The Agreement Hierarchy and Agree

Peter W. Smith

Goethe-Universität, Frankfurt

October 2016

Abstract. Hybrid nouns - nouns whose morphological shape differs from their semantic interpretation - can sometimes control agreement that either reflects the semantic interpretation or their morphological form. When there are multiple targets of agreement with a hybrid noun, mismatches among the targets are tolerated, however, only in certain configurations which appear to reflect Corbett's (1979) Agreement Hierarchy (*the government is embarrassing themselves* vs. **the government are embarrassing itself*). In this paper I argue that these restrictions follow from a two-step agreement model (Benmamoun et al. 2009, Arregi & Nevins 2012) of Agree, composed of Agree-Link, which takes place iteratively in the narrow syntax, and Agree-Copy, which is by-and-large a post-syntactic process. This formulation of Agree is shown to interact with structure building, such that in the relevant configurations, targets that merge earlier into the derivation delimit the potential agreement values that can be shown by targets that merge later by virtue of the fact that they undergo Agree-Link earlier.

1   Introduction

Since its introduction into the Minimalist Programme, the operation of Agree has been the primary means of elements sharing featural information. There has been considerable debate over how exactly should Agree be formulated. What is of interest to this paper is whether agreement is a one-step operation that takes place all in one component (for instance, among many others, Chomsky 2000, 2001, Preminger 2015), or whether it is composed of two operations that are distributed over different components (among others, Arregi & Nevins 2012). Whereas Agree has by and large been formulated in the former terms, various evidence has been given for the latter formulation in recent years, thanks mostly to studies on closest conjunct agreement (Benmamoun et al. 2009, Bhatt & Walkow 2013, Marušič et al. 2015). In this paper I aim to argue in favour of

1

the latter formulation through consideration of the interaction between morphologically motivated agreement and semantically motivated agreement. Semantically motivated agreement - an agreement value taken from a feature that is not morphologically realised, but nonetheless semantically interpreted - has been recently proposed to show different properties than traditional ϕ agreement. Whilst ϕ-agreement on T has been traditionally formulated as the probe looking downwards in the structure (see Preminger 2013, Preminger & Polinsky 2015 for recent defenses of agreement only looking downwards), Zeijlstra (2004, 2012) argues that negative concord motivates a view of agreement where the agreement target (the probe, in Minimalist terms) looks upwards in the structure to find its controller (the goal), a conclusion that Zeijlstra (see also Wurmbrand 2011, 2012, 2014).

The upwards versus downwards debate does not play a big factor in this paper, however, related to this debate, Author (2015), Smith (to appear) have recently proposed that semantically motivated agreement is evaluated off the output of syntactic structure, as it is sensitive to LF structures, such that the controller must c-command the target at the level of LF. Smith bases his discussion on so-called 'hybrid' nouns (see Corbett 1979, Wechsler & Zlatić 2003), which seem to be able to control an agreement value that does not match their morphological shape, but does match their semantic interpretation. For instance, in dialects of English, collective nouns (CNP) like *commitee* are able to control plural agreement on verbs, despite being unambiguously morphologically singular. The plural agreement ostensibly reflects the plurality of the CNP, which are formed of multiple individual members.

(1)   The committee are making a decision now, please try to be patient.

In this paper, I will argue that the interaction between morphological agreement and semantic agreement motivates a two-step agreement (TSA) model, where Agree consists of the sub-operations of Agree-Link and Agree-Copy (see also Arregi & Nevins 2012 in addition to the references cited for closest conjunct agreement). Specifically, by considering evidence from the interaction of semantic agreement and morphological agreement, this paper will argue that Agree-Link takes place throughout the narrow syntactic derivation, applying at the first derivational point that controller and target are in the structure together. However, I will also review data showing that whilst this operation of linking takes place early, the operation of copying (Agree-Copy) takes place comparatively late, on the output of syntactic structure.

The phenomenon used to motivate the early placement of Agree-Link will come from what I term Agreement Hierarchy Effects. The Agreement Hierarchy is a generalisation given by Corbett (1979), who shows that in a language with

2

hybrid nouns, certain targets are more likely to show semantic agreement than others, which are more likely to show morphologically motivated agreement. For instance, as in (1), English CNPs can either control singular or plural agreement. Plural, as noted, appears to be semantically motivated, whilst singular appears to be motivated by the morphological shape.[1] The ability to show both singular and plural agreement is not restricted to verbal agreement, but is also found on pronominal elements.

(2)   a.   The government is embarrassing itself with this strategy.
        b.   The government are embarrassing themselves with constant scandals.

Corbett (1979) (see also Corbett 1983, Levin 2001) shows that across a corpus the level of semantic agreement with hybrid nouns is higher for some targets than others. Specifically, he proposes the following hierarchy, where elements to the right on the scale are more likely to show semantic agreements than elements to the left, which are more likely to show morphological agreement:

(3)   attributive — predicate — relative pronoun — personal pronoun
        ← morphological agreement   semantic agreement →

Corbett also shows that the Agreement Hierarchy is monotonic: if semantic agreement is possible for one slot on the hierarchy, then all slots to the right on the scale will also show semantic agreement in the language, though not necessarily slots to the left. Conversely, if a slot shows morphological agreement, then all slots to the left will also show morphological agreement, but not necessarily those to the right. Crucially, in Corbett's formulation, these generalisations hold across a corpus, and predict relative frequencies of agreement for the targets. However, as will be discussed throughout this paper, the Agreement Hierarchy, superficially at least, appears to reach beyond the level of a 'corpus' generalisation, and apply at a deeper level. For instance, in English, there can be a mismatch of values between multiple targets that agree with a hybrid nouns, such that an anaphor shows semantic agreement whilst the verb shows morphological agreement. So, whilst mismatches are tolerated as shown in (4-a), the converse mismatch is not allowed (4-b):

---

[1] It is important to stress that the singular agreement does not just apply to CNPs that are perceived as a whole, which if so, could simply represent again the semantics. For instance, the following sentence is fine, even though the use of the predicate *gather* ensures that the CNP is not semantically singular, as *gather* must take a semantically non-singular subject.

(i)   The committee is gathering back at Head Office for the crucial decision.

(4)   a.   The government is embarrassing themselves day after day.
       b.   *The government are embarrassing itself day after day.

This seems to reflect the Agreement Hierarchy, since mismatches are allowed when the element to the right in (3) shows semantic agreement and the element to the left morphological, but not conversely. I will term these restrictions on mismatches Agreement Hierarchy Effects (AHEs), and turn to them in the next section, where more patterns will be introduced. I will provide an explanation for these effects in section 3, showing that we can understand them if AGREE-LINK applies iteratively throughout the syntactic derivation. In section 4 I will extend this proposal to restrictions on mismatches that do not easily reflect Corbett's hierarchy, since the two targets come from the same slot. In section 5 I discuss how the current proposal compares to the general appraoch to hybrid nouns given by Wechsler & Zlatić (2003), as well as a recent proposal concering DP-internal agreement in Hebrew by Landau (2016). Before concluding, I consider in section 6 the implications for the general theory of AGREE, discussing recent evidence showing that although the linking between the controller of agreement and its target happens throughout the syntactic derivation, the copying of feature values happens post-syntactically. This will crucially argue for a TSA model, where AGREE is composed of distinct operations that are distributed across different components of the grammar.

2   Agreement Hierarchy Effects

2.1   AHEs

Hybrid nouns, as mentioned above, nouns whose morphological shape does not reflect its semantics. They are seen in languages with grammatical gender. For instance, the German word for 'small girl', *Mädchen* has neuter grammatical gender (as indicated by its combination with the neuter definite article *das*), even though the referent is semantically feminine. When hybrid nouns are able to control either morphological or semantic agreement and there are two targets of agreement, we expect four configurations: two configurations where the agreements on the targets match, and two cases of mismatches.[2] Consider once more CNPs in English, which can control plural agreement or singualr agreement on both verbal elements and anaphors. When a sentence contains both a verb and an anaphor, there are in principle 4 agreement configurations:

---

[2]Not all hybrid nouns are able to control semantic agreement, see section 3 below and Author (2015) for discussion.

(5)

Verb	Anaphor	
Morphological (singular) Semantic (plural)	Morphological (singular) Semantic (plural)	Matching
Morphological (singular) Semantic (plural)	Semantic (plural) Morphological (singular)	Mismatching

Now consider the following data (from Smith to appear).

(6)   a.   The government **has** offered **itself** up for criticism (with this economic policy).
     b.   The government **have** offered **themselves / each other** up for criticism.
     c.   The government **has** offered $^?$**themselves / each other** up for criticism.
     d.   *The government **have** offered **itself** up for criticism.

In (6), we see the following patterns. (6-a) and (6-b) show that, unsurprisingly, sentences where the value on verb and anaphor match are grammatical. It is also possible for there to be a mismatch between the two as in (6-c) where the anaphor agrees for plural and the verb singular. However, it is not possible for the mismatch to go the other way, such that the anaphor is singular and the verb plural (6-d). Since mismatches are in principle allowed, and each element can independently alternate between the two, it is surprising that this sentence is disallowed. Interestingly, the allowed mismatch is the one that reflects Corbett's hierarchy: it has the anaphor (personal pronoun on Corbett's scale) showing semantic agreement, and the verb (the Predicate in (3)) showing morphological agreement. But, the sentence is ungrammatical the pattern goes against the Agreement Hierarchy, with the pronominal element showing morphological agreement and the predicate showing semantic agreement. This 3/4 effect, where 3 out of the available 4 configurations grammatical, is the characteristic of what I term here Agreement Hierarchy Effects (AHE). As we will see below, not all mismatch restrictions obviously reflect the Agreement Hierarchy.

A similar AHE comes from Russian, given in Corbett (1983) (see also Pesetsky 2013). In this instance, the controller of agreement is the noun *vrač* 'doctor'. *Vrač* potentially has a mismatch according to gender. The noun has masculine grammatical gender, however when the referent is a female doctor, gender agreement can reflect the semantic gender, feminine.[3] When there are two targets, here the adjective *novyj* 'new' and the verb *skazal* 'said', we again see three out of the four predicted patterns arising. The two targets can match in gender according to the

---

[3]If the referent of *vrač* is a male, then there is no gender mismatch.

morphological shape (7-a) or match according to the semantics of the referent, (7-b). Similarly, they can mismatch if the adjective shows morphological agreement, and the verb semantic agreement. (7-c). However, the converse mismatch however is not possible, where the adjective shows semantic agreement, and the verb agrees morphological:

(7)   a.   Novyj   vrač   skazal.
            new.MASC doctor said.MASC
            'The new doctor said.'
       b.   Novaja   vrač   skazala.
            new.FEM doctor said.FEM
            'The new doctor said.'
       c.   Novyj   vrač   skazala.
            new.MASC doctor said.FEM
            'The new doctor said.'
       d.   *Novaja   vrač   skazal.
            new.FEM doctor said.MASC
            The new doctor said.'

Turning to Hebrew, (Landau 2016), another 3/4 instance of agreements arises. Landau (2016) shows that the Hebrew noun *be'alim* 'owner(s)' has another mismatch between its morphology and its semantic interpretation. The noun is morphologically plural, showing the *-im* suffix which marks masculine plural nouns. However, semantically it can either mean refer to a singular owner or a plurality of owners.

(8)   a.   hem/hen   hayu   ha-be'al-im   šel
            they.MASC/they.FEM were.3.PL the-owner-MASC.PL of
            ha-dira
            the-apartment
            'They were the owners of the apartment.'
       b.   hu/hi   haya/hayta   ha-be'al-im   šel
            he/she was.3.SG.MASC/was.3.SG.FEM the-owner-MASC.PL of
            ha-dira
            the-apartment
            'He/she was the owner of the apartment.'

In cases where the noun refers to a single owner, there is then a mismatch on the number information of the noun. Similar to the English and Russian data above, when *be'alim* controls agreement on two separate targets, we again see that the full array of agreements fails to arise, with only three out of four possible. In this paradigm in (9), we see that the verb and adjective can match according to

	Attributive	Predicate	Pronoun	
Russian	■	▨		
Hebrew	■	▨		Grammatical
BrE		■	▨	
Russian	▨	■		
Hebrew	▨	■		Ungrammatical
BrE		▨	■	

Table 1: Grammatical and ungrammatical mismatches.

the morphological shape (9-a), or match according to the semantic shape (9-b). With regard to mismatches, if the adjective shows morphological agreement and the verb semantic, then the mismatch is fine (9-c), but the converse is not allowed (9-d):

(9)   a.   ha-be'al-im   ha-kodem   maxar et   ha-makom lifney
            the-owner-PL the-previous.SG ACC.   sold.3.SG the-place   before
            šana
            year
            'The previous owner sold the place a year ago.'
       b.   ha-be'al-im   ha-kodm-im   maxru   et ha-makom lifney šana
            the-owner-PL the-previous-PL sold.3.PL ACC the-place   before year
            'The previous owners sold the place a year ago.'
       c.   ?ha-be'al-im   ha-kodm-im   maxar   et ha-makom lifney šana
            the-owner-PL the-previous-PL sold.3.SG ACC the-place   before year
            'The previous owner sold the place a year ago.'
       d.   *ha-be'al-im   ha-kodem   maxru   et ha-makom lifney šana
            the-owner-PL the-previous.SG sold.3.PL ACC the-place   before year
            'The previous owner(s) sold the place a year ago.'

Up to this point, we have been considering instances where a mismatch seems to be ruled out because it goes against the Agreement Hierarchy. The possible and impossible mismatches are visualised in table 1. Gray cells indicate when a target has semantic agreement and Black cells indicate morphological agreement.

However, it should be noted that I am not proposing here that the explanation is directly related to the Agreement Hierarchy. The explanation could equally be independent of this, and just make it look like the Agreement Hierarchy. We also find instances of mismatch restrictions, which do not reflect the Agreement Hierarchy (though I continue to label them as AHEs, with AHE referring to mismatch restrictions in general). Consider Chichewa, as discussed in Corbett (1991). In the following examples, the hybrid noun is *ngwazi* 'hero', which Corbett claims

to be of gender class 9. However, it can also trigger agreement of gender class 1, which is the default class of animates in Chichewa. Put into current terms, we can analyse class 9 as the morphological shape, with class 1 reflecting its semantics (a hero being semantically animate). When two elements agree with *ngwazi*, we see a familiar pattern. Where the two targets match in value, the sentence is grammatical (10-a) and (10-b). With regard to mismatches only one configuration is possible. If the ordinal shows semantic agreement and the possessive pronoun targets morphological agreement, then they are fine to mismatch (10-c), but not conversely (10-d):

(10) a. ngwazi y-athu y-oyamba
          hero   9-our   9-first
          'Our first hero.'
      b. ngwazi w-athu w-oyamba
          hero   1-our   1-first
          'Our first hero.'
      c. ngwazi y-athu w-oyamba
          hero   9-our   1-first
          'Our first hero.'
      d. *ngwazi w-athu y-oyamba
          hero   1-our   9-first
          INTENDED: 'Our first hero.'

This is a different type of AHE than what we have encountered up to now, however, it is the same general pattern. The noun *ngwazi* can control both semantic and morphological agreement and mismatches between two targets are allowed but only tolerated in one direction. All that differentiates this pattern from the previous ones is that both targets are drawn from the same slot on the Agreement Hierarchy, specifically, attributives. Now, it may be argued at this point that the term 'attributive' on Corbett's hierarchy is too coarse, and should be further subdivided. In fact, in Corbett (1983), the slot 'predicate' is further decomposed to distinguish between finite verbs, participles, (predicate) adjectives and (predicate) nouns. A further decomposition of 'attributive' could be in order, and if so, then Chichewa would be consistent with other AHEs discussed here.

A further subdivision of the 'attributive' slot may or may not be warranted, however, I do not pursue this any further due to evidence from Hebrew. Recall from above that Hebrew that the noun *be'alim* can control either singular or plural agreement. Landau (2016) shows that when there are two adjectives that combine with *be'alim*, the adjectives can show a mismatch between agreements, however only in one way. The adjective further away from the noun is allowed

to show semantic agreement:[4]

(11) a. ?ha-be'alim ha-pratiyim   ha-axaron šel ha-tmuna   haya
          the-owner the-private.PL the-last.SG of the-painting was.3SG
          ha-psixo'analitika'i Jacques Lacan
          the-psychoanalyst   Jacques Lacan
          'The last private owner of the painting was the psychoanalyst Jacques
          Lacan.'
      b. *ha-be'alim ha-prati   ha-axron-im šel ha-tmuna   haya/
          the-owner the-private.SG the-last-PL   of the-painting was.3.SG/
          hayu   ha-psixo'analitika'i Jacques Lacan.
          was.PL the-psychoanalyst   Jacques Lacan
          'The last private owner of the painting was the psychoanalyst Jacques
          Lacan.'

Note here that though there is seemingly an AHE pattern, the targets involved are drawn from the same slot on the hierarchy. A further subdivision of 'attributive' does not help us since both targets are one and the same thing.

3   Deriving AHEs

In this section, I provide an account of AHEs, capturing why some mismatches between targets are allowed, yet the converse mismatches are disallowed. For this section, I restrict attention to patterns where the two targets are taken from different slots on the Agreement Hierarchy. Thus, the paradigms of interest come from English (6), Russian (7) and Hebrew (8). The multiple attributive patterns from Hebrew and Chichewa will be discussed in the following section. From the discussion in section 2, what appears to be the case is that the element to the right on the hierarchy controls what elements to the left can show. I will first discuss some prelimaries regarding semantic and morphological agreement, before laying out an informal schema of the analysis. I then flesh out the details in the subsequent subsections.

3.1   Hybrid Nouns: Theoretical Preliminaries

Generally, there is a predictable correspondence between the morphological shape of a lexical item and its semantic interpretation. For instance, in (12-a) the subject noun phrase is singular, and is interpreted as if there is one (particular) owl who

---

[4]Though I phrase this here in terms of linearly furthest away, this is not intended to be contentful with respect to what I propose below. It is tempting to look for an explanation based on linear order, however, linear relations do not play a role in the analysis I give below.

is sitting on a branch. In (12-b) *owls* is inflected for plural, and the interpretation is that there are multiple owls who are being referred to.[5]

(12)   a.   The owl is sat on a branch.
         b.   The owls are sat on a branch.

Despite such correspondences between morphology and semantics being the normal case, we have seen various exceptions to this above, where the correspondence between morphology and semantics breaks down. So called 'hybrid' nouns then involve a mismatch between the morphological shape and its semantic interpretation. For instance, in a grammatical gender language, a noun may be grouped into a grammatical gender without it having any semantic gender, such as *der Tisch* 'the.MASC table' in German. Alternatively, a noun may have a semantic interpretation and a morphological form that overtly mismatch.

The ability to control either morphological or semantic agreement is not fully free. For instance, in English, CNPs (when morphologically singular) are never allowed to combine with plural demonstratives, even when other agreement targets in the same sentence are plural. This fact is in accordance with Corbett's hierarchy, as demonstratives will qualify as attributive elements, and so the 'cut-off point' for semantic agreement in English is between 'attributive' and 'predicate' in (3):

(13)   a.   This committee are blaming each other for the mess.
         b.   *These committee are blaming each other for the mess.

Following in the spirit of Wechsler & Zlatić (2003), Smith (to appear) proposes that the 'dual' nature of hybrid nouns can be analysed as there being a divergence of information carried on a feature between the morphological value and the semantic value. Smith proposes that all ϕ-features are composed of a morphological half and a semantic half. During the syntax, features consist of both of these halves, before they are separated at the point of spell-out. The morphological half *u*F, which is semantically uninterpretable is sent along the PF-branch for interaction with the morphological component, whilst the semantically interpretable *i*F is sent along the LF branch for interaction with the morphological component. In the normal case, where there is a no discrepancy between the form of a feature and its interpretation, features can be represented as in (14)a, whereas a 'hybrid' feature, will be represented as in (14)b (here the number feature of a CNP):

---

[5]I restrict attention to declarative contexts here, and do not take into account generic statements where the singular is used with multiple reference, such as *the cat is a member of the feline family*, where every cat is referred to.

(14) a.   b.

$$\phi_{number}$$

uF:singular   iF:singular   uF:singular   iF:plural
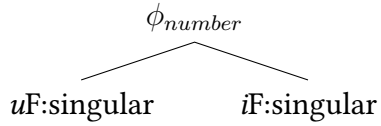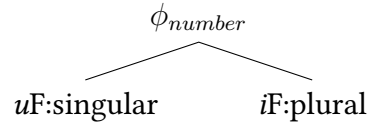
Following Smith (to appear), we can thus characterise semantic and morphological agreement in terms of what type of feature on the controller donates its value to the target. In semantic agreement, the value is drawn from the iF, whilst in morphological agreement, the value is drawn from the uF. As discsussed in section 6, and outline in more detail in Smith (to appear), there are further properties that distinguish semantic agreement from morphological agreement, which do not matter for the time being, though I return to them later on.

It is important to note though that semantic agreement is not able to be controlled by all potential controllers in a language. English does allow for semantic agreement with CNPs, as outlined above, so agreement with iFs is possible in the language. Imposter constructions (Collins & Postal 2012) can be seen as another instance of a hybrid noun, since the person features of the DP does not match how it is interpreted. They are morphologically third person DPs that are interpreted as first person.[6] However, it is not possible for the iF to factor into verbal agreement. In the following, the DP *this reporter* controls 3.sg agreement on the auxiliary, however the referent of the sentence is (on one reading) the speaker. However, 1.sg agreement on the auxiliary is not possible.

(15) a.   This reporter is expecting an announcement shortly.
      b.   *This reporter am expecting an announcement shortly.

That (15-b) is ungrammatical, suggests that Imposter DPs are not able to control semantic agreement, in contrast to CNPs. Therefore, we can conclude that semantic agreement is not something that is manifested at the level of a language, but rather is a property of (natural) classes of DPs.[7] I assume that whilst uFs are by default active for agreement (uF agreement seems to be the general case), certain DPs also make their iFs active for agreement also. For reasons that will become clear presently, I assume that activity is relevant to Agree-Link.[8]

---

[6]Collins & Postal (2012) give a structural account of imposter constructions, whereby there is a covert SPEAKER or ADDRESSEE argument merged above the Impsoter DP. I do not follow this here, and assume that Imposter DPs involve a mismatch on the person ϕ-feature.

[7]An issue that I abstract away from here is that within a class of semantic agreement controllers we see different rates of semantic agreement. *Team* is more likely to occur with plural agreement than *batallion* (Levin 2001).

[8]Note that the effects of iFs will only be seen with hyrbid nouns, since one can then contrast the value against the uF. In a regular non-hybrid noun, iF and uF agreement are indistinguishable.

(16)   An *active* feature is able to be targeted by Agree-Link.

3.2   General Schema of the Analysis for AHEs

Assuming Smith's characterisation of hybrid nouns, mismatches between targets of agreement are seen when one of the targets gets a value from the *i*F on the controller and the other gets its value from the *u*F. If the controller is a hybrid noun, then the result is a mismatch between the value given to Target 1 and that given to Target 2. However, the question that must be answered is why this situation is restricted in the manner that we have seen with AHEs. What I will propose is that an *i*F that is active for agreement can be optionally deactivated in the derivation. Furthermore, I will assume that this is a one way process; a feature that is inactive cannot be reactivated. Now, if there is an ordering between two targets of agreement, such that Target 1 undergoes agreement before Target 2, and the *i*F is active for Target 1 but is deactivated before the point that Target 2 undergoes agreement, then a mismatch will be created: Target 2 can only get a value from the *u*F of the controller as the *i*F is no longer active.

(17)

Step	Controller	Target 1 Value	Target 2 Value	
1.	[uF:▲,$\underline{i}$F:■]	–	–	
2.	[uF:▲,$\underline{i}$F:■]	■	–	Agreement
3.	[uF:▲,$\underline{i}$F:■]	■	–	*i*F deactivation
4.	[uF:▲,iF:■]	■	▲	Agreement

In order to avoid an ungrammatical mismatch being made possible, we must assume that *i*Fs that are active for agreement cannot be ignored. This would prevent a derivation where Target 1 undergoes agreement with the *u*F of the controller, before Target 2 undergoes agreement with the *i*F. Thus, I propose the following:

(18)   An active *i*F cannot be ignored by Agree-Link.

Now, if we superimpose an ordering of agreement targets as in (19) that is reflective of the Agreement Hierarchy, then we are able to account for AHEs, allowing for the attested ones, and disallowing the unattested patterns. Specifically, since anaphors undergo agreement before verbs, there is the possibility of the *i*F of the controller becoming inactive between the point that the anaphor undergoes agreement and the verb undergoes agreement, creating a mismatch between the two targets. Crucially, since the *i*F cannot be ignored for agreement when it is active, and inactive features cannot be made active, then the situation where an anaphor would undergo agreement with the *u*F and the verb the *i*F cannot arise.

(19) **Agreement Ordering Hypothesis**
anaphor $\prec$ verb $\prec$ adjective
("$x \prec y$" indicates x precedes y)

Whilst this gets us the right result for AHEs there are two problems. Firstly, why should (19) be the order, rather than that in (20), below. Simply stipulating (19) as part of UG is uninformative. Better is to try to derive the order from general principles. I turn to this in the next subsection.

(20) *Impossible:*
adjective $\prec$ verb $\prec$ anaphor

The second problem is that such an approach does not generalise to the instances where we find AHEs where the two targets are drawn from the same slot on the Agreement Hierarchy. This will be the focus of section 4.

3.3 Why anaphor $\prec$ verb $\prec$ adjective?

In this paper I argue for the following TSA model, which is based on Arregi & Nevins (2012), though modified in order to account for the phenomenon of semantic agreement (see Author 2015, Smith to appear). As can be seen from the formulation, AGREE is formed of two sub-operations: one operation of linking together controller and target and one operation of copying the value of the controller to the target. These sub-operations are termed AGREE-LINK and AGREE-COPY, respectively.[9]

(21) Agreement by a target with controller proceeds in two steps:
   a. AGREE-LINK: a target has unvalued ϕ-features that trigger Agree with a controller (possibly more than one). The result is a link between target and controller.

---

[9]Throughout this paper I use the informal terms 'controller' to refer to goal the element that donates the agreement value and 'target' to refer to the element that receives the agreement value. Thus, in keeping with the minimalist terminology of 'probe' and 'goal', we can formulate (21) as follows:

(i) Agreement by Probe Γ with Goal γ proceeds in two steps:

   (a.) AGREE-LINK: a Γ has unvalued ϕ-features that trigger Agree with γ (possibly more than one). The result is a link between Γ and γ.

   (b.) AGREE-COPY: the values of the ϕ-features of γ are copied onto Γ linked to it by AGREE-LINK.

      i. if AGREE-COPY happens at transfer, this requires that γ c-command the Γ.

b. AGREE-COPY: the values of the φ-features of controller are copied onto the target linked to it by AGREE-LINK.

i. if AGREE-COPY happens at transfer, this requires that controller c-command the target.

In order to account for the AHEs given above, I propose here that AGREE-LINK happens as soon as possible in the derivation: as soon as the controller and the target can establish an AGREE-LINK relation, they must do so. In practice, this means the first derivational step at which controller and target stand in a configuration where they can be linked. Anaphors canonically merge into the structure before T, and so they will appear in the structure with their controller before T does. Furthermore, I will propose that adjectives merge into the structure after T (assuming the late merger of adjuncts, Stepanov 2001), and will appear in the structure after both anaphors and T, which is discussed in more detail below. Thus, the Agreement Ordering Hypothesis ultimately stems from the Principle of Early AGREE-LINK, similar in spirit to the Earliness Principle of Pesetsky (1989), and the timing of merge.

(22) **Principle of Early AGREE-LINK (PEAL)**
Establish AGREE-LINK between a probe and a goal as soon as possible.

As mentioned, I restrict myself to discussing AGREE-LINK for the time being, and in section 6 return to the question of why PEAL references AGREE-LINK, rather than AGREE as a whole (in other words, why it is PEAL, not PEA).

I propose that we can understand AHEs through PEAL. As soon as controller and target are in the derivation together, they will be linked through AGREE-LINK. Recall from above, that I assume that *i*Fs can be either *active* or *inactive* for agreement. *Active i*Fs cannot be ignored by AGREE-LINK, so an active *i*F will always control agreement. However, if this is true, then it is important that we allow for features that are active to become inactive, otherwise mismatches will not be possible. To this end, I assume that if an *i*F undergoes AGREE-LINK, it potentially is rendered inactive. Once inactive, it will be unable to be seen by further iterations of AGREE-LINK, and only the *u*F value of the feature will be able to be targeted. In a configuration with two targets of agreement, since AGREE-LINK happens as early as possible, then mismatches arise when an active *i*F becomes deactivated by AGREE-LINKwith Target 1, leaving Target 2 only able to agree with the *u*F. Crucially, since *i*Fs that are active cannot be ignored, and inactive features cannot be reactivated, then the possibility of Target 2 undergoing AGREE-LINK with the *i*F, but Target 1 undergoing AGREE-LINK with the *u*F is ruled out.
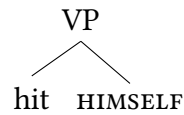
3.3.1 Anaphor ≺ Verb

To begin witConsider the sentence in (23), which has the numeration in (24):[10]
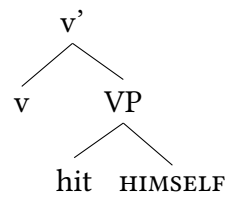
(23)   John hit himself.

(24)   {John, T, v, V, himself}

    Taking the standard assumption that structure is built from bottom-to-top (but see, e.g. Phillips (2003) for a top-down approach), the derivation proceeds as follows. As is standardly assumed, I assume that 'subject' agreement in English is carried on the functional T head. For our purposes it is irrelevant whether the verbal root raises covertly to T, or T lowers postsyntactically, so I ignore verbal movement in the derivation:

(25)   a.   Merge V and HIMSELF

                VP
              ／＼
          hit   HIMSELF

    b.   Merge v

                v'
            ／＼
        v   VP
               ／＼
          hit   HIMSELF

    c.   Merge *John*

               vP
         ／＼
      John   v'
             ／＼
         v   VP
              ／＼
         hit HIMSELF
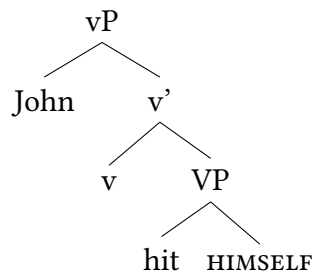
    d.   Merge T

---

[10]In the trees below, I represent targets of agreement with small caps notation to reflect their potential to appear as if they have agreed with the *u*F of the controller, or the *i*F when they undergo agreement with a hybrid controller.

```
                    T'
                  /    \
                 T      vP
                      /    \
                   John     v'
                          /    \
                         v      VP
                              /    \
                            hit   HIMSELF
```

e. Remerge *John*

```
                    TP
                  /    \
              John_i    T'
                      /    \
                     T      vP
                          /    \
                        t_i     v'
                              /    \
                             v      VP
                                  /    \
                                hit   HIMSELF
```
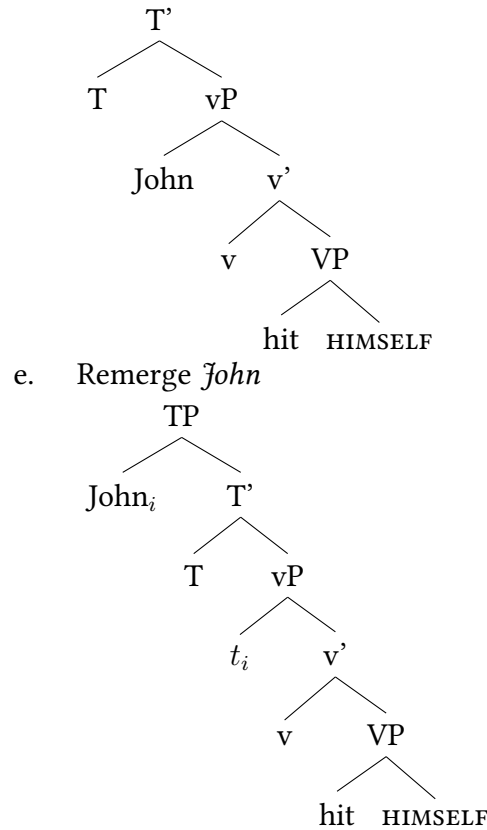
The important stages in the derivation are (25-c) and (25-d). The stage in (25-c) is the first stage in the derivation where *John* and the anaphor appear in the derivation together. (25-d) is the first stage where T and *John* are in the derivation together. Thus, if AGREE-LINK is formed as soon as possible between targets and controllers, it follows that the anaphor will undergo AGREE-LINK with the controller before T does, deriving the first part of the timing issue.[11]

With regards to mismatches, it is possible that the anaphor will undergo AGREE-LINK with the controller before T has merged into the structure (stage c in the above derivation). Thus, if AGREE-LINK deactivates the *i*F at this stage, then T will only be able to agree with the *u*F value of the controller.

---

[11]Note that AGREE-LINK does not copy the features of the goal onto the probe, but only establishes the relation between the two.

(26)

TP

T'

T   vP

Step 2:AGREE-LINK   CNP   v'

v   VP

embarrass   ANAPHOR

Step 1:AGREE-LINK
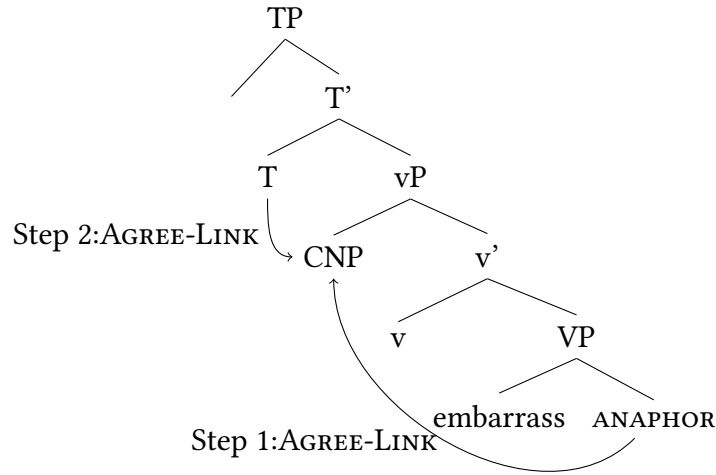
The structure (26) represents the (abbreviated) structure for the English mismatches, with the steps above conflated into one structure. Note that the structure is given before the point at which the CNP controller remerges into Spec,TP, so as to represent the position at which T and the anaphor undergo AGREE-LINK with it. As the derivation proceeds, as soon as the CNP is merged into Spec,VP, the anaphor undergoes AGREE-LINK. If the *i*F on the CNP is active, then the anaphor will be linked to *i*F:plural.[12] Since AGREE-LINK can either leave an active feature active or deactivate it, there are two possibilities now. If the feature remains active, than at Step 2, when T merges into the structure, it will undergo AGREE-LINK with the CNP and the result will be that it too becomes linked to the *i*F of the CNP (and eventually, plural agreement here). However, if at step 1 the *i*F of the CNP is deactivated, then at Step 2, T will only be able to undergo AGREE-LINK with the *u*F of the CNP, and the result will be the mismatch. It is worth repeating that at the point of Step 1, the anaphor cannot ignore an active *i*F, nor can an inactive *i*F become active between Step 1 and 2, and so the unattested mismatch cannot arise.[13]

---

[12] As will be discussed in section 6 this linking is crucially *not* the point at which the value of the controller is copied to the target: AGREE-LINK is a precondition for AGREE-COPY to successfully copy the value.

[13] Anaphors canonically merge into the structure before T, since they are canonically objects. It would be interesting to test whether different patterns of agreement are found in languages where anaphors can merge in the 'subject' position. I do not have data from hybrid nouns in languages that allow for subject anaphors, and so I leave this prediction to future research.

3.4 Verb ≺ Adjective

The second part of the timing issue relates to adjectives, and specifically, why they undergo agreement after verbs and anaphors. Following the above logic, it would seem that adjectives merge into the structure after T. Stepanov (2001) proposes that adjuncts must merge late into the structure. Since attributive adjectives are adjuncts, this means they merge counter-cyclically into the derivation, crucially after both anaphors and T.[14] Late merger, whilst controversial (see Sportiche 2016 for recent discussion) has also been assumed in various places in the literature (Lebeaux 1988, Fox & Nissenbaum 1999, Fox 2002, Landau 2007, Takahashi 2006, Takahashi & Hulsey 2009, Stanton 2016), yet Stepanov proposes that it is not just a possibility for adjuncts, rather, it is obligatory.

Taking this to be the case (though as noted, there is controversy around both the existence of late merger, and whether adjuncts must always merge late), this means that as attributive adjectives are adjuncts, the first point in the derivation where an adjective can undergo AGREE-LINK with its controller, occurs after both T and anaphors will have established AGREE-LINK relations with the controller. To illustrate, consider the sentence in (27), which has the numeration in (27-b). The derivation proceeds as in (28):[15]

(27)   a.   Hungry owls feed themselves.
       b.   {hungry, owls, T, v, feed, themselves}

(28)   a.   Merge *feed* and THEMSELVES

       b.   Merge v

---
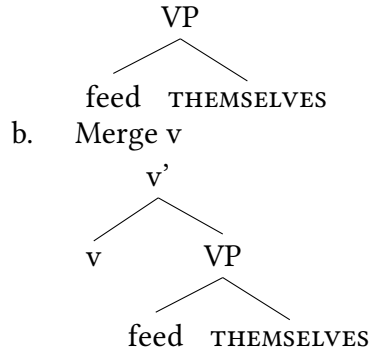
[14]There is controversy over whether adjectives are adjuncts or not. Cinque (2010) argues that adjectives each head their own projection, which would make them unlikely candidates for being adjuncts, however, Bošković (2013) argues that in many languages adjectives are adjuncts. There is potentially certainly scope for both approaches to be correct, and that languages differ parametrically on this point (see also the discussion of Chichewa possessive pronouns below), but here I take the strong view that all attributive adjectives are adjuncts. Note that, this does not mean that adjectives in a predicate position will merge late. In fact, (Corbett 1983) explicitly notes that predicative adjectives behave differently than attributive adjectives with regard to the Agreement Hierarchy.

[15]In the derivation, irrelevant structural details and steps are ignored for reasons of space.
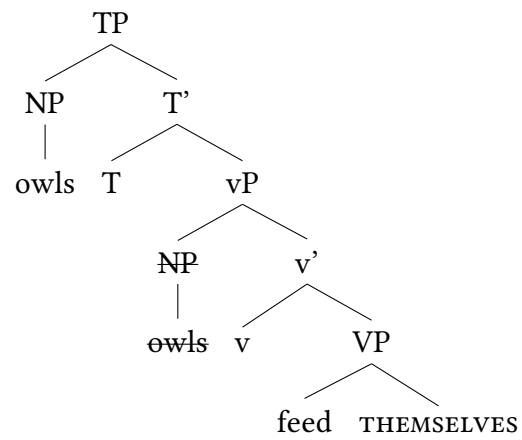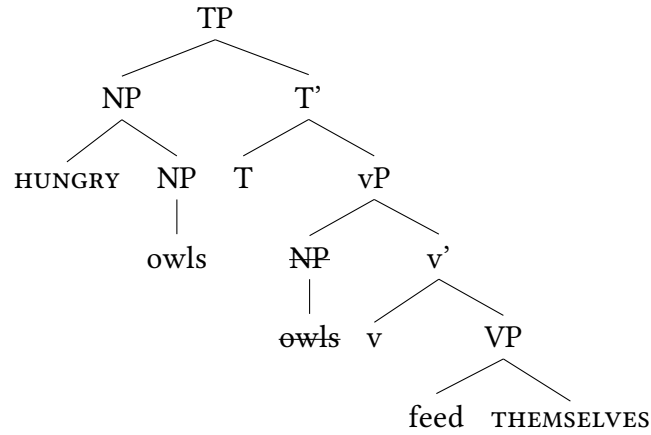
c. Merge *owls*

```
              vP
          ┌────┴────┐
        NP          v'
        │        ┌───┴───┐
       owls     v        VP
                      ┌───┴────┐
                    feed   THEMSELVES
```
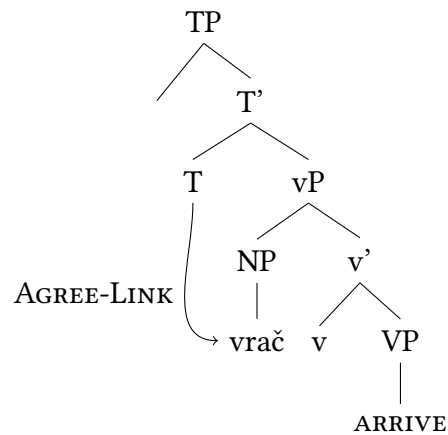
d. Merge T

```
               T'
          ┌────┴────┐
         T          vP
                ┌────┴────┐
              NP          v'
              │        ┌───┴───┐
             owls     v        VP
                            ┌───┴────┐
                          feed   THEMSELVES
```

e. Remerge *owls* in Spec,TP

```
                 TP
           ┌─────┴─────┐
         NP            T'
         │         ┌────┴────┐
        owls      T          vP
                         ┌────┴────┐
                       NP          v'
                       │        ┌───┴───┐
                      owls     v        VP
                                     ┌───┴────┐
                                   feed   THEMSELVES
```

f. Late-Merge HUNGRY

TP
NP   T'
HUNGRY NP   T   vP
owls   ~~NP~~   v'
~~owls~~ v   VP
feed THEMSELVES
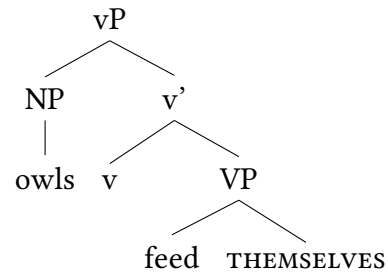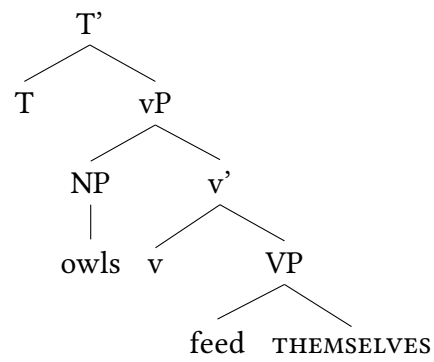
If the adjective were to undergo agreement with the head noun, it would only be able to do so at the point of the derivation in (28-f). Thus, through PEAL, coupled with a bottom-up derivation and late merger of adjuncts, we are able to derive the timing effect, such that anaphors agree before verbs, and verbs in turn before adjectives. The AHE from Russian is shown below. At Step 1, T undergoes AGREE-LINK with *vrač*, at the point that T is merged into the structure. In Step 2, after the subject NP has remerged in Spec,TP, the attributive adjective is merged counter-cyclically, and undergoes AGREE-LINK with *vrač*. In the same manner as in the English derivation in (26), if the iF:gender is active at Step 1, but deactivated there through AGREE-LINK, then a mismatch will arise between verbal and adjectival agreement, as the adjective can only be linked to the uF at Step 2. The same analysis holds for the Hebrew cases where the is a mismatch between the verb and the adjective.

(29)   a.   Step 1

TP
T'
T   vP
    NP   v'
AGREE-LINK
→ vrač v   VP
ARRIVE

    b.   Step 2

```
                          TP
                       /      \
                     NP        T'
                   /    \     /   \
                 NEW    NP   T     vP
            AGREE-LINK    |       /   \
                    ↘   vrač   NP      v'
                                |     /  \
                              vrač   v    VP
                                          |
                                        ARRIVE
```
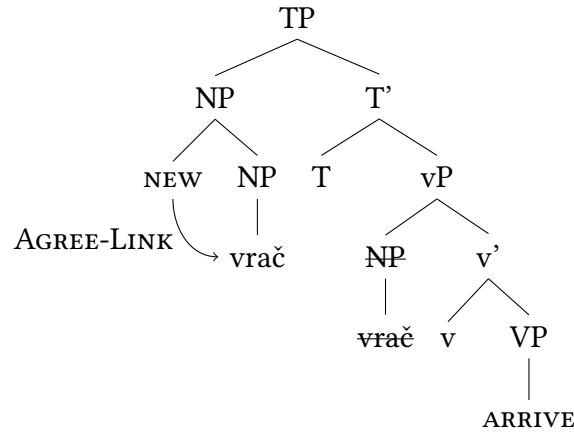
4 Capturing Mismatches from the Same Slot

In the preceding section, I have proposed the AHEs result from the timing of AGREE, coupled with the optional deactivation of *i*Fs. In all the cases up to now, I have been discussing instances where a mismatch is found from two targets that are drawn from different slots on the Agreement Hierarhcy. As they come from different slots, it has been relatively straightforward to see why certain items merge into the structure before others. In this section, I return to discuss the cases of multiple attributive elements from Hebrew and Chichewa, and show that the same logic can be applied there, resulting in a unified approach capturing the mismatches. For ease of reading, I repeat the relevant paradigms from Hebrew (30) and Chichewa (31) here.

(30)   a.   ?ha-be'alim ha-pratiyim   ha-axaron šel ha-tmuna   haya
              the-owner the-private.PL the-last.SG of the-painting was.3SG
              ha-psixo'analitika'i Jacques Lacan
              the-psychoanalyst   Jacques Lacan
              'The last private owner of the painting was the psychoanalyst Jacques
              Lacan.'
        b.   *ha-be'alim ha-prati   ha-axron-im šel ha-tmuna   haya/
              the-owner the-private.SG the-last-PL   of the-painting was.3.SG/
              hayu   ha-psixo'analitika'i Jacques Lacan.
              was.PL the-psychoanalyst   Jacques Lacan
              'The last private owner of the painting was the psychoanalyst Jacques
              Lacan.'

(31)   a.   ngwazi y-athu y-oyamba
              hero   9-our   9-first

21

'Our first hero.'

    b.   ngwazi w-athu w-oyamba
           hero   1-our   1-first
           'Our first hero.'

    c.   ngwazi y-athu w-oyamba
           hero   9-our   1-first
           'Our first hero.'

    d.   *ngwazi w-athu y-oyamba
           hero   1-our   9-first
           INTENDED: 'Our first hero.'

4.1   Differing structures for Hebrew and Chichewa

Before proceeding, it is worthwhile to consider the DP-internal structures of Hebrew and Chichewa before moving on. Both languages are N-initial, leading Landau (2016) to treat them as having the same internal structure, that is, left to right linear order is equal to low to high, with elements further from the noun c-commanding elements that are closer, as is the case below.

(32)

                        Adjective₂

      Noun Adjective₁
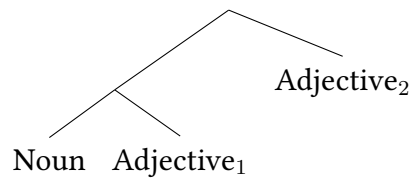
This rightward-is-highest structure is supported for Hebrew by data from Shlonsky (2004), who shows that the order of Hebrew adjectives is consistently the opposite order than English.

(33)

	English		Hebrew
	colour > nationality/origin		**nationality/origin > colour**
a.	a brown Swiss cow	c. *	para xuma švecarit
c. *	a Swiss brown cow	d.	para švecarit xuma
	shape > colour		**colour > shape**
f.	the long black table	g. *	ha-šulxan ha-'arox ha-šaxor
h. *	the black long table	i.	ha-šulxan ha-šaxor ha-'arox
	age > shape		**shape > age**
j.	the old round hat	k. *	ha-kova ha-yašan ha-'agol
l. *	the round old hat	m.	ha-kova ha-'agol hayašan

Now, assuming that adjectives merge in a fixed hierarchy into the tree (Cinque

[2010](), then this implies that in contrast to the rightest-is-lowest order of English adjectives, Hebrew has a leftest-is-lowest order.[16] Along with Landau, I assume that this is the correct structure for Hebrew adjectives.

However, the question is whether this is the way that we should treat Chichewa too. Specifically, given that possessive pronouns are linearly closer to the noun than ordinal numerals, does this mean that there are structurally lower than ordinal numerals? Landau assumes without discussion that Hebrew and Chichewa share a rightest-is-highest structure, however, it is notable that Carstens ([1991](), [1993]()) has argued the opposite for Bantu languages. Specifically, she has argued that the N-initial character of Kiswahili is derived through movement of the Noun to a high head position within the DP. This is supported by the fact that the unmarked order of elements that follows the noun is the same order that is found in English:

(34) picha   hii   yangu nzuri   ya   Busi
      9.picture 9.this   9.my   9.good 9.of
      'this nice picture of mine of Busi

As in Chichewa, the position of postnominal elements is subject to change, however Carstens notes that in Kiswahili, the order Noun-Demonstrative-Adjective-Theme is the *unmarked* order of those elements, suggesting that it is the basic one. Secondly, she notes that the possessive is apparently more limited in where it can appear, and generally appears close to the noun. Assuming once again that English is representative of the basic functional order (Cinque [2005]()) without movements to reorder elements, then we can conclude that the postnominal elements in the unmarked word order have not moved; all that has moved is raising of the noun to D. As the possessive is closer to the noun, then this suggests that the Bantu noun-internal word order branches in a rightwards-is-downwards fashion, precisely the opposite order of Hebrew.

In Chichewa, possessive pronouns generally precede ordinal numerals (Sam Mchombo, p.c.). Assuming a right branching DP structure, they are structurally higher. I thus assume the structure in [(35)]() for the relevant NPs in Chichewa. There are two points of note. I assume the N to D raising of Carstens ([1991]()) mostly for convenience. There is a debate over whether languages that do not have articles have the category D or not, notably, see Bošković ([2005](), [2013]()). This debate is orthogonal to the current point, so I do not wish to take a stand here on the status of 'D' in Chichewa, and assume that there is a D-layer only for conve-

---

[16] Though a fixed hierarchy of adjectives is has been conflated with LCA (Kayne [1994]()) compatible structures (e.g. Cinque [2010]()), a fixed hierarchy does not entail fixed, cartographic positions in the tree, see among others Bobaljik ([1999]()), Abels & Neeleman ([2012]()).

nience.[17] The second point is that Carstens (1991) assumes a structure whereby possessive pronouns appear in the specifier of NumP (having moved there from Specifier of NP), however, I assume that they are adjoined to NP.[18]

(35)

```
              DP
             /  \
                  D'
                 /  \
               D      NP
              / \    /  \
             D  hero Poss  NP
                          /  \
                    Ordinal   NP
                               |
                             hero
```
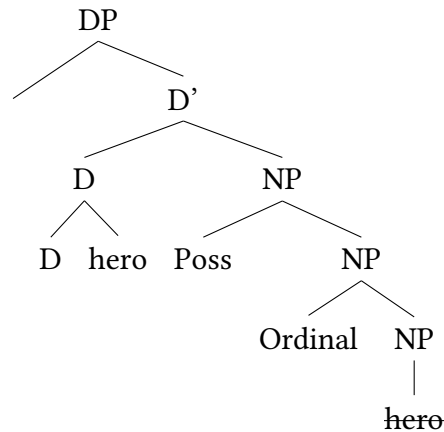
The final point that the reader should note about this discussion regarding the difference between Hebrew and Chichewa is that the languages appear to differ in one key regard. Assuming that the difference between the two languages is correct in that Hebrew DPs branch leftward whilst Chichewa DPs branch rightward, then it is the structurally higher adjunct that shows semantic agreement in Hebrew, but the structurally lower adjunct that shows semantic agreement in Chichewa. This raises an interesting issue: if we are to use the same explanation

---

[17]If D is absent in articleless languages, then there must be some functional projection above NP to host the raised noun in Bantu languages.

[18]Whether or not possessive pronouns are adjuncts or not seems to be a point of cross-linguistic variation. Bošković (2005), Despić (2011) shows that possessives in Serbo-Croatian have the same status of adjectives in the language (see also Bošković & Hsieh (2012) on Chinese, which shows the same). In Serbo-Croatian, possessives are allowed to bind out of the NP (Despić 2011):

  i. *Kusturicin$_i$ najnoviji film ga$_i$ je zaista razočarao
     Kusturica's latest film him is really disappointed

     'Kusturica$_i$'s latest film really disappointed him$_i$.'

Despić takes this to show that possessives do not occupy the same structural position as they do in English. In English, a possessive does not give a binding Principle B violation in the same environment:

  ii. John$_i$'s mother loves him$_i$.

Despić aruges that this contrast between the two languages shows that possessives in Serbo-Croatian are not buried within a PossP underneath a DP, but rather the DP layer is lacking altogether in this language, and that possessives adjoin to the NP layer along with adjectives.

given above in section 3 to account for the possible and impossible mismatches, then it must be the case that the higher adjunct has undergone AGREE-LINK before the lower adjunct in Hebrew. However, in Chichewa, it is the lower adjunct that undergoes AGREE-LINK before the higher one.

4.2   Accounting for the difference between Hebrew and Chichewa

In order to explain why it is the higher adjunct in Hebrew that shows semantic agreement, but the lower one in Chichewa, I propose that languages differ in the order in which they merge adjuncts. Key to the proposal above is that adjuncts merge counter-cyclically (Stepanov 2001), as assumed above. Merge, in the sense of Chomsky (1995), works in a cyclic manner because merge takes place at the root. Yet given that adjuncts merge counter-cyclically, there is no obvious requirement that they do so in a cyclic manner from lowest to highest. There is thus no reason to suspect that 'cyclic' counter-cyclic merger is forced in the case of adjunction. Suppose that it is a parametric choice across languages as to whether when merging multiple adjuncts at the same site, the adjuncts merge in top-down or bottom-up fashion. We can formalise this with the following statement:

(36)   In case of adjunction, adjoin to the {highest/lowest} segment of the adjunction site.

Now depending on which option the language takes with respect to where adjunction takes place, we see a difference in whether structurally higher or structurally lower adjectives have merged first in the derivation. Assuming that there is a fixed hierarchy of adjectives in the style of Scott (2002) and Cinque (2010), but that this hierarchy is enforced independently of the timing of merge (i.e. as long as the final order respects the hierarchy of adjectives, there is no requirement that lower ones merge before higher ones), for languages that merge at the HIGHEST segment, then adjectives which are higher in the structure have merged after ones that are lower. For languages which merge adjuncts at the LOWEST segment, adjectives that are higher in the structure must have merged before the ones that are lower.

To illustrate this, take a simple noun phrase like *big red car* in English. The adjectives *big* and *red* must appear in that order, and the opposite order **red big car* is ungrammatical. Suppose that English is a HIGHEST language, and merges adjectives at the *highest* segment of the adjunction site, there are two possible derivations. In the first, *red* merges before *big*:

(37)   a.   *Build NP*

NP
|
car

b. *Merge* red *with highest segment*

NP
╱╲
red NP
|
car

c. *Merge* big *with highest segment*

NP
╱╲
big   NP
╱╲
red NP
|
car
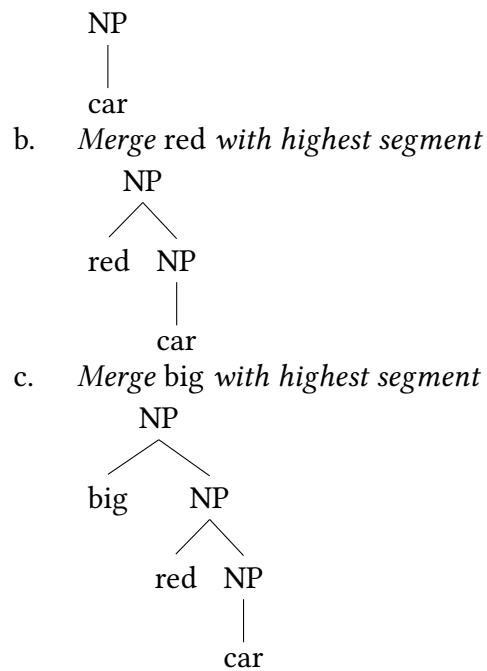
In the above, the correct order surface order is obtained. The first adjunction trivially targets the highest segment of NP, since there is only one segment. In the second adjunction, *big* then merges to the highest segment, and the result is that it is structurally higher than *red*. This gives the observed output of English. In the alternate derivation where *big* would merge into the structure before *red*, we end up with the ungrammatical **red big car*.

(38)   a. *Build NP*

NP
|
car

b. *Merge* big *with highest segment*

NP
╱╲
big NP
|
car

c. *Merge* red *with highest segment*

```
        NP
       ╱╲
    red    NP
          ╱╲
       big   NP
              │
             car
```
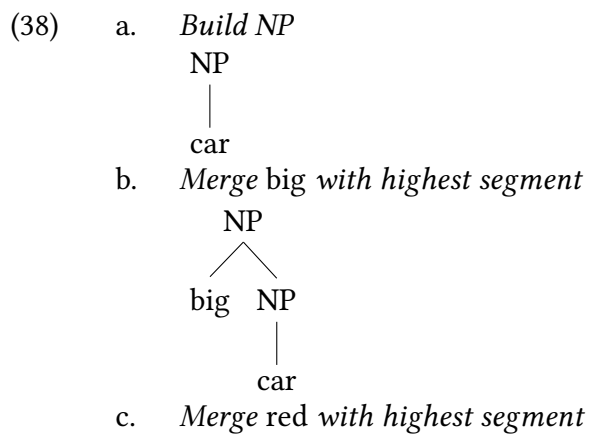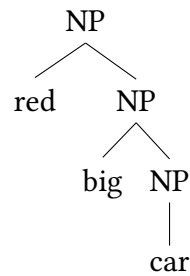
Now consider a Lowest language, where adjuncts target the *lowest* segment of the adjunction site. Again, I will illustrate with an English example, without committing to which option English chooses. If *red* merges first, the segment attached to is trivially the lowest one. However, when *big* comes to merge into the structure, it targets the lowest segment, and so will attach to the bottom segment of NP. The result is the ungrammatical order of adjectives, **red big car*:

(39)   a. *Build NP*

```
   NP
    │
   car
```

      b. *Merge* red *with highest segment*

```
       NP
      ╱╲
   red   NP
          │
         car
```

      c. *Merge* big *with highest segment*

```
       NP
      ╱╲
   red    NP
         ╱╲
      big   NP
             │
            car
```

The grammatical result would be obtained in this type of language if adjectives that are structurally highest have merged before adjectives are structurally lowest. In this derivation 'big' merges before 'red', and the correct output is generated.

(40)   a. *Build NP*

```
                NP
                |
                car
    b.   Merge big with lowest segment
                NP
               / \
            big   NP
                  |
                  car
    c.   Merge red with lowest segment
                NP
               / \
            big   NP
                 / \
              red   NP
                    |
                    car
```
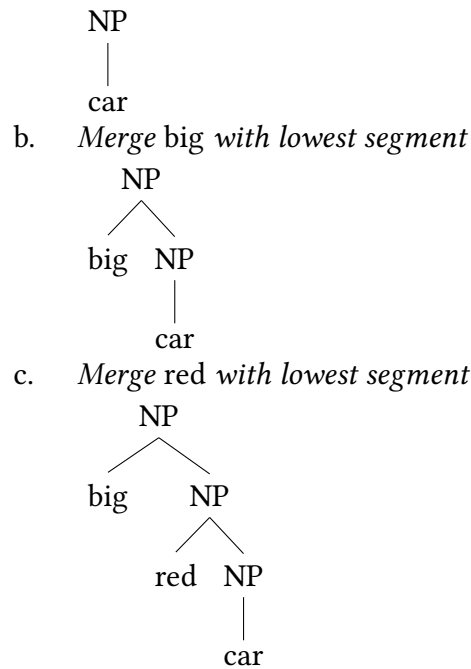
This discussion has served to show that depending on the choice that a language makes regarding the adjunction site (36), this will have an impact on whether structurally higher adjuncts have merged before or after structurally lower ones. In summary:

(41) a. If a language merges adjuncts at the HIGHEST segment of the adjunction site, adjuncts that are structurally higher have merged *after* adjectives that are structurally lower.

 b. If a language merges adjuncts at the LOWEST segment of the adjunction site, adjectives that are structurally higher have merged *before* adjectives that are structurally lower.

As I will show in the next two subsections, this distinction affords us the flexibility to use the same explanation for AHEs given in section 3 to account for the AHEs seen in multiple attributive patterns.

4.3 Hebrew

It is the structurally higher adjective in Hebrew that shows semantic agreement. According to the approach laid out above, this must mean that the structurally higher adjective has undergone AGREE-LINK before the structurally lower one. I assume that Hebrew has the following in effect:

(42)   In case of adjunction, adjoin to the Lowest segment of the adjunction
        site.

The attested mismatch is generated in the following way. As adjectives merge at
the lowest segment of the adjunction site, this means that more peripheral adjec-
tives merge first. In the following derivation, in Step 1, the peripheral adjective
LAST merges and undergoes Agree-Link with the head noun. At this point, the
*i*F is deactivated on the noun, leaving only the *u*F available for further iterations
to Agree-Link. In Step 2, PRIVATE merges and undergoes Agree-Link. Since
adjuncts in Hebrew merge at the lowest segment of the adjunction site, then it
will adjoin in between the head noun and LAST.

(43)   a.   Step 1
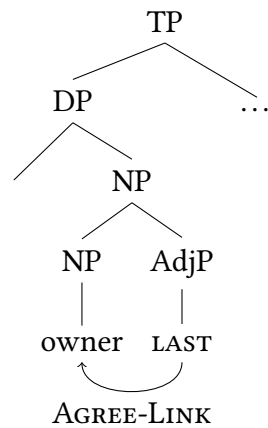
        b.   Step 2
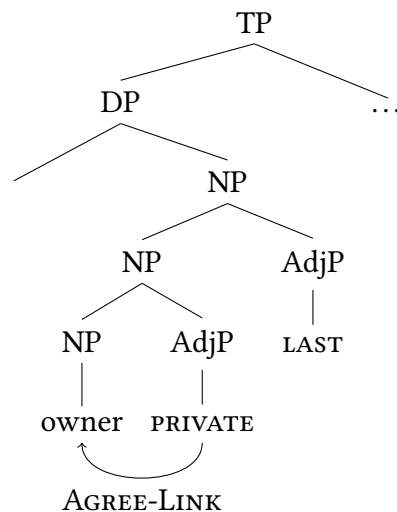
Pesetsky (2013) notes the same pattern as Hebrew for Russian. When there are two adjectives modifying a noun like *vrač*, there can be mismatches between the agreements on the adjectives. However, it must be the case that the structurally higher adjective shows semantic agreement, and the lower one morphological agreement, such as in (44) (I have altered the transliteration of Pesetky's minimally to be consistent with the Russian data given above). I assume the same analysis for Russian as given for Hebrew.[19]

(44)   a.   ?U menja očen' interesn-aja nov-yj vrač
              by me very interesting-FEM.NOM.SG new-MASC.NOM.SG doctor-NOM.SG
              'I have a very interesting new (female) doctor.'
        b.   *U menja očen' interesn-yj nov-aja vrač
              by me very interesting-MASC.NOM.SG new-FEM.NOM.SG doctor
              'I have a very interesting new (female) doctor.'

4.4   Chichewa

In contrast to Hebrew, I propose that Chichewa adjuncts merge at the highest segment of the adjunction site:

(45)   In case of adjunction, adjoin to the HIGHEST segment of the adjunction
        site.

As Chichewa merges adjuncts at the highest segment and the possessive is structurally higher than the ordinal, it must have merged after the ordinal. The attested mismatch then arises in the following derivation. In Step 1 the ordinal merges first, and undergoes AGREE-LINK with the head noun, before the posses-

---

[19]Having the same analysis for Russian is complicated by the fact that there adjectives in Russian that Pesetsky determines are low in the structure that never show semantic agreement.

(i.)   Glavn-yj/*Glavn-aja   vrač   poliklinik-i skazal-a
       head-MASC.NOM.SG/*head-FEM.NOM.SG doctor-NOM.SG clinic-GEN.SG say-PAST.FEM.SG
       čtoby   ...
       that.SUBJ
       'The (female) head doctor of the clinic ordered that ...'

Pesetsky argues that these adjectives have merged into the derivation before a feminising head is merged into the structure (the adjectives are structurally lower), and they undergo agreement before this head is merged. It is not possible to use Pesesky's analysis here, since I assume that all attributive adjectives, high or low, would merge into the structure counter-cyclically, which is presumably after the feminising head. In order to account for the non-semantically agreeing adjectives, I must stipulate here that they as a class are simply unable to enter into any kind of agreement relation with *i*Fs. Essentially, they would be somewhat akin to demonstratives in English, which can never show semantic agreement with CNPs.

sive merges and undergoes Agree-Link in Step 2. If the *i*F remains active after Step 1, then we will have matching *i*F agreement on the two adjuncts. However, a mismatch arises if the *i*F is deactivated at Step 1.

(46)   a.   Step 1
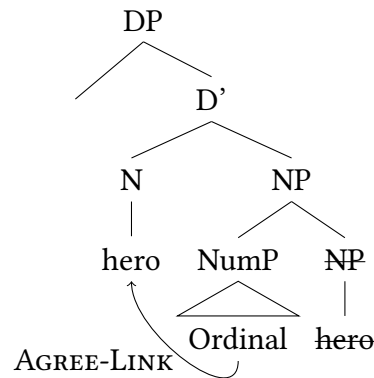
b.   Step 2
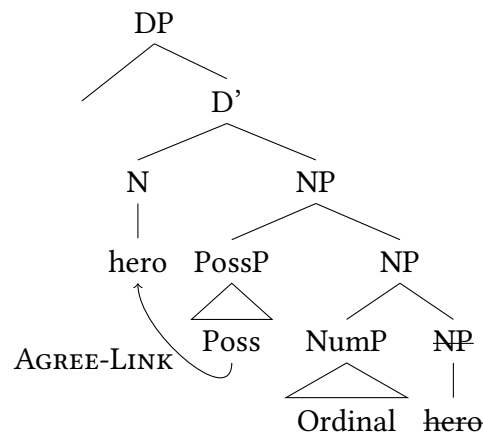
4.5   Summary

In this section I have shown that we can generalise the proposal given in section 3 to account for all AHEs, and not just for those that arise when the two targets of agreement come from different slots on the Agreement Hierarchy. Throughout this section, I have been assuming that Hebrew and Chichewa have different NP-internal structures, such that Hebrew is by and large left branching, whilst Chichewa is right branching. Parameterising the height of adjunction site has allowed for us to fit this difference between the languages with the proposal here. It should be noted that the general approach given here does not suffer if Hebrew and Chichewa are shown to have the same NP-internal structure. If

both are left-branching like Hebrew, then we must assume that adjunction in both languages targets the lowest segment of the adjunction site. However, if both languages are analysed as right-branching, then both must adjoin at the highest segments.
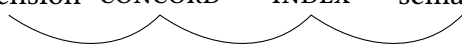
5   Comparison with other approaches

It is not a new observation that 'hybrid nouns' can induce agreement mismatches on different agreement targets. Nor is the observation that certain mismatches are allowed and others disallowed. In this paper, I have offered a novel analysis of this phenomenon that captures a wide range of data, however other approaches have been advanced in the existing literature, to which I now turn.

As pointed out in the introduction, the patterns that have been discussed in this paper are reminiscent of Corbett's Agreement Hierarchy (Corbett 1979, 1983, 2000, 2006, 2012). In Corbett's formulation, the Agreement Hierarchy merely controls the frequency of agreement at the level of a corpus. That is, all else being equal, across a corpus of data, personal pronouns will show a greater frequency of semantic agreement than relative pronouns, will, which in turn will show a higher frequency than predicates, so on and so forth. Such a statement, whilst descriptively true, is insufficient for our purposes here, since it does not mention what happens at the level of individual clauses.

Hybrid nouns have been prominently discussed within HPSG by Wechsler & Zlatić (2003), and later following largely in their footsteps Landau (2016). Wechsler & Zlatić's approach is similar to the one offered here, however couched in different albeit similar terms. They propose a distinction between CONCORD features and INDEX features. INDEX features are more closely related the semantic information of the noun, whilst CONCORD are more closely related to the morphological, or declension, information of the noun. Constraints govern the flow of information throughout the system such that in the general case, the information carried on each feature type matches up with the others. However, in certain instances, these links are broken, and the information does not match across all types of features, which gives rise to hybrid nouns.

Of interest to us here is that these constraints hold among contiguous regions in (47). That means, it is not possible for INDEX to have some value different from the semantics, but for CONCORD to have a value that it gets directly from the semantics.

(47)   a.   Regular case:

Declension   CONCORD   INDEX   semantics

b. Possible:

Declension CONCORD ‖ INDEX semantics

c. Not possible:

Declension CONCORD ‖ INDEX ‖ semantics
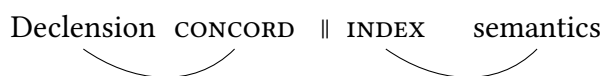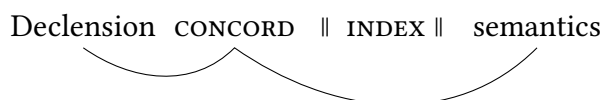
Wechsler & Zlatić discuss the Agreement Hierarchy and how this can be produced by their system. CONCORD and INDEX are on different featural paths (within HPSG). Relative pronouns can differ from personal pronouns, because, whilst relative pronouns always undergo anaphoric binding (agreeing with IN-DEX, it is possible for personal pronouns to also undergo pragmatic binding (taking an agreement value directly from the SEMANTICS in addition to INDEX). Thus, whenever there is a mismatch it will be the personal pronoun that more closely resembles the semantics, rather than the relative pronoun.

Wechsler & Zlatić note that their theory does not direcetly predict that predicates will have a lower frequency of semantic agreement than attributive element, or that attributive elements will have a higher frequency of morphological agreement. In their terms: why should attributive targets more frequently undergo agreement with the CONCORD feature? The answer that they suggest comes from the diachronic development of predicate agreement. They note that agreement markers historically evolve from incorporated pronouns, thus having a closer link to pronoun agreement (and thus being more likely than attributive elements to agree with the INDEX feature). The second point they note is that attributive elements are more local to the head noun than predicate elements, however, as they note, the notion of syntactic locality would need to be relativised to the attributive/predicate distinction, since it does not work with relative pronouns.

Wechsler & Zlatić note that the monotonic effect of Agreement Hierarchy follows from their account, because of the system of contiguous contraints that is built into their system. As pronominals can undergo either agreement directly with the semantics or the INDEX feature, they will always be able to reflect 'semantic' agreement even when INDEX reflects the morphological shape of the work (i.e. where there is a disconnect between INDEX and SEMANTICS). Similarly, where there is a disconnect between INDEX and CONCORD, as predicates can undergo agreement with INDEX or CONCORD, but attributive elements mostly only undergo agreement with the CONCORD feature, this means that predicates will show a higher frequency of agreement that is semantically motivated than attributive elements.

However, there are a couple of problems. Whilst their approach can in princi-

ple work for instances where the mismatching targets come from different slots on the agreement hierarchy, it does not provide any explanation for instances where the mismatches arise on elements from the same slot on the hierarchy, as is the case in Chichewa and Hebrew noted above. One could argue that syntactic locality can arbitrate here - in both Hebrew and Chichewa the target with morphological agreement is linearly and structurally closer to the controller. Given that both elements can in principle agree with the semantic information, there is no clear reason why a more peripheral adjective cannot agree morphologically across a less peripheral attributive element.

Landau (2016) offers an explanation here.Landau proposes that the plural morphology is the result of an inherent plural CONCORD feature, whereas the INDEX feature is open, and can be either plural or singular depending on the context. Landau proposes that since [CONCORD:plural] is inherently specified, it should be located on N (see also Acquaviva 2008, Author 2015), whereas the INDEX feature, not being inherent, is introduced in a NumP higher up in the nominal spine. The difference in adjectives comes from there being multiple places where adjectives are allowed to merge into the structure. Landau proposes that in the general case, adjectives merge lower than NumP, but that they can optionally merge higher than NumP. If there are then two adjectives that merge in different sites, and agreement only looks downwards (and cannot skip over a feature), then we get a mismatch:[20]

(48)
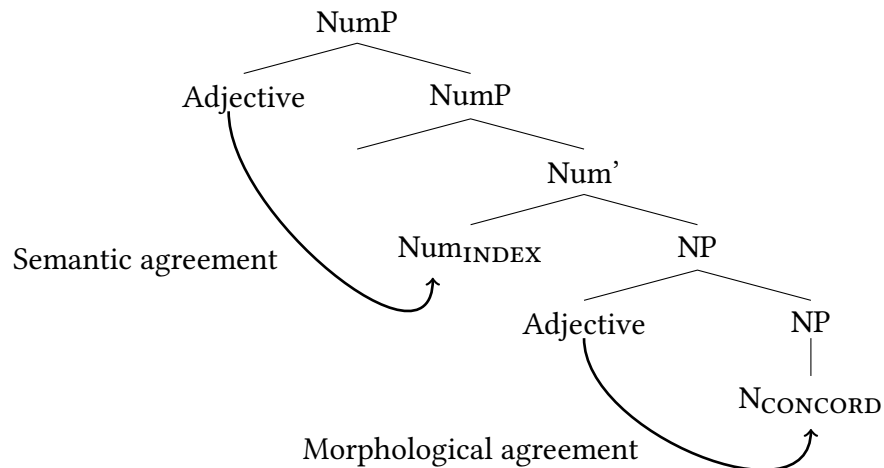
Whilst this does give the correct result for Hebrew multiple adjectives, it should be stressed that it only works for Chichewa as long as all N-initial lan-

---

[20]Linear order is ignored in the tree in (48).

guage involve the same, left-branching NP structure. As discussed above, this conclusion is open for debate. Furthermore, it is reliant on AGREE operating exclusively in a downward manner. As soon, as AGREE can look upwards in the structure, then there is a problem. It could be the case that there are two adjectives merged beneath NumP, and the adjective that is lower undergoes (upward) agreement to NumP, and the higher one undergoes (downward) agreement to N. However, upward agreement does appear to exist. Most notably for the current paper (and discussed below) Author (2015), Smith (to appear) shows that upward agreement is necessary for semantic agreement in English, but there is increasing evidence more generally that AGREE should not be limited to only looking downwards in the structure (see Baker 2008, Zeijlstra 2012, Wurmbrand 2011, 2012, 2014, Bjorkman & Zeijlstra 2015).

Putting aside these problems with Landau's proposal for the moment, it is worth revisiting the general approach to see how it compares with what is offered here. Firstly, though Wechsler & Zlatić's difference between predicates and attributive elements is derived by the development from personal pronoun to predicate marker, it should be noted that this holds only for diachronic development. As far as I can tell, there is nothing that prevents a language from developing beyond this stage into one where predicate targets are more likely to go with CONCORD features. Yet, this is unattested. If we ignore the proposal of diachronic development, all that remains for the approach of Wechlser & Zlatić to differentiate between attributive and predicates is the notion of locality. However, as they note, this could only be used to differentiate between attributive elements and predicates, as it is not consistent with relative pronouns. Thus syntactic locality appears unsuitable to be the ultimate arbitrator between targets in how sensitive they are to either CONCORD agreement or INDEX agreement.

Furthermore, even if we assume that there is some mechanism such that pronouns are more likely INDEX targets than predicates, this explanation is only suitable for sentences like the following:

(49)   *These committee is going to make a decision.

We know that in English, attributive targets do not undergo agreement for the INDEX feature, and so can only agree with CONCORD, whilst the predicate is able to inflect according to the information on INDEX (or CONCORD). However, Wechsler & Zlatić's proposal apparently faces a problem when there are two (or more) targets that can inflect for both INDEX and CONCORD. To see this, consider again Russian *vrač*, which as we have seen above, has the ability to control both masculine and feminine agreement on various targets. Let's assume for the sake of the argument that the featural representation of *vrač* is as follows (simplified from the more elaborate HPSG representation that Wechsler & Zlatić assume):

35

(50)   [CONCORD:masculine, INDEX:feminine]

Now, since adjectives and verbs can each show both singular and plural agreement, it must be the case that both must have the ability to undergo agreement with either CONCORD or INDEX. If this is true, then it is confusing as to why there should be a contrast between the mismatches familiar from the discussion of AHEs.

(51)   a.   Novyj   vrač   skazala.
              new.MASC doctor said.FEM
              'The new doctor said.'
        b.   *Novaja vrač   skazal.
              new.FEM doctor said.MASC
              The new doctor said.'

There doesn't seem to be anything in Wechsler & Zlatić's approach prevents the above situation other than potentially superimposing stipulations into the system, such that when predicates agree with the INDEX feature, then pronouns are preventing from agreeing with the CONCORD feature.

On the other hand, sentences like (51-b) follow from the present account without further stipulation in addition to the other patterns that Wechsler & Zlatić can capture. Additionally, the current account is able to handle the multiple attributive sentences of Landau (2016) under the same analysis.

In truth, it is probably not correct to see the approach offered here as being in opposition to Wechsler & Zlatić (2003). Aside from theoretical orientation (HPSG vs Minimalism), the major point of difference is that in the approach given here (elucidated in greater detail in Author 2015), there are only two potential agreement values - one from the iF and one from the uF. Wechsler & Zlatić allow for there being three - SEMANTICS, INDEX and CONCORD. Yet, the need for three different features or two is orthogonal to this paper, and has been recently discussed elsewhere (Alsina & Arsenijević 2012, Wechsler & Zlatić 2012). Given the similarities between the two approaches, then depending on the readers' disposition toward merging HPSG and Minimalist approaches, it is possible to see the ideas pursued here and in Wechsler & Zlatić as being complementary to one another.

6   AGREE-COPY is not (always) part of narrow syntax

I turn now to the second operation of AGREE, AGREE-COPY. Above, I have argued that AGREE-LINK is part of the narrow syntax, since it happens at the first derivational point where target and controller are in the structure. However,

up to this point, the position of Agree-Copy has not been discussed in any real detail. The only restriction is that it happens after Agree-Link, so it could in principle happen immediately after Agree-Link in the narrow syntax. If it happens immediately after Agree-Link is established, then there is little need to view Agree as decomposing into Agree-Link and Agree-Copy. However, here I review data showing that Agree-Copy is distinct from Agree-Link, and thus the decomposition is meaningful. I discuss two strands of evidence. Firstly, and extremely relevant to this paper is the discussion in Smith (to appear), who argues that semantic agreement is sensitive to LF positions, and so Agree-Copy must happen at the earliest at the point of LF. The second set of data come from closest conjunct agreement, which have been used by various authors to argue that Agree-Copy happens after the syntax, since it is sometimes sensitive to linear, rather than hierarchical position.

6.1 Semantic Agreement is sensitive to LF positions

Smith (to appear) argues at length that there is a crucial difference between semantically motivated agreement and morphologically motivated agreement. On the basis of data from CNPs in English (see also the discussion on Russian QNPs in Author 2015), Smith has argued that in order for semantic agreement to be possible, then the controller must c-command the target at the level of LF, whereas morphological agreement is not subject to any hierarchical restrictions other than either controller or target must c-command the other. Author proposes the following generalisation:

(52)   Semantic agreement is possible only if the controller of agreement c-commands the target at LF.

Since I do not have space to repeat all the arguments here, I show two constructions which eivdence this, and refer the reader to the paper for full argumentation. Firstly, Elbourne (1999), Sauerland & Elbourne (2002) show that in when there is singular agreement in the following, the CNP can be interpreted as taking either wide or low scope relative to *likely*. However, when the verb shows plural (semantic) agreement, only wide scope is possible, and the reconstructed reading is possible.

(53)   a.   A northern team is likely to be in the final.   ∃ >likely / likely >∃
        b.   A northern team are likely to be in the final. ∃ >likely / *likely >∃

Smith claims that the reason why plural agreement fails here is because the CNP does not c-command the target at LF, since it has reconstructed into the lower clause. Importantly, the contrasts also shows that it is not sufficient to

c-command the target of agreement at some point in the derivation: rather it is necessary that the controller remain above the target right the way through to LF.

Secondly, as noted by Elbourne (1999), Sauerland & Elbourne (2002), CNPs in (certain dialects of) English have the ability to control either plural or singular agreement, as noted above. However, plural agreement is not possible in every environment that singular agreement is. For instance, with existential constructions, only singular agreement is grammatical on the copula, whilst plural agreement is ungrammatical.

(54)   a.   There is a committee in there.
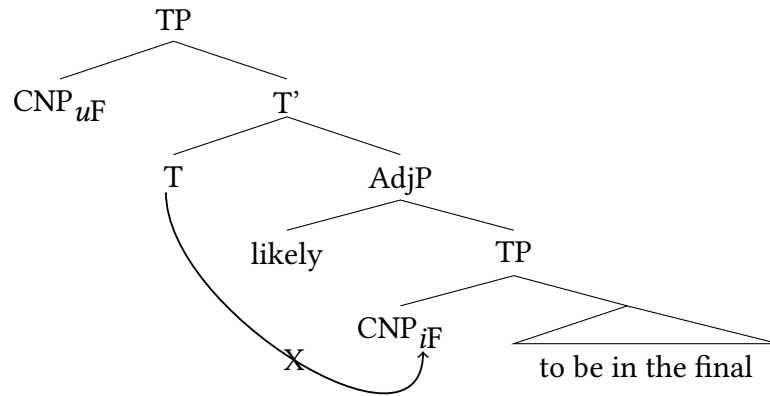   b.   *There are a committee in there.

There have been claims that the reason why plural agreement is not allowed here is due to some restriction that prevents plural-agreeing CNPs from appearing in existential constructions, see den Dikken (2001), Sauerland (2004a,b). Smith (to appear) however shows that CNPs can in principle control plural agreement as plural anaphors are licensed.

(55)   There is a team starting to psych themselves up in there.

Den Dikken (1995) shows that the postverbal DP in existential constructions remains low in the structure, and at LF does not appear in Spec,TP. Thus, the data in (54) are in accordance with (52).
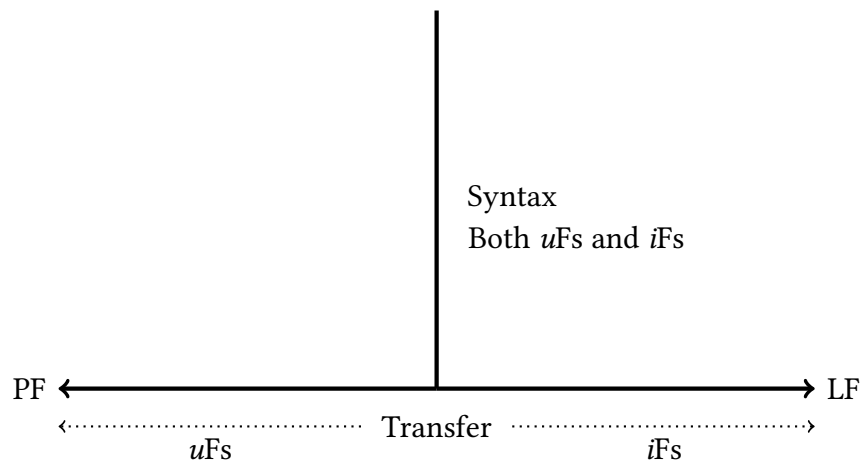
In order to explain the restrictions on semantic agreement, Smith assumes that "LF-position" should be understood as the position where the *i*Fs of an element are interpreted. Though this mostly correlates with the position where the *u*Fs are pronounced, in instances of reconstruction, the *u*Fs are pronounced in the matrix Spec,TP, whilst the *i*Fs are interpreted low in the matrix clause. Thus, the structure of a scope reconstruction sentence would be as follows:

(56)

```
                    TP
          ╱──────────────╲
      CNP_uF              T'
                   ╱──────────────╲
                  T              AdjP
                          ╱──────────────╲
                       likely              TP
                                  ╱──────────────╲
                              CNP_iF        to be in the final
                    X
```

Smith argues that *i*Fs and *u*Fs are together within the narrow syntax, before they are distibuted to the relevant interfaces, *i*Fs going to along the LF branch and *u*Fs going along the PF branch.

(57)

```
                                    │
                                    │
                                    │
                                    │   Syntax
                                    │   Both uFs and iFs
                                    │
                                    │
     PF ◄─────────────────────────────────────────────────► LF
        ◄········································· Transfer ·······························►
                  uFs                                        iFs
```

With regard to AGREE-COPY, given that it is the LF position which is important, Smith argues that we can make sense of this if AGREE-COPY happens at the earliest on LF structures. Since *i*Fs are not present in the PF branch, then in order for them to have morphological expression, it must be the case that the value of the *i*Fis copied from controller to target before the derivation proceeds to PF. Assuming a single-output model of syntax (for example, Bobaljik 1995), Smith proposes that this happens at the point of transfer, which ensures that both *i*Fs are still visible, and that no further syntactic movements can move the controller. However, AGREE-COPY which takes place at the point of transfer can only look upwards in the structure. Thus, there are at least two things that happen at the

point of transfer. Firstly, the positions of *i*Fs and *u*Fs are chosen, which I label here Privilege (*c.f. copy privileging* of Bobaljik 1995, 2002). Secondly, Agree-Copy applies looking upwards in the structure. Agree-Copy can happen in the PF branch, and in this case, it is able to look both upwards and downwards in the structure. This means that *u*Fs, which are still present in the structure in the PF-branch in contrast to *i*Fs, are able to be targeted by Agree-Copy even if they are beneath the target of agreement.[21] If we couple this proposal regarding Agree-Link discussed above in the structure with Smith's proposal, this leads to the following decomposition of Agree:

(58)

The decomposition of Agree
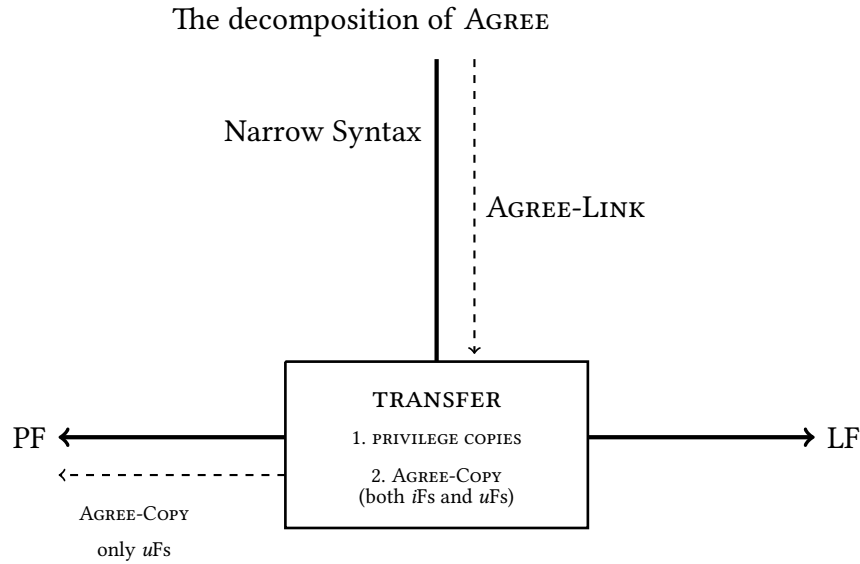
Note that if Agree-Copy were to happen earlier in the derivation, then there is nothing to stop Agree-Copy applying, and then the *i*Fs being privileged beneath the target of agreement, which would be a violation of (52). Thus, Smith concludes that Agree-Copy cannot be part of the narrow syntax. Crucial for our purposes here, it shows that Agree-Link and Agree-Copy are distinct operations that take place at different times. If the discussion about AHEs above is on the right track, then Agree-Link happens iteratively throughout the narrow syntax, at the first derivational point that target and controller are in the structure. However, Agree-Copy operates on the output of syntax, and crucially not iteratively.

---

[21]Smith does not offer an explanation as to why Agree-Copy taking place at the point of transfer can only look upwards in the structure, noting only that that is the conclusion that the data leads us to.

6.2 Closest Conjunct Agreement

Proponents of a TSA model have traditionally appealed to the phenomenon of closest conjunct agreement. Closest conjunct agreement is important in this realm, as it appears as though linear order plays a role for AGREE-COPY. Consider the following from Hindi and Tsez. The data here for both languages are taken from Benmamoun et al. 2009 ((59-a) is simplified). For other discussion of CCA in Hindi see Bhatt & Walkow 2013, for CCA in Tsez see Polinsky & Comrie 1999, and for apparent CCA in Slovenian, Marušič et al. 2015.

(59) a. maiN-ne ek chaataa   aur ek saaRii
          I-ERG   an umbrella.ABS.MASC.SG and a   saaree.ABS.FEM.SG
          khariid-ii
          buy-PERF.FEM.SG
          'I bought an umbrella and a saree.'
      b. us-ne   kharid-ii   kursii   aur sofa
          he-ERG buy.PERF-FEM.SG chair.FEM.SG and sofa.MASC.SG
          'He bought the chair and the sofa.'

(60) a. kid-no   uži-n   Ø-ik'i-s
          girl.ABS.II-and boy.ABS.I-and I-went
          'A girl and a boy went.'
      b. y-ik'i-s kid-no   uži-n
          II-went girl.ABS.II-and bou.ABS.I-and
          'A girl and a boy went.'

As can be seen in the data, verbal elements in Hindi can show agreement with either the last conjunct (59-a) or the first conjunct (59-b) (in both cases, it is the feminine, singular agreement on the verb that shows the agreement marking).[22]

The same pattern arises in Tsez, where the verb can undergo agreement with the last conjunct (60-a) or the first conjunct (60-b), this time marked by a prefix indicating the gender of the controller. Both languages are in general head final languages, in which case agreement can be with the second conjunct (see (59-a) and (60-a)), however, under certain circumstances, the verb can come to precede the coordination, in which case we observe agreement with the first conjunct (59-b) and (60-b).

There are a variety of accounts that have been proposed to account for CCA, see Benmamoun et al. (2009), Bhatt & Walkow (2013) and Marušič et al. (2015), among others. These approaches vary in the details, however, common to the

---

[22]This is a simplification, with clarifications below. CCA in Hindi is only observed if the coordination is a non-subject argument. If the conjunction is a subject, then the verb shows resolved agreement with the entire coordination (see Bhatt & Walkow 2013 for discussion).

cited work is the proposal that during the syntax, the verbal element is linked to the coordination phrase, &P, and that there is a separate operation that happens post-syntactically where the value of agreement is copied from target to controller. This second operation of copying looks for the linearly closest DP within &P. Important for our purposes here is that CCA is sensitive to linear relations, and not hierarchical structure (though see Bošković 2009, Puškar & Murphy 2015 for an alternative view, which does not rely on linear structure).[23] Linear structures are often taken within Minimalism to be a property of PF-structures (see Chomsky 1995, Fox & Pesetsky 2005, Arregi & Nevins 2012), and that syntax operates only on hierarchical structures (a notable exception is work assuming Antisymmetry, *e.g.* Kayne 1994). Supposing that it is true that the narrow sees only hierarchical structures, with linear structures only imposed on PF, then if AGREE-COPY can be sensitive to linear order, then it must at least be able to take place in PF.

Before leaving the discussion of CCA, it is worth considering the status of AGREE-LINK here. Though the discussion around (59) and (60) above presents a simplified picture, the key fact to take away is that linear relations appear to matter for AGREE-COPY, which leads to my following in the footsteps of the previously cited references that the mechanism of feature copying, AGREE-COPY, happens in the PF component. Yet just because AGREE-COPY happens in PF, one might wonder whether this shows that agreement is really composed of distinct operations. Could it be the case that AGREE-LINK also happens at PF in these languages, allowing for agreement to apply purely postsyntactically in one step?[24]

That agreement would be done purely postsyntactically is undermined by other factors that play into agreement in the two languages. Firstly, with regard to Hindi, as noted in various places, agreement in Hindi targets the structurally highest non-case marked argument (Bhatt 2005, Bobaljik 2008), and will 'ignore' a non-case marked object, despite the fact that it intervenes linearly between subject and verb, the following example from Bobaljik (2008).

(61) siitaa kelaa   khaatii   thii
      Sita   banana eat.IMPERF.FEM be.PAST.FEM
      'Sita (habitually) ate bananas.'

Similarly, agreement in Tsez has been shown to be sensitive to other syntac-

---

[23]Benmamoun et al. (2009) argue in detail that CCA cannot be reduced to simply *highest conjunct agreement*, which is attested in many languages, by assuming that coordination phrases are right branching in first conjunct agreement, and left branching in last conjunct agreement. They show that even with last conjunct agreement, there is clear evidence that the conjunct that is linearly first c-commands the linearly second conjunct.

[24]Borsley (2009) proposes that agreement is purely sensitive to linear relations in Welsh, which could be a candidate for this.

tic factors. As is well known, Tsez shows Long Distance Agreement, where a matrix verb can agree with an argument of an embedded finite clause (Polinsky & Potsdam 2001). However, this is only possible when the embedded argument is interpreted as a topic.

(62) eni-r   [už-ā   magalu   b-āc'-ru-ɬi]
     mother-DAT [boy-ERG bread.III.ABS III-eat-PSTPART-NMLZ]
     b-iy-xo
     III-knowPRES
     'The mother knows that the bread, the boy ate.'

For Polinsky & Potsdam (2001), this means that the argument is, at the level of LF, in the left periphery of the embedded clause, which means that it will be visible at the edge of the phase, according to the PIC of Chomsky (2001). Thus, it is the hierarchical position of the embedded absolutive that is crucial, and so there is some aspect of AGREE must take place prior to the conversion of hierarchical structure into linear structure. Note that if the embedded absolutive is not interpreted as a topic, then LDA is not possible, even though the linear position is the same (the lack of topichood indicated in the translation):[25]

(63) eni-r   [už-ā   magalu   b-āc'-ru-ɬi]
     mother-DAT [boy-ERG bread.III.ABS III-eat-PSTPART-NMLZ]
     r-iy-xo
     IV-knowPRES
     'The mother knows that the boy ate bread.'

A final point should be made here. Having AGREE-COPY as part of the PF component does not entail that AGREE-COPY will only be senstive to linear order. Linear order becomes relevant if AGREE-COPY happens after the linearisation of syntactic structure. If it takes place before that, then it will remain sensitive to hierarchical structures (see Arregi & Nevins 2012 for discussion on the timing of post-syntactic operations).

7 Conclusion

The aim of this paper has been two-fold. Firstly, I have provided an account of AHEs, patterns of allowed as disallowed mismatches involving morphological

---

[25]Bjorkman & Zeijlstra (2015) claim that the information structure features of the embedded argument make it visible to agreement. Their account also makes use of a distinction between linking and copying in AGREE, albeit in a slightly different way than here, which is based entirely in the narrow syntactic component.

and semantic agreement that seem to mirror the Agreement Hierarchy of Corbett (1979) *et seq.*. Given that Corbett's hierarchy is formulated to hold at the level of a corpus of data, to see it apparently operative at the level of single clauses is surprising, and raises questions of how different targets of agreement with the same controller can interact with each other. In the account given here, I have proposed that the interaction comes about because AGREE-LINK occurs as soon as controller and target are in the structure together (PEAL). AGREE-LINK cannot ignore *i*Fs if they are active for agreement, however, an instance of AGREE-LINK can deactivate an *i*F, leaving only the *u*F available for further instances of AGREE-LINK. Once this proposal was combined with the assumption that adjuncts are obligatorily merged late into the derivation, the result is that we can understand all AHEs under one mechanism, an apparently welcome result that to my knowledge has not been proposed elsewhere in the literature.

The second aim of this paper was to provide novel arguments that AGREE is composed of two sub-operations, one of linking the controller and target and one of copying features from the controller to the target. Crucially, the distinct operations take place in different components of the grammar. Such a model of agreement has been proposed before, as discussed above, however the majority of work that motivates this comes from the phenomenon of closest conjunct agreement. However, here I have offered an argument for this model of agreement from another phenomenon. I have argued here that AGREE-LINK happens early in the syntactic derivation, whilst structure is still being built. I have further contrasted this with AGREE-COPY, which applies on the output

References

Abels, Klaus & Ad Neeleman (2012) Linear Asymmetries and the LCA. *Syntax* **15**(1): 25–74.

Acquaviva, Paulo (2008) Roots and lexicality in Distributed Morphology. *http://ling.auf.net/lingbuzz/000654* .

Alsina, Alex & Boban Arsenijević (2012) The two faces of agreement. *Language* **88**(2): 369–379.

Arregi, Karlos & Andrew Nevins (2012) *Morphotactics: Basque Auxiliaries and the Structure of Spellout.* Dordrecht: Springer.

Author (2015) *Title.* Ph.D. thesis, University.

Baker, Mark (2008) *The syntax of agreement and concord.* Cambridge: Cambridge University Press.

Benmamoun, Elabbas, Archna Bhatia, & Maria Polinsky (2009) Closest conjunct agreement in head final languages. *Linguistic Variation Yearbook* **9**: 67–88.

Bhatt, Rajesh (2005) Long-distance agreement in Hindi-Urdu. *Natural Language and Linguistic Theory* **23**: 757–807.

Bhatt, Rajesh & Martin Walkow (2013) Locating agreement in grammar: An argument from conjunctions. *Natural Language and Linguistic Theory* **31**: 951–1013.

Bjorkman, Bronwyn & Hedde Zeijlstra (2015) Upward Agree is Superior. Manuscript.

Bobaljik, Jonathan D. (1995) *Morphosyntax: The Syntax of Verbal Inflection*. Ph.D. thesis, MIT.

Bobaljik, Jonathan D. (1999) Adverbs: the hierarchy paradox. *Glot International* **4**: 27–28.

Bobaljik, Jonathan D. (2002) A-Chains at the PF-Interface: Copies and Covert Movement. *Natural Language and Linguistic Theory* : 197–267.

Bobaljik, Jonathan D. (2008) Where's phi? In *Phi-theory: phi-features across modules and interfaces*, Daniel Harbour, David Adger, & Susanna Béjar, eds., Oxford Studies in Theoretical Linguistics, Oxford: Oxford University Press, 295–328.

Borsley, Robert D. (2009) On the superficiality of Welsh agreement. *Natural Language and Linguistic Theory* **27**: 225–265.

Bošković, Željko (2005) On the Locality of Left Branch Extraction and the Structure of NP. *Studia Linguistica* **59**(1): 1–45.

Bošković, Željko (2009) Unifying first and last conjunct agreement. *Natural Language and Linguistic Theory* **27**(3): 455–496.

Bošković, Željko (2013) Phases beyond clauses. In *The Nominal Constructinos in Slavic and Beyond*, Lilia Schürcks, Anastasia Giannakidou, & Urtzi Etxeberria, eds., Berlin: Mouton de Gruyter.

Bošković, Željko & I-Ta Hsieh (2012) On word order, binding relations and plurality within Chinese NPs. In *Proceedings of IsCLL 13*.

Carstens, Vicki (1991) *The morphology and syntax of determiner phrases in Kiswahili*. Ph.D. thesis, University of California Los Angeles.

Carstens, Vicki (1993) On nominal morphology and DP structure. In *Theoretical Aspects of Bantu Grammar*, Sam A. Mchombo, ed., Leland Stanford Junior University.

Chomsky, Noam (1995) *The Minimalist Program.* Cambridge, MA: MIT Press.

Chomsky, Noam (2000) Minimalist inquiries. In *Essays on minimalist syntax in honor of Howard Lasnik*, Roger Martin, David Michaels, & Juan Uriagereka, eds., Cambridge, MA: MIT Press, 89–155.

Chomsky, Noam (2001) Derivation by phase. In *Ken Hale: A life in language*, Michael Kenstowicz, ed., Cambridge, MA: MIT Press, 1–52.

Cinque, G. (2005) Deriving Greenberg's Universal 20 and its exceptions. *Linguistic Inquiry* **36**: 315 – 332.

Cinque, G. (2010) *The syntax of adjectives.* Cambridge, MA: MIT Press.

Collins, Chris & Paul Postal (2012) *Imposters: A study of pronominal agreement.* MIT Press, Cambridge, MA.

Corbett, Greville (1979) The Agreement Hierarchy. *Journal of Linguistics* **15**: 203–225.

Corbett, Greville (1983) *Hierarchies, targets and controllers: Agreement patterns in Slavic.* Penn State Publishing.

Corbett, Greville (1991) *Gender.* Cambridge: Cambridge University Press.

Corbett, Greville (2000) *Number.* Cambridge: Cambridge University Press.

Corbett, Greville (2006) *Agreement.* Cambridge: Cambridge University Press.

Corbett, Greville (2012) *Features.* Cambridge: Cambridge University Press.

den Dikken, Marcel (1995) Binding, expletives and levels. *Linguistic Inquiry* **26**(2): 347–354.

den Dikken, Marcel (2001) Pluringulars, pronouns and quirky agreement. *The Linguistic Review* **18**(1): 19–41.

Despić, Miloje (2011) *Syntax in the absence of Determiner Phrase.* Ph.D. thesis, University of Connecticut.

Elbourne, Paul (1999) Some correlations between semantic plurality and quantifier scope. In *Proceedings of NELS 29*, 81–92.

Fox, Danny (2002) Antecedent-contained deletion and the copy theory of movement. *Linguistic Inquiry* **33**: 63–96.

Fox, Danny & Jon Nissenbaum (1999) Extraposition and scope: A case for overt QR. In *Proceedings of WCCFL 18*, Sonya Bird, Andrew Carnie, Jason D. Haugen, & Peter Norquest, eds., 132–144.

Fox, Danny & David Pesetsky (2005) Cyclic linearization of syntactic structure. *Theoretical Linguistics* **31**: 1–45.

Kayne, Richard (1994) *The Antisymmetry of Syntax*. Cambridge, MA: MIT Press.

Landau, Idan (2007) Contstraints on Partial VP-fronting. *Syntax* **10**(2): 127–164.

Landau, Idan (2016) DP-internal semantic agreement: A configurational analysis. *Natural Language and Linguistic Theory* **34**: 975–1020.

Lebeaux, David (1988) *Language acquisition and the form of the grammar*. Ph.D. thesis, University of Massachusetts Amherst.

Levin, Magnus (2001) *Agreement with collective nouns in English, Lund Studies in English*, vol. 103. Lund University.

Marušič, Franc, Andrew Nevins, & William Badecker (2015) The grammars of conjunctino agreement in Slovenian. *Syntax* **18**(1): 39–77.

Pesetsky, David (1989) Languae Particular Processes and the Earliness Principle. Manuscript, MIT.

Pesetsky, David (2013) *Russian Case Morphology and the Syntactic Categories*. Cambridge, MA: MIT Press.

Phillips, Colin (2003) Linear order and constituency. *Linguistic Inquiry* **34**(1): 37–90.

Polinsky, Maria & Bernard Comrie (1999) Agreement in Tsez. *Folia Linguistica* **33**: 109–130.

Polinsky, Maria & Eric Potsdam (2001) Long-distance agreement and topic in Tsez. *Natural Language and Linguistic Theory* **19**: 583–646.

Preminger, Omer (2013) That's not how you agree: A reply to Zeijlstra. *The Linguistic Review* **30**(3): 491–500.

Preminger, Omer (2015) *Agreement and its failures*. Cambridge, MA: MIT Press.

Preminger, Omer & Maria Polinsky (2015) Agreement and Semantic Concord: A Spurious Unification, manuscript, University of Maryland.

Puškar, Zorica & Andrew Murphy (2015) Closest conjunct agreement in Serbo-Croatian: A rule ordering account. In *Topics at InfL, Linguistische Arbeitsberichte*, vol. 92, Anke Assmann, Sebastian Bank, Doreen Georgi, Timo Klein, Philipp Weisser, & Eva Zimmermann, eds., Universität Leipzig, 441–482.

Sauerland, Uli (2004a) "A team," definitely. *Snippets* **9**: 11–12.

Sauerland, Uli (2004b) *A Comprehensive Semantics for Agreement.* Berlin: Ms. ZAS.

Sauerland, Uli & Paul Elbourne (2002) Total reconstruction, PF-movement, and derivational order. *Linguistic Inquiry* **33**(2): 283–319.

Scott, Gary-John (2002) Stacked Adjectival Modification. In *Functional Structure in DP and IP: The Cartography of Syntactic Structures*, vol. Volume 1, Guglielmo Cinque, ed., Oxford: Oxford University Press, 91–122.

Shlonsky, Ur (2004) The form of semitic noun phrases. *Lingua* **114**(12): 1465–1526.

Smith, Peter W. (to appear) The Syntax of Semantic Agreement in English. *Journal of Linguistics* .

Sportiche, Dominique (2016) Neglect (or doing away with Late merger and Countercyclicity). Manuscript, lingbuzz/002775.

Stanton, Juliet (2016) Wholesale late merger in Ā-Movement: Evidence from preposition stranding. *Linguistic Inquiry* **47**(1): 89–126.

Stepanov, Arthur (2001) Late adjunction and minimalist phrase structure. *Syntax* **4**(2): 94–125.

Takahashi, Shoichi (2006) *Decompositionality and identity.* Ph.D. thesis, MIT, Cambridge, MA.

Takahashi, Shoichi & Sarah Hulsey (2009) Wholesale late merger: Beyond the A/Ā distinction. *Linguistic Inquiry* **40**: 487–526.

Wechsler, Stephen & Larisa Zlatić (2003) *The many faces of agreement.* Stanford.

Wechsler, Stephen & Larisa Zlatić (2012) The wrong two faces. *Language* **88**(2): 380 – 387.

Wurmbrand, Susanne (2011) Reverse Agree. *Manuscript, University of Connecticut* .

Wurmbrand, Susanne (2012) The syntax of valuation in auxiliary-participle constructions. In *Proceedings of the 29th West Coast Conference on Formal Linguistics*, Jaehoon Choi et al, ed., Tucson: University of Arizona, 154–162.

Wurmbrand, Susanne (2014) The merge condition: a syntactic approach to selection. In *Mimimalism and beyond: radicalizing the interfaces*, Peter Kosta, Lilia Schürcks, Steven Franks, & Teodora Radev-Bork, eds., Amsterdam: John Benjamins, 139–177.

Zeijlstra, Hedde (2004) *Sentential negation and negative concord*. Ph.D. thesis, University of Amsterdam.

Zeijlstra, Hedde (2012) There is only one way to agree. *The Linguistic Review* **29**: 491 – 539.