with them. Then, in order to get an insight into the psycholinguistic and neurolinguistic mechanisms by which pronominal reference is established, we will look at some experimental research (Hammer et al., 2005; Dong et al., 2015). Certain common findings can be gleaned from these accounts, both theoretical and experimental. Finally, after discussing these commonalities, they will be placed in a wider context, and we will discuss their implications for theories of language processing, as well as what they say about how pragmatics relates to the syntax-semantics interface.

2.   Literature review

2.1 Theoretical accounts

Corbett & Fedden (2016), using the Canonical Typology approach, describe a *Canonical Gender Principle*, which holds that "in a canonical gender system, each noun has a single gender value" (p. 495). They then go on to state that real languages do not all conform to this principle; it is an abstraction, and gender systems in different languages vary in their canonicity.

Furthermore, they introduce the notion of *hybrid* controllers (p. 517ff.), which take different agreements in different domains (in some cases, even in the same domain). Hybrid nouns, it is said, "typically arise when the gender assignment rules of the language are in conflict" (p. 518). This behaviour is constrained by the Agreement Hierarchy, which had been previously described by Corbett (1979, cited as Fig. 5 in Corbett & Fedden, 2016, p. 519). The Agreement Hierarchy describes which targets are more likely to be permitted alternative agreement based on semantic justification. From left to right, the hierarchy is as follows: attributive adjectives > predicative adjectives > relative pronouns > personal pronouns, where the further right one goes in the hierarchy, the more likely it is that such targets are permitted alternative agreement.

The example they give here is that of the German noun *Mädchen* 'girl'. Being a diminutive, this is syntactically neuter, but it allows feminine agreement (at least in the domain of the personal pronoun) because it has a female referent. This implies that the reason "the gender assignment rules of the language are in conflict" is because they are (morpho-)syntactic on the one hand (assigning the value NEUTER), and semantic on the other hand (assigning the value FEMININE).

Audring (2013), discussing apparent gender agreement mismatches in

Dutch (p. 35ff.), rejects the notion of hybrid controllers and proposes "an alternative account that does not make use of the *hybrid* concept" (p. 40). She mentions the supposed semantic rules for gender assignment in Dutch. According to these rules, countable nouns should be masculine, sometimes manifesting as common) and uncountable nouns should be neuter. She then states that these rules do not exist as such. This is to judge from the fact that the Dutch gender assignment patterns for countable and uncountable nouns are in fact inconsistent (for examples, see p. 41). She also argues that some Dutch nouns do not always exhibit hybrid behaviour; moreover, if one takes the use of a masculine pronoun for a common noun as a mismatch, virtually all Dutch nouns exhibit hybrid behaviour in some circumstances. Audring herself does not do this, by the way, as "masculine or common gender fall under the same semantic rule" (footnote p. 41; I agree with this remark).

If the rule that produces what Corbett and Fedden call "hybrid" behaviour, as described by them, is not correct (at least not for Dutch), what, then, is causing the fact that nouns can sometimes take one anaphoric pronoun and sometimes another? Audring argues that hybridity is not a property of the controller, but of the target (in this case, the pronoun). On this view, the pronouns have a semantic load of their own, which may semantically clash with the noun's syntactic gender, when it is felt to be semantically inappropriate, and it is the semantically appropriate pronoun that is used in case of such a clash. Thus, she dispenses with the notion of hybrid nouns entirely (pp. 42–44).

Josefsson (2006), writing about Swedish, distinguishes two systems for assigning gender: a syntactic[3] one and a semantic one. She illustrates this with predicative adjective agreement, which in Swedish sometimes (puzzlingly) appears to mismatch: common nouns can trigger neuter adjectives, as in the following example (after Josefsson, 2006, p. 1347):

(2) senap   är   gul-t.
     mustard   be.PRES   yellow-NEUT
     "Mustard is yellow."

In this sentence, the noun *senap*, despite being common, triggers the neuter adjective *gult*, instead of the common *gul*, which one would expect if agreement in Swedish were purely syntactic, as in the Latin example above.

Josefsson describes the Swedish syntactic gender system as having two genders: common and neuter, and the semantic gender system as having four: MALE, FEMALE, THING, and SUBSTANCE; inanimate nouns can be variably assigned the THING or the SUBSTANCE gender based on the context, although one of these assignments is often prototypical whereas the other is marked (pp. 1349–1352).

Although Josefsson mainly discusses predicative adjective agreement, she also mentions pronouns, and in fact this becomes crucial later on in her analysis. Pronoun selection shows that the syntactic and semantic gender systems interact: semantic males take the masculine pronoun *han* 'he'; females take the feminine *hon* 'she'; things and substances can both take either *den* 'it [common]' or *det* 'it [neuter]', depending on the syntactic gender of the noun with which they agree (p. 1352)—except in cases where

_____

3 Josefsson uses the term "grammatical gender", but I use "syntactic gender" for the sake of consistency and clarity; in my vocabulary, "grammar" includes semantics.

there is no possible antecedent noun present and *den* and *det* appear to be used deictically (examples on pp. 1352–1353). In these cases, according to Josefsson, the choice of pronoun is determined directly by the semantic gender feature of the pronoun itself: *den* is used for things and *det* for substances, which includes events and clauses (pp. 1353–1355).

On this basis, she goes on to argue that the apparent gender mismatches (or cases of "disagreement") are actually cases of agreement because that with which the target adjective gets matched is not, in fact, the gender feature of the target noun, but that of an (unexpressed) "pronominal element merged higher up in the nominal extended projection" (p. 1355). This is much in the same way that, in sentences with so-called "prenominal apposition", it is the prenominally apposed pronoun and not the noun that triggers agreement on the predicative adjective (examples on pp. 1357–1358). This analysis holds that there is a projection called Semantic Phrase (SemP; this is the aforementioned "pronominal element", which may or may not be overtly expressed) on top of the DP.

Comparing Audring (2013) and Josefsson (2006), we see that both authors discuss languages in which one and the same noun, with the same gender value, is referred to with different pronouns in different situations. Whereas the one (Audring) considers this a case of pronouns exhibiting hybrid properties, the other (Josefsson) claims that the apparent violation of agreement is in fact 'true', i.e. syntactic, agreement with an unexpressed part of the deep structure. Although they disagree on this point, we can see that there is a commonality in their findings, perhaps best summarized as follows: *pronouns have a semantic load of their own, rather than being mere syntactic devices used for anaphoric reference, and this semantic load influences the speaker's choice of pronoun depending on the situation*. In what follows, I will refer to this idea as *independent pronoun semantics*. In light of independent pronoun semantics, we can now review our notion of (semantic) agreement. My view is illustrated in Fig. 1:
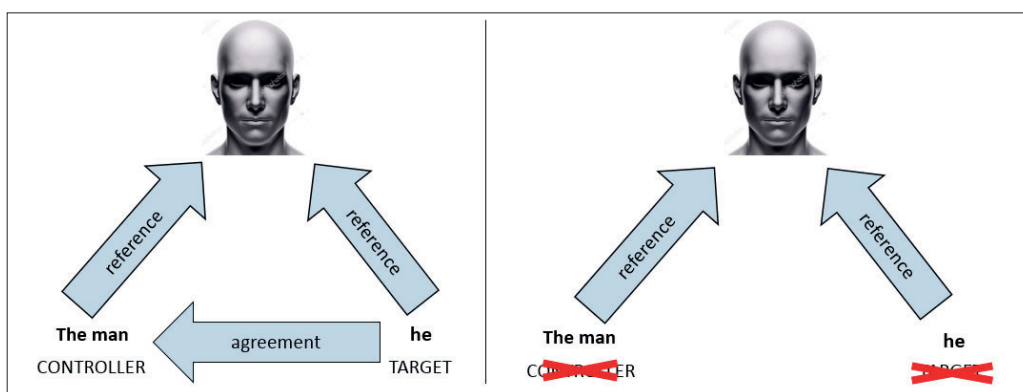

Fig. 1. Syntactic agreement (on the left) vs. 'semantic agreement' (coreference; on the right). The man's head represents the real-world referent, whereas the arrows represent relations between words (syntactic; horizontal) and between words and the real world (semantic; diagonal). The words 'the man' and 'he' should not be taken as the literal English words, but are placeholders for nouns and pronouns in hypothetical languages with full syntactic gender (on the left) and without it (on the right). (Source for man's head: https://depositphotos.com/132561676/stock-photo-generic-human-man-face.html.)

On this view, then, what Josefsson calls semantic agreement is actually simply coreference. Even in languages with syntactic agreement, the pronoun does not merely match its anaphoric referent on syntactic features such as gender, but it also refers by itself, namely to its extralinguistic (real-world) referent. In languages without syntactic agreement, or where syntactic agreement is disappearing, the syntactic relationship between the two is severed, while the semantic relationship is maintained through coreference. Thus, the noun ceases to be a controller, and the pronoun ceases to be a target. For this reason, I suggest we do away with the notion of semantic agreement and use the term agreement only for syntactic agreement.

The fact that personal pronouns have independent gender semantics helps explain why they are on the right in Corbett's Agreement Hierarchy: they are "permitted alternative agreement based on semantic justification" (Corbett & Fedden, 2016, p. 59) more than the elements to their left because, as I have just demonstrated, their semantic gender allows them to refer directly to the referent, rather than via the noun.

2.2 Experimental accounts

So far we have accrued theoretical evidence for the influence of independent pronoun semantics on pronominal reference, but let us now take a more empirical look at what this means for language processing in a psycho- and neurolinguistic sense, by looking at evidence from experiments.

Hammer et al. (2005) investigated event-related potentials (ERPs) in German-speaking subjects presented with German sentences containing anaphoric pronouns that either matched or mismatched their antecedent nouns in gender, where the antecedent noun could refer to either a person or a thing (examples in Table 1 on p. 229). Their results indicate that semantic and syntactic processing interact when listeners (or readers) link a pronoun to its antecedent.

Linking a pronoun to its antecedents proceeds in two stages, according to Garrod & Sanford (1994, cited in Hammer et al. 2005, p. 227): the *bonding* stage, at which the link is initially established, and the *resolution* stage, at which the link is evaluated and ultimately judged to be either successful (the pronoun is congruent with its antecedent) or unsuccessful (the pronoun is incongruent with its antecedent). While this pronoun processing is going on, brain activity may show one of two ERPs: to oversimplify somewhat, an N400 is seen in case of increased semantic processing (indicating difficulty), whereas a P600 is seen in case of increased syntactic processing.

Hammer et al. found both N400 and P600 effects. In their first of two experiments, stimuli consisted of a sentence with one main clause containing the antecedent (such as *Die Jacke ist warm…*) and one subordinate clause containing the pronoun (such as *…weil sie gefüttert ist*). In this experiment, they found a small negativity (not quite an N400) at the pronoun position in case of a mismatch, but only if the antecedent was a thing, contrary to expectations (pp. 229–230); they also found an N400 on the word following the pronoun, again only in the thing condition (pp. 230–232; 235). This indicates that the parser continues to search for an acceptable referent outside the sentence, i.e. either at the discourse level or in the real world, but only in the thing condition (cf. discussion on pp. 232–233; 235–236).

They also found a P600 at pronoun position for both person and things,

but this effect was larger in the person condition, indicating that the preceding negativity 'pulls down' the P600 in the thing condition (p. 230; 235). In light of this finding, they revised their analysis of the P600 effect: it does not reflect purely syntactic processing, but is influenced by the semantic processing that is still going on at the same time (p. 232; cf. also pp. 235–236).

In the second experiment, featuring similar sentences preceded by an additional ('discourse') sentence (such as *Die Frau steht im kalten Schnee*), they looked only at N400 effects (pp. 233–234). This experiment replicated the findings from the first (p. 234), which confirms that semantic and syntactic processing occur together, not separately (cf. p. 236). In other words, the parser uses both semantic and syntactic information at the bonding stage to establish the link between pronoun and noun antecedent. When the noun refers to a person, and thus contains semantic gender information, the pronoun clashes semantically as well as syntactically with the noun, meaning no further analysis is possible. In this case, the parser proceeds to the resolution stage without doing anything, resolving that the pronoun is incongruent. This is expressed in terms of brain activity as a P600.

However, when the noun refers to a thing, and thus does not contain semantic gender information, the possibility remains open that the pronoun refers to something else. Thus the parser is not sure of the link at the bonding stage, expressed by a negative brainwave (indicating difficulty in semantic processing). It then searches outside the sentence, in the discourse, for any other information it could use to establish pronominal reference with something other than the noun inside the sentence. If this information is not found, it resolves that the pronoun is syntactically incongruent at the resolution stage, expressed similarly by a P600, except that the P600 in this case is smaller because it is 'pulled down' by the preceding negativity. But then, crucially, the parser holds out for more information from the following word(s), and only when this information determines once and for all that the pronoun either does or does not have the same referent as the noun does the parser resolve that it is either congruent or incongruent (expressed by an N400 on the word following the pronoun in case of incongruence).

To summarize, although both nouns with inanimate and with animate referents had syntactic gender, it was only the "semantic gender" of the *animate* referents that immediately (at the bonding stage) blocked pronoun integration. This indicates that it is the pronoun's independent semantics that is used to establish pronominal reference, and only if this does not resolve ambiguity do syntactic features come into it.

Another study that suggests semantic information is at least as important for establishing pronominal reference as syntactic information was done by Dong et al. (2016). They did two self-paced reading experiments with Chinese-speaking learners of L2 English. In both experiments, participants were shown sentences with a noun antecedent (such as *Mary goes to the zoo to watch animals every day after work for a good rest*) and an anaphoric pronoun, where the prounoun either matched (*She…*) or mismatched the antecedent in gender (*He considers it the best way to relax and maintain a good mood*; p. 737). In the first experiment, the sentences were preceded by a picture of either a human or a non-human; in the second, the sentences were not preceded by a picture (p. 736).

The authors found that reading times for sentences with a mismatching pronoun were significantly longer—indicating difficulty during processing—only in the condition in the first experiment where the gender of the pictured human matched that of the pronoun (pp. 738–740). In all other circumstances, reading times were the same as for sentences with a matching pronoun (ibid.; pp. 742–743).

They explain this finding as follows (p. 743ff.): in Chinese, 3rd person singular pronouns are not distinguished with regard to gender, except in writing. Thus, when Chinese listeners link a pronoun to its antecedent in Chinese, they do not use gender information (indeed, they cannot, since it is absent). According to Dong et al., Chinese learners of L2 English do the same in English—at least at the bonding stage (see above). At the resolution stage, they use extrasentential information (in this case, the picture) to evaluate the link. This is why the presence of a picture that matches the pronoun, and thus mismatches its antecedent noun, slows down reading times, whereas reading times are unaffected everywhere else (cf. discussion of the mismatch effect on pp. 743–744).

Thus, the studies by Hammer et al. and Dong et al. both indicate that semantic information is important for linking a pronoun to its antecedent; moreover, that extrasentential information (whether from the discourse or from the real world) is used to evaluate this link.

3. Discussion

We have now collected substantial evidence from both theoretical and experimental sources. In section 1 above, we reviewed the definition of agreement, controller and target. In section 2.1, we compared the notion of hybridity as discussed by Corbett and Fedden (2016) and by Audring (2013), respectively; we also discussed Josefsson's (2006) analysis of the Swedish gender system and found that it shares with Audring's analysis of the Dutch gender system that which I refer to as independent pronoun semantics. This idea was then embedded back into Corbett and Fedden's theoretical framework. Finally, in section 2.2, we held these theoretical findings up to empirical scrutiny; we found that independent pronoun semantics directly affects the linking of pronouns to their antecedents and the evaluation of this link, both during comprehension (Hammer et al. 2005) and during productions (Dong et al. 2016).

Building on this evidence, we may now ponder what the *reason* could be that independent pronoun semantics influence the establishment of pronominal reference as seen above. For example, the fact that the parser holds out for extrasentential information in the study by Hammer et al. strongly suggests that pragmatics plays a role here. After all, language is a tool for communication, and in order for communication to be successful it is always necessary to be able to make sense of what the other party is saying. The *Maxim of Relation*, as formulated by Grice (1975, p. 46), states "Be relevant", i.e. (my words) when listening, people expect their conversation partner to say only things that are relevant to the conversation; when speaking, they know that their conversation partner expects them to do the same, so they try to be equally relevant. Thus, when faced with something that seemingly makes no sense (an apparent gender agreement violation), the listener first tries to

make sense of it: "Either this pronoun is incongruent with this noun, but the speaker *meant* to refer to this noun and simply used the wrong pronoun; or the pronoun is in fact congruent, but with another antecedent (found either elsewhere in the discourse or in the real world), not with this one."[4]

This need for 'sense-making' is so inherent to the communicative function of language (and indeed to human cognition more generally) that I would go so far as to suggest pragmatics is part of the mental grammar, i.e. language users have built into their language processing the question *What is proper to say (in this particular situation)? What should be said and what not? What makes sense?*

There is a body of research that supports this suggestion if one makes the necessary connections. On the production side, Antón-Méndez (2010) uses the concept of the *preverbal message* to explain why L2 English speakers whose native language is Spanish make more pronoun gender errors than L2 English speakers with similar native languages, such as Italian. According to her, the information included in the preverbal message is language-specific (pp. 133–135). Spanish has different information requirements than English; in particular, the gender of a third person constituent is often not required in Spanish. Thus, the processing responsible for these errors takes place on the conceptual level, not on the level of syntax. On the comprehension side, this is in accordance with the finding by Dong et al. that Chinese speakers of L2 English process pronoun gender information at the conceptual level. That is, the pronoun is not linked by gender to its antecedent syntactically, but to its real-world referent conceptually; processing difficulty only occurs only if this real-world referent (represented in their experiment by a picture) clashes with the pronoun.

The connection with pragmatics lies in the fact that again, which information is required at the conceptual level depends on what is or is not proper to say in the situation at hand. This differs between languages and is influenced by culture. For example, Pawley (2002) found that inanimate nouns in colloquial Australian English can be animated by referring to them with *he* or *she*. The choice of pronoun depends on semantic properties that, while not fully known, are common to the grammars of the members of this linguistic community, evidenced by the fact that such references are never misunderstood.

An example of culture influencing language comprehension is provided by Hubers et al. (2016), who presented prescriptivist listeners with constructions that are held to be 'improper grammar' by prescriptivist standards, and found that these listeners present such constructions differently from *both* uncontroversially grammatical *and* truly ungrammatical constructions—on a physical, neurolinguistic level, demonstrable through fMRI scans.

Thus, the above findings make a strong case for the psychological reality of Gricean pragmatics. Pragmatics, which is influenced by culture, itself influences language processing—measurably—down to the level of syntax. This appears to support Everett's model of what he calls "cultural constraints on grammar", whereby the grammar of a given language is shaped by the cultural values and attitudes of its speakers as to what is 'proper' to say or not (cf. Everett, 2005, pp. 622–623; 631; 633–634).

---

4 I do not mean to suggest that this attempt at sense-making takes place consciously; on the contrary, my point is that it is as subconscious as the rest of the language processing system.

4.   Conclusion

We can conclude with the following remarks. Although the research question at the beginning of this paper was a single one, the answer turns out to be twofold. Firstly, semantic information indeed appears to be at least as important for establishing pronominal reference as syntactic information, if not more. Returning to the notion of agreement, we see that it is not necessary, even in languages that have it, because independent pronoun semantics ensures that the pronoun-noun link is always maintained simply through coreference. Secondly, this process of establishing pronominal reference is influenced by pragmatics, during both production and comprehension: different languages have different requirements as to what should be expressed, which influences both what speakers choose to say (production) and how listeners make sense of what is said (comprehension).

Unanswered questions still remain. Strikingly, if agreement is not necessary, then why does it exist at all? I suspect the answer may be found in the way the information structure is optimized for easy understanding. Note that Josefsson (2006, pp. 1366–1367) states that pronoun gender facilitates discourse linking. It could be the case that different languages use different forms and combinations of agreement, word order, prosody and perhaps other devices to package information and link discourse together, motivated by the listener's need to make sense of what is said and the speaker's need to make this possible. Further research opportunities lie here.

References

Antón-Méndez, I. 2010. Gender Bender: Gender Errors in L2 Pronoun Production. *Journal of Psycholinguistic Research* 39: 119–139.

Audring, J. 2013. A pronominal view of gender agreement. *Language Sciences* 35: 32–46.

Corbett, G.G. 1979. The Agreement Hierarchy. *Journal of Linguistics* 15: 203–224.

Corbett, G.G. & S. Fedden. 2016. Canonical gender. *Journal of Linguistics* 52: 495–531.

Dong, Y.P., Wen Y., Zeng X.M., & Ji Y.F. 2015. Exploring the Cause of English Pronoun Gender Errors by Chinese Learners of English: Evidence from the Self-paced Reading Paradigm. *Journal of Psycholinguistic Research* 44: 733–747.

Everett, D.L. 2005. Cultural Constraints on Grammar and Cognition in Pirahã. *Current Anthropology* 46 (4): 621–646.

Garrod, S.C. & A.J. Sanford. 1994. Resolving sentences in a discourse context. In: Gernsbacher, M.A. (ed.). *Handbook of Psycholinguistics*, pp. 675–698. New York: Academic Press.

Grice, H.P. 1975. Logic and conversation. In: Cole, P. & J. Morgan (eds.). *Syntax and Semantics: Vol. 3. Speech Acts*, pp. 41–58. New York: Academic Press.

Hammer, A., B.M. Jansma, M. Lamers, & T.F. Münte. 2005. Pronominal Reference about Persons or Things: An Electrophysiological Approach. *Journal of Cognitive Neuroscience* 17 (2): 227–239.

Hubers, F., T.M. Snijders, & H. de Hoop. 2016. How the brain processes violations of the grammatical norm: An fMRI study. *Brain & Language* 163: 22–31.

Josefsson, G. 2006. Semantic and grammatical genders in Swedish—independent but interacting dimensions. *Lingua* 116: 1346–1368.

Pawley, A. 2002. Using *He* and *She* for Inanimate Referents in English: Questions of Grammar and World View. In: Enfield, N.J. (ed.). *Ethnosyntax: Explorations in Grammar and Culture*, pp. 110–137. Oxford: Oxford University Press.