February 2, 2022

Current models of Agree

Amy Rose Deal

1 Aims of syntactic theory

To know a language is to store a set of simplex lexical items and to master the rules for combining them. Accordingly, inquiry in syntax can be roughly divided into two sorts, one concerned with the nature of lexical primitives and one concerned with the nature of combinatoric rules. Questions of interest to syntacticians studying combinatorics include: For a given natural language, what types of combinatoric rules are necessary to describe the structures of this language? For natural language in general, what types of combinatoric rules exist overall? To the extent that there is variation across the combinatoric rules of different languages, how is this variation constrained? And finally, why should this set of combinatoric rules, and not some other set, be characteristic of natural language?[1]

Certain answers to these questions have come into view as syntactic inquiry has unfolded over the past several decades. Major themes that emerge are *asymmetry* and *locality*. On the first count, natural language combinatorics are strongly asymmetric—thus we describe them as *dependencies*, indicating a directional relationship whereby one element depends on another. Some dependencies traditionally of interest to syntacticians include selection/subcategorization, movement, binding, case, and subject-verb agreement, or more broadly, ϕ-agreement. On the second count, dependencies of all these types consistently show locality constraints. They do not hold over unbounded amounts of structure. A question that has thus emerged is whether syntactic dependencies can be productively split into "more local" and "less local" classes. A core hypothesis of the last two decades of generative syntax is that such a two-way division is not only empirically productive but also deeply theoretically grounded. Two basic, abstract, universal operations are involved in syntactic structure building, subject to two very different profiles in terms of locality. These are Merge, the source of extremely local, head to head dependency (e.g. selection), and Agree, the source of long-distance dependency.

The most constrained theory positing Merge and Agree is one that posits no additional operations beyond these two. I will call this proposal the *Merge-Agree Hypothesis*: there are exactly two basic operations of natural language syntax, one that holds extremely locally, head to head (Merge),

[1] Such questions form a hierarchy that reflects the various levels of success for a grammatical description laid out by Chomsky (1964, 1965, 2004). An explanation of the observed, primary data by reference to generalizations which represent a speaker's linguistic competence achieves descriptive adequacy. An explanation of speakers' linguistic competence by reference to generalizations which represent laws of language in general achieves explanatory adequacy. An explanation of laws of language in general by reference to generalizations which represent laws of some distinct, more comprehensive type (cognition, evolution, etc.) reaches what Chomsky has termed "beyond" explanatory adequacy.

versus one that can operate at a distance (Agree). If this hypothesis is correct about language universals, there are ultimately just two general kinds of combinatoric rules in any given natural language. Notably, all dependencies at a distance are to be treated with the same operation—"Agree" is a name for that hypothetical operation, whenever it may apply. This perspective is quite different from the one taken in earlier generative theories, such as Government and Binding Theory or Transformational Grammar. Seen against the backdrop of that tradition, research on Agree has a strongly unificatory flavor; dependencies that were thought of as reflecting separate types of syntactic rules are to be brought together under one umbrella rule type.

The success or failure of a unificatory project of this kind is unlikely to be assessable in broad strokes. It rests on the details of the proposed unification. The key question becomes: how exactly should the Agree operation be formulated, so as to capture the various long-distance dependencies characteristic of natural language syntax? In asking the question, let me emphasize that I would take an answer to be an explicit theory of the operation that underlies all long-distance dependencies. While nearly all current models of such an operation have their origin in Chomsky 2000, 2001, I do not reserve the term 'Agree' for the algorithm outlined in those works. Indeed, in the final section of the paper I present a theory of Agree that differs from Chomsky's in various significant respects. This is the interaction/satisfaction theory, introduced in Deal (2015a, to appear) and further developed in a wide range of crosslinguistic work (Baier 2018, 2019, Halpert 2019, Clem 2019a,c, 2021, to appear, Oxford 2020, Roversi 2020, Branan and Erlewine 2020, Alam and Kumaran 2021, Arregi and Hanink 2021, Deal 2021b, Jenks and Hassen 2021, Scott 2021, Wurmbrand 2021).

This paper is structured as an opinionated survey of issues and perspectives in current models of Agree. On one hand, I aim to highlight the empirical successes of the unificatory project, and on the other, to catalogue the various respects in which these successes have unmoored the enterprise from its original motivations in the work of Chomsky 2000, 2001: the LF deletion of uninterpretable features. I begin with a brief introduction to that work in §2. The three sections following review three strands of literature that have chipped away at the foundation in the course of improving the crosslinguistic empirical adequacy of the theory. These center on valuation and relativized probing, in section 3; defaults and failure to value, in section 4; and the question of whether goals must be made "active" by uninterpretable features, in section 5. In section 6, I review an ongoing debate about the directionality of Agree in light of the issues raised for uninterpretable features in sections 3-5. In section 7, I present what I see as a way forward for the theory of Agree: the interaction/satisfaction theory, which provides a new conceptual grounding for Agree that in various respects makes sense of the empirical landscape uncovered by the past two decades of intensive research on this topic.

2 The Agree algorithm: a starting place

Chomsky (2000, 2001) introduces Agree as motivated by a need to remove elements from a derivation that are incompatible with semantic interpretation:

> "The empirical facts make it clear that there are (LF-)uninterpretable inflectional features that enter into agreement relations with interpretable inflectional features. . . . The obvious conclusion, which we adopt, is that the agreement relation removes the un-

> interpretable features from the narrow syntax, allowing derivations to converge at LF while remaining intact for the phonological component" (Chomsky 2001, 3).

From this perspective, Agree arises from the tension between two properties of language. First, syntactic objects are interpreted by syntax-external systems, and at least one such system (LF) is subject to a principle that Chomsky calls "Full Interpretation": it must be able to assign an interpretation to every symbol that is provided to it. Failure to meet this condition prevents the generation of a grammatical object (or in other terms, "crashes the derivation"). Second, some of the features that lexical items bear are not semantically interpretable. Such items cannot be LF interpreted as is, and in this respect, they are inherently deficient. The motivation for Agree lies in resolving this deficiency by deleting uninterpretable features.

A further proposal concerning the features driving Agree connects their uninterpretability to another property, unvaluedness:

> "The natural principle is that the uninterpretable features, and only these, enter the derivation without values, and are distinguished from interpretable features by virtue of this property." (Chomsky 2001, 5)

While (un)interpretability is clearly intended as a reflection of how a feature relates to LF interpretation, the nature of (un)valuedness is a more complex question, one closely tied to the nature of features themselves. Consider, for instance, the representation of plural. Various versions of an attribute-value schema may be applied, e.g. plural as [NUM:PL] (attribute NUM, value PL) in contrast to singular [NUM:SG], or as [+PL] (attribute PL, value +) in contrast to singular [-PL]. These representations make it straightforward to say what an unvalued feature is: it is an attribute without a value. The notion of (un)valuedness becomes more elusive in light of privative feature theories, e.g. those on which plural is [PL] in contrast to the absence of any number feature in the singular. We return to this issue below. For the time being, it will be helpful to think of features in attribute-value terms.

Let us now consider Chomsky's core algorithm for Agree, presented in terms of the structural description in (1) and the structural change in (2). For ease of reference, I have given each of the subconditions of the structural description and steps of the structural change a descriptive name on the second line. Key concepts for the algorithm overall are Matching, Valuation, and Deletion. Matching is a matter of the structural description: the uninterpretable feature F of the probe (uF) must match with an interpretable feature (iF) of the goal in order for Agree to obtain, (1c)—more properly, the attributes must match (though the values will not).

(1) *Structural description:* Agree holds between a probe and a goal iff all of the following conditions hold.

  a. The probe bears uF: features that are uninterpretable and unvalued.
  (Probe specification)

  b. The probe c-commands the goal.
  (Structural condition)

  c. The uF of the probe matches with iF of the goal.
  (Match condition)

d. The goal is active: it also has uninterpretable features (uF′).
   (*Activity condition*)
e. The goal is the closest element to the probe meeting the conditions above.
   (*Minimality condition*)

The core structural change brought about by Agree is Valuation, (2a); the probe has its uninterpretable unvalued features filled in by matching interpretable valued features of the goal. A side effect of this process, Chomsky suggests, is that the uninterpretable features that made the goal active (as in (1d)) will also be valued in a way determined by the probe (though they do not necessarily match any features of the probe; see Chomsky 2001, 6). For want of a better term, I refer to this as part of the structural change as "Goal flagging", (2b). In addition, some cases of Agree will trigger movement, depending on some further specification of the probe (e.g. an "EPP feature"), as in (2c). Lastly, the motivation for the whole system lies in deletion of uninterpretable features, (2d). Chomsky (2000) presents this step as the conceptual core, if not the very definition, of Agree: "the erasure of uninterpretable features of probe and goal is the operation we called Agree" (p 122).[2]

(2) *Structural change:* Outcome of Agree

a. F's value is copied to the probe from the goal.
   (Valuation)
b. The uninterpretable features of the goal are given values according to the nature of the probe.
   (Goal flagging)
c. Optionally, movement of the goal (re-merge) is triggered, depending on further conditions X on the probe.
   (Movement)
d. Uninterpretable features are deleted.
   (Deletion)

In an application of this algorithm to Agree between T and a DP subject, the probe on T has [uϕ] and c-commands the subject base-generated in *v*P. The subject bears valued, interpretable [ϕ], in virtue of which it matches the probe, and uninterpretable unvalued structural case ([uCase]), in virtue of which it is active. The subject is the closest element to the probe that matches it and is active. As a result of Agree, the ϕ value from the DP is copied to T. Uninterpretable case features on the DP are "assigned a value under agreement. . . the value assigned depends on the probe" (Chomsky 2001, 6); in this case, given that the probe is on T, the value assigned is nominative.[3] In a language like English where subjects move to Spec,TP, Agree triggers movement of the goal. Finally, all uninterpretable features, both on the probe and on the goal, are deleted.

In presenting the algorithm and its motivation in this way, I aim to highlight four separate roles that uninterpretable features have been tasked with. First, they make both probes and goals "active", i.e. able to participate in Agree. Without uF on the probe, condition (1a) is not met.

[2] Chomsky (2001) reformulates slightly, presenting deletion as a step of Spell-out, applying "shortly after" valuation. It is at this step that the Matched, Valued, uninterpretable features are deleted from the probe.

[3] For further discussion of how Case works in this system, see Pesetsky and Torrego (2007, 277).

Without uF on the goal, condition (1d) is not met. Second, uninterpretable features determine what counts as a goal for a given probe and what counts as an intervener with respect to a potential goal, (1e). This is to say that the relativization of a probe (a notion to be elaborated in the following section) is given in terms of uninterpretable features. Third, it is the need to delete uninterpretable features prior to semantic interpretation, combined with the hypothesis that Agree is the primary (if not exclusive) means to do so, that makes Agree obligatory. Fourth—and the reader might be forgiven for finding this conceptually almost an incidental aspect of the system—uninterpretable features, through their connection to unvaluedness, determine what gets copied/valued, (2a).

The following three sections discuss strands of current research on Agree in relation to each of these roles, beginning with seminal work by Béjar (2003) on probe relativization and the nature of valuation.

3 From interpretability to valuation

One of the key roles just reviewed for uninterpretable features lies in determining patterns of minimality and intervention. Consider, for instance, a standard paradigm for superiority effects: in English *wh*-movement (as in many other languages), it is possible to move a *wh*-subject or a *wh*-object, (3a-c). However, it is not possible to move a *wh*-object when the subject of its clause is also *wh*, (3d).

(3) a. Who do you think Lishan gave a cookie to _?
b. Who do you think _ gave a cookie to Basia?
c. Who do you think _ gave a cookie to who?
d. *Who do you think who gave a cookie to _?

Given Agree as in (1)/(2), these data are analyzed as follows. Movement requires the establishment of an Agree relation. In this case, the probe, borne by matrix C, has uninterpretable (and unvalued) [*wh*] features. Each *wh* word has interpretable (and valued) [*wh*] features (along with other features which are uninterpretable, in order to meet the Activity condition, though we will have little to say about these here; see Chomsky 2000, 128). In the well-formed examples (3a-c), the element that enters into Agree with the probe, and therefore moves to the matrix clause, is the highest *wh*-element in the c-command domain of the probe. Thus the locality condition on Agree has been obeyed (see (1e)).[4] In the ill-formed example (3d), the element that enters into Agree with the probe is *not* the highest *wh*-element in the probe's c-command domain. Locality condition (1e) has not been obeyed, and ungrammaticality results. In this analysis, the role for uninterpretable features is to make the correct division between elements that are potential goals and count as interveners vs. those that do not count as either. The crucial comparison is (3a) vs. (3d). In the latter, embedded subject *who* blocks movement of the embedded object, because its interpretable [*wh*] feature matches the uninterpretable [u*wh*] on the probe and it is the closest element to the probe meeting the conditions (1a-d). In the former, embedded subject *Lishan* does not block

[4] The example is biclausal in order to show clearly that subjects do in fact move (a point that is nonobvious in English matrix questions). I set aside the complications related to successive cyclicity that come with this. The concerned reader is invited to imagine that the relevant probe is located not on matrix but on embedded C, and that probing by embedded C feeds a later dependency with matrix C.

movement, because, while it carries features of various generally syntactically relevant sorts, it does not bear an interpretable [*wh*] feature matching the uninterpretable [u*wh*] feature of the probe.

This pattern of locality and intervention is not confined to *wh*-dependencies. Béjar (2003) observed that it seems to hold equally well for a variety of ϕ-features—and indeed that taking this perspective on ϕ-agreement opens up a new line of analysis for agreement systems that previously had resisted syntactic understanding. Consider, for instance, the pattern of ϕ-agreement in Chirag Dargwa (Nakh-Dagestanian).[5] In this system, person marking on the verb is controlled by the subject, if the subject is a local person (1st or 2nd), (4a,b). Otherwise, it is controlled by the object, (4c).

(4) Chirag Dargwa (Sumbatova 2011, 135)

a. dicce { ʕu / it } r-iqqan-da
1SG-ERG 2SG(ABS) 3SG(ABS) F-lead-1
I lead you/her.

b. ʕicce du r-iqqan-de
2SG-ERG 1SG(ABS) F-lead-2
You lead me.

c. ite du r-iqqan-da
3SG-ERG 1SG(ABS) F-lead-1
He/she leads me.

While agreement can be controlled by the object when the subject is 3rd person, it cannot be controlled by the object when the subject is local person. This is parallel to the superiority paradigm in (3): a *wh*-object can enter into a dependency with C, but not when the subject that lies between it and C also bears a [*wh*] feature. The feature in (4) that plays the role corresponding to [*wh*] in (3) is one characteristic of local person arguments, 1st and 2nd person. Béjar (2003) proposes that this is a feature [PART(ICIPANT)]. On this approach, the probe responsible for agreement suffixes in Chirag Dargwa must bear an uninterpretable [uPART] feature. A now-standard consequence of this line of thinking, explored in detail by Béjar (2003) and following work, is that ϕ-features cannot be taken as an undifferentiated bundle as far as Agree is concerned. Probe specifications may make use of individual, fine-grained agreement features.[6]

The gain in empirical understanding bought by a pivot to "relativized" or "articulated" ϕ-probes comes at the cost of some tension with other aspects of the uninterpretable features model. Most notably, capturing patterns like (4) with Agree draws Valuation into a newly central place in the theory—one which a connection to (un)interpretability does not readily capture. To see this, let us consider in more detail the nature of a participant probe. Starting from the proposals in (1)/(2), we might expect to model this probe as [uPART]. A local person argument has interpretable [+PART] (assuming for the sake of argument that this feature is binary—we return to this), and is therefore (modulo any concern about Activity) a possible goal. Valuation of [uPART] has the result that [+PART] from the goal is copied to the probe (or rather, that the + specification of the [PART] feature is so copied). If this were the final result, we would expect that agreement with any participant DP,

[5] While this particular data set is not discussed in Béjar (2003), it matches her predictions for a "high-ϕ" language, i.e. one where the probe originates above both arguments.

[6] For application of this idea, see esp. Béjar and Rezac (2003, 2009), Anagnostopoulou (2003), Nevins (2007, 2011), Rezac (2011), Preminger (2014).

whether 1st or 2nd person, should lead to an identical morphological output. After all, it is only the feature [+PART] itself (or just its + value), in common to all 1st and 2nd person arguments, that is copied to the probe. This, of course, does not match the facts of Chirag Dargwa. Rather, the morphological result of agreement reveals that features specific to first and/or second person have been copied to the probe as well.

Béjar (2003) responds to this issue by distinguishing between the ways that probe specifications determine *Match* versus *Value*: "the probe is matching on the basis of [PARTICIPANT], but it is Agreeing/valuing with more than just the one feature; it doesn't only agree with [PARTICIPANT], it also agrees with [± SPEAKER]" (p 39). This split occurs, she proposes, because a probe that matches a certain feature on a goal will copy back (and in that sense be valued by) that feature *and any more specific feature* found on a goal. Specificity of features is made precise by reference to a feature geometry (Harley and Ritter 2002), which in Béjar's system is given as in (5):

(5) Person representations (Béjar 2003, 45, 48)

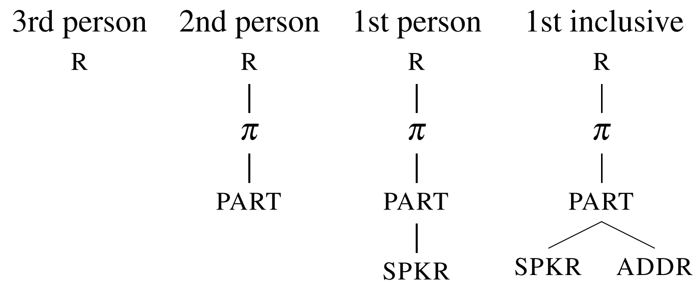

Though the structures in (5) roughly resemble familiar constituent structure trees, it should be emphasized that the graph-theoretic representation stands for something very different in this case. Feature geometries are representations of entailment relations among features: from the fact that R dominates π in (5), we conclude that π may only occur when R occurs. Thus π entails R; entailment is the upward relation in the graph. What the structures in (5) are *not* is syntactic objects. They are not created by Merge. They do not obey familiar principles of headedness and projection.

With this in mind, we should notice the (perhaps surprising) directionality of the entailment condition posited by Béjar (2003) (and subsequently adopted in much work, e.g. Béjar and Rezac 2009, Preminger 2014, Coon and Keine 2021): when [PART] must be copied to the probe, rather than bringing along those features that are entailed by it (higher), Agree brings along those features that entail it (lower). This result is unexpected from the perspective of a minimal Agree theory based on uninterpretability. If deleting [uPART] is the motivation for Agree, copying additional features to the probe beyond those minimally required to trigger deletion of [uPART] should be superfluous. As for feature geometries, because they are not syntactic objects, we cannot (for instance) appeal to constituency in the familiar sense to explain why entailing features should be involved in Valuation.[7] [8]

7 This notion of constituency is perhaps what Preminger (2014) has in mind when he refers to copying of "snippets of ϕ-geometry".

8 It might also be noted that this use of feature geometry differs in a key way from the other primary use to which the notion has been put in the syntax literature, which is explaining patterns of A' intervention. See e.g. Starke 2001, Boeckx and Jeong 2002, Abels 2012, Aravind 2017. For Abels (2012), for instance, feature geometries explain why focus- and *wh*-movement intervene for each other in Italian, whereas focus-movement and topicalization do not. The core idea is that focus- and *wh*-movement are driven by an [OP] feature, and that an [OP] probe cannot overlook a closer [OP] goal in favor of a more distant one (standard relativized minimality). No appeal to copying of entailing

A system that uses relativized probes to capture agreement patterns such as (4) therefore points to an understanding of Valuation that is increasingly autonomous from uninterpretability in terms of conceptual grounding. In Béjar (2003), uninterpretable features do still determine what gets copied/valued, insofar as a probe with uF will copy/value interpretable F and all features that entail F in the relevant feature geometry. This connection is stipulated: it could easily have been the case that a probe with uF would copy/value only F (as Chomsky had proposed) or that it would copy/value only F and features it itself *entails*. The two notions becoming more detached from one another invites the question as to whether both are truly necessary. Notably, it is valuation, rather than interpretability, that plays the key role in accounting for the patterns of ϕ-agreement we see in cases such as (4), as well as numerous other similar cases discussed by Béjar (2003). In this respect, the outcome of Béjar's proposals for the Agree algorithm dovetails with those of authors such as Pesetsky and Torrego (2007), who conclude on rather different grounds that it is not uninterpretability but rather the ability to undergo valuation that is the central characteristic of a probe.

More properly: what is central to the analysis above is not necessarily valuation, in the sense of values and their absence, but *copying*. This brings us to a deferred set of questions concerning exactly what it would mean to value, with an eye to the way that features themselves are represented. We noted above that the notion of copying a value, as in (2a), can be understood most naturally in terms of an attribute-value schema for feature specifications. We also noted that such a schema would lead us to model a participant DP as [+PART]. Such a specification invites the analysis of a 3rd person DP as [-PART]. This, however, will not do for examples such as Chirag Dargwa (4c), where a 3rd person subject fails to intervene in the establishment of an Agree relation with the participant object. This pattern is parallel to (3a), where a non-*wh* DP fails to intervene in the establishment of *wh*-Agree. For the latter case, the natural analysis involves privativity: a non-*wh* subject simply lacks the feature [*wh*] (or [±*wh*]). The participant feature then would need to appear in its positive value on local person DPs, to be absent entirely on 3rd person DPs, and to appear in its negative value. . . nowhere.

This state of affairs is avoided on a privative theory of ϕ-features, such as the one Béjar (2003) defends. There, [PART] is a privative feature borne only by local person DPs. Valuation is not to be understood in terms of supplying a value for a given attribute, but rather in terms of transferring features from a goal to a probe:

> "What does it mean to value given a feature theory in which the category-value distinction does not exist? The answer I propose is that value is structure building, where the features (identified by match) of the goal are 'copied' to the probe." (Béjar 2003, 60)

This conception of Valuation invites a closer look at what makes something a probe, and how probes are different from goals. For Chomsky, both have uninterpretable features, and thus both are LF-deficient. However the probe alone is required to find matching interpretable features in order to fix its deficiency. For Béjar, probes are special in that they can be valued—they have the ability to serve as hosts of structure building; features can be copied to them. This is independent of anything about interpretability. After all, the features, once copied, are promptly removed from LF consideration, as per step (2d). (Similar remarks concerning what makes a probe a probe apply

features is necessary for explanations of this type.

to Pesetsky and Torrego 2007 and other work.) A minimal set of modifications to the structural description and structural change of Chomsky's algorithms could thus be given as in (6), where what it is to be a probe and what it is to value are made more explicit:

(6) Separating interpretability and valuation: revisions to the Chomsky algorithm

a. Probe specification, revised.
The probe has uninterpretable features F (uF) and the ability to serve as the host of structure building.

b. Valuation, revised.
The goal values the probe in terms of F: F and all features that entail it are copied to the probe from the goal.

This type of minimal revision in principle allows uninterpretability to continue to play three of the roles delineated at the end of the previous section—determining what counts as active, determining what counts as a goal/intervener, and making Agree obligatory. But the connection of uninterpretability to the fourth and final function, motivating copying/valuation, has been significantly weakened. Notably, the properties of valuation that are hardest to explain in terms of interpretability—esp. Béjar's entailment condition—are those best supported by the empirical facts (e.g. paradigms as in (4)). Thus the revisions in (6) could be seen as a first step toward an understanding of Agree that parts ways with any appeal to (un)interpretability altogether. The following sections take two additional steps in this direction.

4 Defaults and failure

What happens when there is no goal with iF accessible to a probe with uF? Given the basic setup in (1)/(2), we expect that in this case Agree will not obtain, and therefore uF will remain at the interface. This, by hypothesis, violates some inviolable tenet of interface interpretation, and therefore should cause a derivation to crash. In the terms of Preminger (2014), uF is a "derivational time-bomb", which must be defused (deleted) in connection with Agreement before the derivation concludes. This aspect of the setup for Agree as outlined by Chomsky underpins the explanation in terms of uninterpretability for the obligatoriness of agreement: if agreement does not take place, uF remains in the derivation at LF. This violates Full Interpretation, and ungrammaticality results.

Seeming empirical challenges for this conception have been noted by many authors over many years (e.g. Schütze 1997, Béjar 2003, Bhatt 2005, Preminger 2014). In general, research on agreement paradigms recognizes the existence of a fall-back plan: there is always an elsewhere form for agreement to take when more specific conditions do not apply. In Chirag Dargwa, for instance, when neither argument of a transitive is local person, the verb appears in a form without any overt person affix. This also occurs in 3rd person intransitives. Sentence (7d) completes the previously presented paradigm:

(7) Chirag Dargwa (Sumbatova 2011, 135)

a. dicce { ʕu / it } r-iqqan-da
1SG-ERG 2SG(ABS) 3SG(ABS) F-lead-1
I lead you/her.

b. ʕicce du r-iqqan-de
2SG-ERG 1SG(ABS) F-lead-2
You lead me.

c. ite du r-iqqan-da
3SG-ERG 1SG(ABS) F-lead-1
He/she leads me.

d. ite russe r-iqqle-∅
3SG-ERG girl.ABS F-lead-3
He/she leads the girl.

Similarly, consider the pattern of ϕ-agreement in Hindi-Urdu, (8)-(9). In Hindi-Urdu, like in Chirag Dargwa, agreement is preferentially with the subject, and only targets objects when subjects are not suitable. The suitability condition is different: in Hindi-Urdu, arguments must lack overt case markers in order to Agree (Pandharipande and Kachru 1977). Accordingly, when both arguments bear overt case markers, neither one can Agree. Crucially, in this situation, the sentence remains grammatical, but agreement takes a default form (equivalent to masculine singular). The default nature of agreement is particularly clear in (9c), where agreement is masculine singular despite the absence of any masculine singular DP in the clause.

(8) Hindi-Urdu agreement algorithm (Pandharipande and Kachru 1977, Bhatt 2005, i.a.)

a. If the subject lacks overt case, it controls agreement.
b. Else, if the object lacks overt case, it controls agreement.
c. Else, agreement takes a default form. (M.SG)

(9) Hindi-Urdu agreement (Bhatia 2019)

a. LaRkii ghazalẽ gaa rahii hE
girl.FS songs.FP sing PROG-F be.PRES.3S
The girl is singing songs.

b. LaRkii-ne ghazalẽ gaayiiM thiiM
girl-ERG songs.FP sing-PFV.FP be.PST.FP
The girl sang songs.

c. LaRkii-ne in ghazalõ-ko gaayaa thaa
girl-ERG these songs.FP.OBL-DOM sing-PFV.MS be.PST.MS
The girl sang these songs.

The challenge of "failed agreement" is especially severe for a theory with relativized probes à la Béjar (2003). A [uPART] probe (discussed above for Chirag Dargwa) fails when there is no participant argument. A [uPL] probe (discussed by Béjar for Georgian) fails when there is no plural argument.

One class of responses to this challenge accepts the basic Chomskian Agree algorithm as outlined above, but adds auxiliary mechanisms to prevent unwanted uninterpretable features from reaching the interface. For instance, Schütze (1997) proposes that default agreement results when uF is not present in the numeration—for instance, in Hindi-Urdu, on T. To capture the requirement that agreement take place when it can, Schütze invokes what he calls the *Accord Maximization Principle*, which requires the selection of a numeration that contains as many uninterpretable features as possible. The implementation of this principle involves competition among a set of numerations that differ in the presence or absence of uF. In a case like (9a,b), a numeration including [uϕ] on T will converge, and is thus preferred to a numeration that lacks [uϕ]. For a case like (9c), however, a derivation with [uϕ] cannot converge, and therefore a numeration without [uϕ] is selected. An alternative type of mechanism is outlined by Béjar (2003), who aims not to prevent uF from being present in the numeration, but rather to diffuse unwanted instances of uF by/at the derivation's end. Accordingly, she posits an operation of *Partial Default Agreement*, triggered when Agree has failed to value the features of the probe. Partial Default Agreement removes (certain) uninterpretable features from the probe, avoiding a crash at the interface and producing a distinctive syntactic configuration that default agreement can be connected to. As Béjar puts it: "A probe must be valued and deleted if it can be before the tree is extended. Failure to do so when a controller could have been found in the domain of the probe is what crashes a derivation" (p. 77).

These responses highlight the fact that feature uninterpretability cannot serve all by itself as an explanation for the obligatoriness of Agree. Consider again the case of Chirag Dargwa, where, as above, we might wish to posit a probe on T bearing [uPART]. According to Sumbatova (2011), when there is a participant subject or object, the verb must agree with it. Thus (10b), which uses the same default verb form as in (10a), is expected to be ungrammatical; the only expected form is (10c).[9]

(10) Chirag Dargwa (Sumbatova 2011, 135)

a. ite russe r-iqqle-∅
3SG-ERG girl.ABS F-lead-3
He/she leads the girl.

b. * ite du r-iqqle-∅
3SG-ERG 1SG(ABS) F-lead-3
Intended: He/she leads me.

c. ite du r-iqqan-da
3SG-ERG 1SG(ABS) F-lead-1
He/she leads me.

Assuming this description is correct, can the uninterpretability of [uPART] features explain why default agreement cannot be used in (10b)? Not entirely. On Schütze's (1997) analysis, there are two derivations to consider. In one, T Merges bearing [uPART], but somehow does not Agree. Uninterpretable features are left on the probe, and the derivation crashes, as desired. But there is a second option, too—the T head could be Merged with no [uPART] probe, as in (10a). Appeal to Full Interpretation cannot rule this out; we must appeal to Accord Maximization instead. This is an auxiliary principle whose connection to matters of interpretation at LF is non-obvious. The situation is similar on Béjar's (2003) proposal, where the mechanism of Partial Default Agreement must be stipulated to apply only in certain particular ways, rather than others, so as to appropriately clean up after failed Agree (so to speak) rather than to make agreement less obligatory than it is. To rule out (10b), we must ensure that Partial Default Agreement, which simply deletes uF, cannot apply instead of Agree, which copies features and *then* deletes uF. That is, we must ensure that a more complex operation obtains instead of a simpler one.

9 Sumbatova's description follows Kibrik (2003), published in Russian; she provides no negative data. I provide a * in (10b) based on the generalizations she presents.

A more radical type of response to the challenge of default agreement is articulated by Preminger (2014), who concludes that feature uninterpretability is simply not why Agree is obligatory. If probes bear uninterpretable features, Agree may still under various circumstances fail to assign a value to such features and thus delete them. Such a conclusion entails rejection either of Chomsky's premise of Full Interpretation or his premise that certain features are uninterpretable. The former view is one where the LF interface can simply ignore material not pertinent to it (a claim, incidentally, in keeping with the practice of working semanticists; see also Carstens 2010, Epstein, Kitahara, and Seely 2010). The latter view is one where all features in fact *are* interpretable, though perhaps not in semantically interesting ways (e.g., they are interpreted as identity functions). Whichever choice is made on these questions, the obligatoriness of agreement, Preminger concludes, is not reducible to any interface property of the features on the probe that trigger Agree. On the approach to agreement he sketches, a probe is an element with the capacity to serve as the host of structure-building, following Béjar (2003); in this sense, a probe is "unvalued".[10] The operation that assigns a value is triggered "immediately and obligatorily whenever a head with unvalued features is merged into the derivation" (Preminger 2014, 96). Notably, Preminger rejects the hypothesis that a general operation Agree is implicated in all long-distance dependencies, including both agreement and movement. His proposal, made for (ϕ-)agreement only, is summarized in the directive in (11):

(11) FIND(f): when a head H^0 with an unvalued feature f is merged, look for an XP bearing a value of f, and assign that value to H^0. (Preminger 2014, 96)

This move away from uninterpretability as a source of obligatoriness in agreement has been widely influential, with many subsequent authors adopting Preminger's slogan that "agreement can fail" (see e.g. van Urk 2015, Carstens and Mletshe 2015, Keine 2016, Georgi 2017, Oxford 2017, Kalin 2018, Pfau, Salzmann, and Steinbach 2018, Puškar 2018, Coon and Keine 2021, among many others). Authors adopting this view have typically however *not* also joined Preminger in rejecting the Merge/Agree hypothesis, and thus in seeking to model ϕ-Agreement with an operation distinct from that used in other cases of long-distance dependency. Indeed, quite to the contrary, some authors are very explicit about seeking to use the same type of logic that Preminger discusses for the case of ϕ-agreement for other types of dependencies, e.g. DP movement (Collins 2017) or clitic doubling (Kramer 2014). Such approaches are compatible with (and plausibly understood by their authors as involving) retention of various aspects of the Chomskian Agree algorithm, again with a round of minimal revisions, as outlined in (12). Removed here is any mention of uninterpretable features on the probe; the probe is now strictly unvalued, in Béjar's (2003) sense. Similarly removed is the requirement that uninterpretable features on the probe match with interpretable features on the goal. The match condition is instead simplified, requiring merely that the goal bear the features with respect to which the probe is able to host structure-building. Note that while the revisions in (12) partly build on those discussed in the last section, they are partly independent; for instance, they are certainly compatible with a system where feature geometries play no role in valuation.

[10] As Preminger puts it: "the probe would enter into the derivation with a container for a piece of feature geometry; valuation would consist of copying an appropriate snippet of feature geometry from the goal onto the probe...In such a system, relativized probing amounts to specifying, on this container, what the root of the snippet copied into it must be" (Preminger 2014, 47-48).

(12) "Agreement can fail": revisions to the Chomsky algorithm

a. Probe specification, revised.
The probe has the ability to serve as the host of structure building at least with respect to feature [F].[11]

b. Match condition, revised.
The goal bears [F].

These revisions go quite a bit further than those discussed in the previous section in terms of unmooring Agree from a foundation in lexical uninterpretability and its syntactic elimination. A run through the four roles for uninterpretable features delineated above shows that the elimination of uninterpretable features from any explanatory function is almost complete. First, such features no longer play a role in determining probe specification, or in other words, making probes "active". Whether an element serves as a probe depends on its ability to host a copied feature of whatever type; this notion related to valuation has now deplaced uninterpretability entirely. (We return just below to activity conditions on goals.) Second, relatedly, uninterpretable features no longer determine matching and intervention, which of course are issues related directly to probe specification. Third, the need to remove uninterpretable features no longer underpins the obligatoriness of Agree, as Preminger directly argues. Fourth, giving probe specification explicitly in terms related to valuation removes any need to appeal to uninterpretable features directly or indirectly in determining what gets copied/valued.

5 Activity and the properties of goals

What potentially remains for uninterpretability in the system just outlined is the Activity Condition—the requirement that goals contain an additional uninterpretable feature, uF′, in addition to feature F that matches with the probe—and the related idea that uninterpretable uF′ on the goal receives a value via Agree, thereby making Agree an operation that changes the goal's features. Putting empirical meat on these bones depends on constraining possible choices of uF′. For Chomsky, appeal to an activity condition on goals is motivated by questions about why there should be such a thing as abstract Case (see esp. Chomsky 2000, 127). For Agree holding between T and a c-commanded DP, Agree is driven by uninterpretable ϕ on T (the probe). A-movement to subject position is based on this Agree relation. Hosting an uninterpretable Case feature is, for the goal, essentially the cost of entry to Agree. This cost is paid once, irrevocably, when the goal participates in Agree.[12] This predicts that a DP that has participated in the Agree relation characteristic of A-movement to finite subject position cannot undertake any further ϕ-Agree or movement connected to such agreement (i.e. A-movement): "after structural Case of DP is deleted, the phrase cannot move further to an A position and its ϕ-set cannot induce deletion" (Chomsky 2000, 127).

Serious empirical challenges have been raised for both parts of this prediction. Consider first the question of A-movement from a case-position. This is attempted in the ungrammatical English hyperraising sentence (13), the ill-formedness of which is cited by Chomsky in support of the Activity Condition (2000, 128-129). The Activity Condition rules (13) out because, having

[11] I say 'at least' here to make room for the entailment condition from Béjar (2003) discussed in the last section.

[12] Though note that Chomsky suggests a loophole (which would require modification to (1)/(2)): the goal can maintain its uCase feature upon Agree if the probe is "defective". I return to this idea below.

had its Case feature deleted after ϕ-Agree in the most embedded clause, the subject *John* cannot participate in the further step of ϕ-Agree needed to obtain the subject position of the *seem* clause.

(13) * [John to seem [t is intelligent]] would be surprising

If the Activity Condition is part of Agree itself, and therefore a language universal, we expect the absence of hyperraising also to be universal. Yet there are numerous languages where hyperraising sentences are well-formed (Ura 1996, Rodrigues 2004, Ferreira 2004, Nevins 2004, Zeller 2006, Carstens 2011, Carstens and Diercks 2013, Asarina 2011, Halpert 2016, 2019, Deal 2017, Zyman 2017, Fong 2019, Wurmbrand 2021, i.a.), e.g. Brazilian Portuguese (see (14)) and Zulu (see (15)).

(14) Ninguém$_i$ parece [$_{CP}$ que t$_i$ está doente]
nobody seems.sg that t$_i$ is.sg sick.sg
Nobody seems to be sick. (Brazilian Portuguese; Nevins 2004)

(15) uZinhle$_i$ u- bonakala [$_{CP}$ ukuthi t$_i$ u- zo- xova ujeqe]
Zinhle$_i$ 1S- seem [that t$_i$ 1S- FUT- make steamed.bread]
It seems that Zinhle will make steamed bread. (Zulu; Halpert 2016)

Such sentences are similar to (13) in that a raising predicate embeds a finite clause, where the moving subject participates in visible ϕ-Agree, and presumably therefore also deletes its structural Case, by the reasoning above. If Case is the only uninterpretable feature inherent to nominals (i.e., Case is the only possible choice of uF$'$), the well-formedness of these sentences is unexpected.[13]

Turning to the second part of Chomsky's predictions, there are also numerous languages where a single DP can participate in ϕ-Agreement with more than one head. In the Hindi-Urdu sentences in (16), for instance, both the auxiliary and the main verb Agree in gender and number with the same goal. Recall that this goal is the highest DP without overt case marking, as outlined in the algorithm in (8).

(16) Hindi-Urdu (Bhatt 2005)

a. Rahul kitaab paṛh-taa thaa
Rahul.M book.F read-Hab.MSg be.Pst.MSg
Rahul used to read (a/the) book.

b. Rahul-ne kitaab paṛh-ii thii
Rahul-Erg book.F read-Pfv.F be.Pst.FSg
Rahul had read the book.

With restructuring predicates such as 'want', Hindi-Urdu also allows long-distance agreement by an embedded object. In this case, the same goal Agrees with even more targets. In (17a), the highest DP without overt case marking is the embedded object *ṭehnii* 'branch'. This DP agrees feminine singular with the infinitival verb, the matrix main verb, and the matrix auxiliary. (Contrast (17b), where the matrix subject is unmarked and there is no feminine singular agreement on any of these loci.)

(17) Hindi-Urdu (Bhatt 2005)

[13] Likewise unexpected, for the same reasons, is what Merchant (2006, 2011) calls "polyvalent case", where a single DP receives more than one case assignment. See also Béjar and Massam (1999), Richards (2013), Pesetsky (2013), Assmann, Edygarova, Georgi, Klein, and Weisser (2014) for further discussion of this phenomenon.

a. Shahrukh-ne [ṭehnii kaaṭ-nii] chaah-ii thii
Shahrukh-Erg branch.F cut-Inf.F want-Pfv.F be.Pst.FSg
Shahrukh had wanted to cut the branch.

b. Shahrukh [ṭehnii kaaṭ-naa] chaah-taa thaa
Shahrukh.M branch.F cut-Inf.M want-Pfv.MSg be.Pst.MSg
Shahrukh wants to cut the branch.

Similar patterns of multiple agreement are found in many languages, e.g. Tsez (Polinsky and Potsdam 2001), Innu-aimun (Branigan and Mackenzie 2002), Daga (Baker 2008, 183), Nez Perce (Deal 2015b), and Ojibwe (Oxford 2017), as well as numerous Bantu languages including Swahili, Kilega, and Luusamia (Carstens 2001, 2010, 2011, Carstens and Diercks 2013). The Bantu data, like the Hindi-Urdu data in (17a), are notable in confirming that there is no limit to exactly two instances of ϕ-Agreement for a single goal. In (18), the subject Agrees in person and number with two auxiliary-like elements in addition to the main verb.

(18) (Mimi) ni-li-kuwa ni-ngali ni-ki-fanya kazi.
(1SG.PRON) 1SG-PAST-be 1SG-still 1SG-PERF-do 9.work
I was still working. (Swahili; Carstens 2001, 150)

These data are again unexpected given the Activity Condition together with the thesis that the relevant value of uF′ on the goal is structural Case.

Responses to the challenges for the Activity Condition from hyperraising and multiple agreement come in two general types. A first set of works looks for language-specific solutions to the apparently problematic (non-English) data that leave Activity in place as a language universal. For Ferreira (2004), for instance, Brazilian Portuguese examples such as (14) do not actually involve violations of the Activity Condition. Ferreira builds on a suggestion of Chomsky's (2000, 2001) that Agree only deactivates the goal if the probe is not "defective". For Chomsky, only a probe that is ϕ-complete deactivates its goal; a probe that just bears [uPerson] does not deactivate. Similarly, in Brazilian Portuguese, Ferreira proposes, it is possible for embedded finite T to be defective and thus fail to delete the structural Case feature on the subject. When this happens, the subject remains active, and can undergo raising out of its clause. Ferreira connects the possibility for finite T to be defective in Brazilian Portuguese to its impoverished verbal paradigm.[14] A similar language-specific proposal in terms of defectivity of embedded T is given by Zeller (2006) for Zulu and other Nguni languages (a subgroup of Bantu). Zeller suggests that T in Zulu is defective not because it is ϕ-incomplete, but because it is "not 'verbal' enough"; it has a nominal feature in addition to its tense and [uϕ] specifications. Thus the probe is defective not because of features it lacks, but because of additional features that it has. At a slightly broader level, Carstens (2010, 2011) and Carstens and Diercks (2013) propose for Bantu languages generally that the possibility of hyperraising and multiple agreement reflect the particular ways that gender and Case features are put to use. In particular, Bantu languages lack Case features, but have uninterpretable gender features that participate in ϕ-Agree. These gender features are uninterpretable, but valued. Their valuedness protects them against deletion, and their uninterpretability makes their hosts "hyperactive", that is, able to participate in an unlimited number of ϕ-Agreements.

[14] This requires that English finite T be different, despite the fact that English verbal agreement is similarly impoverished. Similar remarks hold for French, as Rodrigues (2004, 117) notes.

What all these analyses do is preserve an Activity-based explanation for the ungrammaticality of English (13). Variation is reducible to some special other properties of other languages, or certain structures therein, that makes hyperraising and/or multiple agreement possible. Theoretical implementation of the proposals just discussed requires revision to condition (2d) of the Chomsky (2000, 2001) algorithms, ensuring that uninterpretable features only delete if some additional condition is met. (This of course holds for Chomsky's "defectivity" view noted above, as well as the views of Ferreira, Zeller, Carstens and Diercks.) The impact of such a revision is that uninterpretable features retain a role in the Agree algorithm, insofar as they determine the activity of goals. But on a view of Agree as motivated by deletion of uninterpretable features, such a revision of course raises the question of why Agree that does *not* delete uninterpretable features would occur in the first place.

A second type of analysis takes the opposite approach: Activity is removed from the theory entirely, and languages in principle should allow hyperraising and/or multiple agreement, *ceteris paribus*. It is English and similar languages wherein language-specific factors arise that rule these structures out. The central observation here is that appeal to Activity is not actually necessary to rule out hyperraising structures in English. Nevins (2004) observes that movement in (13) plausibly violates the Phase Impenetrability Condition: the embedded finite clause under *seem* is a CP, and cyclic A-movement is not possible through its edge. He proposes that hyperraising is possible in Brazilian Portuguese because subjects move slightly higher in that language than they do in English, which allows them to escape the lower phase. Another possibility is that the Agree relation underpinning movement in (13) violates minimality: English embedded finite CP itself bears ϕ features and is a closer target to matrix T than the embedded subject is (Halpert 2016, 2019). For Halpert, hyperraising is possible in Zulu because CPs lack ϕ-features in that language, which makes them ineligible to intervene between the higher T and the embedded subject. On either of these approaches, the Activity Condition can be removed from the algorithm for Agree with no loss of coverage regarding structures like (13). Thus these approaches are compatible with an algorithm for Agree that removes any reference to uninterpretable features entirely, both as concerns probes and now also as concerns goals.[15]

Theories that remove the requirement that goals bear uninterpretable features cast new light on the way that Agree can change the goal, recorded above as (2b) ("Goal flagging"). In Chomsky's system, the way that Case features are determined by Agree is not made precise. Pesetsky and Torrego (2001) suggest a way forward. They propose that structural Case features are direct signatures of agreement, such that nominative case is the morphological realization of a T feature on DP (and accusative case perhaps the morphological realization of a v feature on DP). Deal (2010a), Clem (2019b) extend similar ideas to ergative morphological case systems. For these proposals, the central idea is of course not that goals bear uninterpretable features, but rather that Agree can change the goal. A more direct recording of this intuition suggests a final round of revisions to the Chomsky algorithm, (19):

[15] I set aside here a third class of analyses, which treats variation with respect to Activity as an irreducible fact, and therefore accords it the status of a parameter (Baker 2008, 2013, Oxford 2017). This view is one where the universal core of Agree does not require reference to Activity, in line with the views just outlined from Nevins and Halpert. However, Activity does play a role in some languages, including (potentially) English. Clearly, an analysis positing a mechanism that cannot be reduced to anything else in the grammar of a language is to be dispreferred to an analysis that reduces either the possibility or the impossibility of hyperraising to other independently motivated facts in a particular language's syntax.

(19) Without activity: revisions to the Chomsky algorithm
a. Activity condition: removed
b. Goal flagging, revised.
The goal is assigned a feature indicating agreement with the probe.

The revisions are again in principle independent of those pursued in the previous sections. However, it is when taken together with them that they make a most striking picture. Uninterpretability has now been removed from the theory of Agreement in its entirety.

6 Questions of directionality

Before leaving uninterpretable features behind entirely, this section considers a consequence of the previous three as concerns ongoing debates surrounding the directionality of Agree. The directionality of movement is well-established: the landing site of movement must c-command the origin site. Accordingly, (upward) topicalization, as in (20a), is well-formed, whereas downward topicalization, as in (20b), is ungrammatical; similar pairs can be provided for movement of numerous types.

(20) a. Anusha will tell Becca that Caleb, she met _.
b. * Anusha will tell _ that Becca, she met Caleb.

Theoretical developments over the past several decades constrain the possible ways in which such directional asymmetries can be explained. Given a copy theory of movement, a ban on lowering as in (20b) cannot be explained purely semantically if there is a possibility of reconstruction via "copy neglect"—i.e., the semantic component can simply ignore the movement (see recent discussion in Keine and Poole 2018). Similarly, the ban cannot be explained phonologically if there is a possibility of pronouncing the lower copy in a movement chain (Bobaljik 2002 and many since). The semantic side of this reasoning suggests that LF interpretation encounters no particular challenge if faced with (20b); it can simply treat *Becca* as occupying the higher position, ignoring the lower copy. The phonological side suggests that PF interpretation encounters no particular challenge, either. Pronouncing the lower copy of a chain violates no core principle of PF. All this points to an analysis where the problem for (20b) and similar cases of attempted downward movement is a core syntactic one.

The nature of this syntactic explanation is further constrained by the Merge/Agree hypothesis. If movement consists of Agree + Merge, as the hypothesis requires, then a syntactic explanation must be in terms of these operations. That is, at least one of the steps in (21) is illicit (where H is the head driving topicalization):

(21) Two steps in a derivation for syntactic lowering, (20b)
a. (Upward) Agree:
[Anusha will tell Becca [that H she met Caleb]]
b. (Downward) Merge:
[Anusha will tell Becca [that **Becca** H she met Caleb]]

Perhaps the most standard view follows Chomsky (2000, 2001) in maintaining a structural condition on Agree whereby probes must c-command goals, as in (1b) (see e.g. Preminger 2013, Preminger and Polinsky 2015, Rudnev 2020, 2021, Diercks, van Koppen, and Putnam 2020, Keine and Dash to appear, Clem to appear). This condition rules out Upward Agree (21a) and therefore, because movement is built on Agree, rules out lowering. Why should Agree be subject to a condition of this type? The constraint plausibly follows from the cyclicity of structure building (Rezac 2003, Béjar and Rezac 2009, Preminger 2014, Carstens 2016; see also Georgi 2014): Agree takes place as soon as the head bearing the probe is Merged. At this point, only the head and its complement have been constructed. Therefore, only material within the complement—i.e., the c-command domain—is available to Agree.

The alternative view is one where Upward Agree as in (21a) is in principle well-formed, available to the grammar—a position on the nature of Agree explored by Zeijlstra (2004, 2008, 2012), Merchant (2006, 2011), Wurmbrand (2012a,b), Bjorkman and Zeijlstra (2014, 2019), Carstens (2016), Baker and Camargo Souza (2020), Arregi and Hanink (to appear), among others. If Upward Agree is possible, and the Merge/Agree hypothesis is correct, then the ban on syntactic lowering must come from a constraint on Merge. Agree is allowed to be countercyclic (in the sense of occurring in a structure long after the probe has been merged) but Merge isn't.

Willingness to accept this conclusion of course depends on the strength of argumentation that can be mustered in favor of Upward Agree in general. And this brings us back to the questions about uninterpretability that occupied the previous three sections. Because of their basis in ideas about interpretability, Chomsky's algorithms suggest that probes and goals can be identified on semantic grounds. While both must bear uninterpretable features in order to be active, it is the probe's features that determine its specification, which in turn determines what counts as a suitable goal and as an intervener. If a long-distance syntactic dependency holds between two elements, where the possibility of participation and intervention in that dependency requires feature F, then the probe in the dependency must bear uninterpretable uF. If we can make a reasonable guess as to the identity of F, and furthermore on the basis of that guess can pinpoint where in the structure F is interpreted, that information allows us to identify probe versus goal. This identification can then be used to determine whether Agree is always indeed downward, as in (1b), or whether Upward Agree is possible (or obligatory) instead.

Zeijlstra (2004, 2008, 2012) follows this reasoning in mounting a series of arguments in favor of Upward Agree. Instead of requiring that the probe c-command the goal, his revised Agree algorithm requires instead that the goal c-command the probe. Central to his arguments is the phenomenon of negative concord. In negative concord, multiple apparently negative elements in a clause give rise to a single negation reading. In Italian, for instance, we can identify items such as *nessuno* 'no one' as negative based on their behavior as negative fragment answers (Zanuttini 1991; for crosslinguistic applications see Giannakidou and Zeijlstra 2017):

(22) Italian (Zanuttini 1991, 109)

Q: Chi ha telefonato?	A: Nessuno.
Who has called?	Nobody.

A clause with two such elements, however, yields a single negation reading, (23a). In Italian, this is also the case when a postverbal negative element cooccurs with sentential negation, (23b). These examples illustrate negative concord.

(23) a. **Nessuno** ha detto **niente**.
nobody has said nothing

b. **Non** ha telefonato **nessuno**.
NEG has called nobody

Nobody said anything. Nobody called. (Zanuttini 1991, 108, 111)

Following earlier syntactic work on negative concord (e.g. Laka 1990, Brown 1999), Zeijlstra proposes that the appearance of multiple negations need not be taken at face value; rather, he suggests, the apparent mismatch between negative morphology and negative semantics reflects the contribution of Agree. Sentential negation semantically contributes negativity (and examples such as (23a) involve a covert sentential negation). Negative concord items (NCIs), such as *nessuno* 'nobody' or *niente* 'nothing', are not themselves semantically negative. Rather, they are indefinites in the scope of the semantic negation. NCIs are a special type of indefinites that Agree with the negation that scopes over them.

Crucial to the argument from negative concord for Upward Agree is the identification of probe versus goal. Zeijlstra's case rests on interpretability: negation is interpretable on the sentence negation, but not on NCIs (which, semantically, are just indefinites). So, if Agree is in terms of a negative feature, then NCIs must bear a [uNeg] feature, and be probes. The semantic negation must bear an [iNeg] feature, and be the goal.[16] An NCI occurs in the scope of negation, thus in its c-command domain, and probes upward to find its goal. Thus NC instantiates Upward Agree—indeed, upward *Multiple* Agree, according to Zeijlstra, as in (24).

(24) 'I haven't said anything to anyone.' (Zanuttini 1991, 147)

Non	ha	detto	niente	a	nessuno.
NEG	have.1sg	said	nothing	to	nobody.
$[\text{iNeg}]_{goal}$			$[\text{uNeg}]_{probe}$		$[\text{uNeg}]_{probe}$

In this way, the analysis is like Béjar's (2003) in revealing a tension between two parts of Chomsky's algorithms as in (1)/(2). In tension here are the role of uninterpretability as providing probe specifications and the structural claim that probes always c-command their goals.

Returning now to the work discussed in the previous three sections, the impact should be clear: if uninterpretable features are not central or even necessary at all for the specification of a probe, the tension dissolves. Questions of where negation is semantically interpreted can no longer diagnose probe versus goal in an Agree relation. There is nothing to prevent the sentential negation, despite its obvious negative semantics, from bearing a probe that Agrees in a negation-related feature with NCIs in its c-command domain. NCIs, for their part, can be treated as special morphological forms that indefinites take when they have Agreed with negation. Such an analysis is purposefully parallel to the treatment, for instance, of nominative case as a set of special morphological forms that DPs take when they have Agreed with finite T. In other words, the morphology of NCIs that Zeijlstra took to reflect valuation of an NCI-hosted probe is treated instead as reflective of goal flagging. A view of aspects of the typology of negative concord that expands along these lines is sketched in Deal (2021b).

Other empirical phenomena central to debates about the directionality of Agree similarly appear in a new light when uninterpretability is taken out of center stage as a means of probe specification. In sequence of tense, for instance, a past tense complement to a past tense attitude verb may receive a simultaneous interpretation (setting aside various additional complications; see e.g. Ogihara 1996), as in (25). Here the embedded clause appears to be interpreted not with a past tense

[16] This reasoning follows earlier work by Brown (1999). Brown uses not Agree but rather its theoretical precursor, feature movement, to account for licensing of Russian NCIs that are not in an overt spec-head relation with negation.

but with a (relative) present tense. This type of reading for embedded past tense is only available when the matrix tense is past, not present, as shown in (26).

(25) Anjali thought it was snowing.
Simultaneous interpretation: Anjali thought, "It is snowing!"

(26) Anjali thinks it was snowing.
Cannot mean: Anjali thinks, "It is snowing!"

Building on Stowell's (1995, 2007) analysis, Zeijlstra (2012) treats structures like (25) as involving only one semantic past tense, located in the matrix clause. The embedded morphological past tense is licensed by upward Agree between a [uPast] feature on the embedded tense and an [iPast] feature on the matrix tense. Like in the case of negative concord, the upwardness of the structural relation is identified on the basis of interpretation: the matrix tense has its ordinary meaning, whereas the embedded tense is semantically empty. Also like in the case of negative concord, therefore, an alternative analysis is readily available once uninterpretability is de-centered. If we wish to maintain that the two T heads Agree with one another, there is no obstacle to treating the higher (semantically meaningful) head as the host of the relevant probe, and the lower (semantically bleached) head as the goal. Past tense morphology on the lower T head results from goal flagging, not valuation of a probe.

Similar remarks apply to the Norwegian and Frisian "parasitic participle" constructions, e.g. (27), discussed by Wurmbrand (2012a,b), where the apparently meaningless material is the participial morphology on the main verb. The perfect auxiliary (here *hadde*) licenses participial forms both for the head of its immediate complement (the modal *villet* 'want') and for the main verb (here *lest* 'read'). With an eye to the long-distance relationship between the main verb and the auxiliary, Wurmbrand posits upward Agree: the auxiliary bears a interpretable (and valued) T feature, which serves as a goal for the lower uninterpretable (and unvalued) T feature on the main verb. The value of the T feature that is transmitted by Agree is *perf*, which produces participial morphology.

(27) Jeg hadde villet lest boka
I had want.PART read.PART book.DEF
'I would have liked to read the book.' (Norwegian; Wurmbrand 2012b, 132)

A simple reanalysis, preserving the main claim that Agree is established between the auxiliary and the main verb, would maintain that the probe is actually borne by the auxiliary, and Agree is again downward. Like in the case of negative concord and nominative case, Agree in this construction has a morphological consequence for the goal, namely perfect morphology. Just as discussed in the previous two types of examples, this type of morphological consequence is in principle neutral as to probe versus goal identification. Instead of indicating valuation on a probe, as Wurmbrand proposes, it can instead be treated as flagging, on a goal.

Notably, the reanalyses just suggested make some headway in understanding two puzzles relevant to current debates about the directionality of Agree. The first is why the phenomena of negative concord, sequence of tense, and participle licensing seemed to contrast with standard ϕ-Agree in their directionality (a contrast emphasized by Preminger 2013, Preminger and Polinsky 2015). The apparent difference is an artifact of particular claims about what features are interpreted where—a matter that turns out to be quite orthogonal to anything about the syntax of Agree, if the previous three sections' conclusions are on the right track. It happens that in ϕ-Agree, ϕ features are generally taken to be interpreted on the DP and not on T, whereas in sequence of tense (for instance), past tense is interpreted on the higher T and not the lower T. A theory where uninterpretability is not a key property of probes faces no obstacles in treating all such cases, despite their internal diversity, with downward Agree. The conclusion is an especially welcome one insofar as this allows for the directionality of Agree to receive a natural grounding in the cyclicity of structure building (Rezac 2003, a.o.).

The second puzzle is why movement is always upward, whereas a morphological change effected by Agree can occur either above or below the element that triggers it. Much literature on the directionality of Agree takes a morphological change triggered by Agree to result from Valuation. This means that an element that changes morphologically as a result of Agree must host a probe (see, e.g. Wurmbrand 2012a,b, Carstens 2016), and suggests that probing can in principle occur in either direction (as in Merchant 2006, Baker 2008, Carstens 2016, Bárány and van der Wal to appear). Why then is movement not similarly free? Crucial to an understanding of this seeming contrast is the nature of Goal flagging. I suggest that the puzzle reflects the ability of Agree to change the goal as well as the probe—both Valuation and Goal flagging are part of Agree, and Goal flagging is not to be reduced to Valuation. There is often (though not always) no *a priori* way to tell whether morphology that results from Agree occurs on the probe versus on the goal. The fact that only the goal is subject to movement, however, rather than the probe, presumably again reflects the cyclicity of structure building. Merger of the probe triggers operations on the probe itself and on the immediate structure wherein it is Merged, including, potentially, internal Merge of the goal in the probe's projection. (Upward) movement of the goal in this way preserves the structure already built up at the point the probe is Merged. The same could not be said for lowering of the probe to the goal, if, hypothetically, the probe were to lower to the goal as the result of Agree.

7 A way forward

Let us now take stock of the picture that has come into view. We have seen that the operation Agree has two types of effects. One is that features are transferred from one locus to another. The other is that movement is triggered. Throughout, Agree is subject to relativized minimality, as we saw in section 3 for superiority effects and Agree in person features. For the cases that do not involve movement, it is clearly feature transfer that is the heart of Agree, most notably (though not exclusively) from the goal to the probe. This transfer has the effect of providing information redundantly in multiple points of a syntactic structure. Movement, seen through the lens of the copy theory, has indeed this same effect.

A fresh look at the foundations of long-distance dependencies, ideas about uninterpretability fully set aside, requires an etiological question about this situation. Why would Agree have this property? Why is redundancy important to language? I have suggested elsewhere that the redundancy created by Agree be recognized as adaptive in view of communicative purposes (Deal to appear). That is, rather than seeing copying (redundancy creation) as a mysterious side effect of LF improvements, I suggest instead a perspective that puts this function at the heart of what Agree is for. A syntactic structure that encodes information redundantly (as the result of Agree) can feed a morphology that *realizes* information redundantly (as the result of Vocabulary Insertion over redundant syntactic structure). Borrowing from functionalist work on redundancy in communication, this redundant morphological realization contributes to "the overcoming of "noise" in the system: multiple cues are better than a single cue when the latter is attenuated, missing, or masked by other material" (Ferguson and Barlow 1988, 17).

Returning now to the main question of this work: what is needed in an explicit statement of the algorithm for Agree, given a foundation in redundancy and in view of the empirical applications we have just reviewed? In the remainder of this paper I sketch what I see as at least the core of an answer. That answer proceeds from a perspective of balancing redundancy against economy. That is, redundancy is useful, but only in moderation. The behavior of agreement in natural language suggests that two types of limits on this redundancy are in place. I will suggest that each limit is connected to a way that probes are relativized.

A first limit is that probes, as elements specified for the capacity to copy features, only attend to certain types of features and not others. In English, the agreement probe on T attends to person and number features of the subject DP, but not to features of the AspP found even more locally. In Chirag Dargwa, the probe attends to features of first and second person DPs, but not to features of third person DPs, or of AspP or other projections in the clausal spine (e.g. vP). Capturing this behavior requires specifying the features that the probe can copy to itself. I will describe these as the features that the probe **interacts** with. These features generally correspond to those described as "unvalued" on the probe in the theories reviewed in prior sections, esp. in valuation-centric approaches such as Béjar (2003), Pesetsky and Torrego (2007), Preminger (2014). I avoid this type of terminology here lest it retain the implication, carried over from the uninterpretable features model, that probes are in some way deficient and Agree is required to cure this deficiency. I suggest instead that the ability to interact with feature F be seen as an ability by the probe to create redundancy with respect to feature F, should structural conditions be conducive. If they are not—generally, because there is no bearer of F present in the appropriate domain—there is simply no Agree, and no penalty incurred.

A second limit on redundancy concerns the conditions under which probing stops. One point of variation among agreement systems concerns those where only a single argument is indexed by a probe versus those where two (or more) arguments are indexed; this division can be found both in systems of verbal agreement and in systems of complementizer agreement.[17] Capturing this behavior requires specifying the features which, when encountered on a target of probing, cause probing to stop. I will describe these as the features that **satisfy** the probe. The notion of satisfaction ties together various threads of the agreement literature treated disparately on the uninterpretable features model. One is the idea that a probe can be specified to participate in Multiple Agree (Hiraiwa 2001, Penka 2011, a.o.; see related ideas in Bošković 1999, Collins 2002) (a proposal which of course would require another round of modification to the algorithms in (1)/(2) if we wished to maintain them). A probe of this type, interacting with F, will create redundancy with respect to every instance of F it encounters. I will describe this type of probe as *insatiable*; it lacks a satisfaction condition. To be contrasted with such a probe is an alternative probe that likewise interacts with F but also is satisfied by F. This second type of probe encounters F, copies it, and thereupon concludes its process of copying.

In presenting an algorithm for Agree that integrates notions of interaction and satisfaction, I

[17] On verbal agreement, see Georgi (2013) for multi-argument indexing (single argument indexing being the more familiar case). On complementizer agreement, see e.g. Haegeman and van Koppen (2012), Diercks et al. (2020) for discussion of the single-argument case, and Deal (2015a) for discussion of the multi-argument case.

will aim for the style that Zobel (2004) calls 'literate code', wherein the algorithm is explained and presented simultaneously. (A much more condensed version written in pseudocode is provided in an appendix.) At a broad level, the Agree algorithm can be described as in (28).

(28) **Agree** searches the domain of an element H and determines whether copying and/or movement will occur. H is specified for interaction features and satisfaction features, denoted respectively as int(H) and sat(H). Each syntactic object in the search space is a target. For a given target, individual features are inspected to determine whether they should be copied to H (interaction) and whether copying of further features to H should halt (satisfaction). A target that interacts with or satisfies H receives a flag feature. Features in int(H) and sat(H) may be specified to trigger movement; a feature F that triggers movement is denoted F^M. If a feature is encountered on a target that corresponds to an ^{M} feature in int(H) or sat(H), the target is re-merged to the projection of H.

The major steps of the algorithm are as follows.

1. Initialize a search through targets.

2. For a given target, initialize a search through that target's features.

3. For a given feature, determine whether it should be copied to the probe and/or cause probing of further targets to halt. If either answer is positive, assign a flag to the target and determine whether movement should also take place. If movement is indicated, re-merge the target to the projection of the element hosting the probe.

Let us now consider these steps in detail.

Agree begins with a search of the domain of the probe. The nature of this search (e.g. in terms of breadth-first versus depth-first search) has recently been discussed in depth by Branan and Erlewine (2020). In using the word 'target', I emphasize the assumption that every syntactic object is considered by the algorithm. Thus a probe on T whose sister is AspP considers AspP, its specifier (if any), the intermediate projection of Asp (if any), the head Asp^0 itself, its sister vP, and so on. None of the syntactic objects just mentioned are likely to serve as a *goal* for ϕ-Agree in the traditional sense. The fact that they do not serve as goals is not a primitive; it comes from an algorithm that inspects them and determines that they do not have appropriate features to interact with or satisfy the probe.

The next two steps of the algorithm involve consideration of the individual features of a target. Continuing our example of a probe on T, consider an Asp^0 head that bears the feature [PERF] and the categorical feature [Asp]. Each of these features is compared to the specification of the probe on T in terms of interaction and satisfaction. This involves asking, for each feature, at least two questions. The interaction question is: should this feature be copied? The satisfaction question is: after encountering a target with this feature, should consideration of any further targets be halted? In the case of an English-style probe on T, the features that interact are, let us say, all ϕ-features; any ϕ-feature satisfies. In considering the features [PERF] and [Asp], the interaction question is therefore answered negatively twice (neither feature should be copied) and the satisfaction question similarly (neither feature causes further probing to halt). These negative answers are what determines that Asp^0 is not a goal of Agree, but merely a target. The search continues.

Suppose now that a DP—say, a first person singular pronoun—is encountered with features [ϕ,PART,SPKR,D]. For these features, the interaction question is answered positively in the case of [ϕ,PART,SPKR], assuming as above that T is specified to interact with all ϕ-features. (That is, the entire ϕ-set is int(T).) Thus each of these features is copied to T. The satisfaction question is also answered positively in the case of these three features. The positive answer to that question means that once the features of this DP are inspected, there will be no further search of the domain of the probe for additional targets. If, for instance, there is a lower DP that bears the feature [PL], the probe on T will not be able to copy it. The two remaining pieces of the derivation are goal flagging and movement. Because the DP has interacted with and satisfies the probe on T, it is flagged as a T goal. In English, this flagging plays a role morphologically in determining that the DP will appear in its nominative morphological form. Lastly, in the case of a language like English, the first DP that T Agrees with—one that is bound both to interact with it and satisfy it, given its probe specification—must also undergo movement. This means that either int(T) or sat(T) bears ^{M} in English. Accordingly, when a target is determined to bear ϕ-features, that target is re-merged to the projection of T.

With these steps in mind, a final piece helpful for application of the theory is some notation for the interaction and satisfaction specification of probes introduced in Deal (to appear). For a probe that interacts with all features grouped under particular node N in a feature geometry, I write the interaction condition as N; thus for a probe that interacts with all ϕ-features, I write [INT:ϕ]. To be clear, this is an abbreviation. There is a non-singleton set of features that will be copied to an [INT:ϕ] probe; such features otherwise behave as a natural class, which motivates their feature-geometric grouping. To get from the abbreviated, feature-geometric form to a non-abbreviated, fully explicit form, it suffices to replace the label of the node N with the exhaustive listing of the nodes dominated by N in the geometry. Thus if the ϕ-set consists of {ϕ,PART,SPKR,PL}, for instance, instead of writing [INT:ϕ] we might write [INT:{ϕ,PART,SPKR,PL}]. For a probe satisfied by any ϕ-feature, I write [SAT:ϕ]. This means that the feature [ϕ], when encountered, terminates further probing of additional targets. If any element that bears any ϕ-feature also bears the feature [ϕ] itself, there is no need to unpack the abbreviation any further. The feature [PART], for instance, does not by itself satisfy a [SAT:ϕ] probe. But [PART] always co-occurs on a DP with [ϕ], which satisfies. A probe that copies all ϕ-features, and stops after the first bearer of ϕ-features, is thus [INT:ϕ, SAT:ϕ]. Such a probe is found on T in a language like German. In a language like English where the bearer of ϕ-features must also move to T, there is an additional ^{M} specification either on the interaction or the satisfaction feature (the two yielding identical outcomes, hence equally good guesses by the learner): either [INT:ϕ^M, SAT:ϕ] or [INT:ϕ, SAT:ϕ^M].

There is much more to say about this algorithm, its applications, and its potential further developments (e.g. in the direction of *dynamic interaction*, as in Deal to appear). For the purposes of this paper I set the modest ambition of demonstrating the response that this theory provides to the issues raised in the previous sections. It should be clear that Béjar's (2003) intuition that Agree is structure-building, i.e. copying, is the central notion. Also central is Béjar's innovation of probes that halt when they encounter a particular feature, but which copy back more than just that feature. For Chirag Dargwa, for instance, the agreement suffix reflects a probe satisfied by [PART] but that copies at least [PART] and [SPKR]. Such a probe is specified [INT:PART, SAT:PART]. Cases of Agree where a probe copies features from multiple targets are likewise handled naturally. For instance, a probe that Agrees with all ϕ features until it encounters the feature [ADDR] is specified as [INT:ϕ, SAT:ADDR] (see Deal 2015a). A probe that Agrees with all ϕ features until it encounters the feature [PART] is specified as [INT:ϕ, SAT:PART] (see Oxford 2020). A probe that Agrees with all *wh* features in its domain, without limit, and moves all targets it interacts with, is specified as [INT:WH^M, SAT:-] (the - indicating the absence of a satisfaction condition). Notably, these last three examples involve specifications where it is not merely the 'halt feature' and everything geometrically entailing it that is copied to the probe, a possibility not recognized in Béjar's (2003) system. Insofar as these specifications prove useful in accounting for the rich typology of agreement systems, the broader range of expected patterns is an advantage of the theory.[18]

It should also be clear how the interaction/satisfaction theory fits within the class of analyses that Preminger (2014) describes as "obligatory operations models". This is not a theory where Agree freely applies, or not, with the outcome assessed against a series of filters, LF or otherwise (e.g. a ban on uninterpretable features present at LF). Agree is instead taken as an operation which is triggered by a lexical item when Merged into a structure, presumably either immediately or in an order determined by lexical specifications of its bearer (as in Georgi 2014). Because Agree is triggered when the probe is Merged, search proceeds downwards, i.e. into the c-command domain; this is the structure that has already been built. Notably, Agree is a process that can result in redundancy iff the structure wherein the probe is Merged is able to support that. In a structure where five targets are local to the probe, none of which can interact with it or satisfy it, all five targets are considered, and then the algorithm exits. Structure building continues, e.g. with further application of Merge. The same holds in a case of "pure satisfaction", where the probe does not encounter any features it can interact with but does encounter its satisfaction feature. These correspond to two cases where Agree "fails", in Preminger's sense, in that no features are copied; either the probe finds no features to copy, or the search is terminated before relevant features are encountered.

A final notable property of the interaction/satisfaction theory, little explored to date, concerns goal flagging. The idea that Agree can change the goal is certainly not a new one, and indeed has been reflected in different strands of the Agree literature in different ways. This includes not only Chomsky's (2000, 2001) proposals regarding uF$'$ on goals, reviewed above, but also the idea that Agree involves feature sharing between probe and goal, such that the goal's features are marked for their connection to the probe (Pesetsky and Torrego 2007), as well as the idea that Agree is decomposed into one step that involves linking probe and goal in this way ("Agree-Link", *per* Arregi and Nevins 2012) versus another that involves actual copying of features ("Agree-Copy", *Ibid.*) (on this type of decomposition, see Deal 2010b, 228-229, 399-400, Arregi and Nevins 2012, Bhatt and Walkow 2013). On the interaction/satisfaction theory, the fact that the goal can be changed by Agree is a general property of the way Agree works. It does not require the goal to have any particular property—e.g. an unvalued uninterpretable feature—which would limit it to participation in only one Agree dependency. Thus cases of multiple agreement, as in the Hindi-Urdu and Swahili data reviewed in section 5, are handled straightforwardly. As for the typology of hyperraising, the theory fits well into the set of approaches reviewed above that connect variation to the theory of locality, whether absolute locality (e.g. phases; Nevins 2004, Deal 2017) or relative locality, i.e. minimality (Halpert 2016, 2019).

A welcome consequence of these developments is a new way of understanding directionality in long-distance dependencies—notably, why movement is always upwards, whereas morphological changes triggered by Agree can occur above or below their trigger. The interaction/satisfaction

[18] This point is discussed in more depth in Deal (2015a, to appear).

theory makes it possible to maintain the view that Agree is always downward, i.e. into a probe's c-command domain. The fact that morphological changes can occur on either side of an Agree dependency simply shows that the operation involves both feature copying to the probe *and* goal flagging. Similarly dissolved on the interaction/satisfaction theory is the apparent puzzle of how Agree can hold between a higher meaningful element (e.g. sentence negation) and a lower "meaning deficient" element (e.g. a negative concord item). On this approach, determination of probe versus goal is simply not a semantic matter. Thus in the analysis of negative concord, there is no obstacle to an analysis that treats sentence negation is a meaningful element that Agrees into its c-command domain (Deal 2021b). The interaction feature is one found on all (and presumably only) negative concord items. As a result of Agree, negative concord items receive a goal flag from negation, which leads in the morphology to their distinctive "negative" morphological forms (e.g. Italian *nessuno* 'nobody').

Appendix: the Agree algorithm once more

In this appendix I briefly present a formalization of the Agree algorithm in pseudocode, a style of presentation intended to foreground the structure of an algorithm while leaving open irrelevant details of implementation for computation.[19] This algorithm goes beyond that presented in the main text in two ways. First, it incorporates dynamic interaction (useful for the analysis of certain types of PCC patterns; Deal to appear). Second, it distinguishes two kinds of goal flags, one connected to interaction and one to satisfaction (a distinction useful for the analysis of certain types of "composite probing" systems, e.g. in Khanty; Deal 2021a). I will briefly review its structure with reference to the line numbers provided to the far left of each line.

The algorithm consists of two loops. The first loop (opened on line 1, concluded on line 19) loops through targets for a probe H, with the order of search presumably determined structurally (see Branan and Erlewine 2020). For each target, two variables (STOP-FLAG and I-STORE) are initialized, line 2. The second loop then commences, looping through the features of the target (lines 3-16). For each feature, three questions are asked.

The first concerns interaction (lines 4-8). If the feature is in the interaction specification, then it is copied to the probe H (line 5) and the target is given an interaction flag, i|H (line 6—the bar in the feature notation mnemonic of a flagpole). If the interaction specification has a movement subfeature ^{M}, and the target is not already in the specifier of H, then the target moves (lines 7-8).

The second question (lines 9-10) concerns dynamic interaction (Deal to appear). If the feature under consideration has a subfeature ↑, then that feature is set aside in I-STORE.

The third and last question concerns satisfaction (lines 11-15). If the feature under consideration is the satisfaction condition of H, then the STOP-FLAG is raised (line 12). This indicates to the algorithm that no further targets should be considered. However, the algorithm does not necessarily halt immediately; remaining features on the current target, if any, *are* considered. The target is given a satisfaction flag, s|H (line 13). If the satisfaction specification has a movement subfeature ^{M}, and the target is not already in the specifier of H, then the target moves (lines 14-15).

Once these questions have been answered for every feature on the current target, the algorithm executes two steps before potentially taking another target into consideration. First, it asks if the STOP-FLAG has been raised—has the satisfaction feature been encountered? If it has (line 17),

[19] For a previous use of pseudocode in grammatical theorizing, see Trommer (1999).

then the algorithm exits. If not, the I-STORE is considered—has a dynamically interacting feature been encountered? If it has (line 18), then the interaction specification on the probe is updated. The new interaction specification is the maximal set of features in the feature geometry that includes the I-STORE feature and all features that it dominates. I refer to this set as the *geometric closure* of I-STORE.

(29) Agree with dynamic interaction

```
1:  for target
2:      let STOP-FLAG = F, I-STORE=∅
3:      for feat(target)
4:          if feat(target) ∈ int(H)  then
5:              copy feat(target) to H
6:              let target have i|H
7:              if int(H) has M and target is not Spec,H then
8:                  remerge target as Spec,H
9:          if feat(target) has ↑ then
10:             let feat(target) → I-STORE
11:         if feat(target) = sat(H) then
12:             let STOP-FLAG = T
13:             let target have s|H
14:             if sat(H) has M and target is not Spec,H then
15:                 remerge target as Spec,H
16:     endfor
17:     if STOP-FLAG = T then exit
18:     if I-STORE ≠ ∅ then let the geometric closure of I-STORE → int(H)
19: endfor
```

References

Abels, Klaus. 2012. The Italian left periphery: a view from locality. *Linguistic Inquiry* 43:229–254.

Alam, Samir, and Elango Kumaran. 2021. Focus-sensitive restrictions on multiargument agreement in Maithili. Presented at WCCFL 39, University of Arizona.

Anagnostopoulou, Elena. 2003. *The syntax of ditransitives: Evidence from clitics*. Walter de Gruyter.

Aravind, Athulya. 2017. A' interactions and feature geometries. In *A Pesky Set: papers for David Pesetsky*, 333–342.

Arregi, Karlos, and Emily Hanink. 2021. Reverse weak PCC in Washo. Talk presented at NELS 52.

Arregi, Karlos, and Emily Hanink. to appear. Switch reference as index agreement. *Natural Language and Linguistic Theory* .

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

Asarina, Alya. 2011. Case in Uyghur and beyond. Doctoral Dissertation, MIT.

Assmann, Anke, Svetlana Edygarova, Doreen Georgi, Timo Klein, and Philipp Weisser. 2014. Case stacking below the surface: On the possessor case alternation in Udmurt. *The Linguistic Review* 31:447–485.

Baier, Nicholas. 2018. Anti-agreement. Doctoral Dissertation, University of California, Berkeley.

Baier, Nico. 2019. The locus of variation in Ā-sensitive agreement. Handout from Parameters workshop for Lisa Travis, McGill University, May 2019.

Baker, Mark, and Livia Camargo Souza. 2020. Agree without agreement: Switch-reference and reflexive voice in two Panoan languages. *Natural Language and Linguistic Theory* 38:1053–1114.

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

Baker, Mark C. 2013. On agreement and its relationship to case: Some generative ideas and results. *Lingua* 130:14–32.

Bárány, András, and Jenneke van der Wal. to appear. We don't agree (only) upwards. *Linguistic Inquiry*.

Béjar, Susana. 2003. Phi-syntax: a theory of agreement. Doctoral Dissertation, Toronto.

Béjar, Susana, and Diane Massam. 1999. Multiple case checking. *Syntax* 2:65–79.

Béjar, Susana, and Milan Rezac. 2003. Person licensing and the derivation of PCC effects. In *Romance linguistics: theory and acquisition*, ed. Ana Teresa Pérez-Leroux and Yves Roberge, 49–62. Amsterdam: John Benjamins.

Béjar, Susana, and Milan Rezac. 2009. Cyclic Agree. *Linguistic Inquiry* 40:35–73.

Bhatia, Sakshi. 2019. Computing agreement in a mixed system. Doctoral Dissertation, University of Massachusetts Amherst.

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, and Hedde Zeijlstra. 2014. Upward agree is superior. Manuscript, lingbuzz.

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

Bobaljik, Jonathan. 2002. A-chains at the PF-interface: copies and 'covert' movement. *Natural Language and Linguistic Theory* 20:197–267.

Boeckx, Cedric, and Youngmi Jeong. 2002. On the nature of syntactic intervention. In *Proceedings of WECOL 14*, ed. Brian Agbayani, Paivi Koskinen, and Vida Samiian, 33–41.

Bošković, Željko. 1999. On multiple feature checking: multiple wh-fronting and multiple head movement. In *Working minimalism*, ed. Samuel Epstein and Norbert Hornstein, 159–187. Cambridge, Massachusetts: MIT Press.

Branan, Kenyon, and Michael Yoshitaka Erlewine. 2020. Locality and (minimal) search. To appear in *Cambridge Handbook of Minimalism*.

Branigan, Phil, and Marguerite Mackenzie. 2002. Altruism, A′-movement, and object agreement in Innu-aimûn. *Linguistic Inquiry* 33:385–407.

Brown, Sue. 1999. *The syntax of negation in Russian : a minimalist approach*. CSLI.

Carstens, Vicki. 2001. Multiple agreement and case deletion: Against *phi*-(in)completeness. *Syntax* 4:147–163.

Carstens, Vicki. 2010. Implications of grammatical gender for the theory of uninterpretable fea-

tures. In *Exploring Crash-Proof Grammars*, ed. Michael Putnam, 31–57. John Benjamins.
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.
Carstens, Vicki, and Michael Diercks. 2013. Parameterizing case and activity: Hyperraising in Bantu. In *NELS 40*, ed. Seda Kan, Claire Moore-Cantwell, and Robert Staubs, 99–118. Amherst: GLSA.
Carstens, Vicki, and Loyiso Mletshe. 2015. Radical defectivity: implications of Xhosa expletive constructions. *Linguistic Inquiry* 46:187–242.
Chomsky, Noam. 1964. The logical basis of linguistic theory. In *Proceedings of the ninth international conference of linguists*, ed. H Lunt. Mouton.
Chomsky, Noam. 1965. *Aspects of the theory of syntax*. Cambridge, Massachusetts: M.I.T. Press.
Chomsky, Noam. 2000. Minimalist inquiries: The framework. In *Step by step: Essays on minimalist syntax in honor of Howard Lasnik*, ed. Roger Martin, David Michaels, and Juan Uriagereka, 89–156. Cambridge, Mass: MIT Press.
Chomsky, Noam. 2001. Derivation by phase. In *Ken Hale: A Life in Language*, ed. Michael Kenstowicz, 1–52. Cambridge, Mass: MIT Press.
Chomsky, Noam. 2004. Beyond explanatory adequacy. In *Structures and beyond: The cartography of syntactic structures*, ed. Adriana Belletti, volume 3, 104–131. Oxford: Oxford University Press.
Clem, Emily. 2019a. Agreement, case, and switch-reference in Amahuaca. Doctoral Dissertation, University of California Berkeley.
Clem, Emily. 2019b. Amahuaca ergative as agreement with multiple heads. *Natural Language and Linguistic Theory* 37:785–823.
Clem, Emily. 2019c. Object-sensitive switch-reference and insatiable probes. In *NELS 49: Proceedings of the Forth-Ninth Annual Meeting of the North East Linguistic Society*, ed. Maggie Baird and Jonathan Pesetsky, volume 1, 173–186. Amherst: GLSA Publications.
Clem, Emily. 2021. Accounting for parallels between inverse marking and the PCC. Talk at NELS 52, Rutgers.
Clem, Emily. to appear. Cyclic expansion in Agree: Maximal projections as probes. *Linguistic Inquiry* .
Collins, Chris. 2002. Multiple verb movement in ǂHoan. *Linguistic Inquiry* 33:1–29.
Collins, James N. 2017. Samoan predicate initial word order and object positions. *Natural Language and Linguistic Theory* 35:1–59.
Coon, Jessica, and Stefan Keine. 2021. Feature gluttony. *Linguistic Inquiry* 52:655–710.
Deal, Amy Rose. 2010a. Ergative case and the transitive subject: a view from Nez Perce. *Natural Language and Linguistic Theory* 28:73–120.
Deal, Amy Rose. 2010b. Topics in the Nez Perce verb. Doctoral Dissertation, University of Massachusetts, Amherst.
Deal, Amy Rose. 2015a. Interaction and satisfaction in ϕ-agreement. In *Proceedings of NELS 45*, ed. Thuy Bui and Deniz Özyıldız, 179–192. Amherst, MA: GLSA Publications.
Deal, Amy Rose. 2015b. A note on Nez Perce verb agreement, with sample paradigms. In *Proceedings of the International Conference on Salish and Neighbouring Languages 50*, ed. Natalie Weber et al., 389–413. Vancouver: UBC Working Papers in Linguistics.
Deal, Amy Rose. 2017. Covert hyperraising to object. In *Proceedings of NELS 47*, ed. Andrew

Lamont and Katerina Tetzloff, 257–270. Amherst, MA: GLSA Publications.
Deal, Amy Rose. 2021a. The logic of agreement: movement, morphology, and composite probes. Talk given at Syntax and Semantics Circle, UC Berkeley, September 2021.
Deal, Amy Rose. 2021b. Negative concord as downward Agree. Manuscript, Berkeley.
Deal, Amy Rose. to appear. Interaction, satisfaction, and the PCC. *Linguistic Inquiry*.
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 program*, ed. Peter W Smith, Johannes Mursell, and Katharina Hartmann, 347–389. Berlin: Language Science Press.
Epstein, Samuel D., Hisatsugu Kitahara, and T. Daniel Seely. 2010. Uninterpretable features: What are they and what do they do? In *Exploring Crash-Proof Grammars*, ed. Michael Putnam, 125–142. John Benjamins.
Ferguson, Charles A, and Michael Barlow. 1988. Introduction. In *Agreement in natural languages: approaches, theories, and descriptions*, ed. Michael Barlow and Charles A Ferguson, 1–22. Stanford: CSLI Publications.
Ferreira, Marcelo. 2004. Hyper-raising and null subjects in Brazilian Portuguese. In *Collected papers on Romance syntax*, volume 47 of *MITWPL*, 57–85.
Fong, Suzana. 2019. Proper movement through Spec-CP: an argument from hyperraising in Mongolian. *Glossa* 4:1–42.
Georgi, Doreen. 2013. Deriving the distribution of person portmanteaux by relativized probing. In *NELS 42*, ed. S. Keine and S. Sloggett, 155–168. Amherst: GLSA.
Georgi, Doreen. 2014. Opaque interactions of Merge and Agree: On the nature and order of elementary operations. Doctoral Dissertation, Universität Leipzig.
Georgi, Doreen. 2017. Patterns of movement reflexes as the result of the order of Merge and Agree. *Linguistic Inquiry* 48:586–626.
Giannakidou, Anastasia, and Hedde Zeijlstra. 2017. The landscape of negative dependencies: Negative concord and n-words. In *SynCom*. Wiley Blackwell.
Haegeman, Liliane, and Marjo van Koppen. 2012. Complementizer agreement and the relation between C^{o} and T^{o}. *Linguistic Inquiry* 43:441–454.
Halpert, Claire. 2016. Raising parameters. In *Proceedings of the 33rd West Coast Conference on Formal Linguistics*, ed. Kyeong-min Kim, 186–195. Cascadilla Proceedings Project.
Halpert, Claire. 2019. Raising, unphased. *Natural Language and Linguistic Theory* 37:123–165.
Harley, Heidi, and Elizabeth Ritter. 2002. Person and number in pronouns: A feature-geometric analysis. *Language* 78:482–526.
Hiraiwa, Ken. 2001. Multiple agree and the defective intervention constraint in Japanese. In *Proceedings of the HUMIT 2000*, ed. O. Matushansky et al., 67–80. Cambridge, MA: MITWPL.
Jenks, Peter, and Himidan Hassen. 2021. Rethinking the distinction between argument- and wh-movement: evidence from Tira. Talk presented at WECOL 2021, November 2021.
Kalin, Laura. 2018. Licensing and differential object marking: The view from Neo-Aramaic. *Syntax* 21:112–159.
Keine, Stefan. 2016. Probes and their horizons. Doctoral Dissertation, University of Massachusetts Amherst.
Keine, Stefan, and Bhamati Dash. to appear. Movement and cyclic Agree. *NLLT*.
Keine, Stefan, and Ethan Poole. 2018. Not all reconstruction effects are syntactic. Unpublished manuscript, USC and UCLA.

Kibrik, A. E. 2003. Materialy k tipologii èrgativnosti. darginskij jazyk (chiragskij dialect). In *Konstanty i peremennye jazyka*, ed. A. E. Kibrik, 482–504. Sankt-Peterburg: Aleteja.

Kramer, Ruth. 2014. Clitic doubling or object agreement: the view from Amharic. *Natural Language and Linguistic Theory* 32:593–634.

Laka, Itziar. 1990. Negation in syntax: on the nature of functional categories and their projections. Doctoral Dissertation, Massachusetts Institute of Technology.

Merchant, Jason. 2006. Polyvalent case, geometric hierarchies and split ergativity. In *Proceedings of the 42nd Annual Meeting of the Chicago Linguistics Society*, ed. J. Bunting et al., volume 2, 57–76. Chicago Linguistics Society.

Merchant, Jason. 2011. Aleut case matters. In *Pragmatics and Autolexical Grammar: In honor of Jerry Sadock*, ed. Etsuyo Yuasa, Tista Bagchi, and Katherine Beals, 193–210.

Nevins, Andrew. 2004. Derivations without the activity condition. In *Proceedings of the EPP/Phase Workshop*. MITWPL.

Nevins, Andrew. 2007. The representation of third person and its consequences for person-case effects. *Natural Language and Linguistic Theory* 25:273–313.

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

Ogihara, Toshiyuki. 1996. *Tense, attitudes, and scope*. Dordrecht: Kluwer Academic.

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

Oxford, Will. 2020. Probe specification and agreement variation: Evidence from the Algonquian inverse. Manuscript, University of Manitoba, September 2020.

Pandharipande, Rajeshwari, and Yamuna Kachru. 1977. Relational grammar, ergativity, and Hindi-Urdu. *Lingua* 41:217–238.

Penka, Doris. 2011. *Negative indefinites*. Oxford University Press.

Pesetsky, David. 2013. *Russian case morphology and the syntactic categories*. MIT Press.

Pesetsky, David, and Esther Torrego. 2001. T-to-C movement: Causes and consequences. In *Ken Hale: A life in language*, ed. Michael Kenstowicz, 355–426. 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. V. Samiian S. Karimi and W. Wilkins, 262–294. Amsterdam: John Benjamins.

Pfau, Roland, Martin Salzmann, and Markus Steinbach. 2018. The syntax of sign language agreement: Common ingredients, but unusual recipe. *Glossa* 3:1–46.

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

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

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

Preminger, Omer, and Maria Polinsky. 2015. Agreement and semantic concord: a spurious unification. Ms., University of Maryland.

Puškar, Zorica. 2018. Interactions of gender and number agreement: Evidence from Bosnian/Croatian/Serbian. *Syntax* 27:275–318.

Rezac, Milan. 2003. The fine structure of Cyclic Agree. *Syntax* 6:156–182.

Rezac, Milan. 2011. *Phi-features and the modular architecture of language*. Springer.

Richards, Norvin. 2013. Lardil "case stacking" and the timing of case assignment. *Syntax* 16:42–76.

Rodrigues, Cilene. 2004. Impoverished morphology and A-movement out of case domains. Doctoral Dissertation, University of Maryland.

Roversi, Giovanni. 2020. How to satisfy probes: person/number hierarchy effects in Äiwoo. In *Proceedings of NELS 50*, ed. Mariam Asatryan, Yixiao Song, and Ayana Whitmal, volume 3, 99–112. Amherst: GLSA Publications.

Rudnev, Pavel. 2020. Agreeing adpositions in Avar and the directionality-of-valuation debate. *Linguistic Inquiry* 51:829–844.

Rudnev, Pavel. 2021. Against Upwards Agree. *The Linguistic Review* 38:65–99.

Schütze, Carson T. 1997. Infl in child and adult language: agreement, case, and licensing. Doctoral Dissertation, MIT.

Scott, Tessa. 2021. Formalizing two types of mixed A/A' movement. Manuscript, UC Berkeley.

Starke, Michal. 2001. Move dissolves into Merge: a theory of locality. Doctoral Dissertation, University of Geneva.

Stowell, Tim. 1995. What is the meaning of the present and past tenses? In *Temporal Reference: Aspect and Actionality Vol. 1: Semantic and Syntactic Perspectives.*, ed. Pier Marco Bertinetto, Valentina Bianchi, and M Squartini, 381–396. Torino: Rosenberg & Sellier.

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

Sumbatova, Nina R. 2011. Person hierarchies and the problem of person marker origin in Dargwa: facts and diachronic problems. In *Tense, aspect, modality, and finiteness in East Caucasian languages*, ed. Gilles Authier and Timur Maisak, 131–160. Bochum: Universitätsverlag Dr. N. Brockmeyer.

Trommer, Jochen. 1999. Morphology consuming syntax's resources: Generation and parsing in a minimalist version of distributed morphology. In *Procceedings of ESSLI Workshop on Resource Logics and Minimalist Grammars*.

Ura, Hiroyuki. 1996. Multiple feature-checking: A theory of grammatical function splitting. Doctoral Dissertation, MIT.

van Urk, Coppe. 2015. A uniform syntax for phrasal movement: A case study of Dinka Bor. Doctoral Dissertation, MIT.

Wurmbrand, Susi. 2012a. Agree(ment): looking up or looking down? Handout from Agreement seminar, MIT.

Wurmbrand, Susi. 2012b. Parasitic participles: Evidence for the theory of verb clusters. *Taal en Tongval* 64:129–156.

Wurmbrand, Susi. 2021. From prolepsis to hyperraising. Handout from Abralin, October 2021.

Zanuttini, Rafaella. 1991. Syntactic properties of sentential negation: A comparative study of Romance languages. Doctoral Dissertation, University of Pennsylvania.

Zeijlstra, Hedde. 2004. Sentential negation and negative concord. Doctoral Dissertation, Amsterdam.

Zeijlstra, Hedde. 2008. Negative concord is syntactic agreement. Ms., Amsterdam.

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

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

Zobel, Justin. 2004. *Writing for computer science*. London: Springer.

Zyman, Erik. 2017. P'urhepecha hyperraising to object: An argument for purely altruistic movement. In *Proceedings of the Linguistic Society of America*, volume 2.