The syntactic limits of probe-goal (a)symmetries

Sandhya Sundaresan & Hedde Zeijlstra
University of Göttingen

June 5, 2021

1 Introduction

In the Minimalist framework of Chomsky (2000, 2001, et seq.), all syntactic dependencies are regulated by the abstract operation **Agree** (for more on which, see the chapters by D'Alessandro, This volume; Preminger, This volume). This chapter serves to provide an overview of the pair of grammatical elements involved in the Agree operation, termed the **probe** and **goal**. We will see that Agree between a probe and goal is possible (or conversely, that two elements qualify as a probe & goal for the purposes of Agree) only if there exists a fundamental featural asymmetry between them. Concretely, the probe is deficient for one or more syntactic features which the goal bears. Agree is then viewed, broadly speaking, as an operation which is triggered to resolve this featural asymmetry. In addition, the probe and goal pair are structurally, and potentially also categorially, asymmetric: e.g. the probe and goal stand in a c-command relationship and typically belong to distinct grammatical categories (e.g. in classic cases of subject ϕ-agreement, we have a DP goal interacting with a T probe). At the same time, for two grammatical elements to qualify as a matching probe-goal pair, in the first place, they must share some core syntactic properties. The central goal of this paper is to explore the upper and lower limits of this syntactic space: i.e. to lay out what featural and structural properties two grammatical elements must share to be identified as a probe-goal pair in the first place, and what properties must minimally distinguish one from the other for the purposes of Agree.

This paper is organized as follows. In Section 2, we set out working definitions of probe and goal which are specific enough to accurately identify them while also being general enough to be accommodate different treatments of these elements. The broadness of these definitions will also be shown to be consistent with different types of Agree relation argued for in the literature. The remaining sections will focus on the nature and range of possible syntactic similarities and differences holding between a probe and goal. Section 3 lays out possible probe-goal featural (a)symmetries, while Section 4 looks at (a)symmetries on a structural level.

2 Working definitions and types of Agree

There is considerable disagreement in the literature about what the precise nature of the asymmetry between the probe & goal should be. What sorts of featural asymmetry are licit: i.e. is featural deficiency on a probe defined in terms of the feature being unvalued or uninterpretable on the probe? Analogously, what sorts of structural asymmetry are licit: e.g. does the probe c-command the goal or vice-versa? An even more basic question has to do with what sorts of features on a probe and goal may drive Agree in the first place? The answers to these questions directly inform the way in which Agree is algorithmically implemented, whether in terms of feature-valuation, -checking or -sharing, or in terms of upward vs. downward probing.

Given the lack of consensus on these issues, we will adopt the following broad working definitions for probe and goal (see Chomsky, 2000, 122, for a formal first definition) going forward:

(1) *Working definitions of probe and goal:*
$\mathcal{G}$, a syntactic element of category α, is a suitable goal for a probe $\mathcal{P}$, another syntactic element of category β (where β might be = α) iff:
  a. $\mathcal{P}$ is deficient for a (potentially unary) set of features $\mathbb{F}$ which $\mathcal{G}$ bears;
  b. $\mathcal{P}$ and $\mathcal{G}$ are syntactically visible to one another: i.e. there is a possibly trivial sequence $\mathcal{P}_1 \ldots \mathcal{P}_n$ where every $\mathcal{P}_i$ is phase-local, minimal and in a c-command relation with $\mathcal{P}_{i+1}$ & $\mathcal{P}_n$ is phase-local, minimal and in a c-command relation with $\mathcal{G}$;

The working definition in (1) is broad enough to cover Agree interactions of differing degrees of complexity. In the simplest scenario, Agree would be a one-step operation between a probe $\mathcal{P}$ and a matching goal $\mathcal{G}$. In this case, the sequence $\mathcal{P}_1 \ldots \mathcal{P}_n$ is trivial, with exactly one probing link (= $\mathcal{P}$) in the chain. $\mathcal{G}$ is both phase-local and minimal to $\mathcal{P}$ and Agree successfully terminates when $\mathcal{G}$ has resolved the featural deficiency of $\mathcal{P}$, as in (2). Classic cases of subject (3) or object agreement would fall under this rubric:

(2) Direct Agree: $[_{phase}\ \mathcal{P} \ldots \mathcal{G}]$

(3) Susan$_{\mathcal{G}}$ make-s$_{\mathcal{P}}$ all her guests drink tea.

In cases of long-distance Agree, however, $\mathcal{P}$ and $\mathcal{G}$ are not necessarily phase-local but are linked to one another transitively, via a sequence of other probe-y elements each of which is phase-local, minimal to, and in a c-command relation with, the next in the chain. The Agree relation between $\mathcal{P}$ and $\mathcal{G}$ is no longer a simple, single-step operation but is cyclic or mediated, as in (4).

(4) Cyclic Agree: $[_{phase}\ \mathcal{P}_1 \ldots [_{phase}\ \mathcal{P}_2 \ldots [_{phase} \ldots [_{phase}\ \mathcal{P}_i \ldots \mathcal{G}]]]]$

Legate (2005, 2) argues that cases of cross-phasal ϕ-agreement warrant such a solution where "agreement applies in a cyclic fashion, through the intermediary of every intervening phase-defining head." Thus, in (5), there is cyclic Agree between the relativized verb *bhuaileas* and a chain of relativized complementizers.

(5) An duine [$_{CP}$ a thuirt e [$_{CP}$ a/*gun bhuaileas e]].
the man C.REL said he C.REL/*that strike.REL he
'That man$_1$ that he said he will hit t$_1$.'

Cases of mediated obligatory control have also been argued to involve cyclic Agree in this sense. The controlled PRO subject first probes its clausemate phasal C head which in turn probes the controller in the higher clause (Landau, 2013; McFadden and Sundaresan, 2018), such that the initial probe (PRO) and ultimate goal (controller) are transitively linked. Cases of Ā-dependency have also been argued to bear the empirical signature of mediated Agree (see e.g. McCloskey, 1979; Chung, 1998; Lahne, 2009; Korsah and Murphy, 2020). We discuss the syntactic status of mediated Agree in greater detail in Section 4.2.

The working definition above also accommodates cases of *multiple* Agree, where a single probe interacts with more than one goal (Béjar and Rezac, 2009; Nevins, 2011; Deal, 2015; Bárány, 2015; Kalin, 2018; Coon and Keine, To appear) or a single goal interacts with multiple goals (Hiraiwa, 2004). The Activity Condition (Chomsky, 2000, 2001), which restricts Agree to probe-goal pairs which are syntactically active by virtue of not having participated in prior Agree relations, is sometimes invoked to prevent such cases of multiple agreement. But the Activity Condition is itself controversial (cf. the discussion in Section 3.1.2): the definition in (1) deliberately refrains from mentioning it for this reason. (1) also remains consciously agnostic about the directionality of probing, i.e. whether $\mathcal{P}$ c-commands $\mathcal{G}$, yielding downward probing (Bošković, 2007, and potentially Chomsky, 2000, 2001), $\mathcal{G}$ c-commands $\mathcal{P}$, yielding upward probing (Wurmbrand, 2011; Zeijlstra, 2012, a.o.), or whether both are systematically allowed (see e.g. Baker, 2008b; Béjar and Rezac, 2009; Preminger, 2014; Puškar, 2018; Bjorkman and Zeijlstra, 2019; Clem, 2019; McFadden and Sundaresan, To Appear). See Section 4.1 for discussion.

Finally, (1) is also flexible enough to accommodate other more complicated instances of Agree that have been proposed in the literature. These include cases of relativized probing (where Agree between the probe and goal may involve a proper subset of the features on the probe or goal, Béjar and Rezac, 2009; Preminger, 2014; Georgi, 2017; Puškar, 2018, cf. Section 3.1.3), dynamic approaches to Agree (where every instance of Agree redefines locality domains for subsequent instances of Agree, Rackowski and Richards, 2005; Puškar, 2018), and Agree relations that are split across an ordered sequence of sub-operations (cf. Agree-Link vs. Agree-Copy in Arregi and Nevins, 2012, Interaction vs. Satisfaction in Deal, 2015 and Match, Value, & Vocabulary Insertion steps in Kalin, 2020).

As mentioned at the outset, we will restrict ourselves to discussion of (a selection of) the Agree variants mentioned here which pertain to the syntactic properties of the probe and goal, like the Activity Condition (Section 3.1.2), relativized probing (Section 3.1.3) and directionality of probing (Section 4.1). But we will have less to say about the other more complicated cases of Agree described here — like multiple or dynamic Agree or Agree that is split across ordered sub-sequences — which we believe ultimately have more to do with the derivational algorithm of Agree than about the properties of the probe and goal.

3 Featural probe-goal (a)symmetries

This section discusses the kinds of featural asymmetry that distinguish a probe from a goal and the kinds of featural parallel that identify two elements as a probe-goal pair for Agree.

3.1 The logical space of featural probe-goal symmetries

This section is devoted to discussing the space of featural symmetries (or commonalities) between a probe and goal. Section 3.1.1 discusses the featural input conditions two grammatical elements must share to qualify as a probe-goal pair. We will then look at particular cases that challenge these basic input conditions in two ways: first, where these input conditions aren't sufficient (cf. Section 3.1.2 on the Activity Condition) and second, where these input conditions aren't fully necessary or are more relativized in various ways (cf. Section 3.1.3 on Relativized Probing).

3.1.1 Feature-matching between probe and goal

The basic input condition for probe-goal selection in classic Agree models (Chomsky, 2000, 2001; Pesetsky and Torrego, 2007, a.o.) is that of feature-matching, as defined in (6):

(6) "The simplest assumptions for the probe-goal system are [...]:
  a. Matching is feature identity [between a probe P and goal G].
  b. D(P) [domain of P] is the sister of P.
  c. Locality reduces to closest c-command [between P and G]." (Chomsky, 2000, 122).

In other words, a grammatical element featurally qualifies as a goal $\mathcal{G}$ for a probe $\mathcal{P}$, iff $\mathcal{G}$ possesses the set of features $\mathbb{F}$ that $\mathcal{P}$ is defective for. For instance, in cases of subject agreement, the T head is typically treated as a ϕ-probe (i.e. deficient for ϕ-features) and the clausal subject as its ϕ-goal. The subject qualifies as the goal both by virtue of being syntactically visible to T and by virtue of bearing ϕ-features that can resolve T's ϕ-featural deficiency. Similarly, in cases of object agreement, the ϕ-probe is typically assumed to be v with its goal being the clausal object. These featural wellformedness conditions regulate that a ϕ-deficient element like T or v will select a ϕ-bearing element (like a pronoun or R-expression) as its choice of goal, rather than an element which bears no ϕ-features at all, like a preposition or negation. Thus, featural wellformedness also regulates categorial asymmetry between the probe and goal, to some extent. It should be mentioned that cases of cyclic or mediated Agree as in (4) do, in theory, allow a ϕ-probe to Agree with another ϕ-probe (e.g. in cases of long-distance ϕ-agreement, reported for Hindi/Urdu in Bhatt, 2005). What is relevant for our purposes, however, is that the *ultimate* choice of goal for a probe or sequence of probes be regulated in terms of feature-matching in the sense discussed here. In other words, probing for a feature $\mathbb{F}$ may not *terminate* with an element that is either also deficient for $\mathbb{F}$, nor may the element be one that itself bears no $\mathbb{F}$-features.[1]

Featural wellformedness conditions have been argued to regulate probe-goal visibility: i.e. whether a set of features $\mathbb{F}$ on an element X is visible to an $\mathbb{F}$-deficient probe $\mathcal{P}$ or not, which in turn conditions X's eligibility as a goal for $\mathcal{P}$. For instance, in many languages, a nominal marked with oblique case cannot trigger co-varying ϕ-agreement on a clausemate verb. Rather, in such cases, agreement tracks either another phase-local nominal which bears non-oblique case[2] or, if the oblique nominal is the sole argument of the clause, we get invariant default agreement. Thus, in the transitive Tamil sentence in (7), the verb tracks the ϕ-features of the nominative object, rather than the dative subject. When the dative subject is the sole argument, as in intransitive (8), we get default agreement on

[1] One source of empirical support for this idea involves the so-called Anaphor Agreement Effect (AAE), an observation going back to Rizzi (1990a) that anaphors cannot trigger co-varying subject or object ϕ-agreement. One line of analyses for this phenomenon (Kratzer, 2009; Tucker, 2011; Sundaresan, 2016; Murugesan, 2019) explores the idea that the AAE stems from the ineligibility of the anaphor as a ϕ-goal. In other words, an anaphor, despite being a nominal, is itself deficient for ϕ-features (Heinat, 2008; Kratzer, 2009; Reuland, 2011; Rooryck and vanden Wyngaerd, 2011, a.o.): as such, it cannot resolve the ϕ-deficiency on a ϕ-probe like T or v, even if it is merged in a structural position that is syntactically visible to the probe. As might be expected, cyclic Agree between a ϕ-probing head and a ϕ-bearing nominal *via* a mediating anaphor does seem to be possible: Sundaresan (2016) argues that this explains cases of putative AAE obviation in Tamil: in such cases, the T head and anaphor together probe for ϕ-features on a ϕ-bearing nominal higher in the CP phase (see Grosz and Patel-Grosz, 2014, for a similar approach to AAE effects in Kutchi Gujarati).

[2] To be more specific, languages vary further in whether any non-oblique case-marked nominal can trigger covarying ϕ-agreement or only nominals with unmarked case may do so (Bobaljik, 2008). Icelandic and Tamil belong to the latter category, Nepali to the former. As discussed in Bobaljik (2008), other types of crosslinguistic variation may be attested as well.

the verb:

(7) Leela-vŭkkŭ naan kaɳɳaaɖi-læ teri-nʤ-een.
Leela-DAT I.NOM mirror-LOC appear-PST-1SG
'Leela saw me in the mirror.'

(8) Leela-vŭkkŭ pasi-čč-ŭdŭ.
Leela-DAT hunger-PST-3NSG
'Leela got hungry.'

A classic explanation for such effects is that the oblique case marking on the nominal renders its ϕ-features invisible to the probe. A featural explanation for such patterns involves either the Activity Condition (AC) or something like Relativized Probing (RP).[3] In general, both the AC and RP challenge the limits of feature-matching between a probe and goal in different ways, so we turn to a discussion of these next.

3.1.2 When feature-matching isn't sufficient: the Activity Condition

The Activity Condition on Agree, defined as in (9), is predicated on the idea that feature-matching between two grammatical elements is a necessary but not a sufficient condition for establishing them as a probe-goal pair:

(9) *Activity Condition (Chomsky, 2001, 6):*
"(3) (i) Probe and goal must both be active for Agree to apply . . . For the Case-agreement systems, the uninterpretable features are ϕ-features of the probe and structural Case of the goal N. ϕ-features of N are interpretable; hence, N is active only when it has structural Case. Once the Case value is determined, N no longer enters into agreement relations and is "frozen in place" (under (3i))."

In other words, the AC has the effect that a head like T can only probe a DP for ϕ-features if that DP still has an "active" Case-feature, i.e. itself is syntactically active by virtue of needing a Case feature. When such a DP resolves the ϕ-deficiency on the T probe, T in turn resolves the Case needs of the DP. An oblique case-marked DP, like in Tamil (7) and (8), will have already had its Case needs dealt with by a more local case-assigner and will thus be inactive for Case at the time T attempts to ϕ-probe it. As a result, the oblique DP is invisible to the T probe and is no longer eligible as a ϕ-goal. In a related vein, Obata and Epstein (2011) argue that case-marking on a DP at the edge of a phase renders it invisible for subsequent ϕ-probing from outside that phase.

The AC was invoked not only to describe the invisibility, in certain languages, of oblique case-marked DPs for ϕ-probing, but also to explain, for languages like

[3]See Section 4.2 for an explanation of these effects in terms of structural wellformedness conditions holding between the probe and goal.

English, the possibility of A-movement in raising contructions like (10) vs. its impossibility in "hyperraising" constructions like (11):

(10) Leela seems [~~Leela~~ to be sleepy].

(11) * Leela seems [~~Leela~~ that is sleepy].

The idea is that the raising of *Leela* in (10) is made possible not only by the ϕ-needs of matrix T but by the Case needs of *Leela*. In embedded subject position, *Leela* cannot be assigned Case because embedded non-finite T is not a licit Case-assigner. As such, *Leela* is "active" and thus syntactically visible to matrix T for ϕ-probing and movement triggered by the EPP. In contrast, "hyperraising" structures like (11) are ungrammatical in English because *Leela* already gets Case from embedded T which, by virtue of being finite, is a licit Case-assigner. *Leela* in (11), being featurally inactive, cannot be ϕ-probed by matrix T or undergo EPP-movement.[4]

But the validity of the AC has since been questioned on a number of grounds. First, the AC is theoretically problematic in that it is unclear whether the active feature on the goal (e.g. the Case feature on a DP ϕ-goal for T/v) has a matching counterpart on the probe. If it does, then it blurs the featural distinction between probe and goal since it entails that a goal can also serve as a probe (for the active feature); if it doesn't, it begs the question why some features (e.g. Case features) should be *sui generis* counterpart-less in this manner while others (e.g. ϕ-features) are not (cf. Section 3.2 for a fuller discussion).

The AC also raises a number of empirical concerns. For instance, the impossibility of hyperraising in (11) has at least two other possible explanations. The first has to do with the Phase Impenetrability Condition (PIC, Chomsky, 2000, 2001) which regulates syntactic dependencies across phases, as in (12):

(12) *Strong PIC (Chomsky, 2000, 108)*:
In phase α with head H, the domain of H is not accessible to operations outside α; only H and its edge are accessible to such operations.

The embedded finite CP in (11) is a phase while the non-finite embedded TP in (10) is not: as a result, the embedded subject *Leela*, which sits in the domain of the phase, is expected to be syntactically invisible to matrix T in the former, but visible to it in the latter. The second solution, which operates in conjunction with the first, involves the Ban on Improper Movement (Chomsky, 1973, 1981; May, 1979) according to which an element cannot move from an A-bar position to an A position. This has the effect of ruling out A-movement of *Leela* in (11) into matrix

[4]Note that a sentence like (i) is, in contrast, perfectly grammatical because the ϕ-needs of matrix T are satisfied by the matrix expletive *it* subject, while the Case needs of the embedded subject *Leela* are satisfied by embedded finite T:

i. It seems [that Leela is sleepy].

subject position (an A position) via the phasal escape hatch in embedded Spec, CP (an A-bar position).

A second empirical challenge comes from constructions like Icelandic (13), which allow raising of oblique DPs (Sigurðsson, 2012, Ex. 2a, 192):

(13) Þeim virtist ekki [hafa verið hjálpað].
them.DAT$_1$ seemed.DFT not [t_1 have.INF been helped.DFT t_1]
Raising *Passive-Movement*
'They$_1$ did not seem [t_1 to have been helped t_1].'

Quite simply, (13) shows that, contrary to what the AC might lead us to expect, a DP like *Þeir*, which is already marked with "quirky" case, can still be syntactically visible for A-movement, first undergoing passive A-movement from embedded object to embedded subject position, and then undergoing raising from embedded to matrix subject positions. In a more recent paper, Keine (2018) presents analogous counter-arguments against the AC from Hindi/Urdu, showing that "all cross-clausal movement in Hindi applies to already case-marked elements" (Keine, 2018, 4). To reconcile the possibility of such sentences with the AC, Freidin and Sprouse (1991) and Chomsky (2000) propose that the embedded subject in (13) may A-move because, while it may bear quirky case, it still remains syntactically active for *structural* Case. But such an account is not just unwieldy, it is also empirically challenged by the observation that non-finite subjects in a number of languages can bear structural nominative Case (McFadden, 2004; Sundaresan and McFadden, 2009; McFadden and Sundaresan, 2011) which can then be visible for further syntactic operations.

Cases of multiple agreement, where multiple probes interact with a single goal, are arguably also problematic for the Activity Condition since, per the AC, a goal that has already participated in Agree with a probe should be rendered inactive for subsequent Agree operations. On the strength of the observation that languages like Icelandic and Turkish disallow multiple probing while others like Georgian, Kinande and Jarawara, allow it, Baker (2008a) argues that the Activity Condition is not language-universal but serves as a macro-parameter that is operational in some languages but not others. More recently, Oxford (2017) has argued that the AC is *micro*-parametric given that such agreement differences are attested even *within* individual languages like Plains Cree, suggesting a more fine-grained parametrization.

Finally, a potentially even more damning challenge to the Activity Condition comes from the fact that hyperraising sentences equivalent to English (11) *are*, in fact, licit in languages like Zulu (Zeller, 2006; Halpert, 2019), Bantu more generally (Carstens, 2011; Carstens and Diercks, 2013) Nez Perce (Deal, 2017) and Greek (but only out of "semi-finite" subjunctive clauses, Iatridou, 1993; Alexiadou and Anagnostopoulou, 1999). Hyperraising in Zulu is illustrated in (14) (Halpert, 2019, 18, Exx. 50a-c):

(14) **uZinhle** **u-/ku-**bonakala [ukuthi **u-**xova ujeqe].
AUG.1Zinhle 1S-/17S-seem that 1S-make AUG.1bread
Literal: 'Zinhle$_i$ seems [$_{CP}$ that t$_i$ is making bread now].'
Intended: 'Zinhle seems to be making steamed bread now.'

To accommodate such patterns, Halpert explicitly proposes eliminating the Activity Condition and phasehood altogether in favor of a model that implements locality in terms of Relativized Minimality (Rizzi, 1990a) and implements ϕ-Agree in terms of relativized probing.

There are thus many reasons to be skeptical of the Activity Condition as an additional feature wellformedness condition on probe-goal pairing, even in its weakened status as a language-internal micro-parameter. At the same time, it is important to keep in mind that the AC is often tacitly assumed in the literature on probe-goal relations and Agree and is often also explicitly used as a strategy to prevent an element from serving as a goal for more than one probe.

3.1.3 When (full) feature-matching isn't necessary: Relativized Probing

Here, we discuss cases that question the extent and nature of core feature-matching between a probe and goal, categorizing these under the general rubric of "relativized probing (RP)". In its most general sense, RP is predicated on the idea that feature-matching between two elements need not be exhaustive to qualify them as a licit probe-goal pair: the goal may have more features than the probe, or vice-versa. We define it in these broadest terms as in (15) below:

(15) *Relativized Probing:*
A grammatical element $\mathcal{G}$ qualifies as a goal for a probe $\mathcal{P}$ which is deficient for a set of features $\mathbb{F}$ iff either:
a. $\mathcal{G}$ bears a superset $\mathbb{F}'$ of features $\mathbb{F}$ on $\mathcal{P}$, or;
b. $\mathcal{G}$ matches $\mathcal{P}$ for a set of features $\mathbb{F}'$, which is a proper subset of $\mathbb{F}$.

In the first logical alternative described in (15a), the goal $\mathcal{G}$ matches $\mathcal{P}$ for the feature-set $\mathbb{F}$ but also bears an additional set of features $\mathbb{H}$. Upon Agree, $\mathcal{G}$ resolves the feature-deficiency $\mathbb{F}$ on $\mathcal{P}$, while $\mathbb{H}$ simply does not featurally interact with $\mathcal{P}$ in any way. Under one variant of this idea, reminiscent of the Activity Condition, the inclusion of $\mathbb{H}$ on $\mathcal{G}$ is a necessary condition for $\mathbb{F}$-probing by $\mathcal{P}$. In other words, simple feature-matching for $\mathbb{F}$ between $\mathcal{P}$ and $\mathcal{G}$ is not enough; $\mathcal{G}$ must also bear an additional dedicated feature-set $\mathbb{H}$. Bobaljik (2008) deploys this idea to explain the unavailability of oblique case-marked nominals for ϕ-Agree in sentences like Tamil (7)-(8). For Bobaljik, such unavailability stems from the condition that only DPs bearing certain cases are accessible for ϕ-Agree. Concretely, the T ϕ-probe in Tamil is *case-discriminating*: i.e. it can only Agree with a nominal which, in addition to bearing ϕ-features, also bears unmarked (nominative) case. The direct object qualifies as such a nominal in the transitive sentence in (7). But in intransitive (8), the oblique case-marked subject is the sole clausal argument: it bears ϕ-features but not unmarked case, thus does not qualify as a licit goal for T.[5]

A logically weaker variant of this idea is captured in the relativized probing approach of Béjar and Rezac (2009) and much subsequent work: $\mathcal{G}$ can bear feature-sets $\mathbb{F}$ and $\mathbb{H}$, as in the previous variant, but the presence of $\mathbb{H}$ is not a necessary input condition for $\mathbb{F}$-probing by $\mathcal{P}$. For Béjar & Rezac, such a model is required to derive variation for person hierarchy effects across languages. Person categories are implicationally ordered, such that (the category of) 1st-person entails 2nd-person which in turn entails the 3rd (i.e. 1st > 2nd > 3rd):

(16) Feature hierarchy for person (Béjar and Rezac, 2009, Table 2A-B, 43):

3rd	2nd	1st
[3] (= [π])	3 [2] (= [participant])	3 [2] [1] (=[speaker])

In a language like Swahili or Choctaw, the person probe is "flat", specified as [u3], deficient only for [3]/[π]. Such a probe is satisfied by a goal of any person category. A 3rd-person DP, specified as [3], will fully satisfy the featural needs of the [u3] probe; a 2nd-person DP, specified as {[3], [2]} will match the probe for [3] alone ([2] will simply not interact with [u3]), while a fully specified 1st-person DP specified as {[3], [2], [1]}, will again match the probe for [3] alone, leaving both [2] and [1] untouched. Such a language thus shows no person hierarchy effects whatsoever, and manifests no preference for 1st vs. 2nd vs. 3rd-person arguments when it comes to person agreeement. Probing successfully terminates after the first cycle of Agree with a person argument.

The second logical option in (15b), where $\mathcal{G}$ bears a proper subset $\mathbb{F}'$ of the feature-set $\mathbb{F}$ $\mathcal{P}$ is deficient for, is exploited to derive person-agreement in languages that *do* show person hierarchy effects. Under this alternative, $\mathcal{P}$ Agrees with $\mathcal{G}$ for $\mathbb{F}'$, rather than $\mathbb{F}$, leaving an "active residue" (to borrow terminology from Béjar and Rezac, 2009) of deficient features on $\mathcal{P}$. This is the featural set-complement $\mathbb{F} - \mathbb{F}'$ which triggers a second cycle of probing by $\mathcal{P}$. Bejar & Rezac propose that languages like Mohawk and Kashmiri involve person probes that are fully articulated for person. This is essentially a 1st-person probe, featurally specified as {[u3], [u2], [u1]}. Such a probe is "picky" in that it can only be *fully* featurally matched by a goal that is also 1st-person, with the feature-specification {[3], [2], [1]}. A 2nd-person goal, specified as {[3], [2]} will match a proper subset of the probe's features, leaving {[u1]} as an active residue; a 3rd-person goal, specified as {[3]} will match the [u3] feature on the probe, leaving as an active residue

[5] In Section 4.2, we turn to a structural treatment of this phenomenon (e.g. Řezáč, 2008) and also discuss the case of ϕ-agreement in Hindi/Urdu, which seems to be regulated both in terms of structural- and featural- (specifically case-) discrimination on the part of the ϕ-probe (Keine, 2018).

the feature-set {[u2], [u1]}.

This yields person hierarchy effects for agreement. Béjar & Rezac assume that probing is downward by default, but the active-residue scenarios trigger a second cycle of probing in an upward direction. Thus, if the internal argument (IA) of a 1st-person probe is either 2nd- or 3rd-person, a second cycle of upward probing is initiated, yielding agreement with the external argument (EA) just in case the EA is higher than the IA on the person hierarchy, i.e. is 1st-person. A second-person probe, in contrast, will allow agreement with a 1st- or 2nd-person EA over a 3rd-person IA, but will yield agreement with a 2nd-person IA, even if the EA is 1st-person. Such a system is attested in Basque. The prefixal agreement slot tracks the EA just in case the EA is 1st/2nd-person and IA is 3rd-person, as in (19), adapted from Béjar and Rezac, 2009, Ex. 2, 37):

(17) Ikusi [z]-in-t-u-da-n.
seen 2-X-PL-have-1-PAST
'I saw you.' ($1 \rightarrow 2 = 2$)

(18) Ikusi [n]-ind-u-en.
seen 1-X-have-PAST
'He saw me.' ($3 \rightarrow 1 = 1$)

(19) Ikusi [n]-u-en.
seen 1-have-PAST
'I saw him.' ($1 \rightarrow 3 = 1$)

In recent work Puškar (2018) derives so-called "hybrid" agreement patterns in Bosnian/Croatian/Serbian, where agreement may reflect either natural or grammatical gender depending on the number feature on the nominal goal. When the nominal is singular, natural gender agreement is enforced, but agreement may alternate between natural and grammatical genders when the nominal is in the plural. Such a pattern requires not only that the probe be able to reliably distinguish natural from grammatical gender, but also that gender features be allowed to interact with number features. Using an analogous system to Béjar and Rezac (2009), Puskar illustrates that such complex patterns may be systematically derived using a model of cyclic Agree driven by relativized probing for number vs. gender in a specific order (number before gender, or gender before number).

Similarly, Halpert (2019) uses a variant of RP to explain the Zulu hyperraising facts in (14). Halpert's solution, similar in spirit to Béjar and Rezac (2009), is to propose that Zulu-hyperraising obtains just in case the embedded CP intervenes for ϕ-Agree between T and the embedded subject, for a proper subset of features on the probe. Relativized Minimality then forces matrix T to Agree with the CP first for this feature-subset. Such an Agree cycle, for Halpert, "unlocks" the CP, allowing matrix T to continue probing in a second cycle of Agree for the remaining features with the next closest candidate, the embedded subject.

In even more recent work, McFadden and Sundaresan (To Appear) argue that

RP can be exploited to systematically derive differences between different types of complementizer agreement. Downward complementizer agreement, attested in many West Germanic languages (Zwart, 1997; Carstens, 2003; van Koppen, 2005; Fuß, 2008; van Alem, 2020, a.o.). But upward complementizer agreement with a higher clausal argument as reported for many Bantu languages (Diercks, 2013; Carstens, 2016; Diercks, van Koppen, and Putnam, 2020), as in Lubukusu or cases of allocutive agreement with an extra-argumental addressee like in Japanese (Miyagawa, 2017), Tamil (McFadden, 2020), or Basque, is due to a second cycle of upward probing by C when the first downward cycle fails to fully satisfy the needs of the probe.

Relativized Probing has thus been very influential in the literature, yielding analytic offshoots that have been used to derive a range of other agreement phenomena beyond just those involved in person hierarchy effects. More recent work in Deal (2015) proposes to implement multiple/cyclic Agree of the kind necessitated by such phenomena using a more radically different model of probe-goal relations. Under this model, a probe is specified more richly for two sets of features: so-called *interaction* features which are the features the probe will copy onto itself from the goal(s), and so-called (optional) *satisfaction* features which determine when the probe's featural needs are satisfied and probing halts. A probe can thus interact with arbitrarily many goals bearing its interaction features (replicating the effects of cyclic Agree under RP) until it finds a goal with a matching satisfaction feature.

3.2 The logical space of featural probe-goal asymmetries

In this section, we will look at the ways in which a probe and goal must be featurally distinguished so as to interact for the purposes of Agree, focusing in particular on the (un)interpretable vs. (un)valued feature distinction between the two. We will see that the somewhat controversial status of the Activity Condition, discussed above, has been instrumental in redefining, in various ways, the limits of the featural asymmetry holding between a probe and goal.

3.2.1 Feature valuation and feature interpretability

Feature-valuation is predicated on the idea that a syntactic feature $\mathbb{F}$ is structured as an attribute-value pair of the form: [$\mathbb{F}$: *val*]. The featural attribute $\mathbb{F}$ categorizes a feature-set (denoting a predicate), like PERSON, NUMBER or GENDER which may also be bundled together more generically as ϕ features. The value *val* characterizes a unique member of the attribute set. Thus, possible values for the PERSON feature-attribute are *1st/2nd/3rd*; analogously, we have *singular/dual/plural* as possible values for NUMBER and *masculine/feminine/neuter* as possible values for GENDER. The classic idea is that some grammatical categories are "born" with the feature valued; others have features which are born unvalued and must be valued (via feature-matching or -copying) in the course of the syntactic derivation. The latter are probes for the unvalued features in question, while the former are eligible to be goals for these same features. To illustrate, ϕ-features are taken to be unvalued on functional heads like T or v which qualify as ϕ-probes on the strength of this; but ϕ-features are taken to be valued on pronouns and R-expressions which can thus serve as ϕ-goals for T and v. The process of ϕ-valuation on T/v by such a nominal yields agreement for ϕ-features.

The advantage of such a system is that it can readily explain why ϕ-agreement on a functional head almost always reflects the ϕ-features on a subject or object: this follows automatically from the idea that the functional head in question has no independent ϕ-values of its own to contribute. At the same time, it is worth noting that there is no *a priori* reason why the ϕ-asymmetry should be modelled in precisely this way. Ackema and Neeleman (2013) argue, in contrast, that a probe and goal are simply not asymmetric for feature-valuedness in this manner; in principle, both a functional head like T/v and the nominal may independently bear valued ϕ-features. Evidence for their position comes from "subset control" agreement phenomena like so-called unagreement in Spanish, where the agreement goal actually seems to bear fewer ϕ-features than the probe.

In contrast, feature-interpretability, strictly speaking, has to do with whether a feature $\mathbb{F}$ is legible at the post-syntactic interfaces of LF (meaning) and/or PF (sound); in practice, it is largely used for semantic legibility at LF. ϕ-features on nominals are considered to be (LF-)interpretable, for the most part, since these then feed differences in the referential interpretation of these nominals. For instance, a 1FSG feature on a nominal is associated with the unique female speaker of a context while a 2MSG is associated with a unique male addressee in that context. In contrast, ϕ-features on functional heads like T and v are held to be LF-uninterpretable.

3.2.2 Early variants of (un)Interpretable and (un)valued features

The earliest definitions of Agree defined featural asymmetries between a probe and goal in terms of (un)interpretable as well as (un)valued features, as elaborated in (20) below (adapting Chomsky, 2000, 2001):

(20) A probe $\mathcal{P}$ can agree with a goal $\mathcal{G}$ iff:
  a. $\mathcal{P}$ carries at least one unvalued and uninterpretable feature and $\mathcal{G}$ carries a matching interpretable and valued feature.
  b. $\mathcal{P}$ c-commands $\mathcal{G}$
  c. $\mathcal{G}$ is the closest goal to $\mathcal{P}$
  d. $\mathcal{G}$ bears an unvalued uninterpretable feature [Activity Condition].

The minimal featural properties of $\mathcal{P}$ and $\mathcal{G}$ before and after Agree may thus be templatically illustrated as in Table 1 (adapted from Zeijlstra, 2012, Exx. 2a-b, 2). Table 2 presents a concrete instantiation of this for ϕ-agreement between a T probe and a subject DP goal in a nominative-accusative language:[6]

Pre-Agree	$\mathcal{P}$	$\mathcal{G}$
	[uF: __]	[iF: *val*] [uK: __]
Post-Agree	$\mathcal{P}$	$\mathcal{G}$
	[~~uF: *val*~~]	[iF: *val*] [~~uK: *val*~~]

Table 1: Template – Probe/Goal asymmetries

Pre-Agree	T	$DP_{subject}$
	[uϕ: __]	[iϕ: *3fsg*] [uCase: __]
Post-Agree	T	$DP_{subject}$
	[~~uϕ: *3fsg*~~]	[iϕ: *3fsg*] [~~uCase: *nom*~~]

Table 2: Illustration – subject ϕ-agreement

Per the template in Table 1, $\mathcal{G}$ both checks the uninterpretable feature F on $\mathcal{P}$ and values it. As a reflex of this, the [uK: __] feature on $\mathcal{G}$ (which makes it syntactically active for probing by $\mathcal{P}$ in the first place) is checked and valued in turn. The uninterpretable features on $\mathcal{P}$ and $\mathcal{G}$ are then assumed to be deleted before they reach LF and PF since, by definition, an uninterpretable feature cannot be parsed at the interfaces. In the case of subject agreement illustrated in Table 2, this essentially means that the DP subject copies its own 3FSG ϕ-values onto the T probe and checks off the uninterpretable ϕ-feature on T. As a reflex of this, the uninterpretable Case feature on the DP is checked and, in languages with nominative-accusative case alignment, valued as nominative.

There are several issues with the classic Chomskyan Agree model illustrated above. First, as already alluded to in Section 3.1.2, it is unclear what checks and values the "active" [uK:__] feature on $\mathcal{G}$. After all, in the classic template in Table 1, $\mathcal{P}$ is not itself specified to have interpretable and valued K features. For Chomsky, the [uK:__] feature on $\mathcal{G}$ is simply *sui generis* uninterpretable and unvalued: i.e. it has no interpretable and valued counterpart. But it is mysterious what higher principles regulate why some features should be *sui generis* in this sense while others, like the [uF: __] feature on $\mathcal{P}$ are not. In addition, as Pesetsky and Torrego (2004) also argue in detail, such a stance violates the idea of radical interpretability (21), originally proposed in Brody (1997):

(21) *Thesis of Radical Interpretability:*
Every feature must receive a semantic interpretation in some syntactic location.

(21) is based on the notion that a syntactic feature cannot be relevant for the narrow computational purposes of the syntax alone, but must also be relevant for legibility at the interfaces. Zeijlstra proposes a learnability argument for (21),

[6]The prefixes "u" and "i" before a feature mark the feature as uninterpretable and interpretable, respectively; the blank line "__" after a feature indicates that the feature is unvalued; a strikethrough across a valued feature indicates that the valued feature has been deleted.

stating that "the cue for acquiring uninterpretable features lies in the fact that language learners are confronted with a morphosyntactic manifestation of a particular semantic property, which is not present on the locus of interpretation. Consequently, no uninterpretable feature can be acquired without there being an interpretable counterpart of it ..." (Zeijlstra, 2012, 2).[7]

Pesetsky and Torrego (2004) deal with the problem of *sui generis* features by proposing that the [uK: __] feature on $\mathcal{G}$ does in fact have an [iK: *val*] counterpart on $\mathcal{P}$. When $\mathcal{G}$ checks and values [uF: __] on $\mathcal{P}$ under Agree, $\mathcal{P}$ in turn checks and values [uK: __] on $\mathcal{G}$, under "reverse Agree". But such a solution comes with its own set of problems. For one thing, it is unclear what higher-order grammatical principles regulate the dependency between Agree and Reverse Agree. In the absence of such systematic principles, it is also mysterious what qualifies a particular operation as the primary Agree operation and the other as a reflex operation parasitic on this. Zeijlstra (2012) in fact argues that cases of ϕ-agreement and Case-licensing actually involve the opposite configuration, with the latter being the primary Agree operation, involving upward probing, and the former its reflex (cf. Section 4.1).

Second, the very notion of feature interpretability is conceptually flawed in some sense. In particular, it is predicated on the notion of "look-ahead": i.e. a feature $\mathbb{F}$ is *syntactically* marked as uninterpretable ([u$\mathbb{F}$]) or interpretable ([i$\mathbb{F}$]), depending on whether it will be deemed legible at the interfaces. But under a modular architecture of grammar where a monotonic syntactic derivation feeds into the meaning and form interfaces of LF and PF, respectively, the legibility status of a feature in the post-syntactic interfaces should not already be determinable in the syntax proper. The look-ahead problem was recognized in Chomsky (2000, 2001). The solution proposed was to treat feature valuation as a lexical representation of interface-interpretability: i.e. the (un)interpretability of a feature $\mathbb{F}$ could be syntactically determined by observing its status with respect to featural valuedness. For this to work, of course, featural uninterpretability must go hand-in-hand with feature unvaluedness; analogously, featural interpretability must go hand-in-hand with feature valuedness. This is explicitly regulated by the bicondition in (22):

(22) A feature F is uninterpretable iff F is unvalued. (Chomsky, 2001, 5)

The only variant of Agree allowed by such a system is between a probe [uF: __] and a goal [iF: *val*] where the goal simultaneously checks off the uninterpretable uF feature on the probe and copies its own feature-value(s) *val* onto it.

[7] At the same time, as discussed in Deal (2015), it is not obvious why the LF-semantics could not just ignore syntactic features that it does not know how to interpret or, alternatively, treat them as involving identity functions.

3.2.3 A four-way distinction of ±interpretable and ± valued features

Pesetsky and Torrego (2007) represents an attempt to address the challenges described above for classic Agree. Pesetsky and Torrego argue that, given that feature interpretability and valuedness ultimately carve out distinct concepts, there is no reason why a feature $\mathbb{F}$ cannot be simultaneously uninterpretable and valued, or interpretable and unvalued. They thus reject the biconditional restriction in (22) above and propose a four-way feature distinction on probes and goals stemming from the free combination of ±interpretable and ± valued features, as shown below:

(23) Four way feature distinction (Pesetsky and Torrego, 2007):

Template	Feature	Category	Description
[uK: __]	[uTns: __]	DP	Case
[iK: __]	[iTns: __]	T	Tense
[uK: *val*]	[uTns: *past*]	V	past-tense on *play-ed*
[iK: *val*]	[iϕ: *3fsg*]	pronoun	*she*

Probing is triggered by unvalued features (which are legible to narrow syntax) rather than by uninterpretable features (which are syntactically unparseable) as a way to avoid the look-ahead problem with feature interpretability. As described in (23), this four-way feature system enables a more nuanced distinction between grammatical categories and a wider range of Agree-based interactions between them. For instance, a probe $\mathcal{P}$ is no longer restricted to be both unvalued and uninterpretable; an element with unvalued and interpretable features will also function as a probe. Pesetsky & Torrego propose that the tense feature on T in a language like English is precisely such a probe. Formalizing the intuition that tense is exponed on the finite verb (in English) but is interpreted on T, it is proposed that T has the tense feature [iTns: __] which probes down to get valued by a matching [uTns: *past*] feature on a finite verb like *played*.

A major advantage of this system is that it can also do away with the *sui generis* activity feature [uK: __] on the goal in the classic Agree model. In the example at hand, the third logical possibility for Tns, namely [uTns: __], is instantiated on an argumental DP like a clausal subject or object. Pesetsky & Torrego propose that this is, in fact, the formalization of Case on such a DP. Case-licensing on a DP thus no longer proceeds as a reverse Agree operation or as an automatic reflex of ϕ-agreement with T/v. Rather, both a DP subject and T function as probes for unvalued Tns. T probes downward first finding the [uTns: __], feature on the subject in Spec, vP. T and the subject form a *shared* probe for unvalued Tns and probe further downward together until they reach the finite verb in V. V values iTns on T and uTns on the subject simultaneously and checks uTns on the DP. Case-licensing on the subject is thus an automatic instantiation of Agree between a subject and the finite verb for Tns.

An attractive consequence of this model, more broadly speaking, is that it allows two elements of the same grammatical category or extended projection to participate in probe-goal relations. For instance T and V belong to the same verbal extended projection and the former probes the latter for Tns. More recent work has extended this idea to derive possible and impossible cases of Agree between two nominals. Work within the Minimalist program (Heinat, 2008; Kratzer, 2009; Hicks, 2009; Rooryck and vanden Wyngaerd, 2011; Reuland, 2011; Sundaresan, 2012) treats the syntactic correlate of anaphoric binding as an Agree relation holding between the antecedent nominal (a pronoun or R-expression) and the anaphoric nominal, with the former serving as a goal for the latter. The four-way feature model in Pesetsky and Torrego (2007) readily lends itself to such an implementation. The antecedent nominal bears interpretable and valued features, e.g. [iϕ: *val*], while the anaphor bears interpretable and unvalued features, e.g. [iϕ: __]. A third logical option, namely [uϕ: __] is instantiated by a functional head like T or *v*.[8] Thus, in addition to allowing the syntactic implementation of anaphora, such a model also allows us to model ϕ-agreement.

The recent proposal in Murugesan (2019) exploits the richness of the Pesetsky/Torrego system to derive the Anaphor Agreement Effect (AAE), described earlier (cf. Fn. 1), namely the observation that anaphors cannot typically trigger covarying ϕ-agreement on T or *v*. For Murugesan, this follows from the idea that the anaphor constitutes a ϕ-probe in its own right (by virtue of bearing unvalued ϕ-features). As such, it cannot value the ϕ-features on a ϕ-probe involved in more conventional agreement like T or *v*, unless the anaphor has itself already been valued by its binder by the relevant point in the derivation. Murugesan convincingly argues that such a state-of-affairs requires a cross-linguistically quite unlikely combination of factors: in most languages, the anaphor is merged first, in object position while its intended binder is merged in subject position in Spec, *v*P. The ϕ-probe responsible for object agreement is merged in an intermediate position between the anaphor and its binder, e.g. *v*. As such, assuming probing happens as soon as possible, the anaphor has no chance to be valued by its binder at the point *v* probes downward for ϕ-values, because its binder will not yet have been merged in the structure. The AAE is the result of the anaphor's failure to satisfy the needs of the *v* probe.

3.2.4 Feature-checking vs. -valuation as structure-building vs. -enrichment

More recent proposals define the distinction between feature-interpretability and feature-valuation in different terms. In Sections 3.2.1-3.2.2, we saw that the solution proposed in Chomsky (2000, 2001) for the look-ahead problem posed by feature-interpretability was to argue that: (i) feature interpretability goes hand-in-hand with feature valuation as per the biconditional in (22), and (ii) the interpretability of a feature on an element should be read off of the valued vs. unvalued

[8] We assume that there is no obvious candidate to fulfill the fourth logical option of [uϕ: *val*].

status of that feature in syntax (since the latter, unlike the former, *can* be determined in the syntactic module). More recent proposals of feature-interpretability avoid the look-ahead problem simply by defining featural interpretability in different terms. For instance, Zeijlstra (2014, 2020) argues that "interpretable" is simply a label for a feature that is purely formal in the sense that it has the ability to check off (a matching) "uninterpretable" feature, but which itself lacks any semantic interpretation. In contrast, "uninterpretable" categorizes a feature that is purely formal in the sense that it cannot survive the derivation unless it gets checked off by a matching interpretable feature. The terms "interpretable" and "uninterpretable" are thus, for Zeijlstra, actually misnomers and perhaps better referred to as "independent" and "dependent" formal features, respectively. Since the intended meaning of these features has nothing to do with interface interpretability, there is also no look-ahead problem associated with them.

Empirical motivation for this more indirect correspondence between independent formal features and semantic features, for Zeijlstra, comes from the observation that mismatches between the two types of features are often attested as, for instance, with deponent verbs (e.g. Latin *loqui* ('talk')), which are formally passive but semantically active. Zeijlstra also argues against Chomsky's argument that an uninterpretable feature must be deleted/erased in the syntactic derivation so as not to feed the LF interface (given Chomsky's principle of Full Interpretation that demands that only semantically active material may end up at LF) on the grounds that it is unclear why the null meaning contribution of an uninterpretable feature could not simply be ignored by the semantic component. In fact, Zeijlstra argues, *not* ignoring the vacuous meaning contribution of an uninterpretable feature might conceivably even result in a logical contradiction.[9]

Zeijlstra (2020, but see also McFadden, Sundaresan, and Zeijlstra, 2021), following Neeleman and van de Koot (2002) and Adger (2003), proposes that for a sentence to be grammatical, every dependent (or uninterpretable) feature must end up in a sisterhood relation with a matching independent (or interpretable) counterpart. If, at a particular derivational stage, such a dependent feature does not stand in a sisterhood with its matching independent feature, it percolates up the the next level. Under this perspective, feature checking reduces to c-selection (or categorial selection) under sisterhood which has been argued to drive structure-building or Merge (see also Adger, 2003). The sole function of a dependent feature is to encode on a particular lexical item what other lexical elements need to be present in a grammatical sentence. A transitive verb thus not only bears an independent V feature, identifying its verbal status, but also a dependent N feature, indicating that it is a predicate looking to merge with a nominal argument (the

[9]Briefly, if the presence of an uninterpretable feature in a structure that is otherwise fully interpretable can make a sentence crash within the semantic component, then, given existing definitions of interpretable features ((Svenonius, 2006)), this uninterpretable feature should be taken to be a semantic feature. But if it is a semantic feature, its presence at LF should not violate the principle of Full Interpretation.

direct object). The noun, analogously, is born with a N feature which may then be additionally specified with valued or unvalued ϕ-features which can drive a secondary operation of feature-valuation, as under standard Agree. Such a model thus distinguishes feature checking operationally from feature valuation. Briefly, and simplifying for now, the former drives Merge, a structure-building operation under sisterhood, while the latter drives Agree, a structure-enrichment process that may apply at a distance.

4 Structural probe-goal (a)symmetries

The discussion so far has focused on the logical space of featural (a)symmetries between probes and goals. But, as we have mentioned, there are also various structural (a)symmetries that apply to probe-goal relations: i.e. a probe and goal must share certain structural properties while also being structurally distinguished.

Here, we explore the limits of these parallels and distinctions, focussing on three main issues. The first concerns the question of whether probe c-commands goal, vice-versa, or both are possible. The second has to do with the observation that c-command is not a sufficient condition for Agree: i.e. not every potential probe-goal pair that stand in a c-command relation need undergo Agree. Probe-goals relations are also subject to syntactic locality, a point that is further complicated by the fact that syntactic locality itself may be defined in different ways. Finally, c-command is not always taken to be necessary for probe-goal relations: some Agree relations between a probe and a goal, e.g. cases of conjunct ϕ-agreement, have been argued to obtain in the absence of one c-commanding the other, with other structural relations, like linear precedence, potentially mattering at least as much as, or even more than, c-command.

4.1 Directionality: what c-commands what

Over the years, consensus has more or less arisen that a core structural relation shared between a probe and goal is that of c-command. At the same time, questions have emerged concerning what c-commands what: does the probe c-command the goal, the goal c-command the probe, or are both licit? This in turn directly affects the directionality of feature-probing — upward, downward or both?

4.1.1 Downward Agree: probe c-commands goal

In the original version in Chomsky (2000), in what has since been termed Downward Agree, an uninterpretable and unvalued feature $\mathbb{F}$ on a probe $\mathcal{P}$ probes down its c-command domain to find a goal $\mathcal{G}$ which bears a matching interpretable and valued feature $\mathbb{F}$.

At the same time, Chomsky's Activity Condition requires that the goal $\mathcal{G}$ also be visible for such probing, by virtue of itself bearing an uninterpretable & unvalued feature $\mathbb{H}$, which must then be satisfied by the probe $\mathcal{P}$. As discussed before, this means that the lower goal in some way also behaves as a probe, and the higher probe as a goal. Classic Downward Agree must thus always be accompanied by an instance of Reverse Agree in the opposite direction: e.g. if Downward Agree is for ϕ-features between T (which bears [uϕ:__]) and the subject in Spec, *v*P (bearing [iϕ: *val*]), then the Activity Condition ensures that this is accompanied by Reverse Agree (in an upward direction) between $\mathcal{P}$ and $\mathcal{G}$ for an unvalued, uninterpretable Case feature ([uCase: __]) on $\mathcal{G}$. Such a model thus introduces bidirectionality for probe-goal relations through the back-door. Bošković (2007) seeks to avoid this problem, proposing that instances of Agree always and only involve downward probing. The only way for a lower goal $\mathcal{G}$, then, to have its case-feature checked against a higher probe $\mathcal{P}$ is by moving to the closest position c-commanding $\mathcal{P}$, namely its specifier. From this new position, $\mathcal{G}$ then probes downward to get its Case checked by $\mathcal{P}$. Bošković's proposal for ϕ- and case-agreement thus also has serious repercussions for the directionality of Agree whereby all operations of Agree obtain under a configuration where a probe c-commands its goal.

Both the classic Chomskyan and more recent Bošković proposals are predicated on the notion that Case-checking is a reflex of ϕ-probing, and differ only with respect to their assumptions about the directionality of the Agree operation for the former. But for Bobaljik (2008) who, as we have discussed, argues that ϕ-probes in many languages can additionally be case-discriminating (i.e. sensitive to the case-properties of a goal $\mathcal{G}$) case-assignment must *precede* probing for ϕ-features. This thus directly undermines the idea that Case-checking is a reverse Agree reflex of ϕ-Agree.

4.1.2 Upward Agree: goal c-commands probe

Zeijlstra (2012) turns Bošković's proposal on its head: like Bošković, Zeijsltra argues that Agree between a probe and goal is always and only uni-directional and that probing for Case and ϕ-features must thus both happen in the same direction. However, for Zeijlstra, Agree is always and only upward, with the goal c-commanding the probe. Case-assignment is the result of a feature checking relation between a higher, potentially ϕ-probing head and a lower DP goal (see also Wurmbrand, 2011). In turn, this DP can raise to the specifier of the probing ϕ-head and check/value the ϕ-features on the probe from this position, again as an instance of upward Agree.

Zeijlstra presents empirical support for Upward Agree from a range of phenomena outside the domain of ϕ-agreement such as negative concord (Haegeman and Lohndal, 2010; Zeijlstra, 2004, 2008), sequence-of-tense (Abusch, 1997; Ogihara, 1989; Kratzer, 2009; Kusumoto, 1999; Stowell, 2007), (strict) NPI licensing (Chierchia, 2013), anaphoric binding (Kratzer, 2009; Tucker, 2011; Sundaresan, 2012; Murugesan, 2019), semantic agreement (Smith, 2017), and inflection doubling (Wurmbrand, 2011; Bjorkman, 2016). Zeijlstra argues that these phenomena are all characterized by a syntactic configuration where the dependent feature, hosted by a probe, is c-commanded by an independent featural counterpart, hosted on a goal. For instance, if negative concord constructions are analyzed in terms of syntactic agreement (as proposed by Haegeman and Lohndal, 2010; Zeijlstra, 2004, 2008, a.o.), negative indefinites should be treated as elements carrying an uninterpretable negative feature [uNeg] that Agrees with a negative operator carrying an interpretable feature [iNeg]. As shown in (24), the negative indefinite in Italian *nessuno* is structurally lower than the semantically negative marker *non*.

(24) Oggi *(non) pretendo che nessuno dica niente.
today NEG ask.1.SG that neg-body says.SUBJ n-thing
'I don't ask today that anybody says anything.'

Zeijlstra's model of Upward Agree has been questioned on a number of grounds, most notably by Preminger (2013), who argues that this proposal is incompatible with cases of long-distance ϕ-agreement where the goal is clearly structurally lower than the probe. For instance, it has been argued (see e.g. Polinsky and Potsdam, 2001) that, in cases of cross-clausal agreement in Tsez (Nakh-Dagestanian), an embedded absolutive topic triggers noun class agreement on the matrix verb ((25), from Polinsky and Potsdam, 2001, Exx. 1a, 584). In the absence of such an embedded absolutive topic, default noun class IV-agreement is triggered instead ((26), from Polinsky and Potsdam, 2001, Exx. 1b, 584):

(25) Enir užā magalu b-āćrułi b-iyxo.
mother boy bread.ABS(III) III-ate III-know
'The mother knows [that (as for the bread) the boy ate it].'

(26) Enir užā magalu b-āćrułi r-iyxo.
mother boy bread.abs(III) III-ate IV-know
'The mother knows [that the boy ate bread].'

Preminger argues that such facts suggest that probe-goal dependencies actually arise from a plurality of sources, with ϕ-probing and negative concord reflecting fundamentally different types of morpho-syntactic dependency. If correct, this then implies that the two phenomena cannot both be handled in terms of Upward Agree.

4.1.3 Bi-directional probe-goal relations

As we have seen above, the classic Chomskyan conception of Agree, though explicitly defined as an operation involving downward probing, actually involves both downward probing (of the probe toward the goal) and reverse upward probing (of the goal toward the probe). As such, it is legitimate to interpret even this classic version of Agree as being bi-directional. More recent theories of Agree are more explicitly bi-directional.

For instance, Baker (2008b) argues that Agree relations for ϕ- and Case are subject to the parametric variation in (27):

(27) *Case Dependence of Agreement Parameter (CDAP):*
$\mathbb{F}$ agrees with DP/NP only if $\mathbb{F}$ values the Case feature of DP/NP (or vice versa)

Yes	No
most Indo-European languages	most Niger-Congo languages

For Baker, the parameter in (27) accounts for the fact that, in languages like Kinande, the finite verb ϕ-Agrees with the closest DP, regardless of its grammatical function (cf. Kinande (28)-(29), taken from Baker, 2008b, Exx. 4a & 6b, 158). It is an open question whether Niger-Congo languages exhibit structural case (Perez, 1985; Diercks, 2013, a.o.) but even if they do, ϕ-agreement is not dependent on it, as is illustrated by the agreement with the locative argument in (29).

(28) Abakali mo-ba-seny-ire olukwi (lw'-omo-mbasa).
Women.2 AFF-2S/T-chop-EXT wood.11 LK11-LOC.18-axe.9
'The woman chopped wood (with an axe).'

(29) Oko-mesa kw-a-hir-aw-a ehilanga.
Leela-DAT hunger-PST-3NSG
'On the table were put peanuts.'

But more needs to be said here as well, as it is not just the closest DP that the finite verb agrees with, but rather the closest higher DP (at least, with respect to surface position). For Baker, this is derived by a second parameter which regulates the directionality of Agree, as in (30):

(30) *Direction of Agreement Parameter (DAP):*
$\mathbb{F}$ Agrees with DP/NP only if DP/NP asymmetrically c-commands $\mathbb{F}$

Yes	No
most Indo-European languages	most Niger-Congo languages

Thus, Baker's model of Agree is neither strictly downward (as it is for Bošković) nor downward and then upward (as under the classic Chomskyan model). Rather, the directionality of Agree is itself hardwired as a parameter which varies across languages.

In even more recent work, Bjorkman and Zeijlstra (2019) propose that feature-checking and feature-valuation carve out distinct syntactic dependencies (cf. also the discussion in Section 3.2.3) which apply in different directions. In particular, they follow (Arregi and Nevins, 2012), who propose that Agree is a composite dependency split across two operational sub-components: Agree-Link (taking place in narrow syntax) and Agree-Copy (taking place in the postsyntactic morphology). For (Bjorkman and Zeijlstra, 2019), uninterpretable features must be checked by their c-commanding matching interpretable counterparts and thus involve Upward Agree; but unvalued features can subsequently be the target of a separate operation of valuation holding between two elements that already stand in a feature-checking relation, which can proceed in a downward direction. This results in a hybrid model where feature-checking always instantiates Upward Agree while feature-valuation, understood as a secondary operation parasitic on feature-checking, may instantiate Downward Agree.

The relativized probing model of Agree (Fernandez and Albizu, 2000; Rezac, 2003; Béjar and Rezac, 2009; Carstens, 2016; Dash and Keine, 2021) discussed in detail in Section 3.1.3 is also bi-directional: here, again, it is proposed that valuation can, in principle, be bidirectional but it is assumed (contra Bjorkman and Zeijlstra, 2019) that Downward Agree is primary and Upward Agree secondary. Other hybrid bi-directional approaches in a similar spirit include proposals in Adger (2003); Merchant (2011); Pesetsky and Torrego (2007); Baker (2008b); Carstens and Diercks (2013); Arregi and Hanink (2018); Clem (2019); McFadden and Sundaresan (To Appear), among others.

4.2 When c-command isn't enough: the role of locality in Agree

Ultimately, c-command between two elements X and Y is not a sufficient condition for Agree to obtain between them, even if X and Y satisfy the featural requirements of being a matching probe-goal pair. Specifically, X and Y also have to be syntactically visible to one another, either by virtue of being directly or cyclically local to one another. Paraphrasing part of our definition in (1), there must be "possibly trivial sequence $X_1 \ldots X_n$ where every X_i is phase-local, minimal and in a c-command relation with X_{i+1} & X_n is phase-local, minimal and in a c-command relation with Y". The issue is further complicated by the fact that syntactic locality itself lends itself to different interpretations (note that our working definition above itself distinguishes minimality from phase-locality). We can articulate (at least) three types of syntactic locality: absolute (or domain-based) locality, relative (or intervention-based) locality, and path-based locality — all of which have been invoked in the syntactic literature — and directly influence how probe-goal structural relations are conceived.

Domain-based locality can be defined as follows: *XP* constitutes a locality domain iff properties inherent to *XP* restrict operations across it. The simplest (domain-)local configuration is one where *X* and *Y* are already in the same locality domain, as in the Direct Agree model in (31), repeated from (2), and can be illustrated by classic cases of subject agreement, as in (32), repeated from (3):

(31) Direct Agree: $[_{phase}\ \mathcal{P} \ldots \mathcal{G}]$

(32) Susan$_{\mathcal{G}}$ make-s$_{\mathcal{P}}$ all her guests drink tea.

A more complex expression of locality is in terms of (successive-)cyclicity. Here, what initially looks like a single unbounded dependency turns out to be comprised of a series of dependencies, each of which is itself local (McCloskey, 1979; Chung, 1998; Lahne, 2009; Korsah and Murphy, 2020), as in (33), repeated from (4):

(33) Cyclic Agree: $[_{phase}\ \mathcal{P}_1 \ldots [_{phase}\ \mathcal{P}_2 \ldots [_{phase} \ldots. [_{phase}\ \mathcal{P}_i \ldots \mathcal{G}]]]]$

The idea of domain-based locality, involved in both direct and mediated dependencies, goes back to *Barriers* (Chomsky, 1986). In Minimalism (Chomsky, 2001, et seq.), they are modelled in terms of *phases*, categorially-defined, semi-permeable locality domains. Under the strong Phase Impenetrability Condition (PIC), " In phase α with head H, the domain of H is not accessible to operations outside α; only H and its edge are accessible to such operations." (cf. Ex. (12) Chomsky, 2000, 108): cyclic locality is thus possible only it is mediated through material at every intervening phase edge.

And yet, disagreement about *what* domains may be phasal abounds. As observed by Polinsky (2003); Bobaljik (2008), long-distance agreement may target the edge of an embedded CP clause but not material that is more deeply embedded (see also Bruening, 2001; Pollock, 1989; Branigan and MacKenzie, 2002). This is quite plausibly a CP-phase effect: once the matrix ϕ-probe enters the structure, only the edge of the embedded clause is accessible, imposing a principled limit on the distance of long-distance agreement (though see Section 3.1.3 for a proposal by Halpert (2019) along the lines of Béjar and Rezac (2009) to "unlock" the phase to see into its domain in cases of hyperraising). At the same time, the phasal status of non-finite, and semi-finite CPs remains unresolved (Boeckx, Hornstein, and Nunes, 2010; Landau, 2012) since these are arguably more transparent for Agree dependencies across them. The phasal status of the *v*P is potentially even more controversial. Chomsky (2001) proposes that all *v*Ps except those that are unaccusative and passive are phasal. In more recent work, (see e.g. van Urk and Richards, 2015; Heck and Himmelreich, 2017) present independent empirical arguments supporting the phasal status of *v*Ps. At the same time, (Keine, 2020), has shown, for Hindi/Urdu, that ϕ-agreement can cross an arbitrary number of vPs which suggests that the vP is never a phase. Also debated is the phasal status of DPs (Legate, 2002; Chomsky, 2005), PPs (Abels, 2003; Svenonius, 2010), and even every phrase (Müller, 2004; Heck and Müller, 2007). These different conceptions of (domain-based) locality have distinct and significant consequences for how probe-goal relations are formalized: in particular, they directly regulate the syntactic conditions under which a potential probe and goal may be visible to one another.

Relative locality is defined, not in terms of domains, but in terms of intervention. Concretely, such locality is *relativized* to the properties of a specific probe, goal, and intervener. (Relativized) Minimality (Rizzi, 1990b), a variant of this, states that, in order for a dependency between X and Y (where X c-commands Y) to obtain for some syntactic feature α, X cannot c-command an element Z marked for α, which in turn c-commands Y (see Keine, 2016, 2019, for a different variant of relative locality, modelled in terms of probes and "horizons" of accessibility):

(34) $[\ldots \mathrm{X}_\alpha \ldots [_{ZP}\ \mathrm{Z}_\alpha \ldots [\ \mathrm{Y}_\alpha]]]$

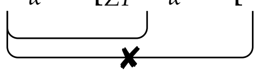

In Minimalism, this idea is modelled in terms of the *Minimal Link Condition* but also surfaces under different guises, such as *Minimal Search* (Chomsky, 1995). Applied to Agree, relative locality boils down to the requirement that a probe can only target a goal if there is no other potential goal that intervenes.

A straightforward structural instantiation of this that, in cases of ϕ-agreement, a T probe in e.g. English cannot ϕ-Agree with an object goal past an intervening subject which is merged higher than the object. But as we have already seen in detail (cf. Sections 3.1.2 and 3.1.3), intervention can be defined in featural as well as structural terms: it is thus not required that the controller of agreement on a probe be the structurally highest potential goal in the search-space of a probe. ϕ-agreement in Hindi/Urdu seems to be regulated by both structural and featural intervention (as described in (35) in Keine, 2018):

(35) If the subject does not bear a case marker N, Agree with the subject;
→ Otherwise: If the object does not bear N, Agree with the object;
→ Otherwise: Use masculine singular default agreement.

Another phenomenon, discussed multiple times here, that has received explanations in terms of both featural and structural intervention has to do with the fact that, in many languages, it is not possible to agree in ϕ-features with a DP that bears inherent case. The solutions proposed for this in terms of the Activity Condition or case-discrimination (Bobaljik, 2008) invoke the general notion of featural intervention. But this phenomenon has also been analyzed in terms of structural intervention. E.g. Řezáč (2008), who refers to this constraint as *Case Opacity*, proposes that oblique case projects a PP phase shell above the DP which typically has the effect of structurally trapping the ϕ-features at the DP level where it is no longer syntactically visible to the probe. It follows straightforwardly that if a subject is no longer acts as a potential goal, the object becomes the the highest possible goal. When both the subject and the object bear oblique case, the ϕ-probe cannot find any accessible goals and the agreement is set to default agreement (or, in terms of Preminger, 2014, an instance of failed agreement emerges).[10] But in certain instances, the P head itself bears a ϕ-probe creating a phasal escape hatch for mediated Agree with the ϕ-features on the DP in the putative lower phase. Rezac argues that this explains the availability of ϕ-covarying agreement with oblique nominals in languages like Basque. In recent work, Preminger (2019) uses an analogous explanation to explain AAE effects crosslinguistically. I.e. for Preminger, an anaphor is prevented from triggering co-varying ϕ-agreement on a clausemate verb because the ϕ-features on the anaphor are trapped under a phasal AnaphP shell. In languages where the AAE effect seems to be obviated, Preminger proposes that the AnaphP is exceptionally non-phasal, rendering the ϕ-features on the anaphor syntactically visible to a ϕ-probe on T or v.

Finally, the third conception of path-based locality is defined, not in terms of the conditions under which probe-goal dependencies are *blocked*, but in terms of when they are *allowed*. Path-based Locality states that two elements X and Y are syntactically visible to each other iff they are connected by an uninterrupted sequence of steps, each of which satisfies the same (syntactic) condition. Such an approach was espoused in certain proposals within the GB framework (see e.g. Pesetsky, 1982, and Kayne, 1984). Analyses in this spirit have also regulated notions of locality in other grammatical frameworks like HPSG/LFG (functional uncertainty in Kaplan and Zaenen, 1989), CCG (Steedman, 1996) and TAG (Kroch, 1989). But path-based locality has not yet found as much currency within Minimalism (though see recent work in McFadden, Sundaresan, and Zeijlstra, 2019; McFadden et al., 2021, for a Minimalist proposal along these lines).

[10] A problem for this line of reasoning concerns cases of so-called *defective intervention*. E.g. while an oblique-case marked DP is invisible for ϕ-Agree, it nevertheless acts as an intervener for Agree past it. Thus, in Icelandic (i), taken from Bobaljik (2008), an dative DP intervene between a ϕ-probe and a nominative object and block agreement between them; at the same time, the dative cannot itself serve as a ϕ-goal. There are thus no viable goals for ϕ-Agree in this sentence, resulting in default agreement:

i. Mér *virð-ast/virð-ist Jóni vera taldir líka hestarnir.
Me.DAT seem-3PL/seem-3SG John.DAT to.be believed.PL to.like horses.NOM.PL
'John seems to me to be believed to like horses.'

Preminger (2014) dubs this the dative paradox: Why is it that a dative that appear to be inert for the purposes of ϕ-feature valuation is simultaneously able to prevent the probe from searching further for a non-inert agreement target? Solutions in terms of both structural and featural intervention have been proposed.

4.3 When c-command isn't necessary: conjunct agreement

Cases of so-called *Closest Conjunct Agreement* (CCA, see e.g. Nevins and Weisser, 2018) as in Welsh (36) and Hindi (37) arguably challenge the idea that c-command is a necessary condition for Agree between a probe and goal:

(36) Gwelais ti a Megan ein hunain.
see.PST.2SG you.SG and Megan 2PL SELF
'You$_i$ and Megan$_j$ saw yourselves$_{i+j}$.'

(37) maiñ-ne ek chaataa aur ek saarii khariid-ii.
I-ERG an umbrella.ABS.MSG and a saaree.ABS.FSG buy-PERF.FSG
'I bought an umbrella and a saaree.'

A fairly standard assumption is that complex, coordinated DPs form Conjunction Phrases (ConjPs) have the templatic structure in (38):

(38)
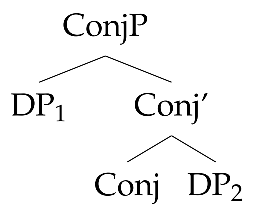

This in turn makes the conjunct agreement in (36)-(37) puzzling for two reasons: (i) the embedded DP inside the ConjP is not in a c-command relation with the ϕ-probe, thus should not be visible to it as a ϕ-goal: (ii) the ConjP as a whole *is* in a c-command relation with the ϕ-probe and should thus bleed ϕ-Agree with a DP inside it.

With respect to the latter, however, it is important to note that CCA is often not a necessary option in a language but alternates with resolved agreement (analyzed as ϕ-Agree with the entire ConjP) or default agreement, as illustrated for Slovenian below (taken from Marusic, Nevins, and Badecker, 2015):

(39) Teleta in krave so odsle/odsla/odsli na paso.
cow.F.PL and calf.N.PL aux.PL went.F.PL/went.N.PL/went.M.PL on graze
'Calves and cows went grazing.'

Such data suggest that it is not the case that the ConjP is always syntactically invisible to the ϕ-probe. As for the first question of how a non-c-commanding DP conjunct can be visible to a ϕ-probe in the first place, it is worth noting that CCA always involves agreement with the conjunct that is *linearly closest* to the ϕ-probe. In the head-initial Welsh example in (36), this is the first conjunct in the head-initial Welsh example (36) and the final conjunct; in the head-final Hindi example in (37), it is the final one. This in turn suggests that linearity may play a greater (or equally significant) role to c-command in regulating the visibility between a probe and potential goal. Various approaches have been formulated to reconcile CCA with the idea that the structural relation between a probe and a goal in narrow syntax is one of c-command and one that is subject to relativized minimality. Below, we briefly outline three such approaches (see also Nevins and Weisser, 2018, for a detailed overview and discussion).

The first approach is based on the notion of *equidistance*. Under equidistance, two elements in nearby structural positions, here ConjP and its DP specifier, count as equally close to a probing head and are therefore both possible agreement goals (Chomsky, 1995). One of the reasons put forward for why both phrases count as equidistant in this sense is that there is no asymmetric c-command relation between the two. For (van Koppen, 2005, 2007), as a result, the probe can agree with both equidistant goals, yielding optionality between resolved and embedded conjunct agreement. A potential drawback of equidistance approaches is that they generally only apply to cases of First Conjunct Agreement (FCA), as it is only the specifier and not the complement that is equidistant to the ConjP (though see Bošković, 2009, for an attempt to treat both First Conjunct Agreement and Final Conjunct Agreement in Bosnian/Serbian/Croatian in such terms).

The second type of approach involves splitting up the agreement process into two sequential operations. The first takes place in narrow syntax, which has only access to structural relations, and the other is implemented at PF, where linear orders are also visible. Unlike equidistance accounts, this type of approach makes explicit reference to linearity. One example of such an approach is Marusic et al. (2015) which follows (Bošković, 2009) in proposing that, in South-Slavic languages, the ConjP is deficient for gender features while its individual conjuncts are featurally specified for gender. As a result, a ϕ-probe can have its number but not its gender features valued when it agrees with ConjP. One option, then, is that the probe remains gender-unvalued which triggers default masculine gender agreement. But another option is that a second operation of Agree applies beyond ConjP, but this time at PF. At PF, the probe can get ϕ-valued either before or after linearization. Before linearization, the probe can only access the structurally higher conjunct; after linearization, by contrast, the probe can only access the linearly closest conjunct. In both cases, the missing gender features will be valued, either by the highest conjunct or by the linearly closest one, yielding the desired optionality.[11]. A similar approach has been developed by (Bhatt and Walkow, 2013) to account for the fact that in Hindi subjects can trigger resolved agreement, but objects trigger CCA.

Finally, Murphy and Puškar (2018) develop a third approach for CCA. Here, the ConjP always serves as a goal for the ϕ-probe but, in particular configurations alone, the ϕ-features of the individual conjuncts may percolate up to the level of

[11]Note that when it comes to preverbal agreement, both cases of PF-based gender agreement will be with the same conjunct and result in First Conjunct Agreement; when it comes to postverbal agreement, pre-linearization PF-agreement will result in First Conjunct Agreement and post-linearization PF-agreement in Final Conjunct Agreement. Hence, this proposal can deal with both types of CCA exemplified before.

ConjP. For Murphy and Puškar (2018), depending on the direction of agreement inside the clause, the features of the highest conjunct may or may not appear in ConjP as well. With post-verbal agreement, the features of the highest conjunct may then be present on the ConjP as well, thus giving rise to the superficial effect of CCA. Crucially, unlike the previous approach, but like the equidistance approaches, ϕ-probing for Murphy and Puškar is insensitive to linearity effects even in cases of CCA and involves nothing but ϕ-Agree with the closest c-commanded goal.

To sum up, CCA appears a challenge for approaches that take c-command and relative locality to underlie the structural relation between probes and goals, but depending on the exact analysis of CCA, this may prove to be a real challenge or only an illusory one.

5 Concluding remarks

In this paper we have attempted to define the featural and structural properties that characterize a probe-goal pair in cases of Agree. We have seen that an Agree relation between a probe and goal is possible only if there exist symmetries as well as asymmetries between the two elements at the featural and structural levels. On the one hand, the two elements must share core featural properties (e.g. matching features) in order to be identified as a matching probe-goal pair; they must also, in the typical case, share structural properties like locality and c-command. On the other hand, the two elements must also be featurally and structurally distinguishable, so they can enter into an Agree relation.

Generally speaking, a probe is deficient for one or more syntactic features which the goal bears; Agree is then viewed as an operation which is triggered to resolve that featural asymmetry. Structurally, either the probe must c-command the goal or vice-versa (with cases of CCA suggesting, as mentioned above, that such c-command perhaps cannot be symmetric), which in turn results in probing proceeding either downward or upward (or both). The definition of probes and goals, repeated in (40) below is, as we have said, deliberately framed in broad terms so as to be compatible with many variants of Agree:

(40) *Working definitions of probe and goal:*
$\mathcal{G}$, a syntactic element of category α, is a suitable goal for a probe $\mathcal{P}$, another syntactic element of category β (where β might be = α) iff:
  a. $\mathcal{P}$ is deficient for a (potentially unary) set of features $\mathbb{F}$ which $\mathcal{G}$ bears;
  b. $\mathcal{P}$ and $\mathcal{G}$ are syntactically visible to one another: i.e. there is a possibly trivial sequence $\mathcal{P}_1 \ldots \mathcal{P}_n$ where every $\mathcal{P}_i$ is phase-local, minimal and in a c-command relation with $\mathcal{P}_{i+1}$ & $\mathcal{P}_n$ is phase-local, minimal and in a c-command relation with $\mathcal{G}$;

Using this broad definition as a springboard, we have explored the syntactic limits of featural and structural (a)symmetries holding between a probe and goal and discussed their empirical consequences. Needless to say, what we have captured here is only a small selection of what has been discussed and described in the literature on Agree. Such a selection is always subjective and other authors most likely would have made different choices. Nevertheless, we hope that our selection will provide readers with a broad and clear impression of what has been, and can be, said about the featural and structural (a)symmetries between probes and goals.

References

Abels, Klaus. 2003. The nature of adposition stranding and universal grammar. Doctoral Dissertation, University of Connecticut.

Abusch, D. 1997. Sequence of tense and temporal de re. *Linguistics and Philosophy* 20:1–50.

Ackema, Peter, and Ad Neeleman. 2013. Subset controllers in agreement relations. *Morphology* 23:291–323. DOI 10.1007/s11525-013-9218-4.

Adger, David. 2003. *Core syntax*. Oxford: Oxford University Press.

van Alem, Astrid. 2020. Complementizer agreement in West-Germanic is clitic doubling. Ms., Leiden University.

Alexiadou, Artemis, and Elena Anagnostopoulou. 1999. Raising without infinitives and the nature of agreement. In *Proceedings of WCCFL*, ed. Sonya Bird, Andrew Carnie, Jason Haugen, and Peter Norquest, volume 18, 15–25. Somerville, MA: Cascadilla Press.

Arregi, Karlos, and Emily Hanink. 2018. Switch reference in Washo as multiple subject agreement. In *Proceedings of NELS 48*.

Arregi, Karlos, and Andrew Nevins. 2012. *Morphotactics: Basque auxiliaries and the structure of spellout*. Dordrecht: Springer.

Baker, Mark. 2008a. The macroparameter in a microparametric world. In *The limits of syntactic variation*, ed. Theresa Biberauer, 351–373. Amsterdam: Benjamins.

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

Bárány, András. 2015. Differential object marking in Hungarian and the morphosyntax of case and agreement. Doctoral Dissertation, University of Cambridge.

Béjar, Susana, and Milan Rezac. 2009. Cyclic agree. *LI* 40:35–73.

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

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

Bjorkman, Bronwyn. 2016. Go get, come see: motion verbs, morphological restrictions, and syncretism. *Natural Language and Linguistic Theory* 34.

Bjorkman, Bronwyn, and Hedde Zeijlstra. 2019. Checking up on (Phi-)Agree. *Linguistic Inquiry* 50:527–569.

Bobaljik, Jonathan David. 2008. Where's ϕ? In *Phi-theory: Phi features across interfaces and modules*, ed. Daniel Harbour, David Adger, and Susana Béjar, 295–328. Oxford: Oxford University Press.

Boeckx, Cedric, Norbert Hornstein, and Jairo Nunes. 2010. *Control as movement*. Cambridge: Cambridge University Press.

Bošković, Željko. 2007. On the locality and motivation of Move and Agree: An even more minimal theory. *Linguistic Inquiry* 38:589–644.

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

Branigan, Phil, and Marguerita MacKenzie. 2002. Altruism, a-movement and object agreement in innu-aimûn. *Linguistic Inquiry* 33:385–407.

Brody, Michael. 1997. Perfect chains. In *Elements of grammar*, ed. Liliane Haegeman, 139–167. Dordrecht: Kluwer Academic Publishers.

Bruening, Benjamin. 2001. Syntax at the edge: Cross-clausal phenomena and the syntax of passamaquoddy. Doctoral Dissertation, MIT, Cambridge, MA.

Carstens, Vicki. 2011. Hyperactivity and hyperagreement in Bantu. *Lingua* 121:721–741.

Carstens, Vicki. 2016. Delayed valuation: A reanalysis of goal features, "upward" complementizer agreement, and the mechanics of case. *Syntax* 19:1–42. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/synt.12116.

Carstens, Vicki, and Michael Diercks. 2013. Parametrizing case and activity: hyperraising in Bantu. In *Proceedings of NELS 40*, ed. Seda Kan, Claire Moore-Cantwell, and Robert Staubs, 99–118. Amherst, MA: GLSA.

Carstens, Vicky. 2003. Rethinking complementizer agreement: Agree with a Case-checked goal. *Linguistic Inquiry* 34:393–412.

Chierchia, Gennaro. 2013. *Logic in grammar: Polarity, free choice, and intervention*. Oxford: Oxford University Press.

Chomsky, Noam. 1973. Conditions on transformations. In *A festschrift for Morris Halle*, ed. Steven Anderson and Paul Kiparsky. New York: Holt, Rinehart, and Winston.

Chomsky, Noam. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.

Chomsky, Noam. 1986. *Knowledge of language*. New York: Praeger.

Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, Mass.: MIT Press.

Chomsky, Noam. 2000. Minimalist inquiries: the framework. In *Step by step: Essays on minimalism in honor of Howard Lasnik*, ed. Roger Martin, David Michaels, and Juan Uriagereka. Cambridge, Mass.: MIT Press.

Chomsky, Noam. 2001. Derivation by phase. In *Ken Hale: A life in language*, ed. Michael Kenstowicz. Cambridge, Mass.: MIT Press.

Chomsky, Noam. 2005. On phases. Ms. MIT.

Chung, Sandra. 1998. *The design of agreement: evidence from Chamorro*. Chicago:

University of Chicago Press.

Clem, Emily. 2019. Agreement, case, and switch-reference in Amahuaca. Doctoral Dissertation, University of California, Berkeley, Berkeley.

Coon, Jessica, and Stefan Keine. To appear. Feature gluttony. *Linguistic Inquiry* .

D'Alessandro, Roberta. This volume. A short history of Agree. In *Cambridge Handbook of Minimalism*, ed. Kleanthes Grohmann and Evelina Leivada. Cambridge: CUP.

Dash, Bhamati, and Stefan Keine. 2021. The ups and downs of agreement. *Ms. USC/UCLA* .

Deal, Amy Rose. 2015. Interaction and satisfaction in $\phi-$agreement. In *Proceedings of NELS 45*, ed. Thuy Bui and Deniz Özyildiz, 179–192. Amherst, MA: GLSA Publications.

Deal, Amy Rose. 2017. Covert hyperraising to object. In *Proceedings of NELS 47*, 256–270. Amherst: GLSA.

Diercks, Michael. 2013. Indirect Agree in Lubukusu complementizer agreement. *Natural Language and Linguistic Theory* 31:357–407.

Diercks, Michael, Marjo van Koppen, and Michael Putnam. 2020. Agree probes down: Anaphoric feature valuation and phase reference. In *Agree to Agree: Agreement in the Minimalist Programme*, ed. Peter W. Smith, Johannes Mursell, and Katharina Hartmann, 347–389. Berlin: Language Science Press. DOI:10.5281/zenodo.354176.

Fernandez, Beatriz, and Pablo Albizu. 2000. Ergative displacement in basque and the division of labor between syntax and morphology. In *Proceedings of the 36th meeting of the chicago linguistic society*, ed. John Boyle and Arika Okrent, 103–117. Amsterdam: Chicago Linguistic Society.

Freidin, Robert, and Rex Sprouse. 1991. Lexical case phenomena. In *Principles and parameters in comparative grammar*, ed. Robert Freidin. Cambridge, Mass.: MIT Press.

Fuß, Eric. 2008. Multiple agreement and the representation of inflection in the C-domain. *Linguistische Berichte* 213:77–106.

Georgi, Doreen. 2017. Patterns of movement reflexes as the result of the order of merge and agree. *Linguistic Inquiry* 48:585–626.

Grosz, Patrick, and Pritty Patel-Grosz. 2014. Agreement and verb-types in Kutchi Gujarati. In *The lexicon-syntax interface: perspectives from South Asian languages*, ed. P. Chandra and R. Srishti, 217–243. Amsterdam: John Benjamins.

Haegeman, Liliane, and Terje Lohndal. 2010. A new approach to negative concord and multiple agree: A case study of west flemish. *Linguistic Inquiry* 41:181–211.

Halpert, Claire. 2019. Raising unphased. *Natural Language and Linguistic Theory* 37:123–165.

Heck, Fabian, and Anke Himmelreich. 2017. Opaque intervention. *LI* 48:47–97.

Heck, Fabian, and Gereon Müller. 2007. Extremely local optimization. In *Proceedings of WECOL 26*, 170–182.

Heinat, Fredrik. 2008. *Probes, pronouns and binding in the Minimalist Program*. VDM

Verlag.

Hicks, Glyn. 2009. *The derivation of anaphoric relations*. Linguistik Aktuell. John Benjamins.

Hiraiwa, Ken. 2004. Dimensions of symmetry in syntax: Agreement and clausal architecture. Doctoral Dissertation, MIT.

Iatridou, Sabine. 1993. On nominative case assignment and a few related things. *MIT Working Papers in Linguistics* 19:175–196.

Kalin, Laura. 2018. Licensing and differential object marking: The view from Neo-Aramaic. *Syntax* 21:112–159.

Kalin, Laura. 2020. Opacity in agreement. In *Agree to Agree: Agreement in the Minimalist Programme*, ed. Peter W. Smith, Johannes Mursell, and Katharina Hartmann, 149–177. Berlin: Language Science Press.

Kaplan, Ronald, and Annie Zaenen. 1989. Long-distance dependencies, constituent structure, and functional uncertainty. In *Alternative conceptions of phrase structure*, ed. Mark Baltin and Anthony Kroch, 17–42. Chicago, IL: Chicago University Press.

Kayne, Richard S. 1984. *Connectedness and binary branching*. Dordrecht: Foris.

Keine, Stefan. 2016. Probes and their horizons. Doctoral Dissertation, University of Massachusetts, Amherst.

Keine, Stefan. 2018. Case vs. positions in the locality of a-movement. *Glossa* 3:1–34. DOI: https://doi.org/10.5334/gjgl.520.

Keine, Stefan. 2019. Selective opacity. *LI* 50:13–62.

Keine, Stefan. 2020. *Probes and their horizons*. Cambridge, MA: MIT Press.

van Koppen, Marjo. 2005. One probe, two goals: Aspects of agreement in Dutch dialects. Doctoral Dissertation, Leiden University.

van Koppen, Marjo. 2007. Agreement with coordinated subjects: A comparative perspective. *Linguistic Variation Yearbook* 7.

Korsah, Sampson, and Andrew Murphy. 2020. Tonal reflexes of movement in Asante Twi. *NLLT* 38:827–885.

Kratzer, Angelika. 2009. Making a pronoun: fake indexicals as windows into the properties of pronouns. *Linguistic Inquiry* 40:187–237.

Kroch, Anthony. 1989. Reflexes of grammar in patterns of language change. *Language Variation and Change* 1:199–244.

Kusumoto, K.İ999. Tense in embedded contexts. Doctoral Dissertation, UMass, Amherst, Amherst, MA.

Lahne, Antje. 2009. Where there is fire there is smoke. local modelling of successive- cyclic movement. Doctoral Dissertation, University of Leipzig, Leipzig.

Landau, Idan. 2012. *Control in generative grammar: A research companion*. Cambridge University Press.

Landau, Idan. 2013. *Control in generative grammar: A research companion*. Cambridge University Press.

Legate, Julie Anne. 2002. Walpiri: Theoretical implications. Doctoral Dissertation, MIT, Cambridge, MA.

Legate, Julie Anne. 2005. Phases and cyclic agreement. *MITWPL* 49:147–156.

Marusic, Franc, Andrew Nevins, and Bill Badecker. 2015. The grammars of conjunction agreement in slovenian. *Syntax* 18:39–77.

May, Robert. 1979. Must Comp-to-Comp movement be stipulated? *Linguistic Inquiry* 10:719–725.

McCloskey, James. 1979. *Transformational syntax and modern theoretical semantics: a case study in Modern Irish*. Dordrecht: Reidel.

McFadden, Thomas. 2004. The position of morphological case in the derivation: a study on the syntax-morphology interface. Doctoral Dissertation, University of Pennsylvania, Philadelphia. `http://ifla.uni-stuttgart.de/~tom/diss.html`.

McFadden, Thomas. 2020. The morphosyntax of allocutive agreement in Tamil. In *Agree to Agree: Agreement in the Minimalist Programme*, ed. Peter W. Smith, Johannes Mursell, and Katharina Hartmann, 391–424. Berlin: Language Science Press.

McFadden, Thomas, and Sandhya Sundaresan. 2011. Nominative case is independent of finiteness and agreement. Presented at BCGL 5: Case at the interfaces.

McFadden, Thomas, and Sandhya Sundaresan. 2018. Reducing *pro* and PRO to a single source. *The Linguistic Review* To appear in GLOW 2016 special issue. lingbuzz/003560.

McFadden, Thomas, and Sandhya Sundaresan. To Appear. Unifying species of C-agreement. In *Proceedings of NELS 51*.

McFadden, Thomas, Sandhya Sundaresan, and Hedde Zeijlstra. 2019. Deriving selective opacity in adjuncts. Presentation at GLOW 42, Oslo.

McFadden, Thomas, Sandhya Sundaresan, and Hedde Zeijlstra. 2021. Towards a path-based theory of syntactic locality and selective opacity. ZAS, Berlin and University of Göttingen.

Merchant, Jason. 2011. Aleut case matters. In *Pragmatics and autolexical grammar: In honor of jerry sadock*, ed. Tista Bagchi Etsuyo Yuasa Yuasa and Katharine P. Beals, 382–411. Amsterdam: John Benjamins.

Miyagawa, Shigeru. 2017. *Agreement Beyond Phi*. Cambridge, Mass.: MIT Press.

Müller, Gereon. 2004. Phrase impenetrability and wh-intervention. In *Minimality effects in syntax*, ed. Arthur Stepanov, Gisbert Fanselow, and Ralph Vogel, 289–325. Berlin: Mouton de Gruyter.

Murphy, Andrew, and Zorica Puškar. 2018. Closest conjunct agreement is an illusion. *Natural Language & Linguistic Theory* 36:1207–1261.

Murugesan, Gurujegan. 2019. Predicting the Anaphor Agreement Effect and its violations. Doctoral Dissertation, University of Leipzig.

Neeleman, Ad, and Hans van de Koot. 2002. The configurational matrix. *Linguistic Inquiry* 33:529–574.

Nevins, Andrew. 2011. Multiple Agree with clitics: person complementarity vs. omnivorous number. *Natural Language and Linguistic Theory* 29:939–971.

Nevins, Andrew, and Philipp Weisser. 2018. Closest conjunct agreement. *Annual Review of Linguistics* 1–25.

Obata, Miki, and Samuel David Epstein. 2011. Feature-splitting Internal Merge: Improper movement, intervention, and the A/A' distinction. *Syntax* 14:122–147. DOI: https:// doi.org/10.1111/j.1467-9612.2010.00149.

Ogihara, T.İ989. Temporal reference in English and Japanese. Doctoral Dissertation, University of Texas, Austin, Austin.

Oxford, Will. 2017. The Activity Condition as a microparameter. *Linguistic Inquiry* 48:711–722.

Perez, C. Harford. 1985. Aspects of complementation in three bantu languages. Doctoral Dissertation, University of Wisconsin.

Pesetsky, David. 1982. Paths and categories. Doctoral Dissertation, MIT, Cambridge, MA.

Pesetsky, David, and Esther Torrego. 2004. Tense, case and the nature of syntactic categories. In *The syntax of time*, ed. J. Guéron and J. Lecarme, 495–538. Cambridge, MA: MIT Press.

Pesetsky, David, and Esther Torrego. 2007. The syntax of valuation and the interpretability of features. In *Phrasal and clausal architecture: syntactic derivation and interpretation*, ed. Simin Karimi, Vida Samiian, and Wendy Wilkins. Amsterdam: Benjamins.

Polinsky, Maria. 2003. Non-canonical agreement is canonical. *Transactions of the Philological Society* 101:279–312.

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

Pollock, Jean-Yves. 1989. Verb movement, Universal Grammar, and the structure of IP. *Linguistic Inquiry* 20:365–424.

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

Preminger, Omer. 2014. *Agreement and its failures*. Cambridge, MA: MIT Press.

Preminger, Omer. 2019. The Anaphor Agreement Effect: further evidence against binding-as-agreement. University of Maryland, https://ling.auf.net/lingbuzz/004401.

Preminger, Omer. This volume. Phi-feature agreement in syntax. In *Cambridge Handbook of Minimalism*, ed. Kleanthes Grohmann and Evelina Leivada. Cambridge: CUP.

Puškar, Zorica. 2018. Interactions of gender and number agreement: evidence from Bosnian/Croatian/Serbian. *Syntax* 21:257–318.

Rackowski, Andrea, and Norvin Richards. 2005. Phase edge and extraction: A Tagalog case study. *Linguistic Inquiry* 36:565–599.

Reuland, Eric. 2011. *Anaphora and language design*. Cambridge, MA: MIT Press.

Rezac, Milan. 2003. The fine structure of cyclic agree. *Syntax* 9:156–182.

Rizzi, Luigi. 1990a. *Relativized minimality*. Cambridge, Mass.: MIT Press.

Rizzi, Luigi. 1990b. *Relativized minimality*. Cambridge, Mass.: MIT Press.

Rooryck, Johan, and Guido vanden Wyngaerd. 2011. *Dissolving binding theory*. Oxford: Oxford University Press.

Sigurðsson, Halldór Ármann. 2012. Minimalist C/case. *Linguistic Inquiry* 43:191–227.

Smith, Peter. 2017. The syntax of semantic agreement in english. *Journal of Linguistics* 53:823–863.

Steedman, Mark. 1996. *Surface structure and interpretation*. LI Monographs. Cambridge, MA: MIT Press.

Stowell, T. 2007. The syntactic expression of tense. *Lingua* 117:437–463.

Sundaresan, Sandhya. 2012. Context and (co)reference in the syntax and its interfaces. Doctoral Dissertation, University of Tromsø (CASTL)/Universität Stuttgart.

Sundaresan, Sandhya. 2016. Anaphora vs. agreement: a new kind of Anaphor Agreement Effect in Tamil. In *The impact of pronominal form on interpretation*, ed. Patrick Grosz and Pritty Patel-Grosz. de Gruyter, Mouton.

Sundaresan, Sandhya, and Thomas McFadden. 2009. Subject distribution in Tamil and other languages: Selection vs. Case. *Journal of South Asian Linguistics* 2:5–34.

Svenonius, Peter. 2006. Interpreting uninterpretable features. *Linguistic Analysis* 33:375–413.

Svenonius, Peter. 2010. Spatial P in English. In *Mapping spatial PPs*, ed. G. Cinque and L. Rizzi, 127–160. OUP.

Tucker, Matthew A. 2011. Even more on the anaphor agreement effect: when binding does not agree. Manuscript written at UCSC.

van Urk, Coppe, and Norvin Richards. 2015. Two components of long-distance extraction: Successive cyclicity in Dinka. *Linguistic Inquiry* 46:113–155.

Řezáč, Milan. 2008. Phi-agree and theta-related case. In *Phi theory: Phi-features across modules and interfaces*, ed. D.H̃arbour et al., 83–129. Oxford: Oxford University Press.

Wurmbrand, Susi. 2011. Tense and aspect in English infinitives. Ms., University of Connecticut.

Zeijlstra, Hedde. 2004. Sentential negation and negative concord. Doctoral Dissertation, University of Amsterdam, Utrecht, LOT Publications.

Zeijlstra, Hedde. 2008. On the syntactic flexibility of formal features. In *The limits of syntactic variation*, ed. Theresa Biberauer, 143–174. Amsterdam: Benjamins.

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

Zeijlstra, Hedde. 2014. On the uninterpretability of interpretable features. In *Minimalism and beyond*, ed. Peter Kosta et al., 109–129. Amsterdam: Benjamins.

Zeijlstra, Hedde. 2020. Labeling, selection, and feature checking. In *Minimalism and beyond*, ed. Peter Smith, Johannes Mursell, and Katharina Hartmann, 137–174. Berlin: Language Science Press.

Zeller, Jochen. 2006. Raising out of finite CP in Nguni: the case of *fanele*. *Southern African Linguistics and Applied Language Studies* 24:255–275.

Zwart, C. Jan-Wouter. 1997. *Morphosyntax of verb movement: A Minimalist approach to the syntax of Dutch*. Dordrecht: Kluwer.