

The linguistics and psycholinguistics of agreement: A tutorial overview

Juan Carlos Acuña-Fariña

*Department of English, Facultad de Filología, Universidad de Santiago de Compostela, Campus Universitario Norte,
15782 Santiago de Compostela, Spain*

Received 10 June 2008; received in revised form 2 September 2008; accepted 3 September 2008
Available online 26 October 2008

Abstract

Agreement lies at the heart of sentence structure in that it usually codifies the formal link between the subject and the predicate. In a way, although not strictly necessary, agreement tends to be what signals that a clause has been created. However, linguists and psycholinguists disagree as to whether this core process is essentially a semantic or a syntactic phenomenon, and there is evidence that suggests that it may be both things at the same time. Berg (1998) suggests that in morphologically impoverished languages like English, low frequency of agreement operations makes for a weak morphosyntactic component that is unable to keep semantic interference at bay, at least in production (where meaning comes first). In his completion study, he manages to show that the strong morphosyntactic defences of German do seem to encapsulate number agreement from non-formal forces. In this paper, I examine the psycholinguistic literature on the processing of agreement in English and Spanish in search of fine-grained evidence for encapsulation in all the relevant *domains* (NP, clausal, supra-clausal). The different grounding of the features of *gender* and *number* is also analysed and evidence for their differentiability in the first cycles of processing is ruled out. In general, Berg's hypothesis is confirmed by existing processing measures (Spanish patterns like German due to its rich morphology), but the fine picture of agreement operations in English and Spanish is also sensitive to the different domains tested and to whether production or comprehension is taken into account. Finally, it is argued that gender classes and the agreement systems based on them constitute a major generative device when constructing the clause. This is because gender is less semantically grounded than number and therefore easier to recruit for phrasal packaging, as this merely requires the co-variance of form.

© 2008 Elsevier B.V. All rights reserved.

Keywords: Agreement; Processing; Psycholinguistics; Features

1. Introduction

Agreement is a challenge for both theories of grammar and theories of comprehension and production. In fact, no theoretical model of any standing can afford not to have a theory of agreement today. There are a number of reasons for this. In the first place, there is the sheer puzzle of it: that is, the fact that it is so pervasive in the world's languages and yet seemingly redundant or even downright useless (Taylor, 2002:332 ff.; Corbett, 2006:XV). In the second place, “agreement is arguably the major interface problem between morphology and syntax, and hence appears particularly difficult when viewed from the heartland of either component” (Corbett, 2006:3; also Eberhard et al., 2005). In the

E-mail addresses: iacanoe@usc.es, carlos.acuna.farina@usc.es.

third place, agreement lies at the heart of sentence structure in that it usually codifies the formal link between the subject and the predicate in many languages. In the fourth place, its nature, source, and functioning display an unsettling mixture of formal and semantic motivations. There are models of grammar and processing that view it as an essentially encapsulated phenomenon whose range of operation is limited to the first cycles of the building of clause structure (Chomsky, 1995, 1999, 2001; Bock and Eberhard, 1993; Eberhard, 1997; Levelt et al., 1999; Carminati, 2005; Franck et al., 2006). Then, there are others which maintain that agreement is either *penetrated* by semantic forces or, more radically, a semantic process in nature (Pollard and Sag, 1988; Barlow, 1999; Vigliocco et al., 1996a; Thornton and MacDonald, 2003; Vigliocco and Hartsuiker, 2002; Haskell and Macdonald, 2003). The different views usually belie the usual differences in where one puts the focus, and since agreement facts span a large portion of linguistic behaviour, it is natural that that focus may rest in quite a large number of different places. For instance, if agreement is seen to be essentially the same phenomenon regardless of the different *domains* where it seems to be at work (phrasal, clausal, supra-clausal, supra-sentential), then as we transit from the smallest domain (NP structure) to the largest one (discourse), we will find plenty of opportunities to see semantics at the wheel, both in grammar and in processing (Corbett, 2006). If, conversely, and precisely because of that, one restricts agreement to the phrase or the clause, then shallow, automatic computations based on rigid formal regulation will be more apparent. However, the frequent intervention of formal processes across large distances and of notional factors in the short distances shows that domains do not have pristine formal vs semantic regulation (Eberhard et al., 2005). Thus, English grammar stipulates that in *a lot of people were there (*is there)* there should be a semantic override of the habitual, formal subject–verb (S–V) agreement pattern, showing neglect of a formal requirement in the computation of the noun phrase. With ‘committee nouns’, the two agreement options (syntactic and semantic) are available (*the committee is . . . ; the committee are . . .*). However, once a form is chosen on the predicate, it tends to be preserved in pronoun choice after the predicate, therefore uniting formal regulation *and* distance, at least in single clauses (Huddleston and Pullum, 2002:495):

- (1) The committee **hasn’t** yet made up **its** mind.
- (2) The committee **haven’t** yet made up **their** mind/minds.¹

And in processing, there is evidence that formal feature matching facilitates the computation of pronoun-antecedent ties as well, that is, even in the long distances (Ehrlich, 1980; Garnham and Oakhill, 1985; Garnham et al., 1995; Arnold et al., 2000).

Therefore, agreement forces one to face interesting questions. Among these, the following will be asked here:

1. Is agreement an essentially formal or an essentially conceptual phenomenon?
2. Are the grammar of agreement and the processing of agreement ‘in good agreement’ with each other?
3. Are there cross-linguistic differences in the way agreement is processed?
4. Do we compute agreement differently for gender and number? For semantic gender and morphosyntactic gender?
5. Does the mind, and the brain, react differently to agreement based on either semantic or formal regulation?
6. Is the processing of agreement the same across domains?

Maybe the greatest puzzle about agreement systems lies in their apparent arbitrariness and uselessness. Consider the facts of Bengali, for instance. According to Ferguson (1964), Bengali’s pronouns have forms for 12 different categories of person, number and respect, and its verbs have syncretized tense/mode/aspect and person suffixes with widespread homophony, but for each tense/mode/aspect set of forms, there is independent marking of five person/number/respect classes. The number of combinations is explosive. For instance, -i is the first person ending in the present but the second person *inferior* in the future, while -o is the second person *ordinary* in the present and first person in the future.² In Banawá, according to Corbett (2006:61), there is subject–verb agreement in person and number. Agreement with the head of the subject noun phrase is marked by a prefix on the verb, but the verb also agrees

¹ As Huddleston and Pullum (2002:495) observe, for cases such as (1) and (2) above, which start with a formally marked singular NP, it “is possible (. . .) to switch from singular verb to plural pronoun—but one would not normally switch from a plural verb to a singular pronoun in close proximity”. This is the habitual relative prevalence of notional agreement when distance between two agreement sites is involved.

² A referee points out that maybe the old number distinction in Bengali has been taken over to code respect, which means that the cross-cutting features are person and respect (corresponding to old number) only. Even so, the number of possible agreement combinations continues to be explosive indeed. I thank the referee for this observation.

in gender with the topic noun phrase via a suffix. This produces an interesting effect with inalienable possession, as in a phrase like *my head*. When this is used in subject position, the verb agrees in person and number with *head*, but also in gender with *I*. On the face of such facts, it is no wonder that gender classes and the complex agreement systems based on them constitute, in principle, one of the quintessential cases of what some have defined as “a clear case of the victory of the indexical aspect of language over its iconic aspect” (Haiman, 1985:162). In other words, on this view agreement “is not only non-iconic but meaningless”. Corbett (2006:274) points out that agreement “often appears to involve a lot of effort for a questionable payoff”. Jespersen (1922:352 ff.) regarded it as “superfluous” and “cumbersome”, and Taylor (2002:332 ff.), a cognitive grammarian for whom admitting the mere existence of form without meaning is problematic, expresses similar concerns:

While the facts of gender (and, more generally, of inflection classes) and of agreement patterns can be represented within a Cognitive Grammar framework (. . .), the very existence of inflection classes is somewhat puzzling, since they contribute little to the symbolization of conceptual structure. English, for example, which lacks noun gender, is thereby not one whit less efficient as a symbolic system than languages that do have gender systems. On the contrary, *elaborate inflection class systems might seem to be dysfunctional*, in that they place a heavy burden on a speaker’s memory. But while inflection classes certainly present the second language learner with severe problems, they are a fact of many languages, and *speakers of these languages show no signs of wanting to give them up. It would even seem that speakers take delight in the formal complexity of their language*. For my part, therefore, I am inclined to see the complexities of inflection classes, and the sometimes elaborate systems of agreement patterns with which they are associated, as manifestations of *humans’ delight in what I called (. . .) form-focused activities*. To the extent that they lack a symbolic-conceptual content, agreement patterns are an example of ‘*pure phonology*’. (emphasis added)

Given the fact that over 70% of the world’s languages (including English) exhibit some form of agreement (Mallinson and Blake, 1981), Taylor’s Anglocentrism sounds extreme in the previous quotation. But indeed, with an English/Spanish perspective in mind, the mere comparison of (3) and (4) below at least partially supports his case:

- (3) [We will keep] the little white candlesticks and the comfortable red chairs.
 (4) [Nos quedamos con] lo-s pequeño-s candelabro-s blanco-s y la-s cómoda-s silla-s roja-s.

Where English uses 2 morphological marks to code the noun phrase in (3), Spanish uses 16 in (4). One really needs to pause to consider if English misses any subtle information that, for some reason, Spanish bothers to code. Indeed, in the face of all the ‘exotic’ agreement systems in the world, like those of Bengali or Banawá mentioned above, and of so many others, an obvious question to ask is: why bother? Why do languages go to such great lengths to implement extremely complex agreement systems across the board if these are “dysfunctional”, “superfluous” and “cumbersome”? Why devote so much attention to the complex interplay of phonological, morphological, and semantic variables that must crystallise in a system which must, at the same time, be quick enough to produce words at a rate of some 15–20 phonemes per second and scan the entirety of a lexicon of about 60,000 words in less than half a second in ordinary word access? With the daunting task of computation of word level, phrase level, sentence level and even discourse level dependencies, why make computational room for the Spanish redundancy in (4) when the minimal coding of English in (2) will suffice?

Notice that at least three related dimensions stem from this paradox. One is simply the function of agreement, if it has any. Alternatively, if it has no important function, the question is what agreement reveals about the faculty of language given that it is not universal but at the same time apparently non-iconic. In other words, if agreement reveals the victory of the indexical over the iconic, then it would be tempting to believe that it is part of universal grammar. That would nip in the bud any need to see ‘substance’ and functional motivation in it, since, as is often the case in the world of matter around us, the microscopic organization of such matter reveals a certain capricious or chaotic emergence and self-organization (like the much-quoted case of the double-helix shape of DNA molecules). But the fact is that not all the world’s languages have agreement. The second aspect of the paradox involves the processing issue. There is evidence in the processing literature that at least in languages with rich inflectional systems like Italian or Spanish, the processing of redundant marks of agreement is fast and shallow, in the sense that it appears to be automatic and relatively immune from semantic interference. Berg (1998) suggests that in morphologically impoverished languages like English, low frequency of agreement operations makes for a weak morphosyntactic

component that is unable to keep semantic interference at bay, at least in production. Since the storing of gender features and the on-line computation of those features in ‘phrase construction’ (Hawkins, 1994, 2004) are two different things, another interesting aspect of the processing of agreement is how word access and phrase construction interact, especially how they do so in different languages like Spanish and English, with different morphosyntax.

The last aspect of the dysfunctional-yet-pervasive puzzle has to do with a Whorfian dimension, which we will not pursue here. One way out of the counter-intuitive dysfunctionality of such a widespread and varied use of agreement systems across languages is to consider the possibility that each agreement system, and especially the gender classes that it is built upon, reflects a particular – if subtle – way to construe reality. Konishi (1993) studied the way in which German and Spanish subjects attributed semantic dimensions to words based on different genders. Speakers of the two languages differed in their ratings for words that had a different gender in the two languages. In a similar test (rating words on a semantic differential scale), Mills (1986:128) found that English and German speakers’ ratings of nouns did not differ on the basis of German gender. Sera et al. (1994) asked Spanish and English speakers to assign a male or a female voice to pictured objects and words. Spanish speakers, but not English speakers, tended to assign voices based on gender differences in the words. Sera et al. (2002) extended the same method to an investigation of French and German. The former behaved like Spanish, while the latter like English. At issue then is whether the gender of words may affect meaning representations. In a recent study of Italian and German, Vigliocco et al. (2005) concluded that that is the case, but only at a lexical (not conceptual) level of representation. Similarly, Bock et al. (2006) have recently investigated whether collective nouns such as *staff* or *committee* in British and American English are represented differently in conceptual structure (“American speakers see a forest where British speakers see trees”, p. 71). In line with the Vigliocco studies, Bock et al. concluded that differences were solid, but only in the lexical structure, not in the way speakers conceptualise reality.

1.1. Linguistic theories of agreement

One way or another, almost all linguistic theories are now enriched by work that tries to demonstrate their psychological adequacy (Dik, 1978, 1991; Kaplan and Bresnan, 1982; Van Valin and Lapolla, 1997; Franck et al., 2006). Given its apparent uselessness, agreement is interesting when contemplated from such a perspective. For instance, generative grammar in the current *minimalist* tradition has spawned a number of issues where agreement plays a central role. Some of these issues have arisen in connection with work on *attraction* (Jespersen, 1924; see especially Eberhard et al., 2005, and Bock et al., 2006 and references therein). Attraction occurs when a complex subject noun phrase is made up of two (or more) nouns with different morphological marking, specifically more often when the combination is singular + plural, as in *the key to the cabinets*. Occasionally, such combinations result in agreement mistakes, as in **the key to the cabinets are in there*. The label *attraction* is meant to imply that the plural feature of the local noun (or ‘interloper’) illegally attracts agreement on the verb. Attraction has been proven strong with the plural feature, not with the gender feature (although this may reflect an English bias in research; see Vigliocco and Franck, 2001). Likewise, Nicol (1988), using a cross-modal priming technique, investigated the processing of English pronouns which were disambiguated by gender or by number and found a significant priming effect in the number matching conditions. Using the same technique, Di Domenico and De Vincenzi (1995) and De Vincenzi (1999) studied Italian sentences where an object clitic pronoun was disambiguated towards one of two antecedents by the gender or the number feature, and obtained significant priming effects only for the latter. On the basis of these facts, the authors argue for a modular theory of sentence processing whereby number is processed in an earlier, syntactic stage, and gender at a later stage when semantic information from the lexicon becomes available. Carminati (2005) has contrasted the disambiguating potential of the features in Greenberg’s (1963) *Feature Hierarchy* (Person > Number > Gender) to argue for a *Feature Strength Hypothesis* according to which there is a correlation between the cognitive significance of a feature and its disambiguating power: the more cognitively important the feature is the better it should be at disambiguating the pronoun that carries it. Thus, even though number and gender are *phi*-features that are supposed to enter the derivation at the same stage and be checked in a more or less uniform way, work of the kind alluded to has served to propose that *number* should have a different status in the grammar from gender and that it should head an independent functional projection. Anderson (2004) reminds us of the usual belief that gender is an inherent property of the lexical N, but number is an inherent property of the syntactic N position (see also Hudson, 1999). I will pay close attention to the evidence for feature differentiability in section 2.4.

The way in which agreement works in connection with movement in early *Minimalism* and in the *Minimalist Enquiries/Derivation by Phase* framework also informs recent work by Franck et al. (2006). Franck et al. reason that rather than having a single hierarchical representation over which all syntactic operations are computed, several intermediate representations need to be postulated which reflect the cyclic derivational steps of the final structure:

Under the assumption of a tight connection between grammar and processor (an assumption which is rarely made ever since the pioneering days of the derivational theory of complexity, but which seems to us to be the null hypothesis), the different derivational steps assumed in linguistics should be traceable in linguistic performance, and for our concerns here, in the way speakers err when producing agreement. (p. 179)

This means that agreement is a derivational and configurational phenomenon which must be computed on tree structures on the basis of hierarchical relations such as c-command and locality. First, in their work, the authors claim that they have proved the superiority of c-command over strict adjacency in *attracting* agreement. At a second stage, departing somewhat from minimalist orthodoxy, Franck et al. point out that agreement involves the two separate components of AGREE and MOVE and that the robustness of agreement in S–V structures (as opposed to V–S structures) is linked to the fact that features are checked twice: (1) through AGREE, following standard assumptions in Chomsky (1995), and (2) in the strictly local Spec-head configuration, after movement of the subject. In V–S structures only one check is conducted (AGREE), with the result that agreement to the right of the verb is generally less consistently – or less syntactically? – marked. In three experiments of elicited spoken production in French and Italian, they manipulate *attraction* errors and strive to show that agreement depends on abstract, hierarchical representations, not on the final surface word order. For instance, the interference effects they claim to have obtained with object clitics in French when they do not intervene between the subject and the verb in the surface structure can only be explained under the assumption that clitics transit via an intermediate position (AgrO) which does intervene on the subject–verb relation. According to the authors, that constitutes “an entirely new illustration of a behavioural reflex of intermediate representations” (p. 210).³

But perhaps the most interesting connection between current generative thinking and experimental work on agreement has to do with issues such as Pesetsky’s (1995) *Earliness Principle*, and especially Chomsky’s (1999, 2001) notions of *feature interpretability*, *legibility* of the core, *phase*, *Phase Impenetrability* and *locality*. The spirit of these proposals is to restrict the domain of application of formal operations like uninterpretable feature deletion, and, despite timid claims to the contrary, it seems clear that they are driven by processing considerations. Thus, for instance, Chomsky (2001:13) insists that probe-goal relations (such as the relation between *controllers* and *targets* in agreement) must be local “in order to minimise search”.⁴ In order to reduce the “computational burden” (1999:9), he proposes that the derivation of expressions proceeds by phase, and that “phases must be as small as possible, to minimise memory” (2001:14; Chomsky seems to identify phases with clausal units). The *Phase Impenetrability Condition*, in particular (any goal in the c-command domain of a phase head is impenetrable to a goal outside the phase), seeks to ensure that, once a phase has been completed, there occurs a transfer of its computational content to the phonological and semantic components for building appropriate phonological and semantic representations, which means that the contents of the *previous* syntactic computation are no longer available. The basic insight is that the ‘language faculty’ can only process limited amounts of structure at one time, so locality and phases do the job of ensuring that everything proceeds by steps and that each step does not involve long searches in memory. This serial thinking, and the obvious derivational cyclicality on which it rests, is precisely the kind of theoretical claim that can be addressed in a laboratory. All one needs to do is to seek to falsify it by showing evidence of semantic *penetration* inside putative phase domains. This is after all what psycholinguistic research is all about. And it is also very much in essence a return to the old days of the derivational theory of complexity. Thus, for instance, in a series of both production and comprehension experiments, Thornton and MacDonald (2003) studied the role of plausibility in subject–verb agreement operations in English and claim to have found robust effects. The plausibility of the verb was manipulated

³ This is a conclusion that I find counterintuitive, but this is not the place to go into that. That said, I am not aware of recent research aimed at either confirming or rejecting it.

⁴ A note on terminology: agreement is generally seen as a matter of systematic covariance (Hudson, 1999), but in a typical case of agreement, say the subject–verb type (S–V), the element that determines the formal features of another element is usually known as the *controller*, while the controlled element is usually called the *target*. There are other terms, like *probe* and *goal*, or *trigger* and *source*, but in general they all reflect the asymmetry that is intuitively strong in at least central cases of agreement: for instance, in S–V agreement in English, the person, and number features of subject pronouns are usually ‘copied’ on to the verb, not the other way round.

so that either the two nouns of a complex subject NP could be plausible passive subjects (*the album by the classical composers . . . BE praised*) or only the head noun could be so (*the album by the classical composers . . . BE played*). There were more agreement errors in production and longer RTs (reaction times) at the verb in comprehension when both nouns were plausible subjects than when only the head noun was plausible. They interpreted their findings as the result of the operation of multiple soft constraints, a typical *constraint satisfaction* view of grammar and processing that rivals Chomsky-inspired, rule-based accounts. Hawkins (2004) provides an account of agreement that makes use of both formal notions like c-command and interactive constraints.

Another interesting issue which arises in connection with agreement is *feature mismatch* (Corbett, 2006:143 ff.). Chomsky (1995:309 ff.) insists that “mismatch of features cancels the derivation”, but certainly, even at the grammatical level mismatches are very often allowed. In *the committee have postponed any decision*, it clearly does not do to invoke some sort of underlying plural feature on *committee*, since instances like **these committee are satisfied* are clearly ungrammatical. This calls for a view of *semantic agreement*, which is no more than a way of sanctioning mismatch (some believe that agreement is *always* semantic: Barlow, 1999; Pollard and Sag, 1988). In *unification* theories like HPSG and LFG (Shieber, 1986; Copestake, 2002), where agreement is seen as a matter of cumulating partial information from the controller and the target, mismatches are as much of a problem. An interesting aspect of mismatches is the processing dimension. For instance, Gernsbacher (1991) showed that people have no difficulty comprehending a plural pronoun when they deal with antecedents which are grammatically singular but conceptually plural. In general, mismatches echo the penetrability or impenetrability of hypothesized components in the grammar and the processing of agreement. This will be an important issue here.

Cognitive Grammar, a theory which explicitly advocates the need to pass the psychological adequacy test, is paradoxically up against maybe the greatest challenge to its foundations when it comes to explaining agreement. Above, we witnessed Taylor’s reaction to the agreement puzzle: agreement is not obviously functional and, therefore, since it is nevertheless very widespread, it must reflect people’s delight in playing with the form of language. I take that conclusion to be naive at best. More aware of the need to face up to the challenge, Langacker (1991:289 ff.) has denied that gender and agreement systems are quintessential instances of the autonomy of form, and has consequently strived to prove their symbolic nature:

Indeed, the apparent arbitrariness of gender assignment over most of the lexicon in European languages is generally the first embarrassing fact to be thrown in the face of anybody with the audacity to suggest that grammar might be semantically based. (p. 304)

Agreement markings are perhaps the archetypal example of sentence “trappings” employed for purely grammatical purposes, and are supposedly inconsistent with any claim that grammar might have a semantic basis. (p. 307)

Against this, Langacker argues that suffixes like *-a* or *-o* in Spanish, when used with inanimates like *mesa* (table) or *cerro* (hill), are noun-forming endings which are schematic for the class of nouns overall. That is, semantically, each is equivalent to the noun class schema: [THING/*-o*] and [THING/*-a*]. The two morphemes are thus seen as polysemous, with each having a prototypical variant which is used with sexed entities and specifies sex, as well as another variant which occurs with unsexed entities and is highly schematic at the semantic pole. As far as agreement goes, Langacker (p. 307) views redundant agreement markings as schematised in [A + *x* . . . B + *x'*] as predications in their own right and, “since it would be counter to both the letter and the spirit of cognitive grammar to describe this situation by a rule that ‘copies’ *x* from A onto B, or in terms of features ‘percolating’ up, down, across, around or through”, *x* and *x'* are both analysed as “meaningful symbolic units”. In a way, this view is reminiscent of declarative, unification-based approaches. In a system where all possible connections constitute constructional *schemas* of higher or lower schematicity, agreement constitutes a case of a higher-level schema that abstracts away from all the actual elaborations (including sexed and unsexed cases). And since the schemas are bound by *the content requirement*, what is the function of so much redundancy?

(. . .) I have no quarrel with the traditional notion that the agreement serves the function of signalling grammatical relationships; it might indicate, for example, that B modifies A, or that A is an argument of B. I would only reiterate in this regard that serving a specifiable grammatical function is perfectly consistent with being meaningful. (p. 308)

One might object that such a curt “traditional” explanation, which may seem right for S–V agreement, still falls short of explaining the intriguing difference between English and Spanish NPs in (1) and (2) above (after all, once we have the article and the noun projecting a noun phrase at the top, agreement on intervening adjectives is not so obviously needed to mark their grammatical function).⁵ Be that as it may, for now, it must be added that there is now a large body of experimental evidence on *attraction* which *might* receive a cognitive explanation. The fact that the attraction effect is robust with the number feature is interesting in that number is much more semantically based than gender (Corbett, 1991, 2000, 2006; Eberhard et al., 2005), which may point to raw cognitive salience. This might be studied in a cognitive framework. Finally, another area where cognitive thinking might be useful is the issue of feature mismatches, referred to above: mismatches between controllers and targets are problematic for any formal theory but may easily be seen as non-prototypical agreement choices which, tied to specific constructions and motivated by various *inheritance* ties, must be listed (Goldberg, 1995). Again, this is an obvious advantage of declarative approaches.

As may be seen, agreement may be “disarmingly simple in appearance” but also a “morass for linguistic and psycholinguistic theories” because, in fact, it “is not only syntactic, not only semantic, and not only pragmatic, but all of these things at the same time” (Eberhard et al., 2005:531). Here, I will use recent linguistic and psycholinguistic evidence from English and Spanish in the main to discuss a few central aspects of its grammar and its processing, in the following way: first, in section 2 the formal and conceptual *sources* of agreement will be shown to cause frequent overrides. This section will also contain a description of the *domains* and kinds of agreement and especially an evaluation of the *Feature Hierarchy Hypothesis* in the light of recent psycholinguistic findings on the processing of gender and number. Section 3 will tackle cross-linguistic differences in agreement with English morphological *attrition* vs Spanish *redundancy* in mind. The evaluation of recent experimental research on the processing of agreement in the two languages (most of it employing the ERP methodology) will also fill most of this section. It will be argued that, unlike Spanish or Italian, English does not have an automatic ‘bonding’ stage (Garrod and Terras, 2000) in the computation of agreement features at the far end of Corbett’s *Agreement Hierarchy*: that is, with pronouns. Finally, section 4 is an epilogue that seeks to re-assess the function of agreement by arguing that grammatical gender is a quintessential generative device. Finally, it should be added at this point that even though several aspects of both the grammar and the processing of other languages, such as Italian or German, are also discussed here, this work has no typological pretensions, although it is hoped that typologists find it inspirational.

2. Formal and conceptual sources of agreement features

2.1. Semantic overrides

As noted, agreement shifts from the formal to the notional with ease. At the grammatical level, and inside the structure of the sentence, formal criteria tend to dominate the scene. Thus, even in English, where agreement morphology is poor and largely restricted to the number feature, whatever little morphology remains is usually formally regulated in the main (e.g. the third person singular is marked by an affix on the verb). In Spanish, there is widespread syntactic agreement in person (verbs), number, and gender (verbs, nouns, adjectives, and determiners). Agreement in gender is particularly arbitrary for non-sexed entities (*mes-a*, *banc-o*), and therefore semantically unmotivated, at least in principle (but see Mills, 1986; Konishi, 1993; Sera et al., 1994; Sera et al., 2002 for qualifications to that).

However, the foregoing merely reflects a typical situation, by no means a fixed rule. In *the hash browns at table nine is getting angry* (Pollard and Sag, 1994), the encyclopaedically-based metonymic conceptualisation of the plural subject phrase allows the verb to be singular. In the ‘nursely we’ construction (*we seem a bit displeased with ourself today*; see Joseph, 1979; also Harley and Ritter, 2002:507), a nurse typically uses a plural pronoun to address a single sick person, so despite the pronoun being plural the co-referential reflexive is singular. In fact, we have already mentioned in passing another two instances where the usual syntactic pattern of agreement is not, or may not, be present in English: the cases of ‘committee nouns’ and number-transparent nouns like *lot*, *number*, *heaps*, and the like:

- (5) The police **is/are** enquiring into that.
- (6) A number of ideas **were/*was** proposed.
- (7) Heaps of time **is/*are** being wasted.

⁵ I owe this point to an anonymous reviewer, to whom I am grateful.

However, there are many more: measure terms, fused relative constructions, coordinated nouns, etc. (see Huddleston and Pullum, 2002:504 ff., from whom many of the following examples are taken):

- (8) Twenty dollars seems a ridiculous amount to pay to go to the movies.
- (9) Three eggs **is** plenty.
- (10) One percent of students **take**/***takes** drugs.
- (11) He withdrew his motion for what **were** obviously very sound reasons.
- (12) What is needed **are** managers with new ideas and the will to apply them.
- (13) What **are** going to be the deciding factors?
- (14) Eggs and bacon **is**/***are** my favourite breakfast.
- (15) The hammer and sickle **was**/***were** flying over the Kremlin.

The same situation obtains in subject predicative agreement across a copula. In principle, given the equative relationship, both constituents (the subject and the subject complement) ought to be feature-matched, and indeed they tend to be, as in (16) and (17). However, mismatches based on semantic interference are common in the features of person (18) and numbers (19) and (20):

- (16) The doctor **is an** honorary member/*honorary members.
- (17) The doctors **are** honorary members/*an honorary member.
- (18) I am a doctor who believes/*believe in euthanasia.
- (19) The only thing we need now **is** some new curtains.
- (20) Our neighbours **are a** nuisance.

Even in Spanish, where morphosyntactic muscle is strong, there are SV agreement mismatches in the cognitively most grounded feature: person:

- (21) *Los hombres somos unos desgraciados.*
the men-3PL are-PRS.1PL a-PL miserable-PL
'We men are miserable'
- (22) *Las abogadas sois muy exigentes.*
the women-lawyers-3PL are-PRS.2PL very demanding-PL
'You women-lawyers are very demanding'

In processing, there are two well-known areas where conceptual overrides are pervasive: so-called *conceptual anaphors* (which involve number) and conceptually-based *gender stereotypes* (gender). These two cases involve pronoun-antecedent ties. In processing, I take semantic override to mean that choice of a mismatched agreement pattern takes no toll in reading times, or is even preferred over its matched formal competitor. In principle, as in grammar, there is evidence that formal cues such as markings for gender and number facilitate antecedent-pronoun coindexation (Ehrlich, 1980; Garnham and Oakhill, 1985; Garnham et al., 1995; Arnold et al., 2000). Thus, when English pronouns are clearly marked for gender and number, subjects typically identify their antecedents more rapidly and read sentences containing such proforms more quickly than otherwise. However, Gernsbacher (1991) demonstrated that conceptual anaphors show the opposite pattern of behaviour: even in the presence of formal cues, the semantically motivated choice of form is often preferred. Gernsbacher used three types of antecedents: collective expressions, such as *a basketball team*, generic types, as in *a book* in general (as opposed to a particular book that someone may be reading), and a multiple item/event noun which refers to an item that everybody is supposed to have many instances of, such as *a plate*. For instance, in:

- (23) After college, my sister went to work for IBM. They/it made her a very good offer,

continuations with *they* were preferred to (faster than) continuations with *it*. Carreiras and Gernsbacher (1992) replicated basically the same findings in Spanish. These examples show that, even though syntax-semantics mappings in number are usually transparent (iconic), when they happen to be opaque, it is semantics that wins out in the end (Corbett, 2000). The *distributivity* of complex NPs like *the label on the bottles* (which must imply that every bottle has

one label, as opposed to, say, *the material of the bottles*) has also been shown to illegally *attract* plural agreement (Eberhard, 1999), and has been interpreted in the (production) *attraction* literature as “evidence that verb number can be notionally controlled” (Eberhard et al., 2005:536).

Although number is the feature that tends to show the greatest sensitivity to semantic interference and gender the one that shows the least, stereotypical gender shows that at times semantics takes its toll there as well. Carreiras et al. (1996) wanted to see how fast pragmatic knowledge is used in pronoun-antecedent ties. In order to do that, they used materials including stereotyped and neutral gender, as in (24) and (25) below:

Stereotyped

- (24) The footballer wanted to play in the match // He/she had been training very hard during the week.

Neutral

- (25) The psychologist studied the student’s test results // He/she was concerned about the low standard.

The reading times clearly indicated not only that the mismatch condition is the hardest to process, but also that, when the gender stereotype is solid, there is a robust facilitatory effect. Similar findings have recently been reported by Kennison and Trofe (2003), and Sturt (2003). In fact, it is noteworthy that in a study by Osterhout et al. (1997) a stereotypical gender agreement violation was perceived as if it were a grammatical gender violation in that both elicited the same ERP brain waveform, the P600 effect normally attributed to syntactic reprocessing (see section 2.4.2).

2.2. Formal overrides

Italian is a language for which quite a large amount of work on the processing of the gender feature has been done, and it offers an interesting contrast with generally gender-impooverished English. For instance, Cacciari et al. (1997, submitted for publication) studied pronoun-antecedent ties involving *epicenes* like *vittima* (‘victim’) and *personaggio* (‘character’). Notice that these nouns are clearly gender-marked but, despite that, they can refer to both male and female referents. In other words, epicenes have one morphosyntactic gender, but two semantic genders. Thus, *la vittima*, for instance, is gender-marked as feminine but it may be used to refer to either a female or a male referent. *Personaggio* presents the opposite case. An interesting question arises as to whether in such cases the conflict between the formal marking and the sex of the referent will impact the processing scene. In order to see that, Cacciari et al. (1997) devised materials like those in (26) and (27):

- (26) La vittima del incidente stradale sbatté violentamente la testa contro il finestrino. Lei (lui), perciò, perse molto sangue e svenne.
‘The victim of the car accident violently slammed the head against the window. She (he), therefore, lost a lot of blood and fainted’.
- (27) L’erede decise di andare in vacanza con I soldi ricevuti dalla zia. Lei (lui), perciò, progettò un lungo viaggio negli USA.
‘The heir decided to go on vacation with the money received from the aunt. She (he), therefore, planned a long trip to the States’.

Erede nouns (those like Italian ‘heir’), which are not gender-marked and in that respect are like ordinary English nouns, were used as a baseline. These bi-gender words do not allow decoders to infer gender based on the suffix, for suffixes like *-e* are gender-opaque in Italian (as opposed to *-a* or *-o*, which are normally clearly associated with feminine and masculine, respectively). In their more recent experiments (Cacciari et al., submitted for publication), the experimenters manipulated the contextual information preceding the antecedent, which was either an epicene or a bi-gender word. They created three contexts: (i) neutral contexts which provided no information on the gender of the referent; (ii) congruent contexts where the referent’s natural gender matched the grammatical gender of the epicene (e.g. a woman referent for *la vittima*, and a man referent for *il personaggio*); and (iii) incongruent contexts, where the opposite situation obtained: syntactic and semantic gender were not matched (e.g. *la vittima . . . lui*):

(28)

L'erede	Lei (Lui)	baseline
La vittima	Lui	non-match
La vittima	Lei	match

(Cacciari et al. 1997, submitted for publication)

Using a moving-window technique, Cacciari et al. sought to find out whether discourse-based gender information prevailed in the earlier phases of pronoun assignment—or whether, alternatively, grammatical cues would do so. If the initial coindexing of antecedent and pronoun is established on the basis of grammatical information, the specific prediction is that there should be faster reading times for grammatically matching anaphors (*la vittima . . . lei*) in the pronoun region, irrespective of contextual biases. In contrast, when the antecedent is a bi-gender word like *erede*, pronoun assignment should be driven from very early on by the only source of information available: the context. If, conversely, context rules the parse in all cases, as predicted by constraint satisfaction models (MacDonald et al., 1994; Tanenhaus et al., 2000), then the same facilitatory effect should be obtained for both epicenes and bi-gender words. In short, what they found was that formal cues ruled the parse for epicenes just as rapidly as semantic cues ruled the parse for bi-gender words. I take that to be a very opportunistic parser, one that knows quite a good deal about grammar, and specifically, about which parts of the grammar of languages are tightly regulated and which are not.⁶ What is interesting is that, when they exist, formal morphosyntactic cues take precedence, and that this occurs in Italian, a richly inflected language like Spanish.

2.3. Domains and kinds

Notice that the Italian data concern pronoun-antecedent ties, that is, a type of agreement that is established at long distances. In a previous section we briefly touched on the notion of derivation by phase in the current *minimalist* tradition. In passing, it was said that Chomsky devised a notion of phase such that it was firmly reined in by strict locality considerations. When it comes to understanding locality in agreement, it is useful to consider Corbett's Agreement Hierarchy first (AH):

The Agreement Hierarchy (Corbett 1979)

Attributive > predicate > relative pronoun > personal pronoun

The AH is intended to illustrate the effect of semantic interference in the different **domains** where agreement takes place. For instance, “*committee* is not uniquely singular or plural: we find both possibilities in the predicate (. . .). However, it is not simply that *committee* is singular or plural. In a different agreement domain there is no choice” (Corbett, 2006:206):

(29) a. this committee b. *these committee

⁶ A similar opportunism was evident in a recent series of experiments on the processing of controlled PROs in Spanish. In Betancort et al. (2006), subject- and object-controlled verbs were compared with a view to examining whether the coindexation of PRO and its controller is faster in object-control verbs like *persuade* than in subject-control verbs like *promise*, due to the proximity of the controller to the PRO gap in the former. The saturation of the PRO gap in control verbs was compared with the gap in adverbial structures containing prepositions *por* and *para* (as in *Juan se enfadó con María por/para PRO marcharse antes*, ‘John got mad at Mary in order (for him) to leave sooner/because she left sooner’). The eye-tracking measures revealed that recency indeed favoured the recovery of the PRO gap in verbs (object-control advantage). However, the fact that no trace of reanalysis could be found in subject-control verbs showed that lexical knowledge was also used very early on with verbs. Interestingly, with prepositions only recency showed its influence: object-control *por*-clauses were faster than subject-control *para*-clauses, and reanalysis was evident in these latter. The reason for the different behaviour of verbs and prepositions seems to be that verbs code PRO control rigidly. However, contextual manipulations do allow one to make a *para*-clause, for instance, a case of object control. The parser seems to ‘know’ that it can trust lexical specification in verbs, so it uses it from the very beginning of its parse blindly. With prepositions it cannot be so sure, so it prefers to fall back upon its usual processing reflex of opting for closer ties over distant ones.

Agreement therefore changes with the domain, even if the controller remains the same:

- (30) a. The committee, which has ... b. The committee, which have ...
 (31) a. The committee ... It ... b. The committee ... They ...

The way the AH works is as follows:

For any controller that permits alternative agreements, as we move rightwards along the Agreement Hierarchy, the likelihood of agreement with greater semantic justification will increase monotonically (that is, with no intervening decrease). (p. 207)

It will be seen that, with the exception of the position taken by relative pronouns, the other three positions in the AH usually involve different distances between controller and targets. Thus, attributive modifiers and determiners are almost always adjacent to their controllers, verbs a little farther away and pronouns even farther. By way of illustration, consider the associative construction in some varieties of Russian (Bogdanov, 1968):

- (32) *moj* *brat* *tam* *toža* *žy-P-i*
 my [M.SG] brother(M) [SG] there also live-PST-PL
 ‘my brother and his family also lived there’

The plural verb shows that it is not just ‘my brother’ who lives there. That plural form on the verb shows semantic agreement, therefore, but in attributive position there is singular (syntactic) agreement between *moj* and *brat*. This is a typical pattern. In Spanish, one typically uses *el fenómeno de* construction, as in *el fenómeno de tu hermana está cansada ahora* ‘the phenomom (masc) of your sister (fem) is tired (fem) now’ (*cansado* would sound odd), with *el* and *fenómeno* entering ordinary formal agreement inside the subject NP, and *cansada* showing semantic agreement across the verb (Casillas Martínez, 2003). Levin (2001) conducted a thorough investigation of ‘committee nouns’ in a corpus of three newspapers written in English in 1995. With a total of 26 nouns, he managed to prove that the predictions of the AH were fulfilled, as agreement with verbs, relative pronouns and personal pronouns showed the expected monotonic increase in semantic motivation.

The AH has some interesting repercussions for linguistic theory. The most obvious one is that, as Corbett himself points out, agreement choices are not a sharply-delimited phenomenon, but rather gradient (Berg, 1998). In particular, the fact that the same controller can control different feature values on different targets at the same time is problematic for almost all theories (Corbett, 2006:227). Another consequence of it is that the sharp distinction that is often established between pronoun targets and the other forms of agreement is undermined. Anderson (2004), for instance, takes this dichotomic position explicitly. However, Barlow (1992:134–152) devotes an extensive discussion to this issue and concludes that the separateness is not linguistically justified (see also Moravcsik, 1978:334; Lehmann, 1982:211; Siewierska, 1999:225; Bock et al., 2004; Eberhard et al., 2005; Corbett, 2006:21). In fact, pronouns typically realise the same features as other targets and are subject to the same kind of *resolution rules*. Corbett (2006:227 ff.) cites Grevisse (1964:405–406) on the following French example, which shows that syntactic agreement is used on the personal pronoun even in the presence of a tempting occasion for semantic override:

- (33) *Votre Majesté (fem) partira quand elle (fem) voudra*
Your majesty will leave whenever she wants

In general, it appears safer to view domains in terms of distance between controllers and targets, therefore, instead of in terms of categorical distinctions, with distance understood structurally, rather than linearly (Corbett, 2006:235–236). The really interesting aspect of the Italian data is that in a morphologically-rich language like Italian (or Spanish), even the processing of long-distance pronoun-antecedent ties seems to be driven by blind, non-strategic computations.

In fact, the issue of whether antecedent-pronoun ties are established by taking into account specific lexicogrammatical properties of the antecedent first or, rather, pragmatic and semantic information divides the processing literature. Indeed, the issue has become a test case for modular vs interactive models of sentence comprehension.

According to the former, agreement is a serial, two-stage process (Garrod and Sanford, 1990; Garrod and Terras, 2000; Rigalleau et al., 2004). Thus, for instance, in the *Bonding and Resolution* model of Garrod and Terras (2000), there is an initial *bonding* stage where a superficial attachment between a pronoun and a possible antecedent takes place, and a subsequent *resolution* stage, when full referential interpretation of the pronoun occurs, including the checking of the initially established link against the discourse scene. During the first component of the process, only formal cues are accessed and semantic penetration is hypothesized to be impossible. By contrast, interactive models (Bates and MacWhinney, 1989; Taraban and McClelland, 1988; Marslen-Wilson et al., 1993; MacDonald et al., 1994; Tanenhaus et al., 2000) claim that agreement is resolved by simultaneously weighing the force of a number of possible sources of information (constraints), notably the information provided by the lexicon (including lexical meaning) and the context, but also that provided by grammatical cues. All forces are supposed to intervene in parallel (hence in one stage), and the ultimate parse is the result of the competition among the different attractors. The Italian data suggest that in morphologically-rich languages at least, a serial route makes more sense.

2.4. The Feature Hierarchy Hypothesis

2.4.1. The grammar of features

The reader may have noticed that there are three different dimensions cutting across the literature briefly reviewed above. On the one hand, there are methodological differences among the various experiments. On the other, there seems to be at least initial motivation for exploring the view that different results are obtained in different languages. This paves the way for a thorough investigation into how agreement systems fit into the overall system of each particular language. In particular, the fact that languages may be either inflection-rich or inflection-poor at large seems particularly relevant to the self-organization of agreement systems. The third dimension alluded to above is feature specificity. Indeed, over the past few years there have been attempts to differentiate gender and number, both at the grammatical level and at the processing level. I focus on features now.

With the exception of gender stereotypes, most semantic overrides affect the number category. This is suggestive of a certain cognitive salience of this feature (in line with Greenberg's (1963) universals 36, 37, and 45, for instance). The typological literature has long ago sanctioned differences in features by capturing their interrelationships in an implicational hierarchy, as in (34):

(34) Feature Hierarchy: Person > Number > Gender

Initially, the motivation for the hierarchy lies mainly in the fact that person is a more firmly established feature in the world's languages than number, and number is more firmly established than gender (Greenberg, 1963). So the implicational hierarchy is meant to imply that if a language has gender, then it must also have number and person, and if it has number, then it must have person, but not necessarily gender. Within each feature subhierarchies have often been proposed (Silverstein, 1985). Very commonly, a split between 1st/2nd, on the one hand, and 3rd, on the other, is proposed for person, for instance (Benveniste, 1971; Tsoulas and Kural, 1998).

However, one thing is typological prevalence, quite another to be able to prove that features are indeed differentiated in terms of degree of cognitive strength—and that this is reflected either in the grammar or in the processing of a given language, or in both. In general, the fact that number is much more prone to semantic interference does suggest that the issue of feature differentiation is worth exploring. In principle, number is indeed more obviously based on a rather neat and objective semantic basis (or 'experiential' basis): singularity vs multiplicity of tokens. Corbett (2000:2) starts off his monograph on number by pointing out that expecting that this feature must be expressed is an "Anglocentric assumption", since there are languages, like Bayso, where there are forms which allow the use of a noun with no number specification at all. However, such typological 'rarities' cannot take one's attention away from the fact that number *is* less likely to be arbitrarily coded than gender. It has been pointed out that number is an *inherent* feature of the noun, that is, one that is not required by the syntactic context (it is, conversely, a *contextual* feature on adjectives or determinatives, for instance; see Booij, 1996). One may very well find a language which uses a gender for women, fire and dangerous things (Lakoff, 1987), or for living creatures, including plants but excluding, say, pigs and monkeys, but it would indeed make front page news to find a language which uses, say, a morphological form for 7 plus or minus 2, and another for

calendar digits only. A distinction between singular, dual, trial, paucal, and plural is usually all one can expect to find at the most (see Harley and Ritter, 2002:497 for the pronominal system). And even all those few distinctions occupy the same unidirectional dimension. Corbett (2000:1) points out that the richest number systems in the world have five number values. The obvious cognitive salience of number thus stands in a marked contrast with the impressive plenitude of gender distinctions in the world's languages (Corbett, 1991:148; see his discussion of Fula's 20 genders, for instance; pp. 190 ff.; see also Harley and Ritter, 2002:514). Indeed, when it comes to gender distinctions, and to the categorizations that underlie them, one wonders if there are any limits at all (there are, of course, also 'exceptions' like Tamil, with highly motivated gender). In Spanish, of course, number is more clearly motivated than gender.

Given the clear difference in semantic motivation for both features, in principle it comes as a surprise that theories of grammar have generally been insensitive to them. In Cognitive Grammar, raw cognitive salience might in principle be seized upon by grammarians keen to establish the cognitive motivations of grammaticalization, but this, to my knowledge, has not happened yet (when it comes to motivating a gender vs number contrast). We have already noted how agreement in general poses a serious problem for cognitive theoreticians due to both the apparent useless redundancy of agreement markings and the apparent arbitrariness of gender. But clearly, number is not so arbitrary, so one would expect cognitive grammarians to have something to say about it. As noted, Langacker (1991:307 ff.) views agreement patterns as predicative constructions where both the controller and the target are symbols in their own right, with the same status in the grammar. This is supposed to apply to the two features, gender and number. In sum, as regards feature differentiation, it seems that Cognitive Grammar is not particularly interested.

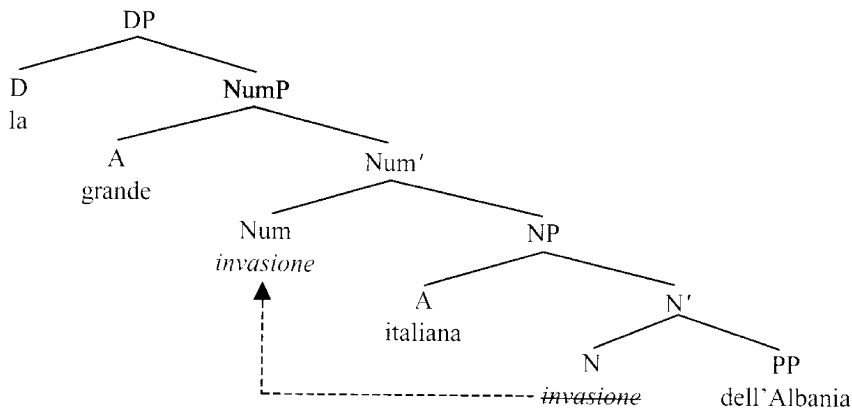
Generative grammar has generally paid more attention to agreement and features. In principle, however, in the current minimalist theory, it is generally assumed that nominal categories access the derivation with all their *phi*-features present. It is also assumed that all the features are 'checked' in the same way. The idea is that subject–verb agreement creates a specific syntactic projection of a functional kind, AgrS, where the morphological specifications are determined in a precise stage of the derivational cycle. The emphasis has always been on the fine delimitation of the cycles⁷: first MERGE fuses together the thematic core of the predication (the verb and its complements); then further MERGE operations fuse the thematic nucleus with the functional structure (specifications of mood, tense, etc.). At this stage, the functional node AgrS enters into an agreement relation with the thematic subject (which is still VP-internal). Since the person and number features of AgrS need to be valued in English and the subject possesses such features, these are 'copied' onto AgrS. This copying occurs under the configurational requirements of locality and c-command: AgrS, the 'probe' of the AGREE operation, looks for a 'goal' inside its c-command domain. Once feature copying takes place, the V moves to AgrS to receive its morphological specification. At an even later stage, the thematic subject moves to the specifier of AgrS leaving a trace behind.

As can be seen, AGREE is at the core of the construction of sentence structure. In particular, AGREE is motivated by the need to erase uninterpretable features (Chomsky, 1999:4) in the sense that it is the presence of these features on a constituent that makes it 'active' and causes it to look for a matching 'goal' within the relevant c-command domain. However, though gender may or may not be semantically transparent, it is supposed to enter the derivation as a valued feature anyway, since, even if it is not transparent to the semantic component, its value is determined – arbitrarily – by lexical specification. *Valuation* (unification of feature values) is what makes features legible at the interfaces. In this way, the two features are assumed to be syntacticized (and, one may surmise, processed) in the same way.

However, the fact that number *is* such a semantically transparent feature has led some authors (Ritter, 1988, 1991, 1993) to argue that NUMBER, as a feature, should head its own functional projection in the way schematised in (35) for the Italian phrase *la grande invasione italiana dell'Albania* (*la-ART.F grande-ADJ.F invasione-NOM.F italiana-ADJ.F della-CONTRACTION.de-la Albania*; 'the great Italian invasion of Albania'), which I take from Radford (2004:179), who toys with that idea:

⁷ In fact, Chomsky (2001) has claimed that there may not be any order of operations inside the phase, but Nunes (1999, 2004) and Epstein and Seely (2002, 2006) have criticised this view on the grounds that it compromises the derivational – cyclic – nature of the whole system of step-by-step structure building, where no simultaneous operations can take place. See Fernández-Salgueiro (2005, 2006) on this (I thank the author for bringing this point to my attention).

(35)



If a feature is a syntactic head, then it is available to the parser for all syntactic operations, including agreement. Number typically has variability (sg/pl) and semantic content, which means that it fulfils the requisites of a syntactic head. By contrast, grammatical gender is not variable and cannot be selected from the lexicon independently of the other nominal features. This is why, according to Ritter (1993), it cannot be the head of its own syntactic projection.⁸ This is a contentious claim, but its proponents often rest on a number of psycholinguistic findings to back it up. To these we turn now.

2.4.2. Psycholinguistic evidence for and against the FH

Using a priming task that combined auditory presentation and lexical decision in pairs of words, Faussart et al. (1999) conducted a series of experiments in French and Spanish and claim to have found differences between gender and number in Spanish, not in French. The authors argue that the result showed up only in Spanish because French no longer marks number phonetically. According to them, the differential treatment that the processor seems to reserve for each feature can be explained in the framework of Bradley and Foster's (1987) model of lexical access. In this model there are three cyclic phases in word retrieval: access, recognition, and integration. In the first stage, the right lexical entry is located and lexical identification takes place. In the second stage, the lexical content of the entry is accessed. This includes semantic information, word category and morphological information. Finally, the third stage accomplishes post-lexical processes of integration with the syntactic context. Grammatical agreement would take place during this last phase. According to Faussart et al. (1999), gender agreement violations would surface in stage 3 and would cause the processor to return to stage 1 to check if the right lexical entry has been chosen. Since number is not a stem-inherent feature, in number agreement violations the processor would only have to check inside the current stage 3 domain without returning to the initial processes of lexical access (stage 1). So the difference in access time when detecting either a gender or a number inconsistency would reflect the additional cost of reinitiating lexical access in the former.

With a similar cross-modal priming technique,⁹ Nicol (1988) investigated the processing of pronouns in English. These could be disambiguated by gender, as in (36), or by number, as in (37):

- (36) a. The ballerina told the skier that the doctor would blame **him** for the injury.
 b. The ballerina told the skier that the doctor would blame **her** for the injury.

⁸ Picallo (1991) also argues for the separateness of number and gender but proposes different phrasal projections for the two features, not just for number.

⁹ The methodology of cross-modal priming consists in mixing modalities of perception. Usually, subjects in experiments must read something on a computer screen (for instance, to decide whether the stimulus is a word or not, that is, a lexical decision task) while something else is heard at the same time.

- (37) a. The landlord told the janitors that the fireman with the gas-mask would protect **him** from getting hurt.
 b. The landlord told the janitors that the fireman with the gas-mask would protect **them** from getting hurt.

The sentences were presented orally and, immediately after the pronoun was heard, a target word was visually presented for a lexical decision. The word could be either related or unrelated to one of the two antecedents, and the assumption was that, if featural information is used early in processing, only the referents with the same feature as that of the pronoun should be re-activated through priming. The author found a significant priming effect for number. For gender something curious happened: masculine pronouns did prime masculine antecedents, but feminine pronouns primed both masculine and feminine antecedents.

A more definitive pattern of results was obtained in Italian by Di Domenico and De Vincenzi (1995) and by De Vincenzi (1999). Using the same priming methodology, these authors investigated Italian sentences containing object clitics that could be disambiguated by gender or number, as in (38) and (39):

- (38) Il lavoratore disse alla cuoca che la padrona de casa non poteva sentir**la**.
 ‘The worker (masc) told the cook (fem) that the landlady could not hear her’.
- (39) Il lavoratore disse alla cuoca che la padrona de casa non poteva sentir**lo**.
 ‘The worker (masc) told the cook (fem) that the landlady could not hear him’.

They found a significant priming effect 1000 ms after the pronoun was heard only for number, but not for gender. De Vincenzi (1999) registered a priming effect for gender only at sentence end. Interestingly, these results were obtained using semantic gender (sexed referents) rather than grammatical gender. The authors used these findings to argue for a modular processor that ‘sees’ number first because number is a syntactic projection. Gender is accessed later because it is bound to lexical processes which an encapsulated processor with a syntax-first filter can only access after mandatory phrase-building operations are under way.

In production, there is evidence that errors are much more numerous for number than for gender (but see Rusakova, 2001, cited in Corbett, 2006:278 for the opposite case in Russian NPs). For instance, in the Vigliocco et al.’s (1996b) completion study of agreement in Spanish and English, out of 2048 responses, only 3 errors in gender occurred, as opposed to 103 in number agreement. It is tempting to associate this to the fact that number is usually a part of the conceptual representation of the message, while gender is much less so, and very often (as with grammatical, arbitrary gender) not at all. Igoa et al. (1999) found that in Spanish exchange errors affected mostly number suffixes, occasionally number and gender, but never gender alone. Igoa et al. also report the results of an experiment in which they tried to elicit such mistakes by asking participants to switch the nouns on complex noun phrase preambles that were presented to them. For instance, the preamble *una niña de los gatos* (‘a girl of the cats (masc)’) should be turned into *unos gatos de la niña* (‘some cats (masc) of the girl’). They found that number stranding was far more likely than gender stranding: *un gato de las niñas* (‘a cat (masc) of the girls’) more likely than *unas gatas del niño* (‘some cats (fem) of the boy’). Again, these differences are often interpreted by assuming that gender is retrieved directly from the lexicon and added to the phrase structure together with the *lemma* of each word, while number is computed supra-lexically. On this logic it follows that the number feature is assigned directly to the phrases as something equivalent to closed class items (Bock and Levelt, 1994; see also Antón-Méndez et al., 2002). Another way of expressing these differences is to say that number is dealt with by *computation*, while gender is dealt with rather parasitically and cost-free when retrieving whole words from the word store. As already noted, Anderson (2004:3) reminds us of the habitual thesis that gender is an inherent feature of the lexical N, while number is an inherent feature of the syntactic N position. Hudson (1999) takes number agreement in English to be a feature that affects subject noun phrases, not nouns in general.

However, despite the logical thinking underlying the interpretation of the previous results, and despite the obvious differences in conceptual motivation, psycholinguistic evidence is much stronger for the view that both features are processed in very much the same way, and that differences emerge only in relatively late stages of processing. In the first place, the very same priming methodology that has been used to provide results that sustain feature differentiability has also produced the opposite pattern of results: no differences. Thus, for instance, Lukatella et al. (1987) compared violations in number, gender, and case in Serbo-Croatian (a richly-inflected language) and found no differences between the various disagreement conditions. Interestingly, these researchers also used a double violation condition

(in gender and case) but the effects obtained did not differ from the single violations, something that allowed them to argue for a processor that behaves in a binary way, that is, one that treats agreement as an on/off switch: if grammatical, processing is fluid; if ungrammatical in *any* way, it sounds the alarm. In the same way, Colé and Seguí (1994) also failed to find differences in a French study with word pairs.

But perhaps the most relevant evidence comes from the methodologies with the finest grain, especially the electrophysiological studies using event related potentials (ERPs). ERP has excellent temporal resolution and is therefore ideal for capturing the millisecond-by-millisecond time course of processing. For instance, the 1000 ms detection threshold that Di Domenico and De Vincenzi (1995) obtained for number using a priming methodology cannot really compete with the array of electrophysiological measures cast by brain waves ‘related to events’. Typically, these range from the circa 200 ms of syntactic violations (the N200 effect, that is, a negative waveform peaking at around 200 ms after the onset of the anomaly), to the P600 measure which usually signals special syntactic difficulty or reanalysis (positive wave form peaking at 600 ms post-anomaly approximately). The N400 range of violations is normally associated with lexical processes and processes of semantic integration. In a pioneer study, Kutas and Hillyard (1983) found an increase in negativity between 200 and 500 ms in anterior areas of the brain for number agreement errors in S–V agreement. Similar effects have been described later and come to be known as left anterior negativity (LAN) effects (Friederici, 1995). LAN is usually synonymous with the malfunction of a syntactic analysis. N400 registers malfunction of a semantic or lexical component. Münte and Heinze (1994) undertook a series of experiments with agreement on word pairs. The experiments mixed languages (German, English and Finnish) and tasks (grammatical judgement and lexical decision). Among other types of violations, gender in article + noun combinations was studied as well. The results showed that agreement in article + noun combinations and S–V agreement produced the same effects: left anterior negativities that seem to reflect automatic processes. The authors themselves point out that such a pattern of results is very similar to the LAN effects found for other classic syntactic anomalies embedded in sentence contexts (Neville et al., 1991; Osterhout and Holcomb, 1992; Hagoort et al., 1993; Osterhout et al., 1994; Friederici et al., 1996). When they tested for semantic and phonological violations they obtained negativities with more posterior distributions, that is, in the N400 range. In a similar study, Barber and Carreiras (2003) did find an N400 type of effect in adjective + noun violations, but this was the same for gender and number. Osterhout and Mobley (1995) focused on English number agreement violations between subject and verb, reflexive pronoun and antecedent, and personal pronoun and antecedent:

- (40) *The elected officials hopes to succeed.
- (41) *The hungry guests helped **himself** to the food.
- (42) *The successful woman congratulated **himself** on the promotion.
- (43) ???The aunt heard that [she/he] had won the lottery.

They obtained effects in the P600 effect time window for the three types of violation. An additional LAN-type effect was found in S–V agreement only. Gender was manipulated in the pronoun condition: no differences with number were observed. Both the gender violation and the number violations produced results different from the semantic violations. Since gender is semantically (or, as some would say, ‘experientially’) based in English pronouns, it could be argued that their semanticity brings them closer to number. However, Hagoort and Brown (1999) reported the same effects in a recent study of Dutch grammatical agreement, Gunter et al. (2000) obtained both a P600 and a LAN effect in German, also with grammatical gender, and Barber et al. (2004) found exactly the same pattern of results in a comparison of Spanish violations like *el faro luminosal* ‘the luminous (fem) lighthouse (masc)’ (with grammatical gender) and *el abuelo delgado* ‘the thin (fem) grandfather’ (with semantic gender). Finally, Deutsch and Bentin (2001) examined S–V agreement in sentences like (44) in Hebrew:

- (44) The woman saw that the boy/diamond_{masc} had fallen_{masc/fem} into the pond

and obtained the classic N400 effect for semantic discrepancies, as well as a P600 effect for syntactic gender violations (although only when this was phonologically transparent).

The previous seems rather conclusive evidence for the view that early syntactic feature checking is the same for both gender and number and that, even for English S–V agreement, agreement is a fast, shallow, automatic process. Although comparisons are always difficult, perhaps the best follow-up of the widely-cited English data in Osterhout

and Mobley (1995) is a recent study by Barber and Carreiras (2005) using Spanish materials. The two Spanish authors conducted two ERP experiments involving word pairs and sentences with violations in number or gender agreement. In the first experiment they manipulated word pairs formed by a noun + an adjective, as shown in (45), as well as pairs formed by an article + noun, as in (46):

(45)

Gender violation:

Arroyo helada [*frozen* (fem-sg) *stream* (masc-sg)]

Number violation:

Arroyo helados [*frozen* (masc/pl) *stream* (masc/sg)]

(46)

Gender violation

La lago [the (fem-sg) lake (masc-sg)]

Number violation

Los lago [the (masc-pl) lake (masc-sg)]

It was hypothesized that the article + noun pairs would involve more syntactic processing because such combinations, unlike noun + adjective pairs, form maximal projections. If so, that should be visible in more anterior effects similar to those found by Münte and Heinze (1994) for this type of anomaly. If De Vincenzi (1999) is right about number and gender differentiability, checking agreement in such minimal and early stages of clause construction should prove informative. The authors found an N400 effect for both gender and number violations. As expected, article + noun sequences also produced an additional LAN effect, but also for both gender and number. No trace of a difference between the two features could be detected. Just as in the Lukatella et al.'s (1987) study on Serbo-Croatian, the double violation condition (in gender and number) did not differ from the single violations. In their second experiment, they embedded the same words used in their first experiment in either sentence initial position (for articles + nouns) or after the verb (for nouns + adjectives). In this way, a more natural reading process was achieved. Violations resulted in a pattern of LAN–P600 for both types of syntactic contexts, although the last segment of the P600 was more pronounced for the middle sentence position (more reanalysis needed). No differences between gender and number were found in the central part of both the LAN and the P600 effects. However, the later part of the 700–900 ms time window showed significant differences in the posterior regions of both hemispheres as well as in the central region of the right side, with larger amplitudes for gender than for number. As the authors point out, these data point to very late differences in the detection of gender and number disagreement during syntactic processing (recall that the 1000 ms effect of Di Domenico and De Vincenzi, 1995 was obtained even later). If we assume that the time window in question reflects the cost of reanalysis, then such late differences bring to mind Bradley and Foster's (1987) model of lexical access (see above), according to which gender is a stem-inherent feature and number not so. In this three-stage model, gender violations force one to reinitiate lexical access in stage 1, so reanalysis is expected to be costlier. But, crucially, what the reported studies suggest is that even though gender and number do differ in **storing**, they do not differ in **computation** when they are both used in agreement processes. This might make the psychological validity of the *Feature Hierarchy* a matter of how one is inclined to see the bottle: half-empty or half-full, with *empty* or *full* correlating with rapidity of processing. In the psycholinguistic literature, however, it is normally assumed that only very early measures reflect the really important aspects of the functional architecture of the linguistic mind: the architectural biases of the mind. Indeed, the whole battle between the different models is ordinarily resolved on the evidence of latencies registered well under half a second post-stimulus. On this view, psychological support for the FH is poor as differences between the two features arise too late in the time course of processing.

One possible counterexample to the previous conclusion could, in principle, be provided by Carminati (2005), working on Italian. Carminati has recently showed that the processing penalty for forcing *pro* to retrieve a referent introduced in object position – against *pro*'s well-established preference for antecedent subjects – is significantly reduced when features higher in the *Feature Hierarchy* disambiguate the pronoun. In particular, she compared (apart

from the person feature) the disambiguating power of gender and number in sentences like (47)–(50) below and used the superior performance of number over gender as evidence for her *Feature Strength Hypothesis* (the more cognitively important a feature is, the better it should be at disambiguating the pronoun that carries it):

- (47) Quando Maria cerca Roberto, diventa ansiosa.
‘When Maria looks for Roberto, she becomes anxious’.
- (48) Quando Maria cerca Roberto, diventa ansioso.
‘When Maria looks for Roberto, he becomes anxious’.
- (49) Quando Maria lo cerca, diventa ansioso.
‘When Maria looks for him, he becomes anxious’.
- (50) Quando **i Rossi** lo cercano, diventa ansioso.
‘When The Rossi look for him, he becomes anxious’.

In particular, Carminati claims that times for (49) are significantly higher than for (50), which includes number disambiguation. However, there are a number of problems of a methodological nature which compromise her conclusions. In the first place, her experiments used self-paced reading times, a methodology with a grain of resolution which cannot rival that of eye-tracking and, especially, ERP. In the second place, the author set the self-paced reading to occur for regions of one clause per bar press (p. 8): that is, one bar press introduced the whole of the subordinate clause, and the next introduced the whole of the main clause. This is clearly too crude a measure, actually diminishing the potency of an already not too powerful methodology. And in the third place, and perhaps more importantly, in her materials number information as in (50) is encountered early, i.e. at the verb, while gender information in (49) appears later, i.e. after the verb. This makes it possible for reanalysis to start earlier for the number-disambiguated condition than for the gender-disambiguated one. Taken together, Carminati’s results cannot seriously undermine the conclusion reached here that gender and number are computed in the same way—at least in languages in which these two features inhabit a rich inflectional paradigm that is put to use every second or so of actual parsing time.

What, then, should one predict for languages with less morphological muscle, like English? Can cross-linguistic differences in the parsing of agreement make sense? We turn to that now.

3. Cross-linguistic differences in agreement: attrition vs redundancy

3.1. Agreement and English attrition

In the previous section we have referred to evidence produced tentatively by [Kutas and Hillyard \(1983\)](#) first, and, more pointedly, by [Münte and Heinze \(1994\)](#) and specially [Osterhout and Mobley \(1995\)](#) later, that various agreement operations in different domains in English are resolved in a way which, electrophysiologically at least, looks very much like a syntactic process. Thus, [Osterhout and Mobley \(1995\)](#) obtained effects in the P600 time window for violations of agreement involving subjects and verbs, reflexive pronouns and antecedents (in both gender and number), and personal pronouns and antecedents. However, only a LAN-type effect was obtained for S–V agreement. This suggests that there are, as it were, degrees of processing speed that are domain-specific, with ‘domain’ meant to designate the place where agreement takes place: the NP, the clause, the sentence or the discourse ([Corbett, 2006](#)). Recall that in Spanish and Italian there is strong evidence that computation of formal agreement is fastest, irrespective of the features involved. In the context of a language like English, so prone to semantic overrides of formal agreement patterns (see section 2.1), it is tempting to suggest that the difference between a P600 response to gender/number agreement violations in pronouns and a LAN–P600 response to S–V agreement incongruities reflects the difference between ‘pure’ syntactic processing in the former and more ‘integrative’ processing in the latter. In short, I suggest that, unlike Spanish and Italian, English agreement is more susceptible to semantic interference, and that that manifests itself both at the level of grammatical sanction and at certain levels of processing at least. I take up the first of these two issues first.

[Corbett \(2006:2\)](#) mentions instances like (51) and (52) below, which he takes from [Morgan \(1984:235\)](#):

- (51) More than one person has failed this exam
- (52) Fewer than two people have failed this exam

to make the point that agreement need not be conceptually based in English. While there are more examples like these, there are also many instances showing that agreement can be conceptually based in English (see Dowty and Jacobson, 1988 for a stern defence of semantic agreement). Normally, both grammatical form and conceptual structure go hand in hand, which makes the individual contribution of semantics opaque. When mismatches arise, however, we have already seen that the final form is more often than not dictated by semantics. In section 2.1 we showed that that is indeed the case for metonymic extensions, ‘committee’ nouns, the *a bunch of* construction, measure terms, fused relative constructions, coordinated nouns, and even predicative agreement across the copula *be*. In fact, as is often pointed out, even ordinary coordination as in *John and his sister have gone out* involves a *resolution rule* resolved semantically, as both the NP *John* and the NP *his sister* are grammatically singular and yet agree with a plural verb (although this is not a peculiarity of English, but, at least, of all Indo-European languages). This has led researchers like Pollard and Sag (1988), Dowty and Jacobson (1988), Reid (1991), or Barlow (1999) to claim that agreement is essentially a semantic process (see also Vigliocco and Hartsuiker, 2002). However, there is also ample evidence that agreement is ruled by syntactic constraints in English. Distance between controller and target is one such powerful constraint. As we have seen, Corbett ties the facts concerning distance to the notion of *domains* of agreement and to his *Agreement Hierarchy*. The idea is that semantic penetration in or interference with agreement forms is more likely the more to the right one moves in the Agreement Hierarchy.

A good testing ground for the penetration vs autonomy of agreement processes is S–V agreement, as this lies somewhere in the middle in terms of distance between controllers and targets and, as already observed here, it is a core process in the building of core clausal structure. One might say that S–V agreement is ‘core agreement’. Berg (1998) compared the penetrability of formal agreement operations in a number of English and German constructions in S–V number agreement. *A priori*, the fact that both languages are members of the “West Germanic branch” might predispose one to expect no great differences in agreement patterns. In fact, what Berg found was a picture of almost perfect opposites. The German author compared the following eight categories:

(53) **Number conflicts in English and German (Berg, 1998)**

1. e.g. *the couple* (formally singular, but semantically plural).
2. e.g. *pancakes* (formally plural, but semantically singular).
3. ART (+ADJ) + N (sg) + *of* (+ADJ) + N (pl): e.g. *a gang of thugs*.
4. e.g. *both he and she* (plural not overtly expressed, must be inferred).
5. e.g. *more than one X* and *fewer than two X* (comparatives).
6. e.g. *many an X*.
7. e.g. *the cause (of the accident) + copula + (to be) + bad brakes*.
8. e.g. *it + (to be) + the politicians who ...; what interests me + (to be) + N (pl)*.

Berg designed a sentence-completion task in which subjects were asked to complete sentences containing the eight categories above. The verb slots were left blank, so the subjects had to fill them in. Forty-two short episodes were used, and 46 German speakers and 57 American English speakers participated. Some cases deserve special mention.¹⁰ The category with the greatest number of constructions, category 3, for instance (with constructions like *the majority of PL*, *a number of PL*, *a gang of PL*, *a series of PL*, *a bunch of PL*, *a set of PL*, *a heap of PL*, *a host of PL*, *a certain percentage of PL*) is particularly interesting. All of the items in that class favour syntactic (singular) agreement in German and almost all of them prefer semantic (plural) agreement in English.¹¹ In this language there are also individual differences due to the salience of the singular meaning associated with the first noun in some specific cases. This is conspicuous in the case of *a bunch of flowers*, for instance. Overall, however, the English data show either direct semanticity or clear trends towards semanticity. By contrast, German behaves syntactically almost without exception. Of special notice is the behaviour of categories 7 and 8, involving predicative constructions, clefts and pseudo-clefts.

¹⁰ There were no conspicuous differences in the responses of German and American subjects in other categories, like categories 4 and 5. There were differences in category 1, with the morphologically singular nouns with plural meaning triggering singular agreement in German and plural agreement in English. The results of category 2 were mixed.

¹¹ Referring to instances like *these sort of things*, Corbett (2006:65) notes that “(t)hese examples are the extreme instance of the semantic head of the noun phrase taking over the syntactic function of head too”.

These continue to separate German from English speakers, but now this happens in the very opposite direction: now the Germans use more plural verbs than the Americans. Category 7 is the clearest in this respect. Agreement choices appear to be influenced by the post-verbal segment in German in a very pronounced way, but not so in English. How can this be? Berg suggests an elaborate explanation, in the following steps: first, the gradient nature of the data suggests that the two forces impacting the agreement stage overlap and compete, therefore producing probabilistic results (“This claim is at odds with all proposals that treat agreement as either a syntactic or a semantic phenomenon”,¹² p. 59). Second, in a production task such as the one Berg used, encoding starts at the message level and, at a second stage, continues when the message is given linguistic form. In principle, the semantic force prevails at the message level, and the syntactic force prevails at the linguistic level. However, “a given force is not totally confined to its ‘home level’. It may spill over to a neighbouring level”. This is because the message representation must be mapped into the linguistic representation and cannot possibly ‘die’ so soon after the mapping is effected. This means that, although semantics will be weak outside of its home domain, it may still exercise some influence in that ‘foreign’ environment. That is, if there is any leakage from semantics to syntax (semantic override), it makes sense to expect it to occur in production, not in comprehension, as the level of representation that is active first in production is conceptual structure (see also Bock’s *marking and morphing* model of the distinct phases of production in agreement processes; Bock et al., 2004; Eberhard et al., 2005). How does this explain the cross-linguistic differences? Here is the explanation:

Given the competitive nature of the system, one major factor determining the strength of the semantic force at the syntactic stage is the strength of the syntactic force in its home domain. If the syntactic influence is relatively weak, it may be overridden by the semantic influence. (...) this appears to be a fair description of the English system. What, then, weakens the syntactic factor in this language? Agreement presupposes a morphology to express the relationship between target and controller. As is well known, the inflectional morphology of English is highly impoverished (...). Thus there is hardly any opportunity for (syntactically based) agreement processes to operate in English. It is a general cognitive principle that **frequency** impacts upon the strength of a phenomenon. The less frequently it occurs, the weaker it is. As a consequence, the limited opportunity that the language provides for expressing syntactically based agreement relationships involves a weakening of the syntactic force. (...) The situation is quite different in German. (p. 60, emphasis added).

So the poverty of English morphology has two interesting consequences: (1) it makes formal linkage more susceptible to semantic interference; and (2) it freezes an SVO order that guarantees that the who-did-what-to-whom message-level mapping to phrasal packages is not compromised despite the loss of morphological marks. As is well known, English turned to SVO the moment it lost its formerly rich morphology (Hawkins, 2004:161 and 166–167 and references therein). One important consequence of that is “**left orientation**”: a sentence arises the moment a leftward NP merges with a VP to its immediate right. In practice, then, as Berg points out, since the subject tends to precede the verb, when it comes to agreement, English speakers rarely take into consideration information that shows up after the verb. This is why in sentences like *the cause of the accident was bad brakes*, English goes syntactic now: that is, it goes to the left. Conversely, German speakers are used to a richer menu of word orders, since after all whatever coding power they lose by renouncing fixed SVO they more than make up for by using their rich morphology. Berg suggests that the conflict between syntax and semantics to determine the form of agreement choices can be predicted in English, in the following way: in cases of mismatch between the preverbal and the post-verbal materials, when the conflict arises within the preverbal domain, the semantic principle is preponderant. This is because English syntactic rails direct the language-user’s attention only to the preverbal NP as such, not to its internal structure (p. 63).

If Berg is right, two consequences should follow:

1. First, if semantics and syntax are likely to interfere with each other depending on their respective moments of ‘ignition’ and on their respective strengths in their home domains, one should observe the very opposite pattern of interference too: in comprehension, as opposed to production, syntactic penetration of form into meaning should occur in languages, like German, Spanish or Italian, with a rich inflectional system, that is, a system well-oiled by frequent and even redundant use of morphological cues.
2. Second, one should observe the same pattern of syntacticity displayed in German in a language like Spanish, as this is also characterised by rich morphology and non-strict word order.

¹² See Bock et al. (2001) for a similar view, reached from the laboratory. Corbett (2006) makes the same point.

We can confirm the first prediction here now. We have seen that Italian epicenes like *la vittima* or *il personaggio* enter agreement relationships by preferentially binding (or ‘bonding with’; cf. Garrod and Terras, 2000) pronouns of the same morphological gender, irrespective of strong contextual manipulations pulling co-indexing in a semantic or a pragmatic direction (Cacciari et al., 1997, submitted for publication). It is noteworthy that this happens at the rightward edge of the AH, that is, with pronouns *and* long distances. The behaviour of bi-gender nouns like *erede* in Italian is particularly relevant: since these nouns are not morphologically marked for gender, they are sensitive to semantic manipulations and thus effectively bond with pronouns that match contextual biases. Interestingly, Italian (or Spanish) bi-gender nouns are to all effects and purposes like ordinary English nouns. We have also surveyed the processing of Spanish gender and number violations when we examined the ERP study of Barber and Carreiras (2005). In all cases (article + noun, adjective + noun, and these pairs embedded in sentence contexts), the authors reported ‘fast’ responses to the violations: an N400 for adjective + noun pairs (recall that, without determiners, these are not NPs), and especially an N400 and LAN for article + noun pairs, and a LAN + P600 when any of the two pairs occurred sentence internally. Manuel Carreiras informs me that fresh ERP experiments in his lab have also obtained LAN + P600 effects in violations of Spanish pronoun-antecedent ties, thus reinforcing the Italian results obtained with a moving-window technique. All this occurred irrespective of the feature tested. In all cases morphological marks were processed as automatic reflexes, suggesting that in Romance languages morphosyntactic cues are privileged cues for phrase construction in the way suggested by Hawkins (1994, 2004). In this connection, I would like to report some incidental evidence of such automatization of agreement processes in Spanish. In an eye-tracking study of Spanish perception verbs currently under way, we aimed to investigate how the missing subjects of infinitives are recovered on-line.¹³ We used structures like (54) below¹⁴:

(54)

A. *Ana_i fue la única de la empresa que consiguió ser nominada para los premios.*⁺

Ana was the only.one-F of the firm who managed to.be nominated-F for the awards
 ‘Ana was the only one in the firm who managed to be nominated for the awards’ [repeated in all conditions]

+Se lo oímos+ \emptyset_{ij} decir+ al mediodía+ a su director_i+ absolutamente+ **pletórico**_i+ de contento.⁺

Se-ASPECT it heard-1PL PRO say-INF at noon to her manager-M absolutely full-M of happy-M
 ‘We heard her manager (masc) say so at lunch time, full (masc) of joy’.

B. *[Ana_i fue la única de la empresa que consiguió ser nominada para los premios.*⁺

Se lo oímos+ \emptyset_{ij} decir+ al mediodía+ a su director_i+ absolutamente+ **pletórica**_j+ de contenta.⁺

+Se-ASPECT it heard-1PL PRO say-INF at noon to her manager-M absolutely full-F of happy-F
 ‘We heard her say so to her manager (mac), full (fem) of joy’.

C. *[Ana_i fue la única de la empresa que consiguió ser nominada para los premios.*⁺

+Hizo + \emptyset_i regresar+ al mediodía+ a su director_i+ absolutamente+ **pletórico**_i+ de contento.⁺

Made-PAST.3S PRO_j return-INF at noon to her manager-M absolutely full-M of happy-M
 ‘(She) made her manager (masc) come back, full (masc) of joy’.

D. *[Ana_i fue la única de la empresa que consiguió ser nominada para los premios.*⁺

+Se lo pro quiso+ \emptyset_j decir+ al mediodía+ a su director_i+ absolutamente+ **pletórica**_i+ de contento.⁺

+Se-ASPECT it wanted-3SG PRO_i say-INF at noon to her manager-M absolutely full-F of happy-F
 ‘(She) wanted to tell her manager (masc), full (fem) of joy’.

¹³ The study in question is part of a research project under way. Apart from the author, Manuel Carreiras and Enrique Meseguer work in it.

¹⁴ Whether perception predicates are best seen as containing something similar to PRO subjects is a moot point. Possibly incorrectly, I simply use PRO to mean ‘infinitival subject gap’ in the materials.

A perception verb like ‘oir’ (*hear*) creates an ambiguous structure when it takes infinitival complementation. Thus, a sentence like *se lo oí decir a su padre* (*I heard her father say so*) is ambiguous in that, even though it would normally be understood to mean that *a su padre* says something, one can grammatically construe *a su padre* as the indirect object of *decir* and *decir* as having an external controller (*Ana* in 54B above: *Ana* is the one saying something to ‘her father’). So we created a condition A which is supposed to reflect the normal bias, a condition B with the opposite bias, where we expected to have large effects at *pletórica* (as the adjective confirms that *a mi padre* is not the agent of *decir* now, contrary to the standard interpretation), and two control (unambiguous) conditions: condition C has a post-verbal PP *a su director* with an agent function relative to the previous infinitive (like A); and condition D has a post-verbal PP which acts as a standard IO of the previous infinitive (like B). In order to prove that the parser experiences a complete breakdown at *pletórica* in condition B, we needed to see no comparable breakdown in the same segment in condition D. As expected, disruption at *pletórica* was obvious in condition B. However, we did find longer reading times at *pletórica* in condition D as well. Notice that that condition is not ambiguous at all. The only explanation is that the juxtaposition of the gender-marked noun *director* (masc) and the gender-marked adjective *pletórica* (fem) creates momentary interference. This is clearly a case of *shallow processing*, but it shows that in Spanish phrasal packaging is built on gender cues that are automatic and use recency. It was obvious from the eye-tracking data that the rapid reaction at *pletórica* in condition D was motivated by the parser not having consulted any relevant meaning representation at that point yet. The automaticity of the initial *bonding* stage proceeded on formal rails. In a Bergian vein, I am inclined to maintain that form spills over into the message level, this time not so much because the message level is weak, but because the form level, in a language like Spanish that may ordinarily use some 10 morphological cues per second, is extraordinarily strong. So even though meaning could have easily salvaged the misanalysis, the autonomy of form clues prevailed.

3.2. English attrition compared with Spanish redundancy

The second prediction that follows from the domain penetrability theory opened up by Berg’s study is that one should observe the same pattern of syntacticity displayed in German in a language like Spanish, the reason being that this is also characterised by a rich morphology and non-strict word order. ‘Syntacticity’ is used here to refer to the strength of morphosyntactic cues, a strength that causes such cues not to be ignored when encoding the message. With a view to gathering some evidence of that, I conducted a little, informal completion test in English and Spanish along the same lines as that of Berg (1998).

The materials used in the study were not identical to Berg’s but essentially very similar. They included the following six kinds of structures, which I illustrate with English materials:

(55) Construction types used in the completion study

1. collectives: *committee, enemy, crew.*
2. *task of the UN’s intervention forces, decision of the board.*
3. *a series of PLURAL, a small number of PL, a number of PL, a bunch of flowers, a heap of PL.*
4. *both he and she, either he or she.*
5. *The cause of the accident BE defective brakes, a case in point BE the revolts in Paris.*
6. *All I could see BE the tree tops in the distance, what worries me BE the reports in the press, what I don’t get BE the shouts of approval.*

The total number of structures tested was 17. In order to divert the readers’ attention, the structures were mixed and embedded in a total of 18 different paragraphs which contained a total of 66 gaps. This created a 17/66 target-to-filler ratio. Subjects who took part in the task were simply required to read “fast but understanding the sentences”. To further distract their attention, they were told that the purpose of the task was to study the effect of context on word selection. These instructions aimed at obtaining a careful pattern of reading for content and at avoiding shallow

processing. The emphasis on speed, however, was meant to prevent the very opposite of shallow processing: over-conscious reflection. There were 26 English speakers and 34 Spanish speakers, with ages ranging from 18 to 63. Results are as shown below:

	English				Spanish			
	<i>N</i>		<i>%</i>		<i>N</i>		<i>%</i>	
	sing	pl	sing	pl	sing	pl	sing	pl
Category 1								
a. Committee	10	11	38.4	42.3	32	2	94.1	5.8
b. Enemy	19	6	73	23	34	0	100	0
c. Crew	11	5	42.3	19.2	34	0	100	0
Category 2								
a. Task of the UN's forces	26	0	100	0	32	1	94.1	3.8
b. Decision of the board	11	0	42.3	0	31	0	91.1	0
Category 3								
a. A series of PLURAL	4	17	5.3	65.3	11	23	32.3	67.6
b. A small number of PL	2	22	7.6	84.6	22	12	64.7	35.2
c. A number of PL	1	22	3.8	84.6	16	18	47	52.9
d. A bunch of flowers	19	6	73	23	31	1	91.1	3.8
e. A heap of PL	5	11	19.2	42.3	7	25	20.5	73.5
Category 4								
a. Both he and she	1	15	3.8	57.6	0	32	0	94.1
b. Either he or she	23	2	88.4	7.6	16	17	47	50
Category 6								
a. Cause of accident BE brakes	23	2	88.4	7.6	8	25	23.5	73.5
b. A case in point BE the revolts	19	2	73	7.6	1	33	3.8	97
Category 7								
a. All I could see BE tree tops	3	22	11.5	84.6	2	32	5.8	94.1
b. What worries me BE reports	2	22	7.6	84.6	2	30	5.8	88.2
c. What I don't get BE shouts	6	20	23	77	3	29	8.8	85.2

Percentages are given only for correct responses. Overall, the data replicate the findings in Berg's study rather closely. Obvious signs of semantics interfering with the morphosyntax of the noun phrase in English can be seen in categories 3b (*a small number of* + PL), and 3c (*a number of* + PL). Obvious signs of orientation to the left can be seen in categories 6a (*the cause of the accident BE brakes*), and 6b (*a case in point BE the revolts*). Additionally, categories 1a, 1b, and 1c also confirm the signs of semanticity for the lexical level (*committee, enemy, crew*): these nouns are robust singulars in Spanish but somewhat ambivalent in English. Although the coordinated structures present shocking results (*either . . . or* has a clear singular bias in English, with Spanish showing no preference at all), overall, it does seem to be the case that semantic overrides affect the internal composition of English NPs much more than is the case for Spanish. And it also seems likely that the stark contrast in the '*the cause of the accident BE PLURAL*' type of structures reflects the preference to establish agreement by reference to an SVO template in English.¹⁵ The only

¹⁵ The study also replicates the variability of Berg's instances of category 3. As the German author suggests, this is probably due to the different lexicalizations that all these forms may have. Such lexicalizations are very likely to reflect habitual *construals* or *profilings* of the same substance of thought. For instance, the *a bunch of flowers* construction is very often conceptualised as a unit (a gift), rather than as a collection of items. In that respect, replacing *flowers* by *kids*, for instance, is very likely to affect agreement choices as it now appears intuitively evident that it is the multiplicity of tokens that receives salience. Such *construals* reflect semantic interference, but one that is resolved at an individual level and is likely to affect all languages. In the same way, the lower-than-expected scores for singular agreement in categories 3a and 3e in Spanish (the *a series of* construction and the *a heap of* construction) may reflect the constructionally driven bleaching of these two highly recurrent expressions in Spanish.

notable exception to the consistency of the opposition between English and German on the one hand and English and Spanish on the other is that both comparisons fail to reveal the expected contrast between left orientation in English and right orientation in Spanish and German in category 7, that of pseudo-clefts: the three languages behave similarly. Berg points out that the different behaviour of English speakers in clefts (with solid singular agreement) and pseudo-clefts (with solid plural agreement) is due to the fact that in cleft constructions the dummy form *it* is unambiguously singular. However, even though that might work in unison with left orientation to establish a rigid singular preference in clefts, it cannot explain why left orientation alone does not cause a singular preference, or at least more singular preferences, in the pseudo-clefts. The only explanation that seems likely is that English speakers have *listed* pseudo-clefts as structures without left orientation, in a *Construction Grammar* sense of the term ‘listing’ (Goldberg, 1995, 2006; Kay and Fillmore, 1999; Culicover, 1999; Croft, 2001). In the same vein, the presence in pseudo-clefts of *what* and the habitual use of both singular and plural agreement with this *wh*-word in questions (e.g. *what is the consequence of this? what are the consequences of this?*) might be seen as either an *inheritance link* or a *taxonomic link* (Goldberg, 1995) that reinforces the distinctiveness of agreement ties in this particular *construction*. Thus, *listing* and *linking* might probably explain the idiosyncratic pseudo-clefts data. But the important thing is that, excluding pseudo-clefts, English and Spanish behaved as predicted in most other categories.

3.3. The processing of English pronouns

In suggesting that there may be cross-linguistic differences in the processing of agreement, I made the point that English is well-known for semantic overrides at the grammatical level, which means that semantics might exert a larger influence in general (i.e. not just in grammar) in English than in other languages. Corbett’s (1979) *Agreement Hierarchy* implies that, at the grammatical level too, overrides are the more likely the more structural distance intervenes between controller and target. This is the reason why one cannot say **these committee* (NP domain), but one can say *this committee have made a proposal* (clause domain), and *my nephew works for a firm called Sonipac. They pay him quite a lot* (discourse domain). Additionally, Berg’s completion test showed that semantic interference manifests itself at the grammatical level too at the moment of encoding the message. And it also showed that German generally manages to keep semantics confined to its home domain by using its strong morphological reflexes. My own little completion study reinforced Berg’s views: Spanish behaves like German due to its strong morphosyntax. Intuition also suggests that. However, when it comes to actual on-line processing, the situation may be an entirely different one. On-line linguistic computations involve processes which occur in very narrow time windows. We have already seen that Osterhout and Mobley (1995) found a LAN-type effect for S–V violations in English in comprehension.¹⁶ This suggests that in comprehension (where form – auditory or visual – is the initial stimulus) at least S–V agreement runs on formal rails even in English. That is, despite frequent grammatical semantic overrides, when syntactic agreement in number occurs in English, this appears to be processed just as fast as in the morphological languages. One should remember that S–V agreement lies at the heart of clause construction, so the principle that practice makes perfect should not *a priori* distinguish English from other languages in excess.¹⁷ It also suggests the *phase impenetrability of a core clause construction process*.

A different issue is whether the same formal bias can be observed in ‘non-core’ agreement processes established over longer distances. The fact that Osterhout and Mobley obtained effects in the P600 time window, but not in the LAN time window, for violations of agreement involving reflexive pronouns and personal pronouns suggests that at the rightward edge of the AH, even in comprehension, automatic formal cues are not quick enough to run the parse in the early stages in English. Recall that the Italian and the Spanish data reviewed in the preceding sections showed that in these two languages not even distance could deter automatic formal *bonding* from occurring at the earliest stages of

¹⁶ See Bock et al. (2004) for differences and similarities between verbs and pronouns in production.

¹⁷ As Eberhard et al. (2005:532–533) note, among verbs, *be* and *have* (with a consistent sg/pl distinction) occur more often than any other predicate in English. Keeping just these two verbs in consideration, English speakers implement S–V agreement at least once every 5 s in running speech. In the *marking and morphing* framework of agreement production that Bock and collaborators have refined over the past few years, verbs get their number in the syntax via a *control* relationship from the subject noun phrase (analogous to a copying rule), while pronouns get their number directly from the semantics and enter *concord* with their antecedents (at least in English, where the poor verbal morphology is not likely – they argue – to code number directly). In principle, this seems reasonable, but it also seems to have the odd result of forcing one to predict that, since the starting point for both verbs and pronouns is the same (meaning), and since pronouns can dispense with the final structural integration process, agreement operations should be faster for pronoun-antecedent ties than for S–V ties.

computation. In the face of strong formal regulation, semantics was made to wait. I would like to show now that the existing evidence on the processing of English pronouns does point to form acting late in this language: that is, it points to the absence of the *bonding* phase that characterises the processing of pronouns in the Romance languages.

In section 2.4.2 we have already made reference to research by Nicol (1988) using the cross-modal priming technique. Nicol investigated the processing of pronouns disambiguated by gender and by number, and found a significant priming effect for number and for masculine gender. We have already argued that those results could not be taken as evidence of the initial phases of processing. The same can be said of the *probe recognition* methodology with which most studies on pronoun resolution have been done. In a typical probe recognition study, participants read passages which are interrupted at various points by the presentation of a written probe word. Participants then have to quickly tell whether they have seen the probe word earlier in the sentence. The rationale is that gender-congruent probes should be recognised faster than gender-incongruent ones. This kind of research has offered conflicting results. For instance, Gernsbacher (1989), Greene et al. (1992), MacDonald and MacWhinney (1990) and McDonald and MacWhinney (1995) showed that immediately after hearing a pronoun the gender-congruent antecedent was no more activated than a preceding gender-incongruent name, even though activation was faster for the congruency condition at a later point. However, other studies have managed to show very early effects (Garnham et al., 1996; Rigalleau and Caplan, 2000). Gordon et al. (2000) have recently surveyed this kind of research and concluded that their results are strongly affected by strategic processes related to the task (see also Arnold et al., 2000; Van Gompel and Liversedge, 2003).

Eye-tracking holds better prospects of providing a finer grain access to pronoun resolution. Ehrlich and Rayner (1983) and Garrod et al. (1994) showed that the earliest effects they could obtain took place at the words following the pronouns. With effects on the pronouns themselves, Vonk (1984) is often cited as evidence of early access to morphological information, but a careful examination of her methodology indicates that that conclusion is in need of revision: her participants had to name the correct reference for the pronoun, which means that strategies not normally used in plain reading may have been used. To date, the best evidence pointing to a rapid use of morphological information in pronoun resolution in English is provided by a recent study by Arnold et al. (2000) in American English. In a series of two experiments, participants viewed a picture with two familiar cartoon characters of either the same or different genders. While they listened to a text describing this picture in which a pronoun referred to either the first or the second character, Arnold et al. recorded the movements of their eyes. A typical example is (56):

- (56) Donald is bringing some mail to [Mickey/Minnie] while a violent storm is beginning. He's carrying an umbrella, and it looks like they are both going to need it.

Arnold et al. created 4 conditions crossing first or second mention of antecedent NP (say *Donald* as subject of the sentence, occurring first, and *Mickey/Minnie* as object, occurring later) and ambiguous pronoun (same gender: *Donald ... Mickey ... he*) vs unambiguous pronoun (different gender: *Donald ... Minnie ... he*). They claim to have found evidence that both gender and *accessibility* have rapid effects on pronoun resolution. It seems that they use the term 'accessibility' to mean 'topical noun phrase'. They observed that within approximately 200 ms after the offset of the pronoun, subjects looked more often at the picture that had the same gender as the pronoun. They concluded that gender information is processed automatically, not strategically.

However, it is not clear that such a conclusion is a solid one. Arnold et al. ran two experiments. In the first one, they used structures like (56) above. When they presented the data of the eye-tracking analysis, they divided the region of the post-pronoun segment into four sub-regions of 200 ms each, as shown in (57):

- (57) Segment 1: pronoun onset to verb onset (about 200 ms).
 Segment 2: next 200 ms.
 Segment 3: next 200 ms.
 Segment 4: next 200 ms.

This is the pattern of referential looks that they found:

- (58) Segment 1: **no effect**.
 Segment 2: **no effect**.
 Segment 3: interaction.
 Segment 4: interaction.

As can be seen, the first two segments yielded no really fast responses. In the second experiment, they manipulated ‘accessibility’ by strengthening the topicality of the first NP, as in (59):

- (59) Donald is bringing some mail to [Mickey/Minnie]. He’s sauntering down the hill, while a violent storm is beginning. He’s carrying an umbrella, and it looks like they are both going to need it.

This time they obtained effects already at the second segment, the one immediately following the pronoun. It seems reasonable to conclude that it is rather the topicality treatment of the experiment, associated with a specific syntactic position (subject), that is running the rapid co-indexation of antecedent and pronoun (see Kazanina et al., 2007). So despite their claims to the contrary, it is not clear that fixations on the target occur as quickly as they claim they do and because of the reasons they point out. Over and above all that, there is another, alternative explanation to their findings: in the visual paradigm they used, participants may have looked at the picture with the same gender as the pronoun without really having established a coreferential relationship between the pronoun and its linguistic antecedent. Overall, it seems premature to conclude that this study conclusively proves that gender is used automatically in antecedent-pronoun coindexation in English.

A more recent eye-tracking study by Van Gompel and Liversedge (2003) does show clear evidence in favour of the opposite conclusion. In this study, both the gender and the number of antecedent noun phrases was manipulated in the establishment of backward anaphora, or cataphora. A sample of the materials used in the gender manipulation is provided in (60):

(60)

a. *gender match*

When **he** was at the party, **the boy** cruelly teased the girl during the party games.

b. *gender mismatch*

When **he** was at the party, **the girl** cruelly teased the boy during the party games.

c. *control*

When **I** was at the party, **the boy** cruelly teased the girl during the party games.

Van Gompel and Liversedge obtained mismatch effects in the early measures of the eye-tracking record in the region immediately following the main clause subject noun phrase. For instance, first-pass times were higher (slower) at the word *cruelly* in (60b), indicating that sentences were easier to read when the cataphoric pronoun and the NP1 were gender-matched than when they were not. There was a firm tendency to assign the cataphoric pronoun to NP1 rather than to NP2. The inclusion of the control condition ruled out the possibility that the ease of the gender-match condition relative to the gender-mismatch condition was due to the fact that the latter contains one extra discourse referent (*he* and *the girl*, instead of *he* and *the boy*, which can co-refer). The feature mismatch effect indicates that the processor has no time to check the gender of NP1 before establishing a coreferential relationship between that NP and the pronoun (contra Cowart and Cairns, 1987). Rather, the mismatch effect reflects the creation of a referential dependency between the pronoun and NP1 *before* relevant bottom-up information about the antecedent becomes available. This results in a processing penalty when the subject NP is not semantically and/or morphosyntactically congruent with the pronoun. For now, I will just use the feature mismatch effect uncovered by Van Gompel and Liversedge (2003) to buttress the view defended here that morphology does *not* rule the initial stages of the processing of pronoun-antecedent ties in comprehension in English.¹⁸ It must be noted that Van Gompel and Liversedge’s experiment 3 also used a *number* condition which cast the same pattern of results as that found for the gender manipulation: number information is also delayed until after the computation of coreferential ties established on independent grounds. Additionally, disruption due to gender and number mismatch arose at exactly the same place in

¹⁸ The fact that pronouns elicit the same grammatically motivated *attraction* errors as verbs (i.e. an increased tendency to be plural after local plural nouns in subject NPs) might suggest that in production, however, some formal routinization is still discernible even in English. This is shocking, in principle, given that one would rather expect more of that routinization in comprehension (where form comes first) and more semantic interference in production (because meaning comes first). There is in fact *also* semantic penetration in *attraction* at the same time, since distributive readings of complex NPs cause more erroneous plural agreement. On balance, it is hard to make sense of the attraction data, not least because as evidence from other languages accumulates, its mechanics seems to hold essentially for all of them, irrespective of their differences. Perhaps there is something echoic in attraction effects that makes it less an architectural hint and more a mistake of the phonological encoding phase. See Eberhard et al. (2005), and references therein, for an authoritative view.

the two experiments, suggesting that, irrespective of **storing** differences between the two features, the two are **computed** in the same way. This further undermines the psychological validity of the Feature Hierarchy Hypothesis (contra Nicol, 1988; De Vincenzi, 1999; Di Domenico and De Vincenzi, 1995; Carminati, 2005; see section 2.4.2).

3.4. Interim summary and a few ramifications

In sum, a review of the event-related brain-potential studies shows that S–V agreement in English appears to be done in a way similar to all other purely syntactic operations. This, at least apparent, encapsulation was not so apparent for the other kinds of agreement, which peaked electrophysiologically in different, and later, brain waves. At the same time, we saw that in Spanish *all* agreement operations exhibited the typical combination of LAN and P600 effects usually associated with the automatization of syntactic processes. I suggest that the difference between English and Spanish is due to the fact that in non S–V types of agreement in English (especially in pronoun-antecedent ties), there is no *bonding* phase (Garrod and Terras, 2000) built on purely automatic reflexes, and that this is due to English being a morphologically impoverished language whose morphosyntax is ‘out of shape’: that is, not well trained through frequent practice.

Corbett’s Agreement Hierarchy predicts that formal agreement processes may be penetrated by semantic constraints the more to the right we move in the hierarchy. That explains why at the grammatical level even the S–V domain (not as far to the right as pronouns) may exhibit semantic interference. However, at the processing level, S–V ties established either on formal co-occurrence or on semantic interference, are processed automatically: that is, like any other syntactic process. This is suggestive of the phase impenetrability of a core clause construction process. It also indicates that the grammar and the processing of agreement are not in perfect agreement with each other. At the processing level, we need to go to the most rightward edge of the *Agreement Hierarchy* to see pronouns not running on formal rails in English. So evidence of semanticity in English varies not just with domains but also with the level of enquiry: grammar or processing. In fact, it varies even more because, as Berg has suggested, domain penetrability may be sensitive to the direction of encoding. This means that semantic interference is more to be expected in production (where meaning comes first) than in comprehension (where form comes first). Berg has demonstrated that German, but not English, generally manages to steer clear of semantic interference due to the strength of its morphosyntactic component. In my own completion study, it was evident that Spanish patterned with German, not with English, on predictable grounds. I made the point that if domain penetrability is sensitive to the force of each home domain (either morphosyntax or semantics), then the opposite pattern of interference should be expected in comprehension, where form happens first. And indeed, the processing of *epicenes* in Romance shows that even in the presence of highly constraining contextual and semantic clues, pronoun-*epicene* ties proceed first on formal rails. The same can be said of recent ERP data on Spanish pronoun-antecedent violations.

In sum, one ramification of the domain penetrability theory is that domain impenetrability will be the more robust the stronger the morphosyntax of the language in question. Another is that morphosyntactic strength is a result of entrenchment: the more often formal cues co-occur, the more entrenched they will be. Still another ramification is the well-known fact that absence of morphological cues results in word order rigidity as a means of phrase and clause construction. In its turn, this means that, when building the clause, English language users show orientation to the left. Such orientation explains why semanticity does not run agreement patterns in clefts: even if something akin to resolution is needed (because the pre-verbal and post-verbal constituents are not feature-matched), the attention of English speakers does not extend to constituents beyond the verb. As an SVO language, English constructs the subject on the left, so it is the left branch of the tree that determines agreement by default.

Finally, another consequence of the domain penetrability theory is that if morphological attrition results in word order fixity and this in left orientation, one should be able to see more signs of such fixity and orientation in other areas of processing. Likewise, if morphologically impoverished languages are more susceptible to semantic interference in agreement, then one should be able to see other effects of interference outside the domain of agreement. Indeed, given that agreement is at the core of sentence structure, it makes sense to expect different rippling patterns spreading from different grammar cores, such as those of English and Spanish, for instance, that is, large repercussions in the rest of the two grammars. Repercussions here relate to the notion of domain penetrability. It seems to me that pursuing this topic is a fascinating enterprise, which cannot of course be paid due justice here, so only a few comments will be made now. In his closing remarks, Berg (1998) himself mentions in passing a number of fairly specific ramifications of the different behaviour of agreement in English and German. The first point Berg mentions concerns the Saxon genitive

(p. 64). Unlike German, English genitive phrases are sensitive to (penetrated by) semantic variables like animacy in the sense that the first noun in the NP must exceed or at least be equal to the second noun in terms of animacy (*Mary's car* vs *the car of Mary*; cf. Hawkins, 1981). The second construction mentioned by Berg is the presentative construction with *there*. In English, this construction can only be instantiated by a small number of predicates, all of which resemble *be* semantically. In German, conversely, all kinds of verbs are allowed: “The syntactic principle is unassailable” (p. 65). The third area is the grammar of NP determination (p. 65). In German, definiteness may be expressed on any noun. In English, however, article-taking nouns are sensitive to the semantic feature abstract vs concrete (*the dishwasher* vs *the *evil*)—yet another interference of meaning on form. Finally Berg mentions punctuation to make the point that this is more syntactically driven in English than it is in German.

I would like to close this section suggesting another ramification of the agreements facts as interpreted throughout the present account. Needless to say, the rigid SVO word order of English has long been recognised. The interesting connections between such rigidity and the coding of basic parameters of information structure are now well known. For instance, *marked focus* in English preserves SVO and thus marks new information in a way that differs from Romance, where *displacement* is usually recruited for the task (see Lambrecht, 1994). What I would like to suggest is that it makes sense to start analysing whether word order rigidity also interacts with parsing routines in different ways cross-linguistically, and the processing of backwards anaphora in English that we reviewed in the previous section is suggestive of that. Recall that in the face of structures like (61) below

(61) While **she** was making pancakes, **Cathy** taught Steve how to carve them into funny shapes.

Van Gompel and Liversedge's (2003) research showed that a dependency is established before all the bottom-up information has been parsed (the *Feature Mismatch Effect*). It is as if when the subordinate clause is being read, the presence of an un-anchored pronoun in it forces the parser to project an antecedent position in the subject slot of the main clause, even before the main clause has actually made its appearance. Thus, before bottom-up information makes its appearance on the processing scene, the parser uses *activeness* to anticipate dependencies, and these are geometrically defined, at least in English. Van Gompel and Liversedge suggest that structural information has “priority” because:

it is often the only type of information that is reliably derivable before bottom-up information about the noun phrase is encountered. If the parser encounters structural evidence for an upcoming clause, then it can immediately and reliably predict that the clause will have a subject noun phrase (. . .) In contrast, other information about the subject noun phrase such as its morphological features and its semantic match to a previously-encountered pronoun cannot be evaluated except via bottom-up information (. . .) Thus, under this view syntactic information plays a crucial role because it enables the parser to make predictions about upcoming material earlier than any other type of information, and hence there is no need to impose architectural constraints that force certain information types to have priority. *Note that this line of reasoning might apply differently in languages that display richer morphological agreement than English*, such that it may be possible to reliably predict morphological properties of an upcoming noun in advance of the noun itself. (p. 23, emphasis added)

Note, importantly, that the expectancy for an antecedent in all the experiments reported is an expectancy for a subject antecedent, that is, for a NP to the left of the verbal agreement target. Although one would need comparable data in a language like Spanish to see that such geometrical aids to activeness are less important, the fact that such geometrical processing has been reported for English makes very good sense given the theoretical principles sketched out here. Kazanina et al. (2007) liken this *activeness* of the processor to gap-filling routines in long-distance filler-gap dependencies (Stowe, 1986; Frazier and Flores D'Arcais, 1989) and, in a series of three off-line and three self-paced reading experiments also on backwards anaphora, offer the best demonstration to date of the suspension of an active search for an antecedent in the face of certain structural domains (like a Principle C violation; see Chomsky, 1981). This apparent geometrical bias of English also echoes the grammar of so-called *misrelated participles*, where an initial participle in a dependent adverbial clause recovers its subject automatically by borrowing it from the main clause subject—in English (GAP *walking home*, *Jane realised she had been fooled*; compare *GAP *walking home*, *the sky seemed ominous*). An interesting sign in the right direction is that misrelated participles do not seem to be a problem for speakers of Spanish.

4. Epilogue: the function of agreement features and the supremacy of gender

There are basically three ways in which phrase and clause construction can take place across languages. The first is via morphology, the second via a consistent word order, and the third is via the use of classifiers (which fall outside the scope of the present work; see Corbett, 1991:136 ff.) A consistent word order is common and useful in clause construction because language users always know where the *who* and the *did* and the *what* are going to be found in the left-to-right development of the message. In a system like this, agreement is less important, or even superfluous at times. English is something like such a system. Its notorious rigidity, the quintessential source of despair for many a second-language learner (my Spanish reflexes hesitate to asterisk **John said yesterday that*), serves the grand purpose of building the structure of the clause. Additionally, as Hawkins (2004:250) points out, “there is no disambiguation motive for rich morphosyntax on SVO languages, since the medial verb position distinguishes the two [main] arguments” of the verb. However, English, like all languages, is no more than the present state of a history of compromises, of innovations and retentions. As a present state, it shows traces of what it used to be: a rich inflectional language like modern Spanish. This means that agreement has not gone entirely by the board. So even today English has some Det–N agreement, some NP–VP agreement, and some antecedent-pronoun agreement. Since word order and agreement are functional means of clause building, it is often not impossible to see functionality driving most of the changes that have taken place during the progressive demise of its once rich morphological component. Take gender, for instance, which English mostly lacks. The standard view on gender is that it belongs with the lexical component. Number, by contrast, is supposed to be a syntactic, computational phenomenon (Ritter, 1993; Harley and Ritter, 2002; Anderson, 2004; Carminati, 2005). That might explain why reanalysis after a gender violation is costlier in Spanish (Barber and Carreiras, 2003, 2005). Thus, as we have seen, in a serial model of lexical access like that of Bradley and Foster (1987), for instance, the difference in access time after detecting either a gender or a number inconsistency would reflect the additional cost of reinitiating lexical access in the former, but not in the latter. However, it is important to understand that the reanalysis data reveal aspects of post-agreement operations. In other words, reanalysis may be costlier for gender anomalies because the processor chooses the wrong lexical feature in the first place, not because that feature, once chosen, is processed (i.e. integrated) differently. One of the most conspicuous conclusions we have reached here is that, when it comes to computing agreement, gender and number are processed in exactly the same way, contrary to several claims made in the last recent years, claims often based on the *Feature Hierarchy* (Carminati, 2005). This distinction between *storing* and *computing* has always been ignored.

The idea that gender is an eminently lexical affair while number is a syntactic affair is also wrong for theoretical reasons. If that were the case, then, given that English has progressively opted for a fixed word order that makes up for the loss of its clause-building morphology, one would have predicted the lexical (i.e. non-clause building) component to be unaffected by that loss and the syntactic component (number) to have gone away with the morphology. Obviously, what happened was the very opposite of that: gender has practically disappeared in English, while number has stayed strong. However, this is a road full of complex turns. Gender does remain precisely in the *domain* where a rigid word order cannot be counted on: the pronominal system. Since pronominal reference usually spans large distances, morphological cues are as necessary there as they have ever been, or as they still are in a richly-inflected language like Spanish. The equation leaps to the eye (at least as far as English and Spanish are concerned): *gender disappears where a fixed word order takes over, and it stays where there cannot be a fixed word order*. It seems, therefore, that the two phenomena may stand in something like a complementary distribution. Can we still maintain the old view that gender is a lexical thing? Maybe not.

Possibly the view that makes more sense is the one that sees gender as the morphosyntactic cue *par excellence*—gender classes arising historically out of capricious cultural categorization. In one culture there might be a specific gender class assigned to inanimates, and another one to animates. Or there may be one assigned to inanimates and women or to inanimates and men (less often). Or one reserved for barnyard animals, groceries and the sun. The situation may be an entirely different one in cultures settled just across the bay, the ocean, or the mountain range. Hence the strong association of gender and lexical idiosyncrasy. Different *chaining principles* (Lakoff, 1987:91–104) are often discernible in such complex systems. Corbett (1991) is a good place to see that every conceivable combination of gender classes is possible worldwide. But Corbett also insists (pp. 8–32) that all gender classes have a semantic basis historically: that is, a certain categorization logic. Since categorizations of an ever changing, large, complex world change in unison with what is (culturally) *profiled* about that world, it makes sense to expect gender to drift quasi capriciously, towards the seemingly arbitrary. Number, by contrast, always remains cognitively motivated,

as *numerosity* does not offer much room for internal construals. All this means that number may be accessed semantically, *directly* as it were, whereas gender can only be captured by *the co-variance of form*. Women can only be categorised alongside fire if both concepts bear the same linguistic tag: in and of themselves, the connection between the two notions is not of the same degree of motivation or *grounding* as that between singular *cat* and plural *cat-s*, for instance. In other words, if there is a morphological cue that is strongly associated with phrase- and clause-building processes, that is gender, not number. Corbett (1991:4) points out that “the determining criterion of gender is agreement”. In that sense, gender is a truly *contextual* (not *inherent*) feature, in Booij’s (1996) sense. In a sample of 50 languages, Bybee (1985) observed the inflectional features marked on verbs and found gender in only 16%, but number and person in 54% and 56%, respectively. This may deceive one into assuming that, since the three main features are presumably equally nominal in origin, gender is precisely the less contextually (i.e. syntactically) motivated. Nothing could be further from the truth. It is important to keep in mind that number and person will always be more prevalent because they are conceptually grounded (Greenberg’s (1963:95) universal 36: “[i]f a language has the category of gender, it always has the category of number”; see also Harley and Ritter, 2002:514 and ff.). The interesting thing is that, when recruited as a clause-building instrument, gender is essential in the languages where it is the main clause-building instrument precisely because it is not so strongly attached to meaning. Additionally, clause building goes well beyond S–V agreement anyway in that the usual redundancy that is typically associated with agreement tends to occur inside the complex structure of the noun phrase. And the interesting count is not how many languages code gender or number, but rather how many code gender if they are of the strict, fixed SVO kind and how many do so which have free word order.

In a way, therefore, the reason why gender is a better clause constructor than number is because number does not tell enough entities of the world apart. Additionally, number classes cannot adapt to the world, as they are too limited in nature – in range and dimension – to capture the huge variability of the whole world. When it comes to coding formal cues that allow one to distinguish the different thematic roles of, say, *sweater* and *blouse* in *the red sweater does not go well with the beige blouse*, for instance, the singularity of the two NPs does not really help much. There are too many singular things in the world or, alternatively, too many amorphously plural ones. Conversely, in a world of continuous variability, arbitrary gender classes allow us to co-classify pieces of the world *ad infinitum*. They provide *creativity*. If a certain language marks, say, ‘the property of being blue’ with, say, an asterisk and the nominal entity ‘supernova’ with an asterisk, it is easy to infer the message to be the synthesis ‘a blue supernova’, however counterintuitive the combination may turn out to be. I can use an asterisk to mark the nominal referent ‘ideas’ and another asterisk to mark the property ‘green’, thus yielding ‘green ideas’. One can conceivably use the same asterisk to code the verb ‘sleep’. The inevitable by-product of punctuated asterisk continuity is a message coded to mean ‘the green ideas sleep’. Thus, gender may be instantiated lexically, but for syntactic reasons. In a way, then, gender classes are a finite means of expressing infinite analyses. As long as they are formally manifested through co-variance, constituents 1 and 4 (or 1 and 13 for that matter) build the superordinate constituent [1–4] by virtue of the mere co-variance of formal cues, irrespective of what constituents 1 and 4 signify.¹⁹ *A priori*, it is no wonder that Taylor (2002; see introduction) expresses his utter bewilderment in the face of a formal machinery that, for him, merely reflects humans’ delight in playing with the form of language. More thoughtfully, Langacker (1991:307) views redundant agreement markings as predications in their own right. These are all “meaningful symbolic units”. However, when it comes to specifying the meaning of such symbols, all we are left with is “the traditional notion that the agreement serves the function of signalling grammatical relationships”.

Both views – the bewilderment and the limited, traditional explanation – have motivation, though. One can see that the same morphosyntactic cue on an adjective and a noun constructs NP in Spanish (Hawkins, 1994, 2004): that is, that agreement serves a grammatical function, as Langacker points out. However, one can also see that the same adjective and the same noun in English may not bear any such cues at all and yet a NP is constructed all the same. To the extent that the purpose of agreement is to liberate clause constituents of strict adjacency, one would expect agreement systems to flourish in languages with widespread discontinuity, like Latin or Warlpiri, and to be practically non-existent in languages which, like English, have frozen a rigid word order. The fact that actual languages are not quite so pristine has led people like Taylor to strike a skeptic note. It is true that the redundancy of Spanish *l-a-s mes-a-s blanc-*

¹⁹ An analogy due to Corbett (2006:19–20) might be useful at this point: “If houses 10 and 12 are both white because it has snowed in both, this is not canonical agreement. If Mrs White paints number 10 white and Mrs Green in number 12 paints her house white as well, then, that is, potentially, canonical agreement”.

a-s y *cort-a-s* (the *while* and *short tables*) cannot be entirely predicted given the previous comments, for after all discontinuity does not run rampant in Spanish. But it is also true that: (1) well-oiled clause-building systems which are functionally motivated (either agreement or word order) tend to be manifested across the board, not just in the specific realisations of the constructions where they are maximally motivated; (2) languages do not have either morphology or attrition like an on/off switch: rather, they have more or less of each, because of historical drift. In English, there are remnants of a once rich morphological system that are still highly productive today. In Spanish, there is more discontinuity than in English, and definitely less word order fixity and left orientation, but constituents *do* tend to be linearised adjacently. Also, its pro-drop tendency is modulated by certain interesting constraints (see Cameron, 1993; Heap, 2000 on Romance in general on this). Siewierska (1998) examined 171 languages and concluded that rigid word order does not allow one to predict the loss of agreement mathematically, but liberal word order does allow one to confidently predict that a language has agreement; Moravcsik (1995:471) posits a case-copying implicational universal to the effect that if agreement through case copying applies to NP constituents that are adjacent, it applies to those that are non-adjacent (see Hawkins, 2004:160); as is well known, in Latin, NP dependents could be separated from the head because of rich formal marking on all constituents, and in the early Germanic languages this was also possible (though not necessary), but in English this ceased to be a possibility precisely after English lost its rich morphology (Sonderegger, 1998); in Kannada, a Dravidian language, accusative-marked NPs can dispense with case marking when they occur adjacent to their transitive verbs but must mark case otherwise (Bhat, 1991:35); and (3) the internal logic of the attrition/redundancy dichotomy is often recognisable in grammatical systems in ways which may not be so clear as one would want them to be, but which are visible anyway. Hawkins (1994, 2004) mentions a number of them:

- a. Discontinuity itself is more a property of German or of Spanish than of English. This cannot be due to chance.
- b. Agreement can expand the construction potential of a node in parsing: thus, for instance, agreement features *-as* in the determiners and the adjectives of Spanish allow one to build NP even before the N makes its appearance on the processing scene:

Agreement may therefore serve to construct a mother or a grandmother node that is categorically distinct from the word that carries the agreement features . . . [T]he adjective (. . .) constructs NP (. . .) by projecting from the nominal agreement features to the noun agreed with, and by constructing the mother of this noun over the adjective. (Hawkins, 1994:368)

- c. Agreement makes it possible for the target to disappear without information loss. Determiners like *una*, *esta*, *cual*, etc. can always supplant the whole NP in Spanish. Their English analogues can only do so in much more restricted circumstances, or not at all. In Spanish the definite determiner merges with adjectives ordinarily to create NPs without nouns (*prefiero los blancos*, *me quedo con los grandes*). There is no special meaning attached to the omission of the noun. In English the combination is so restricted in use and meaning that it is best to see it as *construction* which needs to be *listed* (*the white have voted for a white governor = the white people*). In Spanish, even the indefinite determiner can exhibit the same behaviour (with minimal alterations in the masculine gender: *una blanca estaría mayor*; *yo prefiero uno-o negro*). It is of course well known that verbs can appear without their subjects in pro-drop languages. As Hawkins (1994:369) points out, “their agreement features are sufficient to construct S by AgP, whether or not the subject is overtly present”. That is, “higher node information is either not threatened by the deletion (. . .); or if it is threatened, it can be supplied by AgP”. In German, when the determiner is absent, adjectives take on the strong inflection to signal case and thus construct NP. There are a myriad cases like these in the world’s languages. So, even though the existence of agreement does not guarantee deletion, it obviously makes it easier and makes it available to grammars as an option. Now, grammars may vary in what they delete, but they all delete pieces of semantics because form cannot spell-out all of it. Agreement is extremely functional in that it offers a wide menu of deletion of core elements. This is a precious commodity in a world where the choice between radically different interpretations hinges on analyses conducted in well under half a second.
- d. The worldwide prevalence of subject agreement over object agreement is completely expected given the clause-construction role of agreement (see also Harris, 1981). When a verb agrees with an object NP the merged object is a VP, which adds nothing new to the V alone (since the verb is a VP too). However, when a verb agrees with an S-dominated subject NP then agreement constructs S, hence the greater utility of having a grammaticalised subject agreement system. This explains

why there should be a finiteness category sensitive to the presence or absence of a subject, such that the verb takes one morphological form in clauses that contain a subject, and another in clauses that do not. Finiteness is like subject agreement in this regard, and the reason for their similarity is that both are constructors of S. (Hawkins, 1994:372).

All the experimental evidence reviewed here revealed that agreement features are privileged by parsers, especially in Spanish, where they sometimes even cause processing systems to garden path. Considering that, of all computational devices, parsers are the least likely to pay the best of their attention to “meaningless”, “superfluous”, “cumbersome” and “dysfunctional” features (Jespersen, 1922:352; Haiman, 1985:162; Taylor, 2002:332 ff.), the automaticity that is evident when humans process such features only makes sense if the features in question perform an important job. It does not make sense if all that the features do is reveal humans’ delight in “form-focused activities” (Taylor, 2002). In addition to providing redundancy in noisy channels, the job in question is to build the structure of the sentence. No less. Barring the more restricted mechanism of classifiers, that structure is built in two principal ways: either through a consistent and fixed word order, or through the punctuated presence of co-varying morphological clues. Most languages use the two mechanisms simultaneously, but in different and often unique proportions. Here we have seen a few aspects of the way in which English and Spanish use agreement, and of the way in which English and Spanish language users cope with agreement features in the millisecond-by-millisecond construction of the sentence. The fact that English and Spanish are neither typologically nor historically worlds apart means that radical differences in *all* domains should not be expected as – when building the clause – their respective grammatical cores ripple outwards in the system. Ideally, one would need two languages with either very pronounced or almost missing morphological attrition to see the hints provided by Spanish and English crystallise more firmly across domains and processes. Unfortunately, psycholinguistic evidence is still confined in the main to a few European languages, so that will have to wait.

Acknowledgments

This research was funded by the Spanish Ministry of Science and Technology (grant number SEJ2005-08911/PSI) and the Regional Government of Galicia (grant numbers PGIDIT06PXIC204102PN and INCITE07PXI211031ES). These grants are hereby gratefully acknowledged. I am also indebted to the inspiration, the suggestions and the help provided by Manuel Carreiras, without which this work would not have been possible.

References

- Anderson, S., 2004. Some points of agreement. Internet document.
- Antón-Méndez, I., Nicol, J., Garret, M., 2002. The relation between gender and number agreement processing. *Syntax* 5 (1), 1–25.
- Arnold, J., Eisenband, J.G., Brown-Schmidt, S., Trueswell, J., 2000. The rapid use of gender information: evidence of the time course of pronoun resolution from eye-tracking. *Cognition* 76, B13–B26.
- Barber, H., Carreiras, M., 2003. Integrating gender and number information in Spanish word pairs: an ERP study. *Cortex* 39, 465–482.
- Barber, H., Carreiras, M., 2005. Grammatical gender and number agreement in Spanish: an ERP comparison. *Journal of Cognitive Neuroscience* 17 (1), 137–153.
- Barber, H., Salillas, E., Carreiras, M., 2004. Gender or genders agreement? In: Carreiras, M., Clifton, C. (Eds.), *On-line Study of Sentence Comprehension; Eye-tracking, ERP and Beyond*. Psychology Press, Brighton, UK, pp. 309–327.
- Barlow, M., 1992. *A Situated Theory of Agreement*. Garland, New York.
- Barlow, M., 1999. Agreement as a discourse phenomenon. *Folia Linguistica* 33, 1–24.
- Bates, E., MacWhinney, B., 1989. Functionalism and the competition model. In: MacWhinney, B., Bates, E. (Eds.), *The Cross-Linguistic Study of Sentence Processing*. C.U.P., Cambridge, UK, pp. 1–73.
- Benveniste, E., 1971. The nature of pronouns. In: *Problems in General Linguistics*, University of Miami Press, Coral Gables, Florida, pp. 217–222.
- Berg, T., 1998. The resolution of number agreement conflicts in English and German agreement patterns. *Linguistics* 36, 41–70.
- Betancort, M., Carreiras, M., Acuña-Fariña, J.C., 2006. Processing controlled PROs in Spanish. *Cognition* 100, 217–282.
- Bhat, D.N.S., 1991. Grammatical relations. In: *The Evidence Against Their Necessity and Universality*, Routledge, London.
- Bock, K., Cutler, A., Eberhard, K., Butterfield, S., Cutting, J.C., Humphreys, K., 2006. Number agreement in British and American English. *Language* 82 (1), 64–113.
- Bock, K., Eberhard, K.M., 1993. Meaning, sound, and syntax in English number agreement. *Language and Cognitive Processes* 8, 57–99.
- Bock, K., Eberhard, K.M., Cutting, J.C., 2004. Producing number agreement: how pronouns equal verbs. *Journal of Memory and Language* 51, 251–278.
- Bock, K., Eberhard, K.M., Cutting, J.C., Meyer, A., Schriefers, H., 2001. Some attractions of verb agreement. *Cognitive Psychology* 43, 83–128.

- Bock, K., Levelt, W., 1994. Language production: grammatical encoding. In: Gernsbacher, M.A. (Ed.), *Handbook of Psycholinguistics*. Academic Press, San Diego, CA.
- Bogdanov, V.N., 1968. Osobyj slučaj dialektnogo soglasovanija skazuemogo s podležaščim po smyslu i kategorija predstavitel'nosti. *Naučnye doklady vyssšej školy: filologičeskie nauki* 4, 68–75.
- Booij, G., 1996. Inherent versus contextual inflection and the split morphology hypothesis. In: Booij, G., van Marle, J. (Eds.), *Yearbook of Morphology 1995*. Kluwer, Dordrecht, pp. 1–15.
- Bradley, D.C., Foster, K.L., 1987. A reader's view of listening. In: Frauenfelder, U.H., Tyler, L.K. (Eds.), *Spoken Word Recognition*. C.U.P., Cambridge, pp. 103–133.
- Bybee, J., 1985. *Morphology: A Study of the Relation Between Meaning and Form (Typological Studies in Language 99)*. John Benjamins, Amsterdam.
- Cacciari, C., Carreiras, M., Cionini, C.B., 1997. When words have two genders: anaphor resolution for Italian functionally ambiguous words. *Journal of Memory and Language* 37, 517–532.
- Cacciari, C., Padovani, R., Carreiras, M., submitted for publication. Grammatical and discourse-based gender agreement in Italian pronouns.
- Cameron, R., 1993. Ambiguous agreement, functional compensation, and non-specific tú in the Spanish of San Juan, Puerto Rico, and Madrid, Spain. *Language Variation and Change* 5, 305–334.
- Carminati, M.N., 2005. Processing reflexes of the Feature Hierarchy and implications for linguistic theory. *Lingua* 115, 259–285.
- Carreiras, M., Garham, A., Oakhill, J.V., Cain, K., 1996. The use of stereotypical gender information in constructing a mental model: evidence from English and Spanish. *The Quarterly Journal of Experimental Psychology* 49A-3, 639–663.
- Carreiras, M., Gernsbacher, A.M., 1992. Comprehending conceptual anaphors in Spanish. *Language and Cognitive Processes* 7, 281–299.
- Casillas Martínez, L., 2003. Gender mismatched in Spanish and French N1/A de N2 affective constructions: index agreement vs. morphosyntactic concord. In: Kim, J., Wechsler, S. (Eds.), *Proceedings of the 9th International Conference on Head-Driven Phrase Structure Grammar*, Stanford, pp. 1–17.
- Chomsky, N., 1981. *Lectures on Government and Binding*. Foris, Dordrecht.
- Chomsky, N., 1995. *The Minimalist Program*. M.I.T., Cambridge, MA.
- Chomsky, N., 1999. Minimalist enquiries: the framework. In: Martin, R., Michaels, D., Uriagereka, J. (Eds.), *Step by Step: Essays on Minimalist Syntax: in Honour of Howard Lasnik*. M.I.T., Cambridge, MA.
- Chomsky, N., 2001. *Beyond explanatory adequacy*. M.I.T. Unpublished manuscript.
- Colé, P., Seguí, J., 1994. Grammatical incongruence and vocabulary types. *Memory and Cognition* 22, 387–394.
- Copestake, A., 2002. *Implementing Typed Featured Structure Grammars (CSLI Lecture Notes 110)*. CSLI, Stanford.
- Corbett, G.S., 1979. The Agreement Hierarchy. *Journal of Linguistics* 15, 203–224.
- Corbett, G.S., 1991. *Gender*. C.U.P., Cambridge, UK.
- Corbett, G.S., 2000. *Number*. C.U.P., Cambridge, UK.
- Corbett, G.S., 2006. *Agreement*. C.U.P., Cambridge, UK.
- Cowart, W., Cairns, H.S., 1987. Evidence for an anaphoric mechanism within syntactic processing: some reference relations defy semantic and pragmatic constraints. *Memory & Cognition* 15, 318–331.
- Croft, W., 2001. *Radical Construction Grammar*. O.U.P., New York.
- Culicover, P.W., 1999. *Syntactic Nuts: Hard Cases in Syntax*. O.U.P., Oxford.
- Deutsch, A., Bentin, S., 2001. Syntactic and semantic factors in processing gender agreement in Hebrew: evidence from ERPs and eye movements. *Journal of Memory and Language* 45 (2), 200–224.
- De Vincenzi, M., 1999. Differences between the morphology of gender and number: evidence from establishing coreferences. *Journal of Psycholinguistic Research* 28, 537–553.
- Di Domenico, E., De Vincenzi, M., 1995. Gender and number in the retrieval of pronoun antecedents: differences in use and representation. In: Nash, L., Tsoules, F., Zribi-Herts, A. (Eds.), *Actes du deuxième colloque 'Langues et grammaire'*. Paris, pp. 95–109.
- Dik, S., 1978. *Functional Grammar*. North-Holland, Amsterdam.
- Dik, S., 1991. *Functional grammar*. In: Droste, F., Joseph, J. (Eds.), *Linguistic Theory and Grammatical Description*. John Benjamins, Amsterdam/Philadelphia, pp. 247–274.
- Dowty, D., Jacobson, P., 1988. Agreement as a semantic phenomenon. In: Powers, J., Jong, K.D. (Eds.), *Proceedings of the Fifth Eastern States Conference on Linguistics*. Ohio State University, Columbus, OH, pp. 95–108.
- Eberhard, K.M., 1997. The marked effect of number on subject–verb agreement. *Journal of Memory and Language* 36, 147–164.
- Eberhard, K.M., 1999. The accessibility of conceptual number to the processes of subject-verb agreement in English. *Journal of Memory and Language* 41, 560–578.
- Eberhard, K.M., Cutting, J.C., Bock, K., 2005. Making syntax of sense: number agreement in sentence production. *Psychological Review* 112 (3), 531–559.
- Ehrlich, K., 1980. Comprehension of pronouns. *The Quarterly Journal of Experimental Psychology* 32, 247–255.
- Ehrlich, K., Rayner, K., 1983. Pronoun assignment and semantic integration during reading: eye movements and immediacy of processing. *Journal of Verbal Learning and Verbal Behavior* 22, 75–87.
- Epstein, S., Seely, D., 2002. Rule applications as cycles in a level-free syntax. In: Epstein, S., Seely, D. (Eds.), *Derivation and Explanation in the Minimalist Program*. Blackwell, Oxford.
- Epstein, S., Seely, D., 2006. *Derivations in Minimalism*. C.U.P., Cambridge, UK.
- Faussart, C., Jacobowicz, C., Costes, M., 1999. Gender and number processing in spoken French and Spanish. *Rivista de Linguistica* 11, 75–101.
- Ferguson, C.A., 1964. Baby talk in six languages. *Language* 66, 103–114.

- Fernández-Salgueiro, G., 2005. Agree, the EPP-F and further-raising in Spanish. In: Gess, R., Rubin, E. (Eds.), *Experimental and Theoretical Approaches to Romance Linguistics*. John Benjamins, Philadelphia.
- Fernández-Salgueiro, G., 2006. Romance Null Subjects and the Sensory-Motor Interface. Unpublished manuscript of the University of Michigan.
- Franck, J., Lasso, G., Frauenfelder, U., Rizzi, L., 2006. Agreement and movement: a syntactic analysis of attraction. *Cognition* 101, 173–216.
- Frazier, L., Flores D'Arcais, G., 1989. Filler-driven parsing: a study of gap filling in Dutch. *Journal of Memory and Language* 28, 331–344.
- Friederici, A.D., 1995. The time course of syntactic activation during language processing: a model based on neuropsychological and neurophysiological data. *Brain and Language* 50, 259–281.
- Friederici, A.D., Hahne, A., Mecklinger, A., 1996. Temporal structure of syntactic parsing: early and late event-related brain potential effects. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 22 (5), 1219–1248.
- Garnham, A., Oakhill, J., 1985. On-line resolution of anaphoric pronouns: effects of inference making and verb semantics. *British Journal of Psychology* 76, 385–393.
- Garnham, A., Oakhill, J., Ehrlich, M.-F., Carreiras, M., 1995. Representations and processes in the interpretation of pronouns: new evidence from Spanish and French. *Journal of Memory and Language* 34, 41–62.
- Garnham, A., Traxler, M., Oakhill, J., Gernsbacher, M.A., 1996. The locus of implicit causality effects in comprehension. *Journal of Memory and Language* 35, 517–543.
- Garrod, S., Freudenthal, D., Boyle, E., 1994. The role of different types of anaphors in the on-line resolution of sentences in a discourse. *Journal of Memory and Language* 33, 39–68.
- Garrod, S., Sanford, A.J., 1990. Referential processing in reading. Focusing on roles and individuals. In: Balota, D.A., Flores d'Arcais, G.B., Rayner, K. (Eds.), *Comprehension Processes in Reading*. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 465–480.
- Garrod, S., Terras, M., 2000. The contribution of lexical and situational knowledge to resolving discourse roles: bonding and resolution. *Journal of Memory and Language* 42, 526–544.
- Gernsbacher, M.A., 1989. Mechanisms that improve referential access. *Cognition* 32 (2), 99–156.
- Gernsbacher, M.A., 1991. Comprehending conceptual anaphors. *Language and Cognitive Processes* 6, 81–105.
- Goldberg, A., 1995. *A Construction Grammar Approach to Argument Structure*. Chicago UP, Chicago.
- Goldberg, A., 2006. *Constructions at Work*. O.U.P., Oxford.
- Gordon, P.C., Hendrick, R., Foster, K.L., 2000. Language comprehension and probe-list memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26, 766–775.
- Greenberg, J.H., 1963. Some universals of grammar with particular reference to the order of meaningful elements. In: Greenberg, J.H. (Ed.), *Universals of Grammar*. M.I.T., Cambridge, MA, pp. 73–113.
- Greene, S.B., McKoon, G., Ratcliff, R., 1992. Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory and Cognition* 18, 266–283.
- Grevisse, M., 1964. *Le bon Usage*, 8th ed. J. Duculot, Glemboux.
- Gunter, T.C., Friederici, A.D., Schriefers, H., 2000. Syntactic gender and semantic expectancy: ERPs reveal early autonomy and late interaction. *Journal of Cognitive Neuroscience* 12, 556–568.
- Hagoort, P., Brown, C.M., 1999. Gender electrified: ERP evidence on the syntactic nature of gender processing. *Journal of Psycholinguistic Research* 28 (6), 715–728.
- Hagoort, P., Brown, C., Groothusen, J., 1993. The syntactic positive shift (SPS) as an ERP measure of syntactic processing. *Language and Cognitive Processes* 8, 439–483.
- Harley, H., Ritter, E., 2002. Person and number in pronouns: motivating a feature-geometric analysis. *Language* 78, 482–526.
- Harris, A., 1981. *Georgian Syntax: a Study in Relational Grammar*. C.U.P., Cambridge, UK.
- Haskell, T., Macdonald, M., 2003. Conflicting cues and competition in subject–verb agreement. *Journal of Memory and Language* 48, 760–778.
- Hawkins, R., 1981. Towards an account of the possessive constructions: “NP’s N” and “the N of NP”. *Journal of Linguistics* 17, 247–269.
- Hawkins, J., 1994. *A Performance Theory of Order and Constituency*. C.U.P., Cambridge, UK.
- Hawkins, J., 2004. *Efficiency and Complexity in Grammars*. O.U.P., Oxford.
- Heap, D., 2000. *La variation grammaticale en géolinguistique: les pronoms sujet en roman central (Lincom Studies in Romance Linguistics 11)*. Lincom Europa, Munich.
- Huddleston, R., Pullum, G., 2002. *The Cambridge Grammar of the English Language*. C.U.P., Cambridge, UK.
- Hudson, R., 1999. Subject–verb agreement in English. *Journal of English Language and Linguistics* 3, 173–207.
- Igoa, J.M., García-Albea, J.E., Sánchez-Casas, R., 1999. Gender-number dissociations in sentence production in Spanish. *Rivista di Linguistica* 11, 163–196.
- Jespersen, O., 1922. *Language, Its Nature, Development and Origin*. Henry Holt and Co., New York.
- Jespersen, O., 1924. *The Philosophy of Grammar*. Allen & Unwin, London.
- Joseph, B.D., 1979. On the agreement of reflexive forms in English. *Linguistics* 17, 519–523.
- Kaplan, R., Bresnan, J., 1982. Lexical-functional grammar: a formal system for grammatical representations. In: Bresnan, J. (Ed.), *The Mental Representations of Grammatical Relations*. M.I.T., Cambridge, MA, pp. 173–281.
- Kay, P., Fillmore, C.J., 1999. Grammatical constructions and linguistic generalizations: the *What’s X doing Y?* construction. *Language* 75, 1–33.
- Kazanina, N., Lau, E., Lieberman, M., Yoshida, M., Phillips, C., 2007. The effect of syntactic constraints on the processing of backwards anaphora. *Journal of Memory and Language* 56, 384–409.
- Kennison, S.M., Trofe, J.L., 2003. Comprehending pronouns: a role for word-specific gender stereotype information. *Journal of Psycholinguistic Research* 32 (3), 355–378.
- Konishi, T., 1993. The semantics of grammatical gender: a cross-cultural study. *Journal of Psycholinguistic Research* 22, 519–534.
- Kutas, M., Hillyard, S.A., 1983. Event-related brain potentials to grammatical errors and semantic anomalies. *Memory and Cognition* 11, 539–550.

- Lambrecht, K., 1994. *Information Structure and Sentence Form*. C.U.P., Cambridge, UK.
- Lambrecht, K., 1994. *Information Structure and Sentence Form*. C.U.P., Cambridge, UK.
- Lehmann, C., 1982. Universal and typological aspects of agreement. In: Seiler, H., Stachowiak, F.J. (Eds.), *Apprehension: Das sprachliche Erfassen von Gegenständen II: Die Techniken und ihr Zusammenhang in Einzelsprachen*. Narr, Tübingen, pp. 201–267.
- Levelt, W., Roeloffs, A., Meyer, A.S., 1999. A theory of lexical access in speech production. *Behavioral and Brain Science* 22, 1–75.
- Levin, M., 2001. Agreement with Collective Nouns in English (*Lund Studies in English* 103). Almqvist & Wiksell, Stockholm.
- Lukatella, G., Kostic, A., Todorovic, D., Carello, C., Turvey, M.T., 1987. Type and number of violations and the grammatical congruency effect in lexical decision. *Psychological Research* 49, 37–43.
- MacDonald, M., MacWhinney, B., 1990. Measuring inhibition and facilitation from pronouns. *Journal of Memory and Language* 29, 469–492.
- MacDonald, M.C., Pearlmuter, N., Seidenberg, M.S., 1994. The lexical nature of syntactic ambiguity resolution. *Psychological Review* 101, 676–703.
- Mallinson, G., Blake, B., 1981. *Language Typology: Cross-Linguistic Studies in Syntax*. North Holland, Amsterdam.
- Marslen-Wilson, W., Tyler, L.K., Koster, C., 1993. Integrative processes in utterance resolution. *Journal of Memory and Language* 32, 647–666.
- McDonald, J.L., MacWhinney, J., 1995. The time course of anaphor resolution: effects of implicit verb causality and gender. *Journal of Memory and Language* 34, 543–566.
- Mills, A.E., 1986. *The Acquisition of Gender: a Study of English and German*. Springer-Verlag, Berlin.
- Moravcsik, E., 1978. Agreement. In: Greenberg, H., Ferguson, C.A., Moravcsik, E. (Eds.), *Universals of Human Language. IV. Syntax*. Stanford University Press, Stanford, pp. 331–374.
- Moravcsik, E., 1995. Summing up Suffixaufnahme. In: Plank, F. (Ed.), *Double Case: Agreement by Suffixaufnahme*. Oxford University Press, Oxford, pp. 451–484.
- Morgan, J.J., 1984. Some problems of agreement in English and Albanian. In: Brugman, C., Maccauley, M., Dahlstrom, A., Emanatian, M., Moonwoman, B., O'Connor, C. (Eds.), *Proceedings of the 10th Annual Meeting of the Berkeley Linguistics Society*. Berkeley Linguistics Society, University of California, Berkeley, pp. 233–247.
- Münter, T.F., Heinze, H.J., 1994. ERP negativities during syntactic processing of written words. In: Heinze, H.J., Münte, T.F., Mangun, G.R. (Eds.), *Cognitive Electrophysiology*. Birkhäuser, Boston, pp. 211–238.
- Neville, H., Nicol, J.L., Barss, A., Forster, K.I., Garrett, M.F., 1991. Syntactically based sentence processing classes: evidence from event-related brain potentials. *Journal of Cognitive Neuroscience* 3, 151–165.
- Nicol, J., 1988. *Coreference processing during sentence comprehension*. Unpublished Doctoral Dissertation: M.I.T.
- Nunes, J., 1999. Linearization of chains and phonetic realization of chain links. In: Epstein, S., Hornstein, N. (Eds.), *Working Minimalism*. M.I.T., Cambridge, MA, pp. 217–249.
- Nunes, J., 2004. *Linearization of Chains and Sideward Movement*. M.I.T., Cambridge, MA.
- Osterhout, L., Bersick, M., McLaughlin, J., 1997. Brain potentials reflect violations of gender stereotypes. *Memory and Cognition* 25, 273–275.
- Osterhout, L., Holcomb, P.J., 1992. Event-related brain potentials elicited by syntactic anomaly. *Journal of Memory and Language* 31, 785–806.
- Osterhout, L., Holcomb, P.J., Swinney, D.A., 1994. Brain potentials elicited by garden-path sentences: evidence of the application of verb information during parsing. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20, 786–803.
- Osterhout, L., Mobley, L.A., 1995. Event-related brain potentials elicited by failure to agree. *Journal of Memory and Language* 34, 739–773.
- Pesetsky, D., 1995. *Zero Syntax: Experiencers and Cascades*. M.I.T., Cambridge, MA.
- Picallo, M.C., 1991. Nominals and nominalization in Catalan. *Probus* 3, 279–316.
- Pollard, C., Sag, I.A., 1988. An information-based theory of agreement. In: Brentari, G., Larson, G., McLeod, L. (Eds.), *Papers from the 24th Annual Regional Meeting of the Chicago Linguistics Society*. Chicago Linguistics Society, Chicago, pp. 236–257.
- Pollard, C., Sag, I., 1994. *Head-Driven Phrase Structure Grammar*. University of Chicago Press, Chicago.
- Radford, A., 2004. *Minimalist Syntax*. C.U.P., Cambridge, UK.
- Reid, W., 1991. *Verb and Noun Number in English*. Longman, London.
- Rigalleau, F., Caplan, D., 2000. Effects of gender marking in pronominal coindexation. *The Quarterly Journal of Experimental Psychology: Human Experimental Psychology* 53A, 23–52.
- Rigalleau, F., Caplan, D., Baudiffier, V., 2004. New arguments in favour of an automatic gender pronominal process. *The Quarterly Journal of Experimental Psychology, Section A* 57, 893–933.
- Ritter, E., 1988. A head movement approach to construct state noun phrases. *Linguistics* 26, 909–929.
- Ritter, E., 1991. Two functional categories in noun phrases: evidence from modern Hebrew. In: Rothstein, S. (Ed.), *Syntax and Semantics 25: Perspectives on Phrase Structure*. Academic Press, New York, pp. 37–62.
- Ritter, E., 1993. Where's gender? *Linguistic Inquiry* 24, 795–803.
- Rusakova, M.V., 2001. *Imennaja slovoforma flektivnogo jazyka (soglasovanie v russkom atributivnom slovosocetanii)*. Unpublished PhD dissertation, St Petersburg: Gosudarstvennyj pedagogičeskij universitet im. A.I. Gercena.
- Sera, M., Berge, C., del Castillo, J., 1994. Grammatical and conceptual forces in the attribution of gender by English and Spanish speakers. *Cognitive Development* 9 (3), 261–292.
- Sera, M., Selieff, C., Burch, M., Forbes, J., Rodriguez, W., 2002. When language affects cognition and when it does not: an analysis of grammatical gender and classification. *Journal of Experimental Psychology General* 131, 377–397.
- Shieber, S., 1986. *An Introduction to Unification-based Approaches to Grammar (CSLI Lecture Notes 4)*. CSLI, Stanford.
- Siewierska, A., 1998. Variation in major constituent order: a global and a European perspective. In: Siewierska, A. (Ed.), *Constituent Order in the Languages of Europe*. Mouton de Gruyter, Berlin, pp. 475–551.
- Siewierska, A., 1999. From anaphoric pronoun to grammatical agreement marker: why objects don't make it. *Folia Linguistica* 33 (2), 225–251.
- Silverstein, M., 1985. Hierarchy of features and ergativity. In: Muysken, P., van Riemsdijk, H. (Eds.), *Features and Projections*. Foris, Dordrecht, pp. 163–232.

- Sonderegger, S., 1998. Dichterische Wortstellungstypen im Altgermanischen und ihr Nachleben im älteren Deutsch. In: Askedal, J.O. (Ed.), *Historische Germanische und Deutsche Syntax*. Frankfurt am Main, Lang, pp. 25–47.
- Stowe, L., 1986. *Models of Gap Location in the Human Language Processor*. Indiana University Linguistics Club, Bloomington, IN.
- Sturt, P., 2003. The time course of the application of binding constraints in reference resolution. *Journal of Memory and Language* 48, 542–562.
- Tanenhaus, M.K., Spivey-Knowlton, M., Hanna, J.E., 2000. Modelling thematic and discourse context effects with a multiple constraints approach: implications for the language comprehension system. In: Crocker, M.W., Pickering, M., Clifton, C. (Eds.), *Architectures and Mechanisms for Language Processing*. C.U.P., Cambridge, MA, pp. 90–118.
- Taraban, R., McClelland, J.R., 1988. Constituent attachment and thematic role assignment in sentence processing: influence of content-based expectations. *Journal of Memory and Language* 27, 597–632.
- Taylor, J., 2002. *Cognitive Grammar*. O.U.P., Oxford.
- Thornton, R., MacDonald, M., 2003. Plausibility and grammatical agreement. *Journal of Memory and Language* 48, 740–759.
- Tsoulas, G., Kural, M., 1998. Indexical pronouns as bound variables. In: Bird, S., Tsoulas, G., Kuras, M. (Eds.), *Proceedings of the 18th West Coast Conference on Formal Linguistics*. pp. 545–557.
- Van Gompel, R., Liversedge, S., 2003. The influence of morphological information on cataphoric pronoun assignment. *Journal of Experimental Psychology, Learning, Memory and Cognition* 29 (1), 128–139.
- Van Valin, R., Lapolla, R., 1997. *Syntax*. C.U.P., Cambridge, UK.
- Vigliocco, G., Butterworth, B., Garrett, M.F., 1996b. Subject–verb agreement in Spanish and English: the role of conceptual factors. *Cognition* 51, 261–298.
- Vigliocco, G., Franck, J., 2001. When sex affects syntax: contextual influences in sentence production. *Journal of Memory and Language* 45, 368–390.
- Vigliocco, G., Hartsuiker, R.J., 2002. The interplay of meaning, sound and syntax in language production. *Psychological Bulletin* 128, 442–472.
- Vigliocco, G., Hartsuiker, R.J., Jarema, G., Kolk, H.H.J., 1996a. One or more labels on the bottles? Notional concord in Dutch and French. *Language and Cognitive Processes* 11, 407–442.
- Vigliocco, G., Vinson, D.P., Paganelli, F., Dworzynski, K., 2005. Grammatical gender effects on cognition: implications for language learning and language use. *Journal of Experimental Psychology* 134 (4), 501–520.
- Vonk, W., 1984. Eye movement during the comprehension of pronouns. In: Gale, A.G., Johnson, F. (Eds.), *Theoretical and Applied Aspects of Eye Movement Research*. North Holland, Amsterdam, pp. 203–212.