Possible and Impossible Agreement Mismatches

Peter W. Smith
p.smith@em.uni-frankfurt.de

April 2017

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 one of the two logically possible ways, (*the government is embarrassing themselves* vs. **the government are embarrassing itself*). In this paper I present a number of such patterns, and show that these restrictions follow the order of merge, 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 initiate AGREE earlier than later ones. This gives a unified account of all the patterns, and, among other properties of semantic agreement (Smith to appear), provides further evidence for a two-step agreement model of AGREE (e.g.Benmamoun et al. 2009, Arregi & Nevins 2012), composed of AGREE-LINK, which takes place iteratively in the narrow syntax, and AGREE-COPY, which is by-and-large a post-syntactic process.

1 Possible and Impossible Agreement Mismatches

Corbett (1979) notes for English that nouns that the nouns in (1) are able to license either plural or singular agreement. The singular agreement is reflective of their morphological shape whereas the plural agreement apparently reflects their semantic interpretation.

(1) a. The committee is/are now meeting.
b. The committee yelled at each other for over an hour, and nothing got done.
c. The government set itself up for embarrassment.

These nouns form a class of collective nouns (CNPs) in English that allow for either semantic or morphological agreement. As their semantic interpretation and their morphological shape differ, they have come to be known as 'hybrid nouns', further examples of which we will see from other languages as we proceed. We can see from the examples above that verbs, pronouns and anaphors can all show either plural or singular agreement. Interestingly, when we combine two of these targets in the same sentence, a curious effect arises. Specifically, we see that whilst both sentences with matching

agreements are fine, only one of the mismatches between the targets is allowed (Smith to appear).

(2) 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.

Mismatches between targets are allowed in principle (2-c), therefore it is curious why (2-d) is sharply ungrammatical.

A similar pattern is seen in Russian, this time with gender agreement. Corbett (1979, 1983) notes that profession nouns like *vrač* 'doctor' are morphologically masculine, but have the ability to license feminine agreement when the referent of the noun is a female. Again, we see that mismatches between agreement targets are tolerated, but only in one direction. This time, the verb is allowed to show the semantic agreement whilst the adjective shows the morphologically motivated agreement.

(3) 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.'

The same pattern is seen in Modern Hebrew (Landau 2016). Landau notes that the 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 as shown by the following:

(4) a. hem/hen hayu ha-be'al-im šel ha-dira
they.MASC/they.FEM were.3.PL the-owner-MASC.PL of 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 (5), we see that the verb and adjective can match according to the morphological shape (5-a), or match according to the semantic shape (5-b). With regard to mismatches, if the adjective shows morphological agreement and the verb semantic, then the mismatch is fine (5-c), but the converse is not allowed (5-d):

(5) a. ha-be'al-im ha-kodem maxar et ha-makom lifney šana
the-owner-PL the-previous.SG ACC. sold.3.SG the-place before 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.'

The fact that agreement mismatches are allowed, but only in one way is a surprising result, and it is striking that the patterns are strongly reminiscent of The Agreement Hierarchy (Corbett 1979, and much subsequent work). The Agreement Hierarchy describes how likely a particular element is to show either morphological or semantic agreement. Corbett shows that across a corpus the level of semantic agreement with hybrid nouns is higher for some targets than others.[1] 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:

(6) 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. The Agreement Hierarchy, in Corbett's discussion then describes the effect of differing frequencies of semantic and morphological agreement for different agreement targets, but also describes the implications regarding agreement targets.

What is interesting is that the 3/4 patterns that we have seen in English, Russian and Hebrew all fit with what we would predict if the Agreement Hierarchy were somehow operative at a deeper level, and regulated mismatches within a clause.[2] A mismatch is possible if the element to the right on the hierarchy shows semantic agreement and

[1] See also Levin (2001) for corpus data further cementing Corbett's conclusions.
[2] Indeed, Corbett (1983) considers such an approach, but ultimately rejects it.

the one on the left shows morphological agreement. However, it is possible to find instances of impossible mismatches that are not reflective of the Agreement Hierarchy. Landau (2016) notes that with the same noun *be'alim*, we find agreement mismatches internal to the DP, and hence both agreeing elements come from the 'attributive' slot on the hierarchy. In the following, there are two attributive adjectives that agree with the noun. In the first example, we see that the two adjectives are allowed to mismatch, in that one can show semantic agreement and one can show morphological agreement, however, only in case it is the adjective that is linearly closer to the noun that shows morphological agreement.

(7) 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/ hayu
the-owner the-private.SG the-last-PL of the-painting was.3.SG/ was.PL
ha-psixo'analitika'i Jacques Lacan.
the-psychoanalyst Jacques Lacan
'The last private owner of the painting was the psychoanalyst Jacques Lacan.'

Pesetsky (2013) notes the same pattern as Hebrew for Russian when there are two adjectives. 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 adjective that is further from the noun shows semantic agreement, and the closer one morphological agreement:[3]

(8) 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.'

Superficially, the same pattern is seen in Chichewa (Corbett 1991). Here, the noun *ngwazi* generally takes class 9 agreement, but can also show class 1 agreement. I treat the class 1 agreement as being reflective of semantics, since it is the default class of animates in the language, whilst class 9 reflects the morphlogical classification of the noun. Again, when there are two DP internal targets of agreement, they are allowed to show a mismatch only if it is the element that is closer to the noun that shows morphological agreement:

[3]I have altered the transliteration of Pesetky's minimally to be consistent with the Russian data given above

(9) 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.'

In all of the cases we have seen so far, when there are two targets of agreement, a 3/4 pattern arises. Of the four logical possibilities, 3 are grammatical, with the ungrammatical sentence being one of the mismatches. We can thus formulate the

(10) Condition on (im)possible mismatches (to be revised):
When a hybrid noun controls agreement on multiple elements in a clause, one of the elements can show semantic agreement and the other morphological agreement only if:
a. The element that shows semantic agreement is further right on the Agreement Hierarchy than the element that shows morphological agreement
b. If both elements are attributive, the element that appears linearly further away from hybrid noun shows semantic agreement.

Taken together, both of these statements suffice at a descriptive level, however, it is clear that we are missing something in terms of explanation. One might wonder about the possibility of reducing one of the statements to the other and look for the explanation from there, but this runs into problems. It is clear in the case of the DP-internal mismatches that we cannot appeal to the Agreement Hierarchy, since in both Hebrew and Chichewa we are dealing with multiple attributive elements. Since they come from the same slot on the hierarchy (10) makes no prediction. Conversely, it is not possible to appeal to linear order as an arbitrator for the other cases. Whilst it is true that the element that is linearly further away in English and Hebrew (the anaphor and verb respectively) shows semantic agreement, in the Russian pattern, both elements lie adjacent to the hybrid controller and so are equally near or equally far.

In the remainder of this paper, my aim is to show that the fact that both the Agreement Hierarchy parallels and the linear effects are epiphenomenal, and that all of the patterns are explained by the order in which the elements merge into the structure. As such, I will argue for the following:

(11) Condition on (im)possible mismatches:
When a hybrid noun controls agreement on multiple elements in a clause, mismatches between the targets are allowed only when the element that merged into the structure first shows semantic agreement.

I will return to this in section 3. Before moving onto an explanation of mismatches, it is worth reviewing some properties of semantic agreement, since they will form an important part of the analysis that follows.

2 Semantic Agreement

2.1 Not always available

Semantic agreement is not always possible for all potential controllers that we might expect to see it. Thus, whether semantic agreement is shown in a particular configuration depends on both the target and the controller. In this section I use the term agreement *target* as a pretheoretical term for the element that receives the agreement value, and agreement *controller* as the element that donates the agreement value. When I flesh out the theoretical analysis some more, I switch to the terms probe and goal for target and controller respectively.

Restricting our attention to English, we can see that the properties of the targets are important. As noted above, nouns of collections in English, are able to control singular (morphological) agreement or plural agreement (semantic). This is consistent across anaphoric, pronominal and verbal agreement. As pointed out by Elbourne (1999), demonstratives cannot ever show the semantic plural agreement (12-a), even if other agreement targets in the sentence show plural agreement, and therefore demonstratives must be singular (12-b). Thus, it is a property of demonstratives in English that they cannot show semantic agreement.[4]

(12) a. *These team are congratulating each other.
b. This team are congratulating each other.

Semantic agreement also depends of the type of controller. In order to observe semantic agreement, we must look at hybrid nouns where the morphological shape and semantic interpretation do not neatly line up, otherwise we are unable to tell whether an agreement value stems from the semantics or the morphology. CNPs allow us to observe semantic agreement in English, as discussed above, becuase they are morphologically singular but can be interpreted as plural collections. However, these are not the only class of hybrid nouns in English. Another candidate for hybrid nouns comes from *pluralia tantum* nouns, which are of course morphologically plural but can also refer to singular objects. However, even when referring to a single pair of scissors, (13-a) is not possible, and (13-b) must be used.

(13) a. *This scissors is what I need.
b. These scissors are what I need.

[4]Both den Dikken (2001) and Sauerland (2004) attempt to derive this fact from the internal structure of CNPs in English, however, as shown in Smith (to appear), both of these accounts face problems. It would be preferable of course to derive the restrictions on certain targets showing semantic or morphological agreement from something deeper, however, for the purposes of this article I treat it as a property of the agreement target.

That demonstratives cannot show plural agreement with *pluralia tantum* nouns in English is perhaps not surprising in light of (12-b). Yet, we know that English will allow for semantic agreement on verbs, as the verb can show semantic agreement with CNPs. Thus, the problem seems to be related to the controller of agreement, the pluralia tantum noun, not being able to be a controller of semantic agreement. Similar facts are seen with other hybrid nouns of English, Imposter constructions (Collins & Postal 2012) and *furniture*-mass nouns (Doetjes 1997, Chierchia 1998, Bale & Barner 2009), both of which are classes of hybrid nouns in English, yet neither allow for semantic agreement, see Smith (2015) for discussion. The conclusion to draw from this is that semantic agreement is not a property of a language, but rather is a property of both agreement targets, and classes of controllers (*cf.* Williams 1994).

2.2 Feature mismatches and a syntactic condition on semantic agreement

With this in mind, let us turn to how semantic agreement should be modelled when it does happen. Smith (2015, to appear) shows that the dual nature of hybrid nouns can be captured on the assumption that grammatical features are composed of two halves, one half a morphological *u*F that is legible to the morphology, and one half an *i*F that is legible to the semantics. Having features split into two parts allows us to capture the dual nature of hybrid nouns. Usually, for a given feature, the value that is on the *u*F half of that feature is matched by the value on the *i*F (14)a. However, in the case of hybrid nouns, the value that is carried on the *u*F is not matched by that carried on the *i*F (14)b.

(14) a. b.

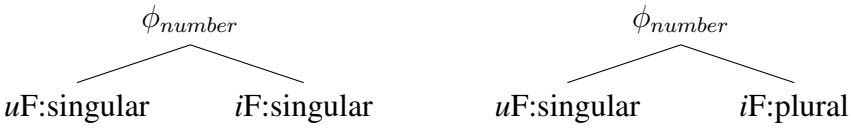

Thus, a hybrid noun like Russian *vrač* in the examples above will have a split on the gender feature as below. The masculine *u*F reflects its masculine morphology and the feminine *i*F reflects that the referent is female:

(15)

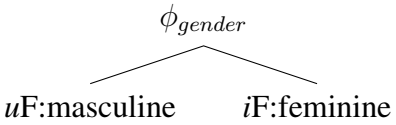

Smith assumes that both halves of the feature are present in the syntax, and then are split apart at the point of transfer. Since *u*Fs are the features legible to the morphology, they proceed along the PF branch and the *i*Fs going along the LF branch, being legible to the semantics.[5]

[5] Smith (2015, to appear) assumes a late insertion model of morphology like Distributed Morphology (Halle & Marantz 1993), where morphology interprets the output of syntax, rather than vice versa.

(16)

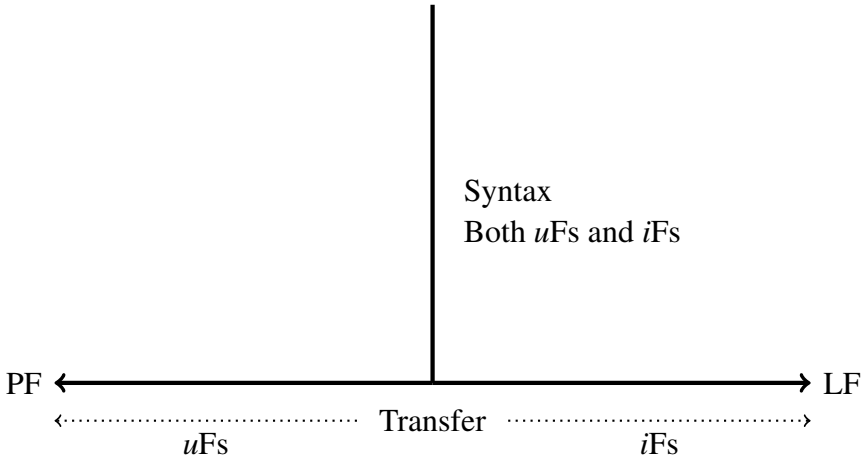

Semantic agreement arises when a controller of agreement has its *i*F active for agreement. Again, recall that this is not a property that holds of a language, but rather (classes of) nouns within a particular language.[6] However, semantic agreement is restricted in a further manner. The generalisation that arises in Smith's discussion of semantic agreement in English is that in order for agreement with an *i*F to be possible, it must be the case that the controller c-commands the probe at the level of LF. In minimalist terms, semantic agreement is restricted to a Reverse Agree configuration (see Wurmbrand 2011, Zeijlstra 2012 on Reverse Agree apporoaches) The relevant descriptive generalisation that Smith provides evidence for is *LF-Visibility*, which covers restrictions on plural agreement in existential constructions and scope reconstruction (Elbourne 1999, Sauerland & Elbourne 2002) and difference to do with whether agreement targets the subject or a predicate of a predication structure (den Dikken 2001).

(17) LF-visibility (descriptive generalisation)
With CNPs, plural (= semantic) agreement requires the controller to c-command the target at LF, but singular (morphological) agreement does not.

Smith derives this generalisation through assuming the following model of agreement. He adopts the two-step agreement approach of Arregi & Nevins (2012), (see also Franck et al. 2006, 2008, Benmamoun et al. 2009, Bhatt & Walkow 2013, Marušič et al. 2015) whereby AGREE is composed of AGREE-LINK, which links the probe and the goal, and AGREE-COPY which copies the feature value from the goal to the probe.[7] The relevant part for Smith (to appear) is the formulation of AGREE-COPY, which in such approaches (especially Arregi & Nevins 2012) is taken to be a post-syntactic process. However, assuming the distribution of features given in (16), then *i*Fs are not present in the PF branch. Smith minimally modifies the proposal of Arregi & Nevins, by adding the clause (b.i.) in (18), which means that AGREE-COPY is largely in the

[6] Indeed, for nouns that do allow for semantic agreement, it is not even the case that they always show semantic agreement. In English, it is an open choice whether agreement with a CNP will be morphological or semantic, with the choice influenced by vary considerations, such as register and individual lexical item (see Levin (2001) for corpus evidence).

[7] My thanks to an anonymous reviewer for pointing me towards Franck et al's work.

PF branch, but can also happen at the point of transfer.[8] This condition ensures that agreement that is still part of the syntax proper can look only upwards, whereas agreement looking both up and down is possible in PF (see Wurmbrand 2011, Zeijlstra 2012, Bjorkman & Zeijlstra 2015 for convergent evidence).[9]

(18) Agreement by Probe with Goal proceeds in two steps:

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

b. AGREE-COPY: After the syntactic derivation, the values of the ϕ-features of the goal are copied onto probe linked to it by AGREE-LINK.

i. if AGREE-COPY happens at the point of transfer, this requires that goal c-command the probe.

This model of agreement, superimposed onto (16) gives the following picture of AGREE:

(19)

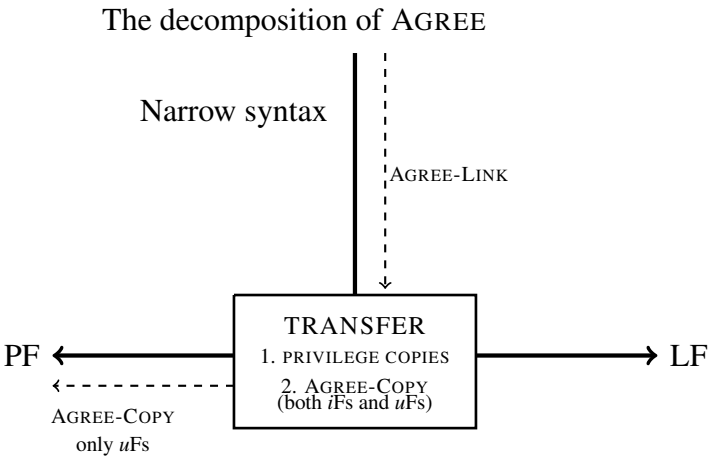

Thus, *LF-Visibility* is derived, albeit in a brute force manner, as semantic agreement is agreement at the point of transfer, and AGREE-COPY at this point can look only upwards in the structure. As will become evident through the rest of the paper, and hinted at in (11), this model of agreement allows us to resolve a tension that semantic agreement raises for models of AGREE. On the one hand, the conclusions of Smith (to appear) shows that semantic is in some sense representational, since it is sensitive to the LF position of the controller. On the other hand, 3/4 patterns will require a derivational view of AGREE whereby it takes place throughout the syntactic derivation (see also Landau 2016 for a derivational approach to some 3/4 patterns). Having a two-step model of agreement as in (18) resolves this tension: AGREE-LINK is the derivational part of AGREE, whilst AGREE-COPY is the representational part, operating on the output of syntax.[10]

[8] In the single-output model that Smith assumes, also corresponds to LF, see also Bobaljik 1995, 2002

[9] I focus in this paper in only using AGREE for morphological agreement, and binding.

[10] That AGREE operates on the output of syntax is also motivated in studies on closest conjunct agreement (Benmamoun et al. 2009, Bhatt & Walkow 2013, Marušič et al. 2015), whereby agreement is

3 Mismatches that reflect The Agreement Hierarchy

I now turn to my account of the 3/4 patterns that were introduced earlier. Firstly, I will account for the patterns that we saw earlier to straightforwardly reflect Corbett's Agreement Hierarchy. Recall that I will argue that this similarity is epiphenomenal and therefore ascribe no formal role to the Agreement Hierarchy to handling the 3/4 patterns under discussion. After these have been discussed I will discuss how my account compares to existing approaches, before turning to the other patterns.

3.1 A bleeding effect on *i*F agreement

Assuming Smith's characterisation of hybrid nouns in terms of a conflict between the values on the *i*F and the *u*F, 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. However, the question that must be answered is why this situation is restricted, such that only one of the mismatches are allowed.

I propose that mismatches happen when an *i*F that is active for agreement is deactivated in the course of the derivation, such that a target undergoing agreement before deactivation shows semantic agreement, whereas targets afterwards do not. Thus, deactivation of an *i*F bleeds semantic agreement on later targets that could otherwise have undergone semantic agreement. Furthermore, I will assume that this deactiviation is a one way process; a *i*F that is inactive for agreement cannot be reactivated. If 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. We can view this schematically in the following derivation:

(20)

Step	Controller	Target 1 Value	Target 2 Value	
1.	[*u*F:▲,*i*F:■]	–	–	
2.	[*u*F:▲,*i*F:■]	■	–	Agreement
3.	[*u*F:▲,*i*F:■]	■	–	*i*F deactivation
4.	[*u*F:▲,*i*F:■]	■	▲	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:

(21) An active *i*F cannot be ignored by AGREE-LINK.

Put together, this gives us exactly the bleeding effect that we need to model 3/4 patterns. If the *i*F on the controller is deactivated before both targets undergo agreement, then the result is matching morphological agreement. If the *i*F remains active throughout, then the result is matching semantic agreement. However, if the *i*F is deactivated after Target 1, but crucially before Target 2, then a mismatch arises such that the first target shows semantic agreement, and the second target morphological agreement. Crucially, it is

sensitive to linear order, and not hierarchical structure.

not possible for Target 2 to show semantic agreement and Target 1 morphological. This would require either Target 1 ignoring an active *i*F (not possible), or an *i*F that is inactive at the point of Target 1 agreement to become active for Target 2 (ruled out by (21)).

Now, this begs the question of how an active feature becomes deactivated. I propose that the deactivation happens through AGREE-LINK. Once an *i*F has entered into an AGREE-LINK operation, then it can become inactive for further operations. Note the use of *can*: this must be an optional process. If it were obligatory, then we would expect there to be maximally one target showing semantic agreement throughout the derivation, which is clearly wrong, given that it is possible for there to be matching semantic agreement across two (or more targets).

Optional deactivation through AGREE-LINK may come across as an *ad hoc* stipulation, and to some extent this is true. However, it should be noted that I am only proposing that this is the case for *i*Fs, but not as a general property of features. This then makes *i*Fs special in comparison to *u*Fs. Though somewhat undesirable that *i*Fs would display properties different from *u*Fs, I believe that it is a burden worth tolerating, given that it allows us to have a unified account of all the patterns in this paper.[11] Furthermore, it may well turn out to be the case that *i*Fs are somewhat special in their behaviour. Semantic features are not the canonical features that are used in agreement and the process seems to be a peripheral one in a language; as noted above, it is only CNPs among the group of hybrid nouns that show semantic agreement in English, as other hybrid nouns like Imposters do not.

That being said, there will doubtless remain readers who are uneasy at the thought of there being optional deactivation of *i*Fs (and only *i*Fs). In section 6, I present an alternative that allows us to have a stronger condition on deactivation of *i*Fs, such that it is obligatorily applies whenever AGREE-LINK targets an *i*F, which would bring *i*F more in line with the behaviour of other morphosyntactic features. However, as I discuss in detail there, this comes with with its own downsides. I defer this discussion until that point however, and proceed with optional deactivation.

3.2 Deriving the bleeding effect

As mentioned above, I will first deal with the mismatches that transparently reflect Corbett's Agreement Hierarchy. Key to the above proposal is that *i*F agreement can be bled by deactivation of the *i*F on the hybrid noun. Thus, in order to capture the pattern that we see in English, it needs to be the case that the anaphor will undergo AGREE-LINK before T.

Consider first English, with the data repeated below:

(22) 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.

[11] It is my hope that further research on the nature of semantic agreement would shed some deeper light on why the different behaviour arises.

The mismatch that is allowed in English is where the anaphor shows semantic agreement and the verb morphological agreement. I propose that probes link to their goals as soon as they can, respecting the ideas of Earliness in Pesetsky (1989), Chomsky (2001). The first derivational point that the anaphor can undergo AGREE-LINK with a CNP controller in the above sentences is when the CNP is merged in Spec,vP (following Sportiche 1988, Koopman & Sportiche 1991). Thus, I assume that AGREE-LINK happens between anaphor and CNP at that point. Notably, this is before the first derivational point that T (the locus of verbal agreement features in English, following standard assumptions) and the CNP share the structure and are able to undergo AGREE-LINK.

The result of this is that it gives us exactly what we need to capture why anaphors are able to undergo semantic agreement and T morphological: the *i*F is deactivated at the point that the anaphor undergoes AGREE with the CNP. If the *i*F is not deactivated at this point, then matching semantic agreement would arise, and if the *i*F is never active, then matching morphological agreement would arise. These are the only three options. Due to a combination of the fact that inactive *i*Fs are not able to become active, and that they cannot be ignored when they are active, then it is not possible for the anaphor to undergo AGREE-LINK to the *u*F of the controller, and then T undergo AGREE-LINK to the *i*F on the CNP. The structure (23) represents the (abbreviated) structure for the English mismatches, with the steps above conflated into one structure.[12]

(23)

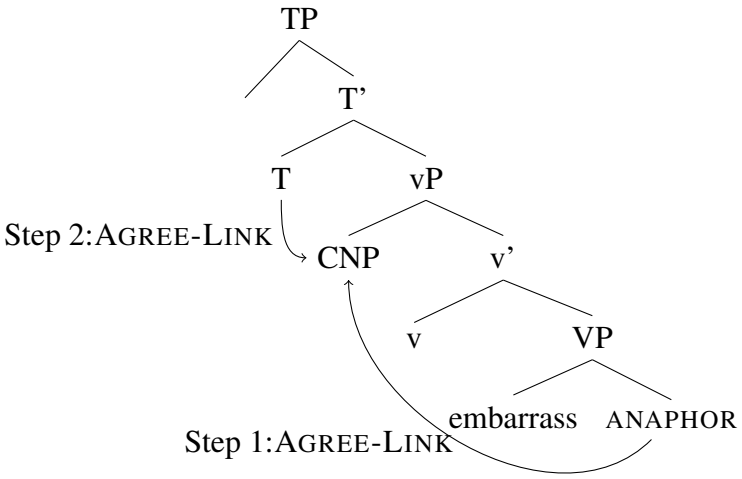

It should be noted that this assumes that binding is a direct process between the antecedent and the anaphor (Hicks 2009, Wurmbrand 2012), rather than one mediated through functional heads, as some prominent accounts have assumed (e.g. Reuland 2001, 2005, 2011, Kratzer 2009). It is crucial here that anaphors establish a binding relationship with their binder directly, since this allows for AGREE-LINK to deactivate the feature on the CNP before the CNP undergoes agreement with T. Mediating the binding relationship through T, or functional heads connected to T would give, if anything, the

[12] 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 to future research.

unattested mismatch. As it happens, the mismatching effects seen with hybrid nouns seem to argue in favour of a direct connection between the anaphor and the antecedent, since the values on T and the anaphor are able to diverge. This is hard to explain if T takes features from the antecedent and passes them onto the anaphor.[13]

3.3 Russian and Hebrew

Turning to Russian and Hebrew, following the above logic, it would seem that adjectives merge into the structure after T. The timing effect was fairly straightforward above, since anaphors canonically appear in object positions. However, the timing is less obvious to derive between attributive adjectives and verbs. In order to answer this, I adopt the proposal of Stepanov (2001), where 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. 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. I will further assume that attributive adjectives are always adjuncts, and hence always merge late.

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 indeed 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.

What this means for our purposes is that adjectives in attributive position will merge into the structure countercyclically, crucially after T has done. I illustrate this with discussion of the Russian *vrač* paradigm, paying attention only to the instances of mismatch. The Hebrew data in (5), receive an identical explanation, and for space I leave out discussion.

(24) 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.'

In step 1, T has merged into the structure, and undergoes AGREE-LINK with *vrač*. In Step 2, the adjective merges into the derivation late, and at this point will undergo

[13]Note that it is also crucial in this approach that the anaphor undergoes agreement first, before T, and is thus at odds with the approach to binding offered in Rooryck & Wyngaerd (2011).

AGREE-LINK with *vrač*. This allows us to capture the fact that a mismatch is allowed when it is the verb that shows semantic agreement and the attributive adjective morphological agreement, since AGREE-LINK happens between the verb and the hybrid noun before it does between the adjective and the hybrid noun. There is therefore the potential for the *i*F on the hybrid noun to be deactivated between these two steps, and for the attested mismatch to arise. Again, it is important to stress that this will not be able to generate the unattested mismatch.

(25) a. Step 1

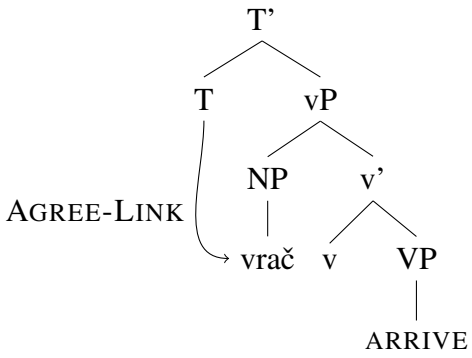

b. Step 2

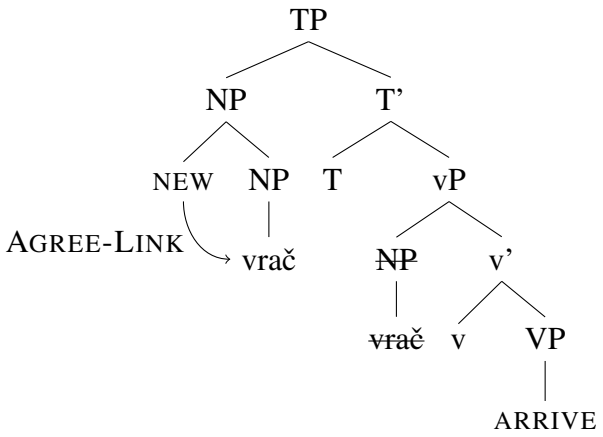

Throughout the discussion to this point, the crucial aspect of ruling in the attested mismatches, and ruling out the unattested mismatches has been that the *i*F on the goal is deactiviated after the first probe has undergone AGREE-LINK. The fact that the cases to this point have all reflected the Agreement Hierarchy has in fact been epiphenomenal. In the anaphor-verb case of Engllish, the anaphor merges into the structure before the verb does, and in the adjective-verb cases of Russian and Hebrew, the adjective merges after the verb by virtue of being an adjunct. This is in no way regulated by the Agreement Hierarchy. Before moving onto the DP-internal cases of Hebrew and Chichewa, it is worth contrasting this approach with other ones that have given a more direct role to the Agreement Hierarchy.

4 Other approaches to 3/4 patterns

4.1 Wechsler & Zlatić (2003): CONCORD and INDEX features

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. Both were noted as early as Corbett (1979) and have been discussed in works since then.

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 some 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 (2000, 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 declensional 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 (26). 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.

(26) a. Regular case:

b. Possible:

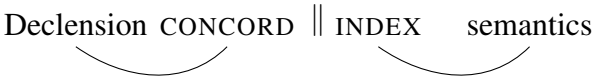

c. Not possible:

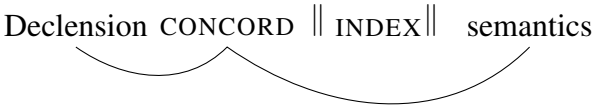

Wechsler & Zlatić discuss the Agreement Hierarchy and how this can be produced at a corpus level by their system. CONCORD and INDEX are on different featural paths (within HPSG). Relative pronouns can differ from personal pronouns, because, whilst relative pronouns always undergo anaphoric binding (agreeing with INDEX), 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. The same holds for predicate agreement, which can comes from the INDEX feature.

Wechsler & Zlatić note that their theory does not directly 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. Put in their terms, the question is 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 only to the attributive/predicate distinction, since it does not work with relative pronouns. Thus through a mixture of diachronic development, and synchronic locality restrictions they are able to model the relative frequencies per category of semantic and morphological agreement.

Wechsler & Zlatić note that the monotonic effect of Agreement Hierarchy follows from their account, because of the system of contiguous constraints 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.

Despite these successes, there are a couple of problems. Whilst their approach can in principle work for instances where the mismatching targets come from different slots on the agreement hierarchy as shown above, it faces a couple of problems with the full range of patterns.

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.

Putting these quibbles aside for the time being, let's grant that there is some mechanism that makes pronouns more likely INDEX targets than predicates, whatever that may

be. Wechsler & Zlatić's explanation is then suitable for sentences like the following:

(27) *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). Thus, (27) is not able to be generated. Wechsler & Zlatić's approach works here precisely because attributive elements in English can only undergo agreement with CONCORD features, whereas verbs can agree with either CONCORD or INDEX. 3/4 patterns of the type in (27) then follow.

However, Wechsler & Zlatić's proposal apparently faces difficulties when there are two (or more) targets that can inflect for either INDEX or 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 featural structure):

(28) [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 3/4 patterns arise in the first place. Why in (29-b), is the attributive element not able to agree with the INDEX value, and the predicate with the CONCORD value.

(29) 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. Yet this simply rebuilds the Agreement Hierarchy into speaker's grammatical knowledge.

On the other hand, sentences like (29-b) follow from the present account without further stipulation in addition to the other patterns that Wechsler & Zlatić can capture. 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 there are only two potential agreement values - one from the *i*F and one from the *u*F. 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 and it should be possible to transpose the account here based on derivational timing into HPSG terms.

A second problem however for Wechsler & Zlatić's approach is that 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. One could argue that syntactic locality can arbitrate here - in both Hebrew and Chichewa the target with morphological agreement is linearly and closer to the controller. Yet, 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.

4.2 Configurational approaches

Pesetsky (2013) and Landau (2016) offer an explanation here and a way to reconcile the DP-internal cases (8) and (7) with Wechsler & Zlatić's approach. I will explicate with the explanation of Landau, but the logic of Pesetsky's analysis runs the same way. Recall that for Hebrew, a mismatch is allowed where the peripheral adjective can show semantic (singular) agreement and the inner adjective morphological agreement. 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, Kramer 2014, Smith 2015 for proposals where inherent features are located on category defining nodes), 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 happens as soon as elements are merged into the structure, then the lower adjective will show morphological agreement, as it has merged before the semantic information is introduced, and the higher element will show semantic agreement. The unattested mismatch cannot be generated by intervention, once NumP is in the structure, then it intervenes betwen higher targets and N, preventing anything from looking past it and taking a value from the CONCORD feature on N. The mismatch case is illustrated below:[14]

[14]Linear order is ignored in the tree in (30).

(30)

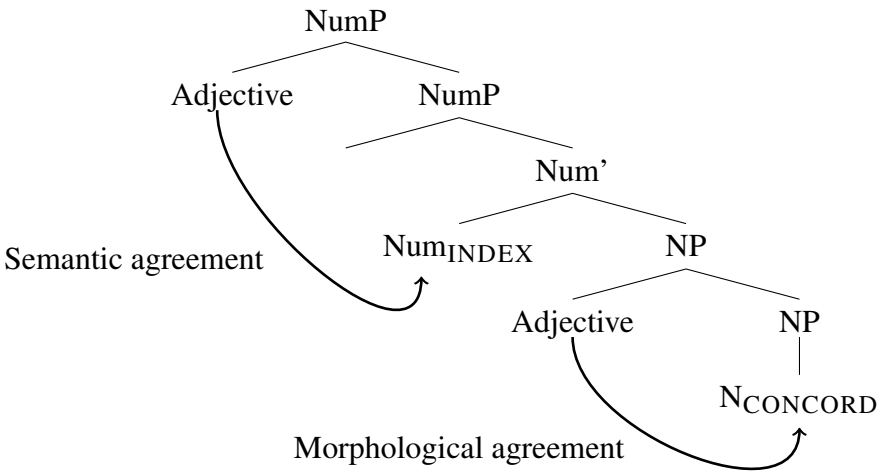

Landau extends this account to the Chichewa data with the same explanation. It is the higher element that shows semantic agreement, and the lower morphological. This relies on the conclusion that Chichewa shares the same DP-internal structure as Hebrew, that it is leftward branching and elements closer to the noun are structurally lower than elements that are linearly further away. As I will shortly discuss, this conclusion has been argued to be incorrect, and as such causes a significant problem with adopting the configurational approach for Chichewa.

Before moving on I would like to make one final point regarding the configurational approach that Pesetsky (2013) and Landau (2016) offer. Though they offer elegant solutions to the 3/4 patterns that exist in Russian and Hebrew, it is not immediately clear to me how the proposals can be extended to account for the 3/4 pattern of English. A configurational account can be extended to instances where one probe is DP-internal and one probe is DP-external. In a 3/4 pattern where the predicate shows semantic agreement, and the attributive element morphological (the Russian case of (24) and the Hebrew case of (5)), the logic of both Pesetsky's and Landau's is that the DP-internal element has merged into the structure lower than the semantic specification, and thus can only agree with the morphological information inherent to the noun. On the other hand, semantic information is transmitted up to DP and is visible to DP-external agreement. Effectively then, elements that show morphological agreement are in the wrong position to show semantic agreement, as they have merged beneath either the NumP in Hebrew, or beneath Pesetsky's feminising head. The problem of English is that we have two instances of DP-external agreement, and so it is not clear how a configurational account would rule in favour of the verb showing morphological agreement and the anaphor showing semantic agreement, and not for instance the other way around. Landau acknowledges this, and suggests that different 3/4 patterns may have different explanations, though it is my contention in this paper that they can all be unified.[15]

[15] As noted in footnote **??**, a configurational account has been offered by den Dikken (2001), Sauerland (2004) to a subpart of the English paradigm, notably why demonstratives can only show morphological agreement with a CNP even when the verb shows semantic agreement. Again, however, it is not clear how these accounts can help if the agreements in question are both DP-external (see also Smith to appear for independent problems with such accounts in English).

5 Different sites of late merge

In section 3, I outlined an approach which allows us to explain why certain mismatches between semantic and morphological agreement are allowed, and others are not. Key to the approach is that it is the element that merges first into the structure that shows semantic agreement, and the latter one morphological, but the system is such that the opposite is unable to be generated.

3/4 patterns that arise with elements that come from the same slot on the Agreement Hierarchy present a different type of challenge, but as I outline below, they fall under the same approach. The descriptive generalisation for these cases identified in section 1 was the (31-b):

(31) Condition on (im)possible mismatches (to be revised):
When a hybrid noun controls agreement on multiple elements in a clause, one of the elements can show semantic agreement and the other morphological agreement only if:
   a. The element that shows semantic agreement is further right on the Agreement Hierarchy than the element that shows morphological agreement
   b. If both elements are attributive, the element that appears linearly further away from hybrid noun shows semantic agreement.

It should come as no surprise to the reader that what I will propose that underlies this effect is that the DP-internal element in these examples has been merged into the structure first, not through any linear preference for closer elements to receive morphological agreement.[16]

However, before proceeding with the explanation, it is important to look into more detail at the DP-internal structure of Hebrew and Chichewa. 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)

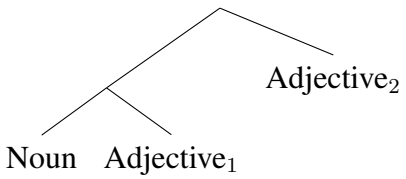

This rightward-is-highest structure is supported for Hebrew by the data in table 1 (from Shlonsky 2004) who shows that the order of Hebrew adjectives is consistently the opposite order than English.

—

Table 1

—

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 ad-

[16] I have shown above that linear effects cannot underlie the general account of 3/4 patterns.

jectives, Hebrew has a leftest-is-lowest order.[17] 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:

(33) picha hii yangu nzuri ya Busi
9.picture 9.this 9.my 9.good 9.of
'this nice picture of mine of Busi [KiSwahili]

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 languages that have the DP-internal word order of English are representative of the basic functional order (Cinque 2005) without movements to reorder elements, then we can conclude from the fact that they show the same order as English, that the postnominal elements in the unmarked word order in KiSwahili 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.

Chichewa possessive pronouns, like in KiSwahili, generally precede ordinal numerals (Sam Mchombo, p.c.). Assuming a right branching DP structure, they are then structurally higher. I thus assume the structure in (34) 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 convenience.[18] 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.[19]

[17] 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).

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

[19] 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

(34)

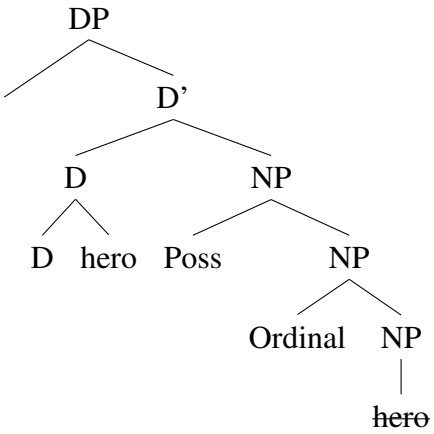

Hebrew and Chichewa then differ in one crucial respect. Assuming that the difference between the two languages is correct in that Hebrew DPs branch leftward whilst Chichewa DPs branch rightward, then it is the structurally higher adjunct that shows semantic agreement in Hebrew, but the structurally lower adjunct that shows semantic agreement in Chichewa. This raises two issues. Firstly, the higher modifier showing morphological agreement and the lower semantic does not sit well with Landau's approach summarised in the above section. since that rested on the idea that elements merged lower in the structure are the ones to show morphological agreement, since their merge preceded the merge of NumP, which carried semantic information.[20] Secondly, if we are to use the same explanation that I am putting forward in this paper, 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. In the current proposal, AGREE-LINK happens as soon as the probe and goal are in the same structure. This, the higher adjunct must have merged before the lower one in Hebrew, whereas in Chichewa, the lower adjunct must have merged before the higher one.

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.

[20]In order to use the same configurational explanation, it would need to be the case that information that is inherent to the Chichewa noun is introduced higher than non-inherent information, however, this runs against the general approach that inherent information is either carried on the lexical root, or introduced on category defining nodes, see Acquaviva (2008), Kramer (2014), Smith (2015) for discussion.

5.1 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, in the sense that there seems no reason to assume that all adjunction mirrors a bottom-up syntactic derivation and merges from lowest to highest.

I propose here that this is a point of variation 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.

(35) In case of adjunction, adjoin to the {HIGHEST/LOWEST} segment of the adjunction site.

An anonymous reviewer of an earlier version of this paper points out that this is effectively giving (35) a parametric status in language and questions what kind of parameter it is, given that parameters are usually conceived of being properties of functional heads since Borer (1984). (35) does not fit into that characterisation, as it is hard to see how it would be related to a functional head. Rather, I propose that (35) is a *grammatical parameter* in the sense of Baker (2008), who writes "[s]omething deserves to be called a grammatical parameter, [...] if it is a relatively general feature of the language, not one that is tied to a particular head or construction (Baker 2008:155)."

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, 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. Assuming that English linearises its DP structure as higher elements being on the left, then it must be the case that *red* merges first into the structure, before *big*, which later merges above *red*. IThe 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.

(36) a. *Build NP*

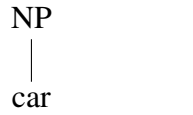

b. *Merge* red *with highest segment*

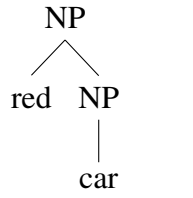

c. *Merge* big *with highest segment*

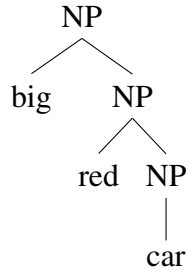

With the HIGHEST option selected, if *big* were to merge first, then a further merge of *red* would lead it to being structurally higher, since it merges at the top segment. This is show in (37), where *big* would merge into the structure before *red*, we end up with the ungrammatical **red big car*.

(37) a. *Build NP*

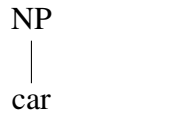

b. *Merge* big *with highest segment*

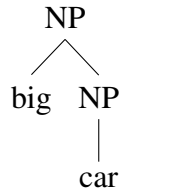

c. *Merge* red *with highest segment*

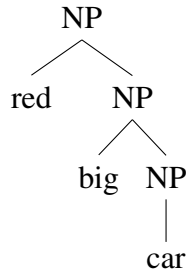

Now consider if Englilsh were instead a LOWEST language, where adjuncts target the *lowest* segment of the adjunction site.[21] If *red* merges first, the segment attached

[21] I do not commit in this paper to what choice English makes, since it is not important for our purposes. I use English examples solely for expository convenience.

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*:

(38) a. *Build NP*

b. *Merge* red *with highest segment*

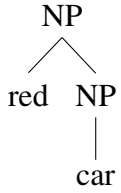

c. *Merge* big *with highest segment*

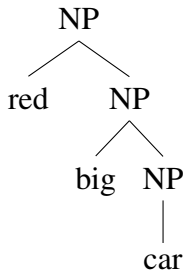

The right result would be obtained in a LOWEST language if adjectives that are structurally highest have merged before adjectives are structurally lowest. In this derivation *big* merges before *red*, and the correct output is generated.

(39) a. *Build NP*

b. *Merge* big *with lowest segment*

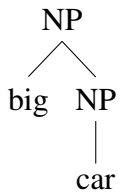

c. *Merge* red *with lowest segment*

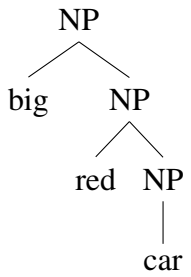

Thus, depending on the choice that a language makes regarding the adjunction site (35), this will have an impact on whether structurally higher adjuncts have merged before or after structurally lower ones. In summary:

(40) 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 that we need in order to account for the 3/4 patterns in Hebrew and Chichewa even considering their differing DP-internal structure. In order to rule out the unattested patterns, it must be the case that languages make a choice with respect to (35) such that the opposite order of merge is not possible in the language.

5.2 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:

(41) 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 adjectives 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 can be deactivated on the noun, leaving only the *u*F available for further iterations of 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. Thus only mismatch that can be generated is where the structurally higher one shows semantic agreement, and the lower one morphological. Matching agreements are derived in the same way as above, where the *i*F remains active for both targets (matching semantic agreement, or is inactive for both (matching morphological agreement).

(42) a. Step 1

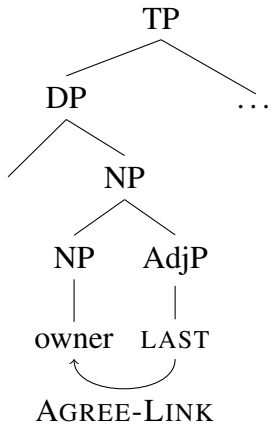

b. Step 2

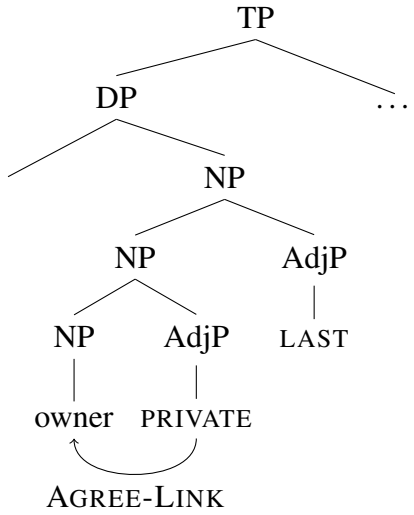

I assume that the same analysis holds for the Russian data in (8) as in Hebrew.[22]

[22]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. As to why these adjectives show this property, I leave to future research.

5.3 Chichewa

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

(43) 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 possessive 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.

(44) a. Step 1

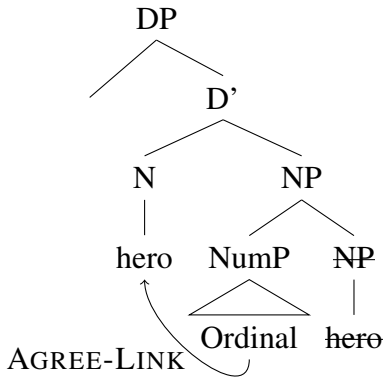

b. Step 2

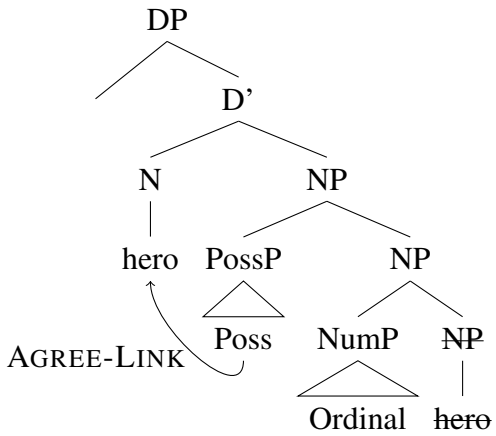

This enables a derivation whereby (i) the ordinal numeral merges first into the structure and showing semantic agreement, (ii) the *i*F on the controller is deactivated, and (iii) the possessive pronoun merges and undergoes morphological agreement.

5.4 Summary

In this section I have shown that we can generalise the proposal given in section 3 to account for all 3/4 patterns, 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 DP-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 DP-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.

6 The deactivation of an *i*F

Throughout this paper I have assumed that 3/4 patterns arise through a bleeding effect, whereby deactivation of an *i*F in the derivation renders further semantic agreement impossible. This allowed for mismatches to be derived through the proposal that elements that merge first into the structure show semantic agreement, and latter ones show morphological agreement.

As I have noted, this has allowed for a unified account of all of the 3/4 patterns, and a principled account of why certain mismatches in agreement are allowed and certain others are not. However, as I also noted in section 3, this has also been predicated on the optional deactivation of an *i*F through the process of AGREE-LINK. I noted earlier the controversy over assuming that an *i*F is optionally deleted, arguing that it need not be thought of as a general property of features (and hence smuggling optional operations into Minimalist syntax through the back door) but rather something that is limited to *i*Fs, which may well simply have special behaviour with regard to agreement.

However, it remains true that not all readers will be comfortable with such a scenario; optional operations have become frowned upon in Minimalist syntax (to put things mildly) a viewpoint that has come about due to the desire for features and satisfaction of their requirements to be the driving force of syntax. Allowing optional operations does not sit well at all with such a model. Thus, in light of this, it would be preferable to remodel my condition on deactivation to something like the following:

(45) An *i*F targeted by AGREE-LINK is deactivated.

Before concluding the paper, I wish to briefly explore this proposal and show that we can hold (45) to be true, whilst capturing the 3/4 patterns in the same manner as I have done in the rest of the paper. However, the conclusion will be that we *can* adopt (45), but it's not so clear that we *should*.

Optionality can be done away with if we assume that AGREE-LINK can take place whenever it wants to, but does not have to take place immediately when controller and target are in the derivation. Rather, at some point in the derivation, an operation

of AGREE-LINK applies and prompts probes to link to their goals. Effectively, this conceives of AGREE-LINK (and by proxy, AGREE more generally) not as an operation that takes place at the immediate behest of features that need a value, but rather as an independent operation that applies in order to link goals to their probes. The difference between the two is subtle, but the latter perspective frees AGREE-LINK from applying immediately.[23]

Now, let's suppose that this is the case, assume that when it does apply, every probe that is in the structure will be linked to its goal. *Per* (45) any *i*F on a goal that is targeted by AGREE-LINK will then become inactive for future iterations of AGREE-LINK. Now, if we further assume that AGREE-LINK has the ability to operate as many times as it needs to so that all goals have had the opportunity to undergo agreement, then this will have the result that elements that merge after the point of (45) will have to show morphological agreement, since the probes can only link to the *u*F of the goal. Thus, the 3/4 patterns are derived in the same way: earlier elements can and must show semantic agreement and all elements merged after (45) can only show morphological agreement.[24]

With these assumptions in mind, let us then revisit the 3/4 patterns that are of interest to us. I will use English as an example for exposition, but will not outline all of 3/4 patterns here for reasons of space. I invite sceptical readers to go through the other derivations themselves and check that the same facts are accounted for.

Using the same, bottom-up derivation assumed above, then the anaphor merges into the derivation before the CNP, which in turn merges before T. There are thus two relevant points of AGREE-LINK to consider.

Firstly, suppose that AGREE-LINK takes place at the point that only the CNP and the anaphor are in the structure. Assuming that the *i*F on the CNP is active, this will mean that the anaphor is linked to the *i*F and will show semantic agreement once AGREE-COPY takes place. The *i*F is then deactivated *per* (45). Now, once T merges into the structure, it will undergo AGREE-LINK (potentially not initially, but at some point before the end of the syntactic derivation). However, due to the earlier inactiviation of the *i*F on the CNP, then this time it can only link to the *u*F of the CNP and will show morphological agreement. Thus, we derive the attested mismatch.

[23] AGREE then is an operation that is not driven by features, but rather affects them.

[24] Note that I am not claiming that AGREE-LINK necessarily links all features to a target, but rather each probe has undergone AGREE-LINK. Preminger (2011, 2015) has shown convincingly that agreement can fail. What is relevant for my purposes is that AGREE-LINK happens wherever it can.

(46)

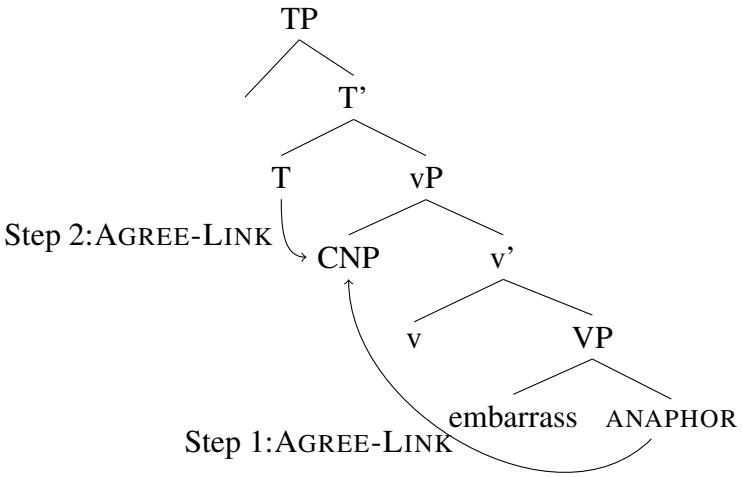

Another option is that the anaphor does not undergo AGREE-LINK before T has merged into the structure. Because AGREE-LINK requires every probe that is in the structure to link to their goal, this means that the anaphor and T will simultaneously link to the CNP. Assuming the *i*F to be active, this will have the effect that both of these are linked to the *i*F and will both show semantic agreement once AGREE-COPY applies. Matching morphological agreement will apply in case the CNP enters the derivation without its *i*F active.[25] This derivation is shown below, and note that there is only one step of AGREE-LINK, reflecting the fact that the anaphor and T undergo AGREE-LINK simultaneously.

(47)

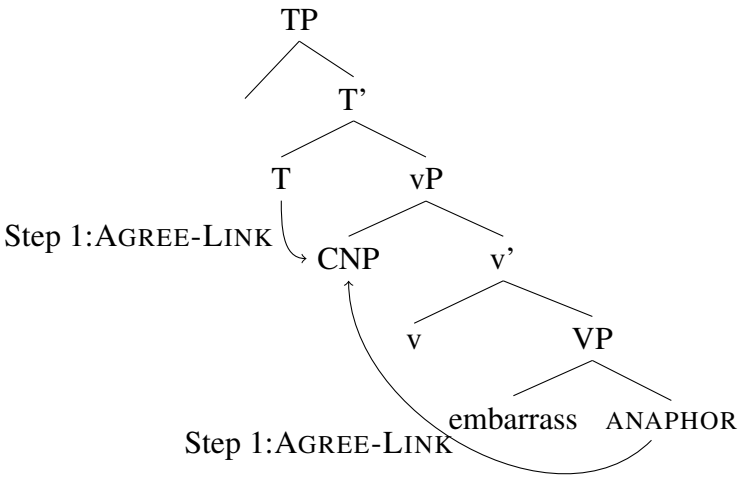

Importantly, such an approach will retain the insight of the rest of the paper that the fact that the anaphor merges before T, and this is what allows for the attested mismatch, but rules out the unattested one. Because AGREE-LINK must apply to all probes that

[25]Recall from above that it is neither necessary nor desirable for all items that potentially control semantic agreement to make their *i*Fs active at the beginning of the derivation, as this would have the effect that there would necessarily be always one element in the structure that shows semantic agreement.

are currently in the structure, then it is not possible for elements that have merged later to undergo semantic agreement, and elements that have merged earlier to undergo morphological agreement. Thus 3/4 patterns, remain derived through the same perspective as we have been exploring throughout the rest of the paper.

However, despite the fact that we are able to do away with the optional deactivation of an *i*F, we are left with having to assume that AGREE-LINK operates somewhat independently from the immediate needs of probes, and allows them to link to a goal only when it decides to apply, effectively giving up on the idea of Earliness (Pesetsky 1989, Chomsky 2001). It has been shown that there is some indeterminacy with respect to Earliness, and that AGREE may not apply directly at the moment that an unvalued feature is inserted into the derivation and there is a body of work emerging that explores the idea that the order of operations is open to some parameterisation, which has produced very interesting results (Müller 2009, Georgi 2014, Assmann et al. 2015, Puškar & Murphy 2015). This proposal is a more extreme version of what is given in these references however, since these works assume that whilst the order of operations is open to some flexibility, the operations that are triggered are still triggered item by item throughout the derivation. The account considered here treats AGREE as an operation independent from each individual features and items and runs the risk of being far too permissive. For instance, it then becomes unclear how to handle instances of opacity that are related to agreement, such as defective intervention effects (Chomsky 2001), if there is the potential for AGREE to apply to multiple elements all at the same time. For this worry I do not wish to advocate for the approach considered in this section, but I offer it as a potential way to alleviate worries concerning *i*Fs being optionally deactivated.

7 Conclusion

This paper has been concerned with various 3/4 patterns that arise with mismatches between morphological and semantic agreement. We began the paper with the observation that not all mismatches in agreement are equal, with some allowed and others not allowed. Throughout, I have outlined an approach which unifies all the 3/4 patterns under the same account, and thus provides us with a general account of impossible agreement mismatches, a welcome result. It remains somewhat mysterious why semantic agreement should act in the way that it does, especially in regards to the discussion surrounding optional deactivation. However, we can draw some more general conclusions from the aforementioned discussion.

Firstly, semantic features have a somewhat special behaviour when they figure in to agreement, in that if this approach is on the right lines, they cannot be ignored for agreement. Other approaches such as Pesetsky (2013) and Landau (2016) have implemented this in different ways, notably a configurational approach where the head that carries semantic information intervenes between probe and the inherent morphological information inherent to the hybrid noun. Though these approaches are perhaps more elegant in explaining the exceptionality of semantic agreement than that presented here, for reasons discussed in section **??**, such a proposal does not allow for the unification of patterns that the present account does. It remains to be seen whether the exceptionality of semantic agreement can be deduced from deeper principles, and a reviewer suggests

that feature geometries may provide some explanation. For reasons of space I have not explored that approach in this paper, and I leave it open, but it could be an interesting direction to pursue.

A second general conclusion is that semantic agreement provides further support for a two-step model of AGREE, bifurcated into AGREE-LINK and AGREE-COPY. This two-step model has allowed us to capture the observation of Smith (to appear) that semantic agreement is shows representational characteristics, combined with the clearly derivational nature of semantic agreement discussed in this paper. Given that a lot of previous support (though not all) for the two-step model of AGREE has come from closest conjunct agreement, this paper has offered a novel perspective on the debate.

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.

Assmann, Anke, Doreen Georgi, Fabian Heck, Timo Klein, & Philipp Weisser (2015) Ergatives move too early: On an instance of opacity in syntax. *Syntax* **18**(4): 343–387.

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

Bale, Alan C. & David Barner (2009) The Interpretation of Functional Heads: Using Comparatives to Explore the Mass/Count Distinction. *Journal of Semantics* **26**(3): 217–252.

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

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.

Borer, Hagit (1984) *Parametric syntax: case studies in Semitic and Romance languages*. Dordrecht: Foris.

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

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.

Chierchia, Gennaro (1998) Plurality of mass nouns and the notion of 'semantic parameter'. In *Events and grammar*, Susan Rothstein, ed., Dordrecht: Kluwer, 53–103.

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

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 (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.

Doetjes, Jenny (1997) Mass and Count: syntax or semantics? *Proceedings of meaning on the HIL* : 34–52.

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.

Franck, Julie, Glenda Lassi, Ulrich H. Frauenfelder, & Luigi Rizzi (2006) Agreement and movement: a syntactic analysis of attraction. *Cognition* **101**: 173–216.

Franck, Julie, Gabriella Vigliocco, Inés Antón-Méndez, Simona Collina, & Ulrich H. Frauenfelder (2008) The interplay of syntax and form in sentence production: a cross-linguistic study of form effects on agreement. *Language and Cognitive Processes* **23**(3): 329–374.

Georgi, Doreen (2014) *Opaque interactions of Merge and Agree: On the nature and order of elementary operations*. Ph.D. thesis, University of Leipzig, Leipzig.

Halle, Morris & Alec Marantz (1993) Distributed Morphology and the Pieces of Inflection. In *The View from Building 20*, Ken Hale & Samuel Jay Keyser, eds., Cambridge, MA: MIT Press, 111–176.

Hicks, Glyn (2009) *The derivation of anaphoric relations*. Amsterdam/Philiadelphia: John Benjamins.

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

Koopman, Hilda & Dominique Sportiche (1991) The position of subjects. *Lingua* **85**(2): 211–258.

Kramer, Ruth (2014) Gender in Amharic: A Morphosyntactic Approach to Natural and Grammatical Gender. *Language Sciences* **43**: 102–115.

Kratzer, Angelika (2009) Making a Pronoun: Fake Indexicals as Windows into the Properties of Pronouns. *Linguistic Inquiry* **40**(2): 187–237.

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 conjunction agreement in Slovenian. *Syntax* **18**(1): 39–77.

Müller, Gereon (2009) Ergativity, Accusativity, and the order of Merge and Agree. In *Explorations of Phase Theory. Features and Arguments*, Kleanthes Grohmann, ed., Berlin: Mouton de Gruyter, 269–308.

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.

Preminger, Omer (2011) *Agreement as a fallible operation*. Ph.D. thesis, MIT, Cambridge, MA.

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

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.

Reuland, Eric (2001) Primitives of binding. *Linguistic Inquiry* **32**(3): 439–492.

Reuland, Eric (2005) Agreeing to Bind. In *Organizing Grammar: Linguistic Studies in Honor of Henk van Riemsdijk*, Hans Broekhuis, Norbert Corver, Riny Huybregts, Ursula Kleinhenz, & Jan Koster, eds., Mouton de Gruyter, manuscript.

Reuland, Eric (2011) *Anaphora and Language Design*. Cambridge, MA: MIT Press.

Rooryck, Johan & Guido Vanden Wyngaerd (2011) *Dissolving Binding Theory*. Oxford Studies in Theoretical Linguistics, Oxford: Oxford University Press.

Sauerland, Uli (2004) *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. (2015) *Feature mismatches: Consequences for syntax, morphology and semantics*. Ph.D. thesis, University of Connecticut, Storrs, CT.

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

Sportiche, Dominique (1988) A theory of floating quantifiers and its corollaries for constituent structure. *Linguistic Inquiry* **19**(2): 425–451.

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 (2000) A theory of agreement and its application to Serbo-Croatian. *Language* **76**(4): 799–832.

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.

Williams, Edwin (1994) Remarks on Lexical Knowledge. *Lingua* **92**: 7–34.

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

Wurmbrand, Susanne (2012) Agreement: Looking up or down? Talk given at MIT, March 2012.

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

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

Table 1: Order of adjectives in Hebrew and English, from Shlonsky (2004).