THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Subset controllers in agreement relations

Citation for published version:

Ackema, P & Neeleman, A 2013, 'Subset controllers in agreement relations', *Morphology*, vol. 23, no. 2, pp. 291-323. <https://doi.org/10.1007/s11525-013-9218-4>

Digital Object Identifier (DOI):

[10.1007/s11525-013-9218-4](https://doi.org/10.1007/s11525-013-9218-4)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Morphology

Publisher Rights Statement:

© Ackema, P., & Neeleman, A. (2013). Subset controllers in agreement relations. *Morphology*, 23(2): 291-323. doi: 10.1007/s11525-013-9218-4

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

SUBSET CONTROLLERS IN AGREEMENT RELATIONS¹

Peter Ackema (University of Edinburgh) and Ad Neeleman (UCL)

1. THE ISSUE

Language could be organized in such a way that features are realized phonologically where they are interpreted. However, this hypothesis is incorrect. One type of exception has to do with PF exponence. It is possible for features to be spelled out in linear positions that do not match their position in morpho-syntax (as in the case of *passer-by*; see Sproat 1985; Ackema and Neeleman 2003). In this paper we argue that a second type of exception exists at the LF interface. That is, features can be spelled out in a position that matches their morpho-syntactic position, but interpreted elsewhere. Our argument is based on ϕ -features in certain agreement relations.

There is a basic asymmetry between arguments and verbs when it comes to the semantics of ϕ -features. Person, number and gender are notions that bear on the interpretation of arguments, but not on the interpretation of verbs. Hence in agreement relations the nominal element is said to be the ‘controller’ of the relation, while the verb is said to be the ‘target’ (see Corbett 2006 for an overview and further references). The question that this phenomenon poses is whether the interpretive asymmetry is reflected in syntax. Many theories of agreement assume that this is indeed the case. The presence of verbal ϕ -features or feature values is supposed to be dependent, in some way or other, on the presence of identical nominal ϕ -features or feature values. Chomsky (2000, 2001), for example, argues that ϕ -features of arguments start out valued, while ϕ -features of verbs are unvalued initially. They acquire a value through the operation of agreement. Frampton and Gutmann (2000) and Pesetsky and Torrego (2007) propose a subtle adjustment of Chomsky’s proposal, suggesting that agreement is an operation of feature sharing. However, they continue to assume that verbal ϕ -features have no initial value and acquire one through the sharing operation. Like Chomsky’s proposal, this alternative is therefore asymmetrical in the sense intended above. Finally, Bobaljik (2007) treats agreement as the copying of ϕ -features from an argument to a verbal head, suggesting an even sharper syntactic asymmetry than Chomsky: verbal ϕ -features are not just unvalued initially, but absent.

Other theories of agreement (in particular, unification-based theories of agreement) are compatible with the view that there is no syntactic asymmetry between nominal and verbal ϕ -features. We will argue in favour of this alternative view, basing our argument on the phenomenon of what we will call subset controllers. These are arguments that at least at face value are specified for fewer ϕ -features than the verb that agrees with them. An example involving number is found in Inari Sami, where dual can

¹ We thank Claudia Borgonovo, Antonio Fábregas, Luisa Martí, Àngels Molina and Martha Robinson for help with the Spanish data in section 4. Examples of quantificational unagreement, unless indicated otherwise, are based on Molina’s MA Thesis (UCL, 2010). Earlier versions of this paper were presented at the *Words don’t come easy* workshop at the University of Verona (2008), at NESS at the University of York (2010) and at the City University of Hong Kong (2011). We thank the respective audiences for useful questions and comments. We would also like to thank Valentina Bianchi, Antonio Fábregas, David Lobina, Andrew Nevins, Hans van de Koot, and Edwin Williams for discussion, as well as two anonymous reviewers and the editor of this special issue, Jochen Trommer.

be expressed in the verb, but not in non-pronominal subjects (pronominal subjects do express the dual):

- (1) a. Alma-h kuálást-ava onne. (Inari Sami; Toivonen 2007)
man-PL.NOM fish-3DU today
 ‘The two men are fishing today.’
 b. Alma-h kuálást-eh onne.
man-PL.NOM fish-3.PL today
 ‘The men are fishing today.’

An example involving gender can be found in (2). First and second person pronouns in Serbo-Croatian are not specified for gender. Nevertheless, the gender of speaker and hearer is identified through past participle agreement:

- (2) a. Ja sam otisla na posao (Serbo-Croatian; Amela Camdzic, p.c.)
I am gone-FEM.SG to work
 ‘I have gone to work’ (said by a woman)
 b. Ja sam otisao na posao
I am gone-MASC.SG to work
 ‘I have gone to work’ (said by a man)

The problem raised by these kinds of data is that a feature appears to be interpreted in a position different from the one in which it is spelled out. If ϕ -features in the verb are dependent on the presence of identical features in the subject, as in an asymmetric theory, then examples like (1a) must involve a lack of exponence of the relevant features in the subject. However, we will argue that, at least in some cases of subset control, there is no evidence for the syntactic presence of these features in the subject. Therefore, asymmetric theories of agreement cannot account for the data.

We also show that symmetric theories of agreement *can* deal with the existence of subset controllers, because they allow fully specified features to be generated in the verb, independently of the ϕ -feature specification of the subject. A simple mechanism of feature spreading (which constitutes a particular type of unification) will ensure the feature is interpreted in the right position at LF, but as feature spreading happens in covert syntax, there is nothing remarkable about the lack of exponence in the subject. What we propose, then, is that data like (1) should be captured through ‘LF exponence’ rather than ‘PF exponence’.

We will formalize our proposal in section 3, but before we do so, we need to introduce a more explicit theory of ϕ -features based on the idea of feature geometry (section 2). The main evidence for the existence of proper subset controllers will be presented in section 4, where we discuss so-called ‘unagreement’ in Spanish.

2. Φ -FEATURES

Our starting point is the assumption that ϕ -features are privative and organised in a feature geometry. When a feature restricts the interpretation of another feature, it is

represented as a daughter of that feature. The idea goes back at least to Gazdar and Pullum 1982; a detailed proposal has more recently been developed in Harley and Ritter 2002. For reasons of presentation we will use a simplified version of the geometry proposed by the latter authors. It is given in (3).

The root of the tree in (3) is occupied by a ϕ -node, which may host features bearing on the interpretation of nominal categories, in particular, person, number and gender features. The gender features [CG] for ‘common gender’ and [FEM] for ‘feminine’ will play only a minor role in this paper. The person system in many languages can be characterized using two features: [PAR] and [ADD] for ‘participant in the speech act’ and ‘addressee’, respectively (see Kerstens 1993, Harley and Ritter 2002, and Ackema and Neeleman 2004). The latter restricts the interpretation of [PAR], and is therefore taken to be dependent on this feature.² Two features are enough to characterize the number systems of most languages, namely [PL] and [MIN] for ‘plural’ and ‘minimal’, respectively (to cover all possible number systems, the geometry has to be slightly more complex; see Harley and Ritter 2002). The [MIN] feature restricts [PL], giving rise to a dual interpretation (two is the minimal number larger than one). Consequently, [MIN] is dominated by [PL] in the feature geometry. A system with three persons and singular, plural and dual can thus be analyzed as follows:

(4)

	1 ST	2 ND	3 RD
SG	a. <pre> φ PAR </pre>	b. <pre> φ PAR ADD </pre>	c. <pre> φ </pre>
PL	d. <pre> φ / \ PAR PL </pre>	e. <pre> φ / \ PAR PL ADD </pre>	f. <pre> φ PL </pre>
DU	g. <pre> φ / \ PAR PL MIN </pre>	h. <pre> φ / \ PAR PL ADD MIN </pre>	i. <pre> φ PL MIN </pre>

² In Ackema & Neeleman 2011, we argue for a variant of this proposal in which second person shares a feature with third person. This change does not affect our argumentation in this paper.

As Harley and Ritter point out, one of the advantages of a feature-geometric analysis of ϕ -features is that it can explain certain linguistic universals. For example, Greenberg (1963:94) suggests that there is no language in which the plural does not have some non-zero allomorphs, whereas there are (many) languages in which the singular remains unexpressed. This follows if plural is represented by a feature [PL], while singular corresponds to the absence of this feature. After all, absence of a feature is not the kind of thing affixes tend to mark. Greenberg further suggests that no language has a dual unless it has a plural. This follows because [MIN] is a feature that restricts the interpretation of [PL]. Therefore, if a language lacks [PL] it cannot have [MIN].

Feature geometric representations of ϕ -features can also capture patterns of agreement found in individual languages. In Jingulu, for example, subjects may be indexed in the verb through a variety of agreement endings. The variation is governed by the hierarchy in (5a): there may be full agreement, or an ending may be replaced by any ending to its left on the hierarchy (Pensalfini 2003:173-174, as quoted in Corbett 2006:153).

- (5) a. singular < plural < dual
 b. Nyama-baji imimikin-bili-rni-rni ardalakbi-wurru-ju. (PL instead of DU)
DEM-PL old.woman-ANIM.DU-F-ERG hot-3PL-do
 ‘The two old women feel hot.’
 c. Kunyirrini dij bila-nya-mi kandirri! (SG instead of DU)
2DU.ERG PRV divide-2SG-IRR bread
 ‘You two cut up the bread.’
 d. Nginda-rni ngaja-mi jurliji-rdarra diyim ka-rdu. (SG instead of PL)
DEM(M)-FOC see-IRR bird-PL fly 3SG-go
 ‘Look at all the birds flying.’

The Jingulu data can be explained in terms of impoverishment (post-syntactic feature deletion; see Williams 1981, Bonet 1991, 1995, among others). We assume that Jingulu always has full agreement in the syntax and that impoverishment optionally affects number features. The geometry in (3) then captures the variation illustrated in (5): deletion of [MIN] in (4g,h,i) yields (4d,e,f), whereas deletion of [PL] in (4d,e,f) yields (4a,b,c). It also explains why it is not possible for an ending to be replaced by an ending to its right in the hierarchy in (5a). How this pattern could be captured in a system that treats singular, plural and dual as values of a number feature is not obvious.

Feature geometry is partly based on the idea of underspecification. The logic behind that concept has the consequence that a given semantic interpretation may be compatible with more than one geometry. For example, reference to the addressee is compatible not only with (4b) (which includes [ADD] and must therefore refer to the addressee), but also with (4a) (since the hearer is a participant) and with (4c) (since (4c) can in principle refer to anything). There must therefore be a system that regulates which constellation of features is chosen if more than one is available.

This issue is comparable with an issue that arises at the PF interface. Given a feature geometry presented to the spell-out system, there is often more than one morpheme available that could realize this geometry. Consider the spell-out rules for pronouns in (6). The geometry in (4b) could in principle be associated with a form

through any of these rules, since all of them are compatible with this geometry. Of course, in practice, (6b) must block (6a) and (6c)

- (6)
- a. $[\varphi \text{ PAR}] \Leftrightarrow /I/$
 - b. $[\varphi \text{ PAR ADD}] \Leftrightarrow /you/$
 - c. $[\varphi] \Leftrightarrow /he/; /she/; /it/$

To enforce this blocking effect, often an appeal is made to the general idea that where there is a choice the most specific spell-out rule must be used. This idea goes back to Panini and was introduced into generative grammar in the form of Kiparsky's (1973) Elsewhere Condition. It can be found in a range of models. Distributed Morphology, for example, implements it as part of the Subset Principle (see Halle 1997). The conceptual motivation behind the use of the Elsewhere Condition for spell-out is that it guarantees that as much of the syntactic input as possible is encoded phonologically.

The issue that arises at the LF interface can be seen in a similar light, as long the syntax is considered an encoding of semantics. The feature geometry licensed in the syntax is the one that encodes as much as possible of the information present in the semantics. We will implement this idea as a very general principle that restricts mappings in the model of grammar in (7) (which is intended to depict a parallel architecture in the sense of Jackendoff 1997; see also Ackema and Neeleman 2004 for motivation). Before we formulate the principle, however, we will say a bit more about the proposed model.

We take the mappings between semantics and LF, between LF and PF, and between PF and phonology to be non-directional. But this does not imply that the principles that define well-formed mappings cannot be asymmetric in that they take one representation as a given and then impose a particular restriction on a second representation.

A clear example of a principle that operates ‘from right to left’ in (7) is Compositionality. After all, Compositionality takes the LF representation as given and restricts the semantics in such a way that the interpretation associated with a syntactic category is determined by the material that that category contains and the way that material is combined.

An example of a principle that operates ‘from left to right’ is the Minimal Link Condition as interpreted by Fanselow (2004) and others. The starting point of Fanselow’s discussion is the observation that superiority violations are tolerated when they result in a meaning that cannot be expressed by alternative syntactic structures (see also Golan 1993, Kitahara 1994, Sternefeld 1997 and Reinhart 2006). Fanselow concludes from this that the Minimal Link Condition must be a principle that is sensitive to interpretation. Given a particular semantics, a syntactic representation must be chosen that has the shortest possible movement links.

When we say that the mapping between LF and semantics is non-directional, what we mean is that this mapping can be characterized jointly by principles that operate from left to right and by principles that operate from right to left. In other words, there is no incompatibility between Compositionality and the Minimal Link Condition as interpreted by Fanselow. In this respect, there is a clear affinity between the model in (7) and work in bi-directional optimality theory (see Blutner 2000, Jäger 2002, De Hoop & Malchukov 2007, among others).

With this in mind, let us return to problem at hand, namely how the interpretation of feature geometries, including underspecified ones, can be regulated. The principle we propose governs all rightward mappings between representations in (7), requiring that as little information as possible is lost in these mappings:

(8) *Maximal Encoding*

A mapping $R \rightarrow R^*$ is licit only if R^* is the maximal expression of R at the relevant level of representation.

We should be explicit about the conditions under which R^* counts as the maximal expression of R , and therefore adopt the definition below:

(9) R^* expresses R maximally if there is no alternative R' that encodes more properties of R or encodes these properties in more locations.

The effects of Maximal Encoding for the mapping between PF and the phonology are uncontroversial. As will be clear, the condition when applied at PF is no more than a restatement of the Elsewhere Condition (or the Subset Principle of Distributed Morphology).

Maximal Encoding also regulates the interpretative effects of underspecification. It implies that given a particular semantics, a feature geometry is licensed at LF only if there is no feature geometry that encodes that semantics more explicitly. The easiest way to understand this effect is to consider a series of syntactic representations that can all be mapped to the same interpretation: $\{\langle R_1, I \rangle, \langle R_2, I \rangle, \dots, \langle R_n, I \rangle\}$. Maximal Encoding demands that given the interpretation I , the syntactic representation chosen is the one with the richest feature specification among $\{R_1, R_2, \dots, R_n\}$.

As an example, consider pronouns that can be interpreted as second person dual. In Inari Sami, the richest feature geometry compatible with this interpretation is the one in (4h). However, as all other feature geometries available in the language are subsets of (4h), all in principle allow an interpretation as second person dual. The reason that there is no optionality in the choice of the pronoun is because (4h), being the most richly specified form, blocks the other feature geometries as a consequence of Maximal Encoding. In Dutch, for instance, there is no expression of dual: the representations in (4g-i) are not available. Therefore, the most highly specified feature geometry compatible with a second person dual interpretation is the one in (4e), which will be selected by Maximal Encoding, blocking (4a-d) and (4f). In other words, (4e) allows a dual interpretation in Dutch, but not in Inari Sami, as in that language a feature geometry encoding dual is available. (For a fuller analysis of Inari Sami based on very similar ideas, see Toivonen 2007.)

Thus, Maximal Encoding constrains the mapping between semantics and syntax (LF), as well as the mapping between syntax (PF) and phonology. In addition to this, we will argue in the next section, it applies internally to the syntax, in that the representation presented at PF must maximally express the representation involved in the mapping at LF.

Depending on one's view of what the grammar is a theory of, the above may be confusing. If the grammar is taken to be a description of an actual procedure that builds linguistic representations in real time, it is hard to see how mapping principles could co-exist that operate in different directions. For example, Maximal Encoding requires look-ahead if the derivation is taken to proceed from the syntax to semantics, and look-ahead seems suspicious in procedural terms. However, in our view the grammar is an abstract function that connects sound and meaning. It determines whether or not a pairing of a meaning representation and a sound representation is well-formed, but it has nothing to say about the algorithm used to compute this function when people produce or understand language (compare Marr's 1982 discussion of levels of description as they apply to computational systems). Thus it is a separate empirical question which algorithm is chosen as the implementation of the grammar in the human language faculty. In our view, it is a category error to impose on the grammar requirements that have to do with efficient computation, such as a ban on look-ahead. For the same reason, consistency of direction is not a reasonable requirement of grammars. All that grammars have to do is provide an explicit formal account of possible associations of sound and meaning. This is, in fact, a very conservative view, going back at least to the competence-performance distinction in Chomsky 1965 (see Neeleman & Van de Koot 2010 for further discussion).

3. PROPOSAL

As we will now argue, the notion of ϕ -feature geometry, in combination with a symmetric theory of agreement, gives us a handle on the analysis of subset controllers. Our starting point is the assumption that the syntax defines a mapping between an LF representation and a PF representation. As discussed, this mapping is not directional. This means that, without fear of inconsistency, we can combine the Minimalist model of grammar, which derives an LF and a PF from a shared underlying structure, with the principle of Maximal Encoding, which works from left to right in the model in (7). In Minimalism, the grammar is typically described as having a direction from the lexicon to the interfaces. However, in our view, this directional account is in essence a claim about the best way to describe the mapping between LF and PF, and should not be seen as a psycholinguistic procedure.

3.1 Regular Agreement as Feature Identification

We first consider aspects of the syntactic derivation relevant to agreement.

The theory of agreement we adopt is symmetrical: ϕ -features are generated on the verb independently of the ϕ -features of the argument. The theory of feature checking in Chomsky 1993 is a recent precursor of this idea. The implementation we adopt is based on the three assumptions in (10), (11) and (12) below. The first of these states that at LF each occurrence of a ϕ -feature must be licensed. Licensing is not the same as interpretation, as not every ϕ -feature has semantic import. For example, French *table*

‘table’ is a feminine noun, but it is not feminine in its interpretation. The condition must therefore accommodate ϕ -features inherent in nouns:

- (10) *Φ -feature licensing*: At LF, each ϕ -feature F must be licensed in each position L with which it is associated. F is licensed in L iff (i) F is inherent in L’s lexical specification, or (ii) F receives a semantic interpretation in L.

Φ -features can be licensed in nominal, but not verbal, locations.³ This implies that any verbal ϕ -feature poses a potential problem that must be dealt with before the syntactic derivation reaches LF. The solution consists of two operations.

As a first step, the verbal ϕ -feature is identified with a nominal ϕ -feature (this process is our version of the rule of agreement; compare Brody 1997). Our analysis of identification is based on the insights of autosegmental phonology (see Williams 1976 and Goldsmith 1976, among others). In autosegmental phonology, phonemes are not unstructured feature bundles, but their features are organised in a hierarchy, just like the ϕ -features in (4). Crucially, dependent features are represented on a separate tier, and can spread from one segment to another (leading to multiple association) or be reassociated with different segments after deletion of their initial segmental host.

We borrow the idea of separate tiers for dependent features and apply it to ϕ -feature geometries. Thus, the root nodes of the various feature geometries in the syntactic representation mark positions comparable to the segmental positions in phonology. The dependent person, number and gender features are represented on separate tiers and can be associated with a ϕ -node on the ‘segmental’ tier. If agreement is identification of ϕ -features, the most transparent representation of this relation is one in which the single ϕ -feature that results from identification is associated with two ϕ -nodes, as in (11). This rule is of course a version of unification: the feature specification of one ϕ -node is unified with the feature specification of another ϕ -node. The output representation in (11) parallels the representation of harmony phenomena in autosegmental phonology.

- (11) *Φ -feature identification*
- $$\begin{array}{c} [\text{DP } \phi] \dots [\text{V } \phi] \rightarrow [\text{DP } \phi] \dots [\text{V } \phi] \\ | \qquad \qquad | \qquad \qquad \diagdown \diagup \\ \text{F} \qquad \qquad \text{F} \qquad \qquad \text{F} \end{array}$$

Subsequent to identification, a rule of disassociation must apply that removes the link between the verb’s ϕ -node and the features it hosts, since these feature are in danger of violating (10) when associated with a verb.

- (12) *Dissociation*
- $$\begin{array}{c} \phi \rightarrow \phi \\ | \\ \text{F} \quad \text{F} \end{array}$$

³ A possible exception is what Corbett (2006:131) calls ‘verbal number’. This is irrelevant to our argumentation, as verbal number does not seem to partake in agreement relations, as far as we know.

This simple formulation of the rule suffices. There is no need to add a structural description that limits dissociation to agreeing verbs. If an association line anchored in the ϕ -node of the verb is deleted while the one linked to the subject is spared, as in (13a), an LF is generated that can be interpreted. However, if the association line anchored in the ϕ -node of the subject is deleted, a representation results that violates (10), as shown in (13b).

- (13) a.
$$\begin{array}{c} [\text{DP } \phi] \dots [\text{v } \phi] \rightarrow [\text{DP } \phi] \dots [\text{v } \phi] \\ \swarrow \quad \searrow \quad \quad \quad | \\ \text{F} \quad \quad \quad \text{F} \end{array}$$
- b.
$$\begin{array}{c} *[\text{DP } \phi] \dots [\text{v } \phi] \rightarrow [\text{DP } \phi] \dots [\text{v } \phi] \\ \swarrow \quad \searrow \quad \quad \quad | \\ \text{F} \quad \quad \quad \text{F} \end{array}$$
- c.
$$\begin{array}{c} *[\text{DP } \phi] \dots [\text{v } \phi] \rightarrow [\text{DP } \phi] \dots [\text{v } \phi] \\ \swarrow \quad \searrow \quad \quad \quad \text{F} \\ \text{F} \quad \quad \quad \end{array}$$
- d.
$$\begin{array}{c} *[\text{DP } \phi] \dots [\text{v } \phi] \rightarrow [\text{DP } \phi] \dots [\text{v } \phi] \\ | \quad \quad | \quad \quad | \quad \quad | \\ \text{F} \quad \quad \text{F} \quad \quad \text{F} \quad \quad \text{F} \end{array}$$

If both association lines are removed, as in (13c), we are left with a feature F not associated with any ϕ -node. We assume that the grammar does not allow such floating features at LF, for reasons discussed below. This ban on floating features also explains why dissociation of features anchored in the verb leads to an ungrammatical result in the absence of agreement. In (13d), the feature associated with the verb is left floating after application of (12). Notice that if any of the output representations in (13b-d) were ruled in, this would make it possible to generate sentences containing inflected but non-agreeing verbs, contrary to fact.

The ban on floating features at LF is reminiscent of similar conditions in early versions of autosegmental phonology. Goldsmith (1976), for example, assumes that every autosegmental feature must be anchored to a segmental position. In subsequent work in phonology, this condition was relaxed in some instances, with authors arguing for deletion or tolerance of certain unanchored autosegmental features at the end of the derivation (see for example Halle and Vergnaud 1981 and Pulleyblank 1986; for an overview, see Odden 1995).

Be this as it may, it is a standard assumption in the syntactic literature that ϕ -features are subject to Full Interpretation: they must be interpreted in the mapping from LF to semantics (modulo our earlier remarks on inherent features). Since it is hard to see what interpretation could be assigned to a floating ϕ -feature, given that ϕ -features restrict the reference of nominal categories, floating ϕ -features violate Full Interpretation and are therefore banned at LF. Note that it would be highly undesirable to deal with this situation by allowing simple LF deletion of floating ϕ -features. If features could be freely deleted at LF, Full Interpretation would be voided of content (ruling in, amongst other

things, theta criterion violations like **John seems that he is ill*, simply by allowing deletion of the entire feature content of the offending argument).

To sum up, the rule in (12) is freely applicable. However, if applied to a subject or to a verb that does not agree with the subject, the output violates Full Interpretation.

We should also stress that both ϕ -feature identification and dissociation operate in the LF branch of the grammar. This means they cannot affect the phonological realization of the categories they apply to. If they did, no language could show overt verbal agreement.

The proposal can be illustrated using the following Dutch examples:

- (14) a. Zij lopen. b. *Hij lopen. c. *Zij loopt.
 they walk-PL *he walk-PL* *they walk-3SG*
 ‘They walk.’ ‘He walks.’ ‘They walk.’

The example in (14a) satisfies the condition on ϕ -feature licensing after application of feature identification, leading to (15a), and dissociation in the verb, leading to (15b).

- (15) a. $\begin{array}{c} [\text{DP } \phi] \dots [\text{V } \phi] \\ \diagdown \quad \diagup \\ \text{PL} \end{array}$ b. $\begin{array}{c} [\text{DP } \phi] \dots [\text{V } \phi] \\ | \\ \text{PL} \end{array}$

In (14b), the verb is specified as plural, while the subject is singular (in form and interpretation). This means that [PL] is associated with the verb, but not with the subject, as in (16a). This representation violates the condition on ϕ -feature licensing and it cannot be rescued by dissociation in the verb, as this results in a floating feature (see (16b)).

- (16) a. $\begin{array}{c} *[\text{DP } \phi] \dots [\text{V } \phi] \\ | \\ \text{PL} \end{array}$ b. $\begin{array}{c} *[\text{DP } \phi] \dots [\text{V } \phi] \\ \text{PL} \end{array}$

Finally, (14c) shows that agreement is obligatory where it is available. It is not clear from the mechanisms involved in the syntactic derivation why this should be so. The representation of (14c), given in (17), does not violate any PF conditions that we are aware of. Since it is identical to the output of dissociation as applied to (14c) (see the output in (15b)), it also does violate any conditions that hold at LF.

- (17) $\begin{array}{c} [\text{DP } \phi] \dots [\text{V } \phi] \\ | \\ \text{PL} \end{array}$

This is not just a quirk of our model. The same is the case in earlier theories of feature checking and in the current probe-goal model of agreement. In both cases, there is nothing inherent in the mechanism of syntax that rules out verbs that contain fewer ϕ -features than can be licensed by the subject.

However, we have already introduced a principle that captures the obligatoriness of agreement. This is Maximal Encoding. We have seen applications of this principle in

the mapping between semantics and LF and in the mapping between PF and phonology. It stands to reason that it should also restrict the mapping between LF and PF. In other words, a given LF must be associated with a PF that encodes as many features of this LF as possible in as many locations as possible. But (17) is not the richest PF representation compatible with an LF that contains a plural subject. That is one in which the verb has a [PL] feature licensed through feature identification, as in (15a). Therefore, by Maximal Encoding, the representation in (15a) will block (17).

That something like Maximal Encoding might regulate the relationship between LF and PF is not a new idea. Bobaljik and Wurmbrand (to appear) argue that the relationship between PF and LF is such that PF is as iconic a reflection of LF as permitted by the syntax of a given language and derive a range of data concerning the encoding of semantic scope from this. (The related idea that child language tends to minimize PF/LF discrepancy to an even greater degree can be found in Klein 1982:195f, Hyams 1986:162f, and Van Kampen 1996.)

In sum, asymmetric and symmetric theories of agreement can both capture the ‘normal’ situation: a full match between the ϕ -features of the verb and the subject. We will argue, however, that symmetric theories are better placed to deal with subset controllers.

3.2 Subset Control as Feature Spreading

As mentioned in the introduction, we use the term ‘subset controller’ to refer to situations in which the agreeing argument expresses fewer ϕ -features than the verb, as in the syntactic representation in (18).

$$(18) \quad \begin{array}{c} [\text{DP } \phi] \dots [\text{V } \phi] \\ \quad \quad \quad | \\ \quad \quad \quad \text{F} \end{array}$$

As noted, asymmetric theories of agreement cannot be reconciled with representations like (18). This is because they are based on the assumption that ϕ -features in the verb must reflect the presence of identical ϕ -features in the agreeing argument. In (18) there would be no source for the feature present in the verb, but absent in the argument. So, if there are genuine subset controllers, this disqualifies asymmetric theories of agreement.

In contrast, a theory of agreement that does not adopt a morphosyntactic asymmetry between controller and target is better placed to deal with subset controllers, provided we adopt the ϕ -feature geometry in (3) and the hypothesis that dependent ϕ -features are placed on a separate tier. The analysis can then be modelled on the way tonal spread and related phenomena are dealt with in autosegmental phonology. Such spreading involves association of a feature anchored in one segment with other segments. In some cases, the segment with which a particular tonal feature was initially associated is deleted.

Suppose the same can happen with dependent ϕ -features, in this case in the LF branch of the syntax. In addition to regular feature identification, as in (11), we can allow feature spreading, as in (19).

$$(19) \quad \Phi\text{-feature spreading}$$

The output of feature spreading is identical to the output of feature identification. Therefore, dissociation in the verb results in the same well-formed LF, the one in (20), where F has the subject as its sole location.

There is a clear affinity between φ -feature identification and φ -feature spreading on the one hand and the process of feature set unification used in theories like HPSG. In both (11) and (19), the feature specification of one φ -node is unified with the feature specification of another φ -node. We will keep φ -feature identification and φ -feature spreading separate, however, because the fact that a language has φ -feature identification (agreement) does not imply that it also allows φ -feature spreading (subset controllers).

Since operations like φ -feature spreading and dissociation take place on the LF-branch of the grammar, they cannot have an effect on the spell-out of either the verb or the subject. Although the input for interpretation, after feature spreading and dissociation, is (20), the input for spell-out is the structure in (18).

For an example of how this works, we return to the Inari Sami data in (1), assuming for the sake of the argument that our description in terms of subset controllers is correct. Two questions need to be addressed. First, how can we capture the interpretive effect of the verbal inflection in (1a)? Second, how can the distribution of subset controllers be restricted in the right way? In particular, we need to rule out subset control involving pronominal subjects.

The dual form of the verb has a feature [MIN] that is dependent on [PL] according to the feature geometry in (3). The presence of [MIN] expresses that the plural is minimal. On the view adopted here, [MIN] is represented on a separate tier. If the subject is a non-pronominal DP, its feature specification will not contain [MIN] (note that there are no dual forms for non-pronominal DPs). If there is indeed such an asymmetry between verb and subject in (1a), agreement must take the form of feature spreading, followed by feature identification, as schematized in (21b,c).

MIN

The representation in (21d), which is derived by dissociation, is input to interpretation, with the desired outcome that the subject is interpreted as a dual.

Note that the class of interpretive effects that can result from feature spreading is quite limited. This is because the interpretation acquired by a DP must be one independently admitted by its original feature specification. As explained in section 2, the absence of a dependent feature in principle allows a range of interpretations including the one encoded by the presence of that feature. This latter interpretation is only blocked (by Maximal Encoding) if a paradigmatically related form is available that expresses the relevant feature. Therefore, the interpretive effect of the agreement pattern is one of selection rather than imposition.

Consider again (1a). A subject specified as [PL] admits any plural reading. This includes a dual reading, unless there is a competing form that explicitly encodes dual. But as we have seen, this is not the case for non-pronominal DPs in Inari Sami. Thus, the effect of feature spreading is that a more specific dual reading is selected from the range of possible meanings for the subject. Notice that the opposite situation cannot exist. There cannot be a language in which subjects that are marked as dual take on a general plural interpretation through the use of a general plural agreement ending on the verb. The mechanism of feature spreading simply cannot remove the [MIN] feature from the subject. Thus, we make the following general prediction:

- (22) If an agreement mismatch has semantic effects, these must always be towards a more specific interpretation of the controller.

Our proposal also explains why subset control is not available for pronominal subjects in Inari Sami. As already mentioned, pronouns in this language differ from non-pronominal arguments in that they can express dual. Maximal Encoding therefore blocks the use of a regular plural pronoun if the interpretation in semantics is dual. Similarly, the Dutch example in (14b) is ruled out not only on an interpretation as ‘he walks’ (as discussed above), but also on an interpretation as ‘they walk’. This second reading is not available,

because Maximal Encoding dictates that the plural pronoun *zij* ‘they’ be used in that case. In general, Maximal Encoding has the following consequence:

- (23) If F^* can be generated on both the controller and the target, giving rise to interpretation I, then generation of F^* on the target only must give rise to an interpretation distinct from I. If such an interpretation is not available, the structure is ruled out as semantically incoherent.

Note that the rule of feature spreading itself does not mention that the target of spreading must not have a more fully specified counterpart. The rule can apply regardless of the paradigmatic relations in which the feature specification of the subject stands. However, in practice, Maximal Encoding rules out the use of an underspecified subject when a more highly specified subject associated with the same interpretation is available. Maximal Encoding does this entirely independently of whether the language allows for feature spreading or not.

Our proposal makes two further predictions. Our basic assumption is that the agreeing DP in relations of subset control is genuinely poorer in syntactic feature content than the verbal head it agrees with. We therefore do not expect agreement phenomena internally to the DP that reflect the relevant verbal feature:

- (24) If F^* is generated on the target only, agreement within the controller will not be for F^* .

Recall that feature spreading is an LF process. Therefore, features acquired by this rule will never have effects on surface forms. Even if further spreading within the DP were allowed, it would not undermine the generalization in (24), since such spreading would not feed PF – hence it would never be visible.

Finally, we predict that subset control must always have an interpretive effect. As already mentioned, ϕ -features need not be interpreted if inherent in a noun (see (10)). However, in the case of subset controllers, the extra feature present in the agreeing head cannot be licensed in this way (otherwise it would have been generated in the subject as well). This implies that the feature in question must have semantic import, giving rise to ‘semantic agreement’:

- (25) If F^* is generated on the target only, it must be interpreted.

This interpretive effect is of course a core observation about examples like (1a).

Our argument for a symmetric theory of agreement only holds water if genuine examples of subset controllers can be found. Proponents of the idea that the presence of ϕ -features in the verb must reflect identical ϕ -features in the noun phrase must argue that the proposed analysis of the Inari Sami data in (1) is incorrect. There are two ways in which such an argument can be developed. First, one may hypothesize that non-pronominal subjects can be specified for dual after all, but that this is not expressed morpho-phonologically, either as a result of impoverishment or because there simply is no overt affix expressing dual that can be attached to nonpronominal nouns. We will call this type of approach, schematized in (26a), a hidden-feature analysis. Second, one may

hypothesize that the apparent controller of agreement is in fact a dislocated or appositional category, while the real controller is a covert pronoun that fully agrees with the verb. We call this kind of approach, schematized in (26b,c), a hidden-controller analysis (or hidden-subject analysis in case the controller is a subject).

- (26) a. [IP [DP PL, MIN] [VP ... [V PL, MIN] ...] → (feature ‘hidden’
[IP [DP PL] [VP ... [V PL, MIN] ...] by impoverishment)
b. [DP PL]₁ [IP [_{pro} PL, MIN]₁ [VP ... [V PL, MIN] ...] (hidden subject,
dislocation of visible DP)
c. [IP [[_{pro} PL, MIN] [DP PL]] [VP ... [V PL, MIN] ...] (hidden subject,
apposition of visible DP)

To be clear, we are not confident that a hidden-feature analysis or a hidden-subject analysis of the Inari Sami data can be excluded. We have used the language only to illustrate the issue. However, there are examples of subset controllers that do not lend themselves to either alternative analysis, as we will argue in the next section.

4. SPANISH ‘UNAGREEMENT’

A strong case can be made for the existence of subset controllers in Spanish. In this language, third person subjects may appear with verbs that carry first or second person plural agreement. This phenomenon has been widely discussed and goes by the name of ‘unagreement’ (see Hurtado 1985 and Jaeggli 1986, among others). It comes in two flavours: the third person subject may be a regular R-expression, as in (27) (from Corbett 2006:132), or a quantifier, as in (28) (compare Moravcsik 1978:351; see also Rivero 2008). We refer to the two sub-types of unagreement as referential and quantificational.

- (27) a. ¡Qué desgraciad-as somos las mujer-es!
how unfortunate-F.PL be.1PL DEF.F.PL women(F)-PL
‘How unfortunate we women are!’
b. ¡Qué desgraciad-as sois las mujer-es!
how unfortunate-F.PL be.2PL DEF.F.PL women(F)-PL
‘How unfortunate you women are!’
(28) a. Cada alumno hablamos diferente.
each student speak-1PL differently
‘Each of us students speaks differently.’
b. Cada alumno habláis diferente.
each student speak-2PL differently
‘Each of you students speaks differently.’

As indicated by the translations, the interpretive effects of referential and quantificational unagreement are different. This difference corresponds to the two ways in which ϕ -features can be interpreted. They may restrict the reference of a DP, as in the case of coreferential pronouns. Alternatively, they can act as a restriction on a variable, as in the case of pronouns bound by a quantifier (see Heim and Kratzer 1998). The former mode

of interpretation is relevant to referential unagreement, while the latter applies to quantificational unagreement. We discuss them in turn.

4.1 Referential Unagreement: Analysis and Predictions

If we take the data in (27) at face value, it appears that the verbal head contains a richer constellation of ϕ -features than the subject. In particular, the verb contains a ‘surplus’ person specification ([PAR] in (27a) and [PAR, ADD] in (27b)):

(29)

<i>Subject</i> (third person plural)	<i>Agreeing head</i> (first/second person plural)
$\begin{array}{c} \phi \\ \\ PL \end{array}$	$\begin{array}{c} \phi \\ / \quad \backslash \\ PAR \quad PL \\ \\ (ADD) \end{array}$

If so, the interpretive effect in (27) comes about through feature spreading (see section 3.2). Below, we give the derivation that connects the input in (29) to the relevant LF representation (we represent [PL] in (30c) above the ϕ -tier to keep things visually clear.)

- (30) a.
$$\begin{array}{c} [DP \ \phi] \dots [V \ \phi] \\ | \quad \swarrow \quad \searrow \\ PL \quad PAR \quad PL \\ | \\ (ADD) \end{array} \quad \text{(Input)}$$
- b.
$$\begin{array}{c} [DP \ \phi] \dots [V \ \phi] \\ \swarrow \quad \searrow \quad \swarrow \quad \searrow \\ PL \quad PAR \quad PL \\ | \\ (ADD) \end{array} \quad \text{(Feature spreading of [PAR (ADD)])}$$
- c.
$$\begin{array}{c} PL \\ \swarrow \quad \searrow \\ [DP \ \phi] \dots [V \ \phi] \\ \swarrow \quad \searrow \\ PAR \\ | \\ (ADD) \end{array} \quad \text{(Feature identification)}$$
- d.
$$\begin{array}{c} [DP \ \phi] \dots [V \ \emptyset] \\ \swarrow \quad \searrow \\ PAR \quad PL \end{array} \quad \text{(Dissociation in V)}$$

Our analysis predicts that feature spreading leads to a narrowing down of the possible readings of the controller, rather than that it imposes a reading on the subject that it could not have otherwise (see (22)). This is correct for the case at hand. A third person plural DP can refer to a set of individuals that includes speaker or hearer. Thus, while *the Dutch* in (31a) is most naturally taken to exclude speaker and hearer, the possibility of coreference in (31b) and (31c) implies that *the Dutch* may have a reference comparable to ‘we the Dutch’ or ‘you the Dutch’.

- (31) a. Anyone who knows the Dutch realizes they no longer wear wooden shoes.
 b. Anyone who knows the Dutch realizes we no longer wear wooden shoes.
 c. Anyone who knows the Dutch realizes you no longer wear wooden shoes.

These interpretational possibilities follow from the feature geometry in (3)/(4). An R-expression like *the Dutch* receives a third person interpretation by default, because it lacks the features [PAR] or [ADD]. However, given that R-expressions do not have counterparts that do encode [PAR] or [ADD], Maximal Encoding does not block non-default readings in contexts like (31b,c). The interpretive effects of agreement in (27) now follow. The default third-person interpretation of the subject is no longer available after feature spreading. Instead, the subject receives a first-person or second-person plural interpretation: its reference must include the speaker or the addressee (‘we/you women’).

Conversely, our account predicts that unagreement cannot occur if the subject does not allow non-default first or second person plural interpretations to begin with. One such case is that of coordinated singular R-expressions. A singular R-expression cannot normally be used to refer to the speaker or the hearer (possibly for pragmatic reasons). This has the consequence that (32b) and (32c) are ungrammatical. Given that neither Jan nor Piet can be understood to be the speaker or the hearer, and given that these individuals make up the plural referent of the coordinate structure, it is impossible to refer back to that plural referent using a first- or second-person plural pronoun. In line with this, unagreement is impossible when the subject consists of coordinated singular R-expressions (see (33)).

- (32) a. Anyone who knows Jan and Piet realizes they no longer wear wooden shoes.
 b. *Anyone who knows Jan and Piet realizes we no longer wear wooden shoes.
 c. *Anyone who knows Jan and Piet realizes you no longer wear wooden shoes.

- (33) a. *¡Qué desgraciad-as somos Juanita y Carmen!
 how unfortunate-F.PL be.1PL Juanita and Carmen
 ‘How unfortunate we, Juanita and Carmen, are!’
 b. *¡Qué desgraciad-as sois Juanita y Carmen!
 how unfortunate-F.PL be.2PL Juanita y Carmen
 ‘How unfortunate you, Juanita and Carmen, are!’

Non-default first- or second-person plural readings can also be blocked by Maximal Encoding, namely in the case of third-person plural pronouns. These cannot receive such readings, as shown in (34), because of the availability of other pronouns that specifically encode them (first- and second-person plural pronouns); see the discussion around (22) and (23).

- (34) a. Anyone who knows them realizes they no longer wear wooden shoes.
 b. *Anyone who knows them realizes we no longer wear wooden shoes.
 c. *Anyone who knows them realizes you no longer wear wooden shoes.

Therefore, unagreement with pronominal subjects should be impossible. This is correct. If the verb carries first-person plural inflection, the only permissible pronominal subject is *nosotras* ‘we’. It is impossible to insert a less specified pronoun and use unagreement to force a first-person plural interpretation (see (35a)). The same is true, *mutatis mutandis*, for verbs inflected for second-person plural (see (35b)).

- (35) a. ¡Qué desgraciad-as somos nosotras/*yo/*ellas/*ella!
how unfortunate-F.PL be.1PL we/I/they/she
 ‘How unfortunate we (women) are!’
 b. ¡Qué desgraciad-as sois vosotras/*nosotras/*tú/*ellas/*ella!
how unfortunate-F.PL be.2PL you.PL/we.PL/you.SG/they/she
 ‘How unfortunate you (women) are!’

Thus, the data in (33) and (35) confirm the prediction in (23). A further aspect of this prediction is that it should be impossible for the subject in cases of Spanish referential unagreement to be singular. After all, plural can be expressed on R-expressions just as well as on pronouns. This blocks the examples in (36).

- (36) a. *¡Qué desgraciad-as somos la mujer!
how unfortunate-F.PL be.1PL DEF.F woman(F)
 ‘How unfortunate we women are!’
 b. *¡Qué desgraciad-as sois la mujer!
how unfortunate-F.PL be.2PL DEF.F woman(F)
 ‘How unfortunate you women are!’

The examples in (27) differ from those in (35) and (36) in that the ϕ -features expressed exclusively in the verb are not possible features of R-expressions (at least in Spanish, but probably universally; see section 4.2 below). Therefore, there are no more highly specified forms of the subjects in (27) that would have to be used because of Maximal Encoding (which would void the possibility of unagreement).

Referential unagreement has three further properties predicted by our analysis. First, as stated in (22), the semantic effects of agreement mismatches must always be towards a more specific interpretation of the controller. This is because agreement mismatches result from the interpretation of an additional feature in the verb with respect to the subject, whereas less specific interpretations require the *absence* of features. The

absence of a feature in the verb cannot cancel the interpretation of features present in the subject. Thus, a third-person verb cannot impose a third-person interpretation on a first or second person subject. Neither can first-person verbs impose a first-person interpretation on a second person subject. The data are in line with this:⁴

- (37) a. *¡Qué desgraciad-as son nosotras/vosotras!
how unfortunate-F.PL be.3PL we/you.PL
 ‘How unfortunate they (women) are!’
 b. *¡Qué desgraciad-as somos vosotras!
how unfortunate-F.PL be.1PL you.PL
 ‘How unfortunate we (women) are!’

Second, there should be no evidence internally to the subject for the presence of [PAR] or [PAR, ADD]. These features are absent in the syntactic representation of the DP, other than at LF. Since feature spreading is an LF process, it will not have PF effects, and it should therefore not affect the morpho-phonology of the DP. Indeed, although Spanish has DP-internal gender and number agreement, material within the controller is never inflected for person, not even in cases of unagreement.⁵

Third, referential unagreement should be ‘semantic agreement’: it must always involve a feature that is interpreted, rather than licensed by being inherent in the subject (see (10) and (25)). This is the case in Spanish. The language has many nouns with an inherent gender feature, as well as some *plurale tantum* nouns, which have an inherent number feature. However, there are no nouns that trigger first or second person agreement obligatorily. This means that apparently no noun has an inherent specification for first or second person.

4.2 A First Alternative: Hidden Features

We now turn to alternative analyses of referential unagreement that maintain the standard view of agreement as syntactically asymmetric (where verbal ϕ -features depend on matching features in the subject). An obvious asymmetric alternative of the Spanish data would be to assume that R-expressions can bear person features and hence trigger first or second person agreement. Of course, the analysis would have to assume that these features are never realised on the R-expression itself, possibly as a result of a rule of impoverishment. This is what we have called a hidden-feature analysis in the discussion around (26) above.

There are several facts that are difficult to reconcile with this type of analysis. First, as we have seen, there are no R-expressions in Spanish that have inherent person features, while there are R-expressions that have inherent number and/or gender features. In a hidden-feature analysis, it remains mysterious why there should be such an

⁴ The contrast between (27) and (37) follows from the feature geometry in (3)/(4). In theories that treat first, second and third person as independent features, it remains unclear why first- and second-person features in the verb can overrule a third-person feature on the subject, while a third-person feature in the verb cannot overrule first- and second-person features in the subject.

⁵ Note that even if there is further feature spreading within the DP at LF, this cannot have interpretive effects either, since these features are not interpretable on the other elements within the extended projection of the noun. The relevant ϕ -features would therefore be dissociated subsequently in all elements except DP itself in order to guarantee compliance with (10).

asymmetry between the features for which an R-expression can be lexically specified according to that analysis. The asymmetry follows from our analysis, as R-expressions are not lexically specified for person features at all.

Second, although it might be a coincidence that the features postulated to characterize R-expressions under the hidden-feature analysis are not overtly realized in Spanish, it would require a conspiracy to account for the apparent universal absence of a spell-out of such features on R-expressions. As far as we know, this universal absence is a fact, however. We do not know of any language in which *we linguists* is expressed using a first person inflectional ending attached to the noun *linguists*, or in which a noun like *linguists* has an irregular first person form.

Third, if it is the case that R-expressions in fact bear person features, it is hard to explain why some languages allow unagreement, while other languages do not. This question is of course also awaiting a satisfactory answer in our approach to the problem, but at least we can identify a difference in the grammar of Spanish and the grammar of, say, Italian: only Spanish has a spreading rule for person features. The hidden-feature analysis appears to lack anything that allows one to even approach the question of cross-linguistic variation in this respect.

Fourth, in Mancini et al.'s (in press) ERP study, a three-way difference was found between regular agreement, unagreement and agreement errors. That is to say, the neurophysiological processing correlates of unagreement could not be identified as errors (reflecting the fact that the relevant examples are grammatical), but could also not be unified with the processing correlates of regular agreement. These results have a natural interpretation in our theory, where unagreement involves an additional grammatical process (namely feature spreading), compared to regular agreement. Indeed Mancini et al. propose a comparable analysis of unagreement, which they dub 'reverse agreement'. But such results seem much harder to reconcile with a hidden-feature analysis, which after all aims to equate unagreement to regular agreement.⁶

These considerations lead us to conclude that a hidden-feature analysis is not the way to approach the Spanish data. Notice that this conclusion is based on empirical findings; it leaves open the question what it is in the theory of grammar that rules out such an approach. In other words, why should it be that R-expressions do not bear person features? This is an important question, but one that current theories of grammar have little to say about. Given that R-expressions can refer to groups that may or may not include speaker and hearer (see (31)), it is unlikely that the answer should be semantic. Rather, one would imagine that there is something about the morpho-syntax of R-expressions that blocks the presence of person features. This does of course fit the idea of feature spreading, as that rule enriches the semantic representation of R-expressions at the LF interface.

4.3 A Second Alternative: Dislocation

⁶ Notice that very many instances of regular agreement presumably involve impoverishment, namely the majority of cases involving syncretism. Hence, non-spell-out of features as such cannot be the factor that triggers the effect found with unagreement (left posterior negativity followed by a more central negativity; absence of the P600 effect). One would have to argue it is specifically impoverishment in, or underspecification of, the controller that is associated with this effect. This could be tested using English examples like *you should be ashamed of yourself/yourselfs*, where the form of the reflexive disambiguates *you* as singular or plural.

The second type of alternative analysis that maintains a syntactically asymmetric theory of agreement is what we have called the hidden-controller analysis. In such an analysis, the apparent subject is assumed not to be the subject at all. Instead, there is a null subject that is fully specified for all ϕ -features realized on the verb. What seems to be the overt subject is either claimed to be in dislocation (or some other type of non-argument position; compare Hurtado 1985)), or taken to be adjoined to the null pronominal subject, parallel to apposition in examples like *we the boys* (compare Suñer 1988).⁷

The hidden-subject analysis seems more promising than the hidden-feature analysis for the Spanish data, given that this language allows empty subjects in general. Moreover, as argued by Barbosa (1996, 2000), Alexiadou and Anagnostopoulou (1998) and Ordoñez and Treviño (1999), apparent preverbal subjects in pro-drop languages may be in dislocation, rather than occupy a true argument position. The subject position could then be occupied by a silent pronoun. Thus, the kind of structure required by the hidden-subject analysis has been argued to be available independently.

Nonetheless, the success of a hidden-subject analysis depends on the extent to which the properties of unagreement are compatible with the syntax of null subjects and the syntax of dislocation or apposition. Our assessment is that it is not. We begin by discussing the analysis of unagreement in which the apparent subject occupies a dislocated position, while clause-internally the subject position is occupied by a null pronoun that agrees with the verb in the standard way (compare Olarrea 1996).

An immediate problem for this analysis is that the distribution of unagreeing subjects matches that of regular subjects. For example, Torrego (1996) points out that they need not appear in a peripheral position, but may also occur sandwiched between the verb and an object (see (38); translation slightly adjusted). Even theories that take preverbal subjects to be in dislocation recognize this position as a regular argument position.

- (38) Firmamos los lingüistas la carta
signed-IPL the linguists the letter
 ‘We linguists signed the letter.’

Moreover, the dislocation analysis only shifts the problem, because in other dislocation structures the sentence-internal pronoun and the associated dislocated category must have matching ϕ -features. This is of course best illustrated in languages that lack pro drop. In Dutch, for example, there is a strict requirement of feature matching between pronoun and dislocated element:

- (39) a. De jongens, ze zijn aan elkaar gewaagd.
the boys, they are to each other weighed
 ‘The boys, they are well matched.’

⁷ Some proposals treat verbal inflection as the subject (see Ackema et al. 2006 and references mentioned there). That is, the agreement morpheme is seen as the category that receives the verb’s external theta-role. We abstract away from such proposals, because they do not bear on the agreement *relation*. If anything, they fit in better with symmetric theories, as pro drop can then be seen as the most extreme case of feature spreading: the agreeing category receives its entire feature content from the verb. An asymmetric theory must assume six distinct *pros*, each with a full feature specification.

- b. *De jongens, we zijn aan elkaar gewaagd.
the boys, we are to each other weighed
- c. *De jongens, jullie zijn aan elkaar gewaagd
the boys, you are to each other weighed

Assuming that apparent subjects in unagreement structures are in dislocation thus begs the question why no feature matching is necessary between the hypothesized null subject and the dislocated category associated with it.

One way out would be to assume that Spanish has an additional type of construction in which a null subject pronoun is doubled by a full DP. This hypothetical construction, which we will call ‘low dislocation’, must be distinguished from both hanging-topic left dislocation and clitic left dislocation in two respects. First, it should not require feature matching. Second, the category associated with the null pronoun should appear clause-internally, rather than in a peripheral position.

This may appear ad hoc, but clitic doubling constructions could provide some independent evidence for low dislocation. Clitic doubling is attested with certain types of objects in Spanish. In these cases, the object behaves as if it is clause-internal. Moreover, mismatches between the ϕ -features of the clitic and its double can be observed, parallel to those found with subjects and unagreeing verbs:

- (40) a. Nos denunciaron a las mujeres (Hurtado 1985:201)
us denounced-3PL to the women
‘They denounced us women.’
- b. Vos denunciaron a las mujeres
you.ACC denounced-3PL to the women
‘They denounced you women.’

Thus, clitic doubling and unagreement with subjects could be grouped together as cases of low dislocation: both structures would involve a pronominal argument (silent in the case of subjects), doubled by a DP in an A’-position internally to the clause (compare the analysis of clitic doubling in Kayne 1975).

Although there is a parallel between the two structures, we do not think that an analysis in terms of low dislocation is feasible. Clitic doubling is usually no longer taken to involve a pronominal argument coindexed with a DP in A’-position, partly because the relevant DP behaves like a regular object in certain respects. Instead, it is regarded as a case of agreement between a DP-argument and a functional category (the clitic) generated either within the verbal extended projection (Sportiche 1998) or within the nominal extended projection (Uriagereka 1995, Papangeli 2000 and Kayne 2002).

However, if clitic doubling involves agreement rather than low dislocation, all we can conclude from the Spanish data in (40) is that the problem of unagreement extends to those cases of object agreement (or DP-internal agreement if the clitic starts out DP-internally) that are attested in Spanish. The problem remains one of lack of feature matching between otherwise agreeing categories.

4.4 A Third Alternative: Apposition

There is a second incarnation of the hidden-subject analysis for referential unagreement. The apparent subject could be assumed to be in apposition with a null pronominal subject specified for all features relevant to verbal agreement (Suñer 1988 suggests a related analysis for quantificational unagreement; see below). Thus, (41a) would have a structure much like its English translation in (41b), except that the pronoun is not spelled out.

- (41) a. ¡Qué desgraciad-as somos [DP [DP *pro*] [DP las mujer-es]]!
how unfortunate-F.PL be.1PL DEF.F.PL women(F)-PL
 b. How unfortunate [DP [DP *we*] [DP women]] are!

Notice that overt apposition in Spanish is allowed in the relevant contexts. Therefore the examples in (27) could be related to the structures in (42) through pro drop.

- (42) a. ¡Qué desgraciad-as somos nosotras las mujer-es!
how unfortunate-F.PL be.1PL we DEF.F.PL women(F)-PL
 ‘How unfortunate we women are!’
 b. ¡Qué desgraciad-as sois vosotras las mujer-es!
how unfortunate-F.PL be.2PL DEF.F.PL women(F)-PL
 ‘How unfortunate you women are!’

This version of the hidden-subject analysis has two obvious advantages. First, the semantics of referential unagreement is similar, if not identical, to that of close apposition. Second, apposition does not require feature matching for person between the argument and the category adjoined to it. Examples like *we the people*, *you boys*, and so on, are unobjectionable. This means that the analysis may provide a genuine way of reconciling referential unagreement with standard views of agreement.

Nonetheless, the analysis faces a serious difficulty, which has to do with the distinction between weak and strong pronouns (see Cardinaletti and Starke 1996 for an overview). Weak pronouns differ from strong pronouns in their syntactic distribution, their prosody and their interpretation. The distinction is relevant here because apposition is only possible with strong pronouns. In Dutch, for example, there is a formal opposition between the strong first-person plural pronoun *wij* and its weak counterpart *we* (see (43a)). That *we* is indeed a weak pronoun is shown by the fact that it cannot be coordinated, in contrast to *wij*, as illustrated in (43b). As it turns out, only the strong form can be inserted in appositional structures like (43c). Similar data exist in English (with *you/ye*). The pattern repeats itself for the German first-person plural pronoun. The strong form *wir* ‘we’ has a weak counterpart *ma* (44a,b) in certain varieties. Only the former can host an appositional DP (see (44c)).

- (43) a. Wij/we gaan dan uit eten.
we go then out eating
 ‘We will go for dinner then.’
 b. Wij/*we en de studenten gaan dan uit eten.
we and the students go then out eating
 ‘We and the students will go for dinner then.’
 c. Wij/*we studenten gaan dan uit eten.

we students go then out eating
 ‘We students go for dinner then.’

- (44) a. Ich hoffe daß wir/ma uns dann amüsieren können. (Klaus Abels, p.c.)
I hope that we us then amuse can
 ‘I hope that we can amuse ourselves then.’
 b. Ich hoffe daß wir/*ma und die Studenten uns dann amüsieren können.
I hope that we and the students us then amuse can
 ‘I hope that we and the students can amuse ourselves then.’
 c. Ich hoffe daß wir/*ma Studenten uns dann amüsieren können.
I hope that we students us then amuse can
 ‘I hope that we students can amuse ourselves then.’

Cardinaletti and Starke (1996) argue in some detail that null pronouns must be classified as weak. Indications of this are that, in contrast to strong pronouns, they can have non-human referents, cannot be coordinated and can be used as subjects of weather verbs. We illustrate these properties for Spanish below:

- (45) a. *pro* es bonito.
pro is pretty
 ‘He/it is pretty’
 b. **pro* y el de Juan son bellos.
pro and that of John are beautiful
 ‘This one and that of John are beautiful.’
 c. *pro* hace mucho frío.
pro makes much cold
 ‘It is very cold.’

But if silent pronouns must be classified as weak, and if weak pronouns cannot appear in appositional structures, then the Spanish examples in (27) cannot involve hidden apposition.

4.5 Quantificational Unagreement

We now turn to quantificational unagreement, as exemplified in (28) above. In order to understand this phenomenon, we will first need to explore the interpretive effects of ϕ -features in a little more detail. One obvious interpretive contribution that ϕ -features can make is in restricting the reference of a DP. The best example of this is cross-sentential coreference between a DP and a pronoun. Thus, the gender features of the pronoun in (46a) restrict its reference in such a way that it must be coreferent with *John* rather than *Mary*. The gender features of the pronoun in (46b) give the opposite result.

- (46) a. John met Mary. He was wearing a helmet.
 b. John met Mary. She was wearing a helmet.

This interpretive contribution is relevant to the analysis of referential unagreement, as discussed in the previous subsection.

However, ϕ -features can also appear on elements that are not referential. This is what we find in the case of pronouns bound by a quantifier. The possessive pronoun in *every boy loves his mother* does not refer to any particular boy, and yet its ϕ -features must make a semantic contribution, given the ungrammaticality of **every boy loves her mother*. Heim and Kratzer 1998 argue that in this case the ϕ -features act as a restriction on the variable introduced by the pronoun. Since the pronoun is bound by the quantifier, and the quantifier operates on the set of boys, there is an interpretational clash in the latter example, as x is simultaneously presupposed to be a boy and to be feminine:

- (47) Every x , x = a boy, x loves x 's mother, x = female

The interpretation of ϕ -features as restrictions on variables is relevant to quantificational unagreement, as we will now explain.

Our analysis of examples like (48) (where (48a) repeats (28)) runs largely parallel to that of referential unagreement. The only difference comes from the interpretation of ϕ -features in the context of quantification. The ϕ -features of the verb in (48a,b) cannot narrow down the reference of the subject, simply because the subject is not referential. Instead, ϕ -features in this case serve to determine the interpretation of the set on which the quantifier operates.

- (48) a. Cada alumno hablamos/habláis diferente.
each student speak-1PL/speak-2PL differently
 'Each of us students speaks differently.'
 b. La mayoría podemos/podéis conducir con una mano.
the majority can-1PL/can-2PL drive with one hand
 'Most of us/you can drive with one hand.'

In general, the ϕ -features of a QP originate in the NP complement of the Q-head. In other words, although it is QP that agrees with the verb, it is NP in which the ϕ -features are interpreted.⁸ Specifically, the ϕ -features present in the subject contribute to the content of the restrictor of the quantifier. As before, unagreement is an instance of feature spreading. This results, starting from (49), in the representation in (50a) (for reasons of readability, in (50) we abstract away from the feature percolation from NP to QP and we place number and person features on opposite sides of the ϕ -tier; this has no theoretical significance). After dissociation in the verb takes place, we are left with (50b), with the result that the features generated in the verb pertain to the interpretation of the subject: the operator now applies to a plural set including speaker or addressee.

(49)

Subject	Agreeing head
<pre> QP [ϕ_1] / \ Q NP [ϕ_1] </pre>	<pre> ϕ_1 / \ PAR PL </pre>

⁸ The operation of feature percolation probably also takes place in regular referential DPs, which inherit their ϕ -features from the complement to the determiner. However, in this case the interpretation of ϕ -features affects the reference of the DP.

	<div style="text-align: center;"> (ADD)</div>
--	--

(50) a. (Feature spreading)

b. (Dissociation in V)

This yields the right interpretation for examples like (48). For instance, (48a) will be assigned the following representation (where S is the set of students):

(51) Each x , $x \in S$ & $|S| > 1$ & speaker/hearer $\in S$, x speaks differently.

Note that in the examples in (48) unagreement does not only involve person features, but also a number feature. As discussed, this situation is blocked by Maximal Encoding in the case of referential unagreement, given that referential DPs have a plural counterpart. Quantifiers like *cada* ‘each’, however, lack a plural counterpart, so that no blocking effects can obtain. For those quantificational expressions that have both a singular and plural form, however, blocking effects *are* expected, on a par with those found with referential unagreement. In other words, for such QPs unagreement should be impossible in the singular, but permitted (for person features) in the plural. This appears to be correct, as illustrated for *algún* ‘some’ and *todo* ‘every’ below:

- (52) a. *Algún paciente hemos/habéis llamado a la doctora.
some patient have-1PL/have-2PL called to the doctor
‘Some of us/you patients have called the doctor.’
b. Algunos pacientes hemos/habéis llamado a la doctora.
some patient have-1PL/have-2PL called to the doctor
‘Some of us/you patients have called the doctor.’
c. *Todo niño creemos/creéis en los Reyes Magos.
every kid believe-1PL/believe-2PL in the Reyes Magos
‘All of us/you kids believe in the Magi.’
d. Todos niños creemos/creéis en los Reyes Magos.
every kid believe-1PL/believe-2PL in the Reyes Magos

‘All of us/you kids believe in the Magi.’

Of course, quantifiers that *only* have a plural form (such as *ambos* ‘both’ or *tres* ‘three’) also allow for unagreement in person features. In other words, the generalization that emerges is that quantificational unagreement is allowed with plural quantifiers, and with singular quantifiers as long as they do not have a plural counterpart.

If our analysis is correct, we would expect our predictions about the subject in referential unagreement (see section 4.1) to carry over to the complement of the Q-head in quantificational unagreement. We cannot test this for the full range of DPs discussed above, because not all DPs are possible restrictors of quantifiers. However, there is a sharp contrast between NP restrictors and pronominal restrictors. When QPs have non-pronominal restrictors, this does not interfere with unagreement, whether the restrictor is a PP (see (53a)) or an NP (see examples in (48a) and (52) above). But just as referential unagreement is blocked by Maximal Encoding if the subject is a pronoun, quantificational unagreement is blocked if the quantified subject has an overt pronominal restrictor (see (53b/c); note that pronominal restrictors must be contained in a PP).

- (53) a. Todos de los que dejaban lo vimos/visteis.
all of those that left him saw-1PL/saw-2PL
‘All of us/you that left saw him.’
b. Todos de nosotros/vosotros lo vimos/visteis.
all of us/you.PL him saw-1PL/saw-2PL
‘All of us/you saw him.’
c. Todos de ellos lo vieron/*vimos/*visteis.
all of them him saw-3PL/saw-1PL/saw-2PL
‘All of them/us/you saw him.’

As expected, these data parallel interpretive restrictions observed in English:

- (54) a. All of the boys could have done it, because they were all in the schoolyard.
b. All of the boys could have done it, because we were all in the schoolyard.
c. All of the boys could have done it, because you were all in the schoolyard.
(55) a. All of them could have done it, because they were all in the schoolyard.
b. *All of them could have done it, because we were all in the schoolyard.
c. *All of them could have done it, because you were all in the schoolyard.

4.6 Alternatives: Dislocation, ‘High’ Apposition and ‘Low’ Apposition

We now turn to alternative analyses of quantificational unagreement. As with referential unagreement, a hidden-feature analysis is a non-starter. Person is never expressed on quantifiers, making it unlikely that these elements can carry either [PAR] or [PAR, ADD]. When a quantifier does not take an overt complement, as in (56), there might be a silent pronominal complement with the relevant person specification (this is essentially the analysis in Suñer 1988). But this analysis does not carry over to cases of unagreement in

which the quantifier takes an overt non-pronominal complement (as in (52b,d)), as non-pronominal NPs do not carry first or second person features either.⁹

- (56) Todos hemos ido de excursión alguna vez.
all have-1PL gone of excursion some time
 ‘Every one of us has gone on a trip some time.’

A hidden-subject analysis based on dislocation of the apparent subject is also unlikely to work for quantificational unagreement, quite apart from the fact that the problems observed in section 4.3 carry over. There is cross-linguistic evidence that at least certain quantifiers resist dislocation. We illustrate this below with data from Dutch, but the effect seems universal. For example, Jelinek 1984, 2006 and Baker 1996, 2006 suggest that in polysynthetic languages all NPs are in dislocation and derive from this that such languages lack NP-quantifiers.

- (57) a. *Niemand₁, die₁ kwam *t*₁ gisteren thuis.
no-one, that came yesterday home
 ‘No-one came home yesterday.’
 b. *Niemand₁, die₁ heb ik *t*₁ gezien.
no-one, that have I seen
 ‘I saw no-one.’

The same restriction seems to hold in Romance. Rizzi (1986) shows that in Italian negative quantifiers, among others, resist clitic left dislocation. Arregi (2003) confirms that in Spanish, too, certain quantifiers cannot appear in dislocation. There is a sharp contrast between the examples in (58a) and (58b), but as (59) shows, negative quantifiers *can* function as subset controllers:¹⁰

- (58) a. Juan₁, nosotros lo₁ vimos.
Juan, we him saw-1PL
 ‘As for Juan, we saw him.’
 b. *Nadie₁, nosotros lo₁ vimos.
No-one, we him saw-1PL

- (59) a. ?Nadie fuimos a la boda de Fran.

⁹ It is highly unlikely that in such cases there is a second, hidden, complement specified for these features, given that quantifiers typically take only a single restrictor. Indeed, adding an overt pronominal restrictor to a quantificational phrase like *todos de los niños* ‘all of the boys’ or *todos niños* is ungrammatical (**todos (de los) niños de nosotros* ‘all (of the) boys of us’). It is possible to add a PP-adjunct, as in *todos de los niños entre nosotros* ‘all of the boys among us’. Hypothesizing that the adjunct can remain silent, however, cannot provide a basis for the analysis of examples like (52b), as it would yield an incorrect interpretation: ‘some of the patients among us/you have called the doctor’, rather than ‘some of us/you patients have called the doctor’.

¹⁰ There is some variation in judgments for (59a), which might in this case indicate variation in idiolects. Molina’s (2010) findings are as follows. Out of fifteen speakers, five found the example perfect (scoring it 5 on an ordinal scale from 1 to 5). Three speakers scored it as good (4). Six speakers considered it marginal (3). One speaker completely rejected it (scoring it 1).

- nobody went-1PL to the wedding of Fran*
 ‘Nobody of us went to Fran’s wedding.’
- b. *Niguno pensamos como tú.*
None think-1PL like you
 ‘None of us thinks like you.’

Given the ungrammaticality of (58b), it is unlikely that Spanish negative quantifiers can appear in dislocation when related to null subjects. But if negative quantifiers resist dislocation, dislocation cannot be the phenomenon responsible for quantificational unagreement in examples like (59). It is of course still possible to analyze the phenomenon in terms of a special kind of ‘low’ dislocation, but as argued in section 4.3, this does not provide a solution either.

An analysis of quantificational unagreement in terms of apposition of the quantifier to a silent pronoun is equally problematic, simply because quantificational categories cannot function as appositional phrases. So, even if we ignore the fact that weak pronominal forms do not allow apposition (see above), the analysis must be rejected in view of examples like (60). This example shows that the universally quantified phrase *todos niños* cannot even stand in apposition to a strong pronoun. (The same conclusion can be drawn with respect to other quantifiers that allow unagreement.)

- (60) a. *[Nosotros [todos niños]] creemos en los Reyes Magos.
We all boys believe-1PL in the Reyes Magos
 ‘All of us boys believe in the Magi.’
- b. *[Vosotros [todos niños]] creéis en los Reyes Magos.
You_{PL} all boys believe-2PL in the Reyes Magos
 ‘All of you boys believe in the Magi.’

A more plausible alternative would be to say that it is not the entire QP that stands in apposition to a silent pronoun, but just its complement (on a par with *each of us boys*, *all of you boys*, etc.). Of course, this suggestion faces the by now familiar difficulty that weak pronominal forms do not permit apposition to begin with. But irrespective of that, it appears to be insufficiently general. For example, apposition with an overt pronoun is not possible with the NP complement of *cada* ‘each’ (see (61)). However, we have already seen that *cada*+NP partakes in unagreement (see (28/48a)).

- (61) a. *[Cada [nosotros alumno(s)]] hablamos diferente.
each we student speak-1PL differently
 ‘Each of us students speaks differently.’
- b. *[Cada [vosotros alumno(s)]] habláis diferente.
each you student speak-2PL differently
 ‘Each of you students speaks differently.’

6. CONCLUSION

In this paper, we have investigated to what extent there is an asymmetry in agreement relations, beyond the semantic asymmetry that limits the interpretation of ϕ -features to

nominals. According to a number of theories, agreement is also asymmetric morpho-syntactically, in that features of the target are in some way dependent on the presence of identical features in the controlling argument.

We have argued that the data discussed in this paper can be analysed more successfully if the presence of the features in the verb does not depend on the presence of features in the subject. The two ϕ -feature sets are generated independently. It is the role of the agreement process to associate target and controller, in a way that is familiar from theories based on unification.

Our case was based on the existence of subset control: structures in which the controller is specified for fewer features than the target. The existence of such structures would make it impossible that features in the target always depend on identical features in the controller. Proponents of asymmetric theories of agreement have dealt with apparent examples of subset control by arguing for a hidden feature in the controller, or a fully specified hidden controller. We have shown that such analyses must be abandoned at least in the case of Spanish unagreement. To account for the full range of data, they would have to rely on assumptions that are difficult to square with what is known about the syntax of subjects and related issues.

The alternative we propose is to allow that features are generated freely in target and controller, as long as their distribution at the interface with semantics is such that all ϕ -features can be licensed. We have argued that this is established by two distinct processes. One is feature identification. The other is a process of feature spreading akin to the spreading of autosegmental features in phonology.

It is an open question how widespread the distribution of subset control is. Although we cannot go into this here, there are many other instances of agreement that we think are most fruitfully analysed in these terms. A well-known case is British English plural agreement with collective nouns. Other cases include the Tamil and Talisk Russian data discussed in Corbett 2006:161,209. It is possible that subset controllers are actually quite widespread, given that the phenomenon of agreement-related pro-drop can be seen as the most extreme instance of subset control: all ϕ -features that determine the interpretation of the controller are present in the target only.

REFERENCES

- Ackema, Peter, and Ad Neeleman. 2003. Context-Sensitive Spell-Out. *Natural Language and Linguistic Theory* 21:681-735.
- Ackema, Peter, and Ad Neeleman. 2004. *Beyond Morphology: Interface Conditions on Word Formation*. Oxford: Oxford University Press.
- Ackema, Peter, and Ad Neeleman. 2011. Person Syncretism, Agreement Alternations and Feature Geometry. Ms. University of Edinburgh/UCL.
- Ackema, Peter, Patrick Brandt, Maaïke Schoorlemmer and Fred Weerman (eds.). 2006. *Arguments and Agreement*. Oxford: Oxford University Press.
- Alexiadou, Artemis, and Elena Anagnostopoulou. 1998. Parametrizing Agr: Word Order, V-Movement and EPP-Checking. *Natural Language and Linguistic Theory* 16: 491-539.

- Arregi, Karlos. 2003. Clitic Left Dislocation is Contrastive Topicalization. In Proceedings of the 26th Annual Penn Linguistic Colloquium. *UPenn Working Papers in Linguistics* 9.1:31-44.
- Baker, Mark. 1996. *The Polysynthesis Parameter*. Oxford: Oxford University Press.
- Baker, Mark. 2006. On Zero Agreement and Polysynthesis. In: Ackema et al. 2006, pp. 289-320.
- Barbosa, Pilar. 1995. *Null Subjects*. PhD dissertation. MIT
- Barbosa, Pilar. 2000. Clitics: A Window into the Null Subject Property. In J. Costa (ed.) *Portuguese Syntax*. Oxford: Oxford University Press. pp. 31-93.
- Blutner, Reinhard. 2000. Some Aspects of Optimality in Natural language Interpretation. *Journal of Semantics* 17:189-216.
- Bobaljik, Jonathan. 2007. Where's Phi? Agreement as a Post Syntactic Operation. In Daniel Harbour, David Adger and Susana Béjar (eds.) *Phi-Theory: Phi features across Interfaces and Modules*. Oxford: Oxford University Press. pp. 295-328.
- Bobaljik, Jonathan, and Susi Wurmbrand. To appear. Word Order and Scope: Transparent Interfaces and the $\frac{3}{4}$ Signature. *Linguistic Inquiry*.
- Bonet, Eulalia. 1991. *Morphology after Syntax: Pronominal Clitics in Romance*. PhD dissertation. MIT.
- Bonet, Eulalia. 1995. Feature Structure of Romance Clitics. *Natural Language and Linguistic Theory* 13: 607-647.
- Brody, Michael. 1997. Perfect Chains. In Liliane Haegeman (ed.) *Elements of grammar*. Dordrecht: Kluwer. pp. 139-167.
- Cardinaletti, Anna, and Michal Starke. 1996. The Typology of Structural Deficiency. In Henk van Riemsdijk (Ed.) *Clitics in the Languages of Europe*. Berlin: Mouton de Gruyter. pp. 145-233.
- Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge: MIT Press.
- Chomsky, Noam. 1993. A Minimalist Program for Linguistic Theory. In Kenneth Hale and Samuel J. Keyser (eds.) *The View from Building 20*. Cambridge, MA: MIT Press. pp. 1-52.
- Chomsky, Noam. 2000. Minimalist Enquiries: The Framework. In R. Martin, D. Michaels and J. Uriagereka (eds.) *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*. Cambridge: MIT Press. pp. 1-59.
- Chomsky, Noam. 2001. Derivation by Phase. In M. Kenstowicz (ed.) *Ken Hale. A Life in Language*. Cambridge: MIT Press. pp 1-52.
- Corbett, Greville. 2006. *Agreement*. Cambridge: Cambridge University Press.
- De Hoop, Helen, and Andrej Malchukov. 2007. On Fluid Differential Case Marking: A Bidirectional OT Approach. *Lingua* 117:1636-1656.
- Fanselow, Gisbert. 2004. The MLC and Derivational Economy. In A. Stepanov, G. Fanselow and R. Vogel (eds.) *Minimality Effects in Syntax*. Berlin: Mouton de Gruyter. pp. 73-123.
- Frampton, John and Sam Gutmann. 2000. Agreement is Feature Sharing. Ms. Northeastern University.
- Gazdar, Gerald and Geoffrey Pullum. 1982. *Generalized Phrase Structure Grammar: A Theoretical Synopsis*. Bloomington: Indiana University Linguistics Club.
- Golan, Yael. 1993. Node-Crossing Economy, Superiority and D-Linking. Ms. Tel Aviv University.

- Goldsmith, John. 1976. *Autosegmental Phonology*. PhD dissertation. MIT.
- Greenberg, Joseph. 1963. Some Universals of Grammar with Particular Reference to the Order of Meaningful Elements. In J. Greenberg (ed.) *Universals of Language*. Cambridge: MIT Press. pp. 73-113.
- Halle, Morris. 1997. Distributed Morphology: Impoverishment and Fission. In: B. Bruening, Y. Kang and M. McGinnis (eds.) *MIT Working Papers in Linguistics 30: Papers at the Interface*. MIT. pp. 425-449.
- Halle, Morris, and Jean-Roger Vergnaud. 1981. Harmony Processes. In M. Klein and W. Levelt (eds.) *Crossing the Boundaries in Linguistics*. Dordrecht: Reidel. pp. 1-22.
- Harley, Heidi, and Elizabeth Ritter. 2002. Person and Number in Pronouns: A Feature-Geometric Analysis. *Language* 78: 482-526.
- Heim, Irene, and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.
- Hurtado, Alfredo. 1985. The Unagreement Hypothesis. In L. King and C. Maley (eds.) *Selected Papers from the Thirteenth Linguistic Symposium on Romance Languages*. Amsterdam: John Benjamins. pp. 187-211.
- Hyams, Nina. 1986. *Language Acquisition and the Theory of Parameters*. Dordrecht: Reidel.
- Jackendoff, Ray. 1997. *The Architecture of the Language Faculty*. Cambridge: MIT Press.
- Jäger, Gerhard. 2002. Some Notes on the Formal Properties of Bidirectional Optimality Theory. *Journal of Logic, Language and Information* 11:427-451.
- Jaeggli, Osvaldo. 1986. Three Issues in the Theory of Clitics: Case, Doubled NPs, and Extraction. In H. Borer (ed.) *The Syntax of Pronominal Clitics*. Orlando: Academic Press. pp. 15-42.
- Jelinek, Eloise. 1984. Empty Categories, Case and Configurationality. *Natural Language and Linguistic Theory* 2: 39-76.
- Jelinek, Eloise. 2006. The Pronominal Argument Parameter. In Ackema et al. 2006, pp. 261-288
- Kato, Mary, and Eduardo Raposo. 2005. Topicalization in European and Brazilian Portuguese. In José Camacho (ed.) *Romance Linguistics 2006: Selected Papers from the 36th Linguistic Symposium on Romance Languages (LSRL), New Brunswick, March-April 2006*. Amsterdam: John Benjamins. pp. 205-218.
- Kayne, Richard. 1975. *French Syntax: The Transformational Cycle*. Cambridge, MA: MIT Press.
- Kayne, Richard. 2002. Pronouns and their Antecedents. In Samuel Epstein and Daniel Seely (eds.) *Derivation and Explanation in the Minimalist Program*. Oxford: Blackwell. pp. 133-166.
- Kerstens, Johan. 1993. *The Syntax of Person, Number and Gender*. Berlin: Mouton de Gruyter.
- Kiparsky, Paul. 1973. "Elsewhere" in Phonology. In Stephen Anderson and Paul Kiparsky (eds.) *A Festschrift for Morris Halle*. New York: Holt, Rinehart and Winston. pp. 93-106.
- Kitahara, Hisatsugu. 1994. *Target α* . PhD Dissertation. Harvard University.

- Klein, Sharon. 1982. *Syntactic Theory and the Developing Grammar: Reestablishing the Relationship between Linguistic Theory and Data from Language Acquisition*. PhD dissertation. UCLA.
- Marr, David. 1982. *Vision*. New York: W.H. Freeman & Co.
- Mancini, Simona, Nicola Molinaro, Luigi Rizzi and Manuel Carreiras. In press. When Persons Disagree: An ERP Study of Unagreement in Spanish. *Psychophysiology*.
- Moravcsik, Edith. 1978. *Agreement*. In J. Greenberg (ed.) *Universals of Human Language; Volume 4 (Syntax)*. Stanford: Stanford University Press. pp. 331-374.
- Neeleman, Ad, and Kriszta Szendroi. 2007. Radical Pro Drop and the Morphology of Pronouns. *Linguistic Inquiry* 38: 671-714.
- Neeleman, Ad, and Hans van de Koot. 2010. Theoretical Validity and Psychological Reality of the Grammatical Code. In M. Everaert, T. Lentz, H. de Mulder, Ø. Nilsen and A. Zondervan (eds.) *The Linguistics Enterprise: From Knowledge of Language to Knowledge in Linguistics*. Amsterdam: John Benjamins. pp. 183-212.
- Odden, David. 1995. Tone: African Languages. In J. Goldsmith (ed) *The handbook of Phonological Theory*. Oxford: Blackwell. pp. 444-475.
- Olarrea, Antxon. 1996. *Pre and Postverbal Subjects in Spanish: A Minimalist Account*. PhD dissertation, University of Washington.
- Ordóñez, Francisco and Esthela Treviño. 1999. Left Dislocated Subjects and the Pro-Drop Parameter: A Case Study of Spanish. *Lingua* 107:39-68.
- Papangeli, Dimitra. 2000. Clitic Doubling in Modern Greek: A Head-Complement Relation. *UCL Working Papers in Linguistics* 12:473-498.
- Pensalfini, Robert. 2003. *A Grammar of Jingulu: An Aboriginal Language of the Northern Territory* (Pacific Linguistics 536). Canberra: Pacific Linguistics.
- Pesetsky, David, and Esther Torrego. 2007. The Syntax of Valuation and the Interpretability of Features. In S. Karimi, V. Samiian and W. Wilkins (eds.) *Phrasal and Clausal Architecture: Syntactic Derivation and Interpretation*. Amsterdam: Benjamins. pp. 262-294.
- Pulleyblank, Douglas. 1986. *Tone in Lexical Phonology*. Dordrecht: Reidel.
- Raposo, Eduardo. 1997. Definite/Zero Alternations in Portuguese: Towards a Unified Theory of Topic Constructions. In A. Schwegler, B. Tranel and M. Uribe-Etxebarria (eds.) *Romance Linguistics: Theoretical Perspectives*. Amsterdam: John Benjamins. pp. 197-212.
- Reinhart, Tanya. 2006. *Interface Strategies*. Cambridge: MIT Press.
- Rivero, María Luisa. 2008. Oblique Subjects and Person Restrictions in Spanish: A Morphological Approach. In R. D'Alessandro, S. Fischer and G.H. Hrafnbjargarson (eds.) *Agreement Restrictions*. Berlin: Mouton de Gruyter. pp. 215-250.
- Rizzi, Luigi. 1986. On the Status of Subject Clitics in Romance. In Osvaldo Jaeggli and Carmen Silva-Corvalán (eds) *Studies in Romance Linguistics*. Dordrecht: Foris. pp. 391-419.
- Sportiche, Dominique. 1998. Pronominal Clitic Dependencies. In Henk van Riemsdijk (ed) *Clitics in the languages of Europe*. Berlin: Mouton de Gruyter. pp. 679–708.
- Sproat, Richard. 1985. *On Deriving the Lexicon*. PhD dissertation, MIT.
- Sternefeld, Wolfgang. 1997. Comparing Reference Sets. In C. Wilder, H.-M. Gärtner and M. Bierwisch (eds.) *The Role of Economy Principles in Linguistic Theory*. Berlin: Akademie Verlag. pp. 81-114.

- Suñer, Margarita. 1988. The Role of Agreement in Clitic-Doubled Constructions. *Natural Language and Linguistic Theory* 6:391-434.
- Toivonen, Ida. 2007. Verbal Agreement in Inari Sami. In Ida Toivonen and Diane Nelson (eds.) *Sami Linguistics*. Amsterdam: John Benjamins. pp. 227-258.
- Torrego, Esther. 1996. On Quantifier Float in Control Clauses. *Linguistic Inquiry* 27:111-126.
- Van Kampen, Jacqueline. 1996. PF/LF Convergence in Acquisition. In K. Kusumoto (ed.) *Proceedings of NELS 26*. Amherst: GLSA. pp. 149-163.
- Uriagereka, Juan. 1995. Aspects of the Syntax of Clitic Placement in Western Romance. *Linguistic Inquiry* 26, 79-123.
- Williams, Edwin. 1976. Underlying Tone in Margi and Igbo. *Linguistic Inquiry* 7:463-484.
- Williams, Edwin. 1981. On the Notions 'Lexically Related' and 'Head of a Word.' *Linguistic Inquiry* 12: 245-274