The featural life of nominals

Ivy Sichel & Maziar Toosarvandani
University of California, Santa Cruz

Abstract We introduce a novel locality violation and its repair in Sierra Zapotec: an object pronoun cannot cliticize when the subject is a lexical DP. This locality effect differs from more familiar ones (e.g., superiority) because the lexical DP does not move. We argue that it is nonetheless able to Agree, consistent with the idea that locality conditions apply to Agree, rather than to a separate movement component. We develop an account in which pronouns and lexical DPs interact with the same probe because they share featural content. In particular, we suggest that the person domain extends to include non-pronominal DPs, so that all nominals are specified for a feature we call δ (to resonate with DP); all and only personal pronouns are specified for π. This account requires a departure from Chomsky's (2000, 2001) classical system of featural co-variation (Agree). A functional head must be able to participate in *overprobing*, interacting with a goal even though its requirements would appear to be met. We introduce a *probe activation* model for Agree, in which, after applying once, the operation can but not need apply again. We also consider two other mechanisms recently proposed to derive overprobing — Deal's (2015, 2020) "insatiable probes" and Coon & Keine's (to appear) "feature gluttony" — though neither is able to account for the locality pattern.

1. Pronouns and lexical DPs

How does natural language differentiate pronouns from non-pronominal DPs? And how does it unify them? Within formal semantics, a lively debate continues to examine the referential trajectories of pronouns and definite descriptions, and whether they are categorically different in this respect or not. But within syntax, discussions of how these nominal classes might overlap or diverge have been surprisingly narrow. They have focused almost exclusively on the typology of nominal structure, and the degree to which pronouns should be distinguished from *lexical DPs* (also called full DPs) in terms of the hierarchical structure they contain: whether pronouns differ from lexical DPs in how much external structure is present — if they are a DP or some smaller constituent (e.g., Cardinaletti & Starke 1999) — and how much (silent) internal structure they possess, such as an elided NP (e.g., Patel-Grosz & Grosz 2017).

There exists, however, a family of quintessentially syntactic phenomena, involving pronominal displacement and cliticization, which sharply distinguish pronouns from lexical DPs. This includes scrambling and object shift, which in many languages affect only pronouns; clitic doubling, which also in many languages targets only pronouns; restrictions on pronominal cliticization like the Person–Case Constraint (PCC); and, differential object marking, which in some languages differentiate between pronouns and lexical DPs. The relevant distinction here does not appear to be directly related to any semantic notions involving referentiality because it places all non-pronominal DPs on the same side. Nor does the hierarchical structure inside nominals obviously seem to be the key for understanding why such different syntactic

phenomena make the same cut. It seems, rather, that since these phenomena involve displacement, natural language distinguishes pronominal from non-pronominal DPs in a currency that is legible to the syntactic operations underlying movement: namely, features. Are there features that all nominals possess? And are there morphosyntactic features that only personal pronouns or lexical DPs possess? There is a rich literature on the featural representation of pronouns — What features distinguish first person from third person? — as well as on the features that are shared across nominal classes, such as number. But hardly any work considers the possibility that pronouns might be distinguished, in featural terms, and as a class, from non-pronouns.

In this paper, we introduce a novel phenomenon related to the movement of pronouns, which provides a window into the featural life of nominals and suggests that just such a distinction is needed. This phenomenon involves a locality violation and its repair in Sierra Zapotec.[1] The locality effect arises when an object clitic pronoun occurs with a subject that is a lexical DP. While subject pronouns can cliticize when the object is a lexical DP, as in (1a), an object pronoun cannot cliticize across a subject lexical DP, as in (1b).

(1)　　a.　*pro > DP*
　　　　　Blenh^4=**ba'^{3}**$_1$　　　　　t$_1$　　　be^{2}ku'^{4}=nh.
　　　　　carry.COMP=**3.HU**　　　　　　dog=DEF
　　　　　'S/he carried the dog.'　　　　　　　　　　　(FSR, SLZ1051-s, 1)
　　　　b.　*DP > pro*
　　　　　*Blenh^4=**eb^4**$_1$　　　　Xwanh^{24}=a'^{4}　　t$_1$.
　　　　　carry.COMP=**3.AN**　　Juana=DEF
　　　　　Intended: 'Juana carried it (an animal).'　　　(FSR, SLZ1051, 7:30)

What makes this pattern so puzzling is that the subject blocks movement of the object even though it does not itself move. In this respect, it is quite different from more familiar interactions between multiple wh-phrases, where a subject that moves prohibits an object from moving.

We will suggest that, although the subject does not move, it is able to Agree, and we will develop an account in which pronouns and lexical DPs can interact with the same probe because they share a certain feature, which we name δ (to resonate with the functional head D). While pronouns have δ plus additional person features, lexical DPs only have δ. Laying out the ingredients for an account of this effect opens up a new perspective on the relationship between the system responsible for featural co-variation (that is, Agree) and the system responsible for displacement itself. In particular, it supports the decomposition of movement into Agree, which is subject to an intervention-based locality constraint, plus an (internal) Merge operation that is otherwise free (Chomsky 2000).

We start, in Section 2, with a discussion of locality and its role in deriving the restriction in (1). As we will argue, only a notion of locality that is syntactic and based on the logic of intervention will suffice. We incorporate this locality into our proposal, which has two components:

---

[1] What we call Sierra Zapotec is a group of closely related Northern Zapotec varieties from the southeastern Sierra Norte. We include data from two of these varieties: (Santiago) Laxopa (for which we report our own fieldwork data) and (Hidalgo) Yalálag (Avelino Becerra 2004, Lopez and Newberg 2005).

1. **Extended person:**
   Lexical DPs and pronouns share a feature (δ) that is part of the structure of person; only personal pronouns are specified for additional person features.
2. **Probe activation:**
   A functional head must Agree with a goal, subject to Attract Closest. Additional Agree relations with other goals are possible, though not required, just in case they are not featurally distinct from the probe (in the requisite way).

With the first component, which we discuss in Section 3, we aim to draw a parallel between the restriction involving lexical DPs, illustrated in (1), and another restriction on cliticization in Sierra Zapotec involving combinations of pronouns, akin to the Person–Case Constraint (PCC) found in Romance and many other languages (Perlmutter 1971, Bonet 1991).

    With the second component, introduced in Section 4, we advance a unified mechanism for deriving both restrictions on pronominal cliticization. This requires a departure from Chomsky's (2000, 2001) classical Agree system. We advance a probing mechanism in which a functional head is "activated" through Agree with one goal, so that it can — but need not — Agree with subsequent goals. We compare this *probe activation* model to two recent theories designed to account for similar *overprobing* patterns: Deal's (2015, 2020) "insatiable probing" and Coon & Keine's (to appear) "feature gluttony." Neither is able to derive the locality effect in (1) involving the interaction between pronominal cliticization and lexical DPs.

2. What is locality?

We can start by completing the paradigm in (1). Both the subject and object can be clitic pronouns (2a) or lexical DPs (2d). While the subject can cliticize when the object is a lexical DP (2b), pronominal cliticization is not allowed across a lexical DP (2c).

(2)   a. *pro > pro*
          Blenh^4=ba'$^3_1$=b$_2{}^4$ t$_1$ t$_2$.
          carry.COMP=3.HU=3.AN
          'S/he carried it (an animal).'   (FSR, SLZ1051-s, 3)
    b. *pro > DP*
          Blenh^4=ba'$^3_1$ t$_1$ be^{2}ku'^{4}=nh.
          carry.COMP=3.HU   dog=DEF
          'S/he carried the dog.'   (FSR, SLZ1051-s, 1)
    c. *DP > pro*
          *Blenh^4=eb$^4_1$ Xwanh^{24}=a'^{4} t$_1$.
          carry.COMP=3.AN   Juana=DEF
          Intended: 'Juana carried it (an animal).'   (FSR, SLZ1051, 7:30)
    d. *DP > DP*
          Blenh^4 Xwanh^{24}=a'^{4} be^{2}ku'^{4}=nh.
          carry.COMP   Juana=DEF   dog=DEF
          'Juana carried the dog.'   (FSR, SLZ1051-s, 2)

Some kind of locality is at play here, since subject and object cliticization are not equally available: it is only the higher of two arguments in (2a–b) that can cliticize. We call this the *Lexical DP Blocking* effect:

(3)   *Lexical DP Blocking (LDB)*
       An object pronoun cannot cliticize when the subject is a lexical DP.

It is important to point out that the LDB is only relevant for third person object pronouns, as local person pronouns do not permit cliticization in object position even when the subject cliticizes. This is the well-known (Strong) Person–Case Constraint (PCC), found in Romance and many other languages, which prohibits first and second person object clitics even when a higher argument cliticizes (Perlmutter 1971, Bonet 1991) (see Section 4.1 for additional details).

There is a structural parallel between the LDB and the well-known superiority effect observed in multiple wh-questions. In the simplest type of multiple wh-question, only the highest wh-phrase can move (though this simple generalization is complicated by numerous interacting factors, including D-linking and transitivity; see Pesetsky 2001 for discussion).

(4)   a.   Who t_1 held what$_2$?
       b.   *What$_2$ did who$_1$ hold t_2?

Superiority was taken by Rizzi (1990) to follow from Relativized Minimality, an intervention-based locality constraint on syntactic representations. It prohibits movement chains in which a moved element (X) and its trace (Z) are separated by an element (Y) that shares properties with X. In subsequent work, Relativized Minimality has been formalized in featural terms (Rizzi 2004).

(5)   *Relativized Minimality (Rizzi 2004)*
       *X ... Y ... Z
       [F]   [F]

The simplest conception of Relativized Minimality prohibits X from moving past Y if they both possess the same feature F. For multiple wh-questions, this could be a [wh] feature.

This concept of locality presumes that movement is a primitive operation in the syntax. More recent theorizing has sought, however, to understand how movement can be reduced to more basic operations. If it can, then the locality effects typically associated with movement might be derived from those operations. Under Chomsky's (2000:135–137) decomposition of movement into Agree and (internal) Merge, superiority effects can be traced to a condition which applies solely to the Agree component. If Merge is the operation actually responsible for displacement (via copying) and is freely available (applying at no cost and free of inherent conditions), then this derivational step may not be subject to any specific locality considerations beyond those deriving from cyclicity. Instead, it is Agree which, as a precondition for Merge, would be subject to locality. One such constraint, Attract Closest, prohibits Agreement between a probe (X) and a goal (Z) when there is a closer potential goal (Y). What counts as a goal is defined in featural terms: both Y and Z are potential goals because they share a feature (F) the probe is looking for.

(6) *Attract Closest (Chomsky 2000:122)*
     *X ... Y ... Z
     [uF] [F]   [F]
       |_____|

This derives Relativized Minimality, along with the superiority effect exhibited by wh-movement. Movement of the lower wh-phrase is never a derivational option, since only the higher wh-phrase can be targeted by Agree.

   Returning to the LDB, it differs from more familiar superiority effects in an obvious way. While an object cannot move past a lexical DP subject, this argument never itself moves. This is consistent with something like Attract Closest in (6), which applies to Agree and not to Merge directly, and may ultimately support the shift to Agree-based locality. But it might also suggest that the logic of attraction is irrelevant, and that the ill-formedness of object cliticization in (2c) is produced through a different locality calculus. Indeed, the paradigm in (2) can also be characterized in a way that only makes reference to movement: if an object is a pronoun, it will cliticize just in case the subject also cliticizes. We should thus consider the possibility that the source of the asymmetry lies in conditions on movement. As we will argue below, however, movement is in fact orthogonal to the problem. The pattern in Sierra Zapotec is best understood as the product of an intervention-based locality constraint on Agree. This, in turn, urges us to seek a feature that lexical DPs and pronouns have in common.

2.1. Locality at the interfaces

If movement is reduced to Merge, little room is left for locality conditions on movement in the syntax, beyond those constraining the operations that are preconditions for Merge. While some properly syntactic constraints have been proposed, such as Shortest Move (Murasugi 1992, Kitahara 1997, Ochi 1997, Richards 1997:109–112, 2002), these have not been incorporated into the decompositional view of movement. And Shortest Move, specifically, does not obviously help with the paradigm in (2). This economy constraint mandates the shortest well-formed movement path, relative to other well-formed movement paths, as determined by some metric (e.g., counting dominating nodes). To rule out object cliticization in (2c), when the higher argument is a lexical DP, there would have to be an alternate well-formed movement dependency that is shorter. But this is precisely what is absent in this paradigm: the lexical DP does not move.

   Under the decompositional approach, movement dependencies can be constrained directly, but only at an interface after they have been formed. A family of *shape conservation* constraints have been proposed to do this, applying at PF to preserve the linear order of elements across levels of syntactic representation (Müller 2000, Fox & Pesetsky 2005; see also Sells 2001, Williams 2003). These have been argued to underlie the blocking effects on pronoun movement in Scandiavian, known as Holmberg's Generalization (Holmberg 1986, 1999). Object shift of a weak pronoun is, for instance, blocked across a verb; it can move only if the verb also moves.

(7) *Swedish*
     a. Jag   kysste$_1$ henne$_2$ inte t_1 t_2.
         I   kissed   her   not
     b. *Jag har   henne$_2$ inte   kysst t_2.
          I   have   her   not   kissed   (Holmberg 1999:1)

While shape conservation theories differ in their details, they all trace the ill-formedness of (7b) to a linear misalignment across syntactic domains (which might be defined cyclically or in some other fashion). If the verb precedes an object pronoun inside, say, the VP, then it must also precede it in all larger domains. Shape conservation is violated if a weak object pronoun moves without the verb also moving.

(8) a. $[_{VP}$ kysst henne]   VP: kysst < henne
     b. $[_{TP}$ jag har henne$_2$ inte $[_{VP}$ kysst t$_2$]]   TP: henne < kysst

By contrast, in the well-formed (7a), the additional verb movement preserves the linear order of the lower domain in the higher domain.

It is clear how shape conservation could be extended to account for the ungrammaticality of (2c). Assuming that at the initial level of representation, the subject and object occur within the same domain (see Fox & Pesetsky 2005:35–38), then their order will have to be preserved in larger domains. This is not possible if the object pronoun moves across the lexical DP subject (which itself would not leave the lower domain).

(9) a. $[_{vP}$ **Xwanha'** blenh =eb]   vP: **Xwanha'** < blenh < =eb
     b. $[_{TP}$ blenh$_1$=eb$_2$ $[_{vP}$ **Xwanha'** t$_1$ t$_2$]]   TP: blenh < =eb < **Xwanha'**

So there is no doubt that shape conservation could derive the ungrammaticality of (2c), if it had the derivation in (9a–b).

But the syntax is not, in fact, able to produce such a representation in the first place. Further shape-conserving movement of the subject to a position that precedes the landing site of the moved object clitic does not remediate the problem caused by cliticizing the object. For instance, A′-moving the subject, either for focus-related reasons (10a) or through wh-movement (11a), does not repair the LDB violation. In these configurations, the object can only surface in its base position as an independent pronoun, as in (10b) and (11b).

(10) a. *Xwanh^{24}=a'$^4{}_1$ blenh^4=**b**$^4{}_2$ t$_1$ t$_2$.
         Juana=DEF   carry.COMP=**3.AN**
         Intended: 'JUANA carried it.'   (FSR, SLZ1051, 14:54)
     b. Xwanh^{24}=a'$^4{}_1$ blenh^4 t$_1$ **leb^{24}**.
         Juana=DEF   carry.COMP   **3.AN**
         'JUANA carried it.'   (FSR, SLZ1051-s, 7)
(11) a. *Nhu$^{42}{}_1$ blenh^4=**eb**$^4{}_2$ t$_1$ t$_2$?
         who   COMP.carry=**3.AN**
         Intended: 'Who carried it?'   (FSR, SLZ1051, 16:31)
     b. Nhu$^{42}{}_1$ blenh^4 t$_1$ **leb^{24}**?
         who   COMP.carry   **3.AN**
         'Who carried it?'   (FSR, SLZ1051-s, 8)

Even if shape conservation is active at PF in Sierra Zapotec, it cannot repair a violation earlier in the derivation. This implies that the violation itself is syntactic. We thus conclude that the locality restriction underlying the LDB is part of syntax, and more specifically part of Agree.[2]

2.2. Another constraint on object cliticization

Our primary analytical claim is that the LDB is parallel to a different constraint, which also makes reference to the grammatical properties of multiple arguments.

In Sierra Zapotec, object pronouns are unable to cliticize even when the subject is another pronoun, depending on the animacy of both arguments. The language realizes a four-way animacy distinction in third person pronouns: elder humans (EL) vs. non-elder humans (HU) vs. animals (AN) vs. inanimates (IN). An elder object pronoun cannot cliticize when the subject is human (12b), though a human object pronoun can cliticize when the subject is elder (12a).

(12)　　a.　*3.EL > 3.HU*
　　　　　　Wkwell=e'$_1$=be'$_2$　　　　　　t$_1$　　　t$_2$.
　　　　　　kick.COMP=3.EL=3.HU
　　　　　　'He kicked him.'
　　　　b.　*3.HU > 3.EL*
　　　　　　*Wkwell=be'$_1$=e'$_2$　　　　　　t$_1$　　　t$_2$.
　　　　　　kick.COMP=3.HU=3.EL
　　　　　　Intended: 'He kicked him.'　　　　　　　　(Avelino Becerra 2004:33–35)

This contrast parallels (2–b), except that the subject pronoun that blocks movement of the object is itself a candidate for movement. In other words, a human pronoun in subject position can block cliticization of an object pronoun, just as a lexical DP does: compare (12b) to (2c). But unlike a lexical DP, this human pronoun can itself undergo movement: compare (12a) to (2b).

Shape conservation is silent about the ill-formedness of (12b). It penalizes linear misalignments across syntactic domains, but here there is no such problem. The higher argument moves, plausibly to a position outside the lower domain, just like the lower argument.

---

[2] In Scandinavian, of course, A′-movement *can* repair a violation of Holmberg's Generalization (Holmberg 1999). The illicit object shift across an indirect object in (i) is repaired by extracting the indirect object (ii).

(i)　*Jag　　gav　　**den**$_1$　　inte　　Elsa　　t$_1$.
　　　I　　　gave　　it　　　not　　Elsa
(ii)　Henne$_1$　visar　　jag　　**den**$_2$　　helst　　inte　　t$_1$ t$_2$.
　　　her　　　show　　I　　　it　　　rather　　not
　　　'I'd rather not show it to HER.'　　　　　　　　(Holmberg 1999: 17)

This is amenable to a shape conservation account like Fox & Pesetsky's (2005). Anagnostopoulou (2004) argues that a syntactic account is also possible by enabling the apparent counter-cyclic extraction via an intermediate projection with multiple specifiers.

If the Swedish intervention effect in (i) is indeed syntactic, as we argue for Sierra Zapotec, the larger question is why it is indirect objects that intervene in Scandinavian and not subjects. Here, we cannot help but observe that subjects surface in different positions: in Sierra Zapotec the subject is lower than in Swedish, where it is in specTP. This places it in a position that is lower than the T probe and higher than the object (somehow, A-movement of the subject in Swedish must remove the potential for intervention).

(13) a. $[_{vP}$ wkwell$_1$ $=$**be'**$_2$ $=$**e'**$_3]$ vP: wkwell $<$ $=$**be'** $<$ $=$**e'**
      b. $[_{TP}$ wkwell$_1=$**be'**$_2=$**e'**$_3$ $[_{vP}$ t$_1$ t$_2$ t$_3]]$ TP: wkwell $<$ $=$**be'** $<$ $=$**e'**

Foley & Toosarvandani (to appear) propose that the contrast in (12a–b) arises from the same source as the *Person–Case Constraints (PCCs)*, a family of restrictions on clitic and other weak pronouns based on their structural position and location on a person hierarchy (Perlmutter 1971, Bonet 1991). It is not just the combination of clitics in (12b) that is prohibited: object cliticization is forbidden whenever the animacy of the object exceeds the subject's on an intuitive hierarchy of animacy.[3] By analogy to the PCC, they call this restriction a *Gender–Case Constraint (GCC)*, following traditional categorizations of animacy as part of gender.

(14) *Gender–Case Constraint (GCC)*
      An object clitic pronoun cannot exceed a subject clitic pronoun on the animacy
      hierarchy, i.e., EL > HU > AN > IN.

The effects of the GCC were shown just for combinations of human pronouns in (12a–b) above, and it is demonstrated pairwise for the other animacy categories in (15) and (16).

(15) a. *3.HU > 3.AN*
          Bchew=be'=ba'.
          kick.COMP=3.HU=3.AN
          'S/he kicked it.'
      b. *3.AN > 3.HU*
          *Bdinn=ba'=be'.
          bite.COMP=3.AN=3.HU
          'It bick her/him.'
(16) a. *3.AN > 3.IN*
          Bchochj=ba'=en.
          hit.COMP=3.AN=3.IN
          'It hit it.'
      b. *3.IN > 3.AN*
          *Bchochj=en=ba'.
          hit.COMP=3.IN=3.AN
          'It hit it.'   (after Avelino Becerra 2004: 34)

Thus, object clicization is systematically excluded even when the subject does cliticize and the original order is preserved. This requires a grammatical source independent of interface conditions tied to shape conservation.

Foley & Toosarvandani derive this hierarchy sensitivity from a grammatical mechanism sensitive to the feature structure of animacy. To encode the four-way animacy distinction in Sierra Zapotec, they posit three privative features — EL(DER), HU(MAN), and AN(IMATE) — that are organized into a geometry based on the semantic entailment relations amongst them. Elder

---

[3] This simplifies things somewhat, as there is variation within Sierra Zapotec in how strictly the hierarchy is obeyed. We have presented data here from the Yalálag variety, which conforms to it most completely.

pronouns are fully specified, bearing all three features; non-elder humans and animal pronouns are partially specified; inanimates bear none of these features.

(17) a. *3.EL*   b. *3.HU*   c. *3.AN*   d. *3.IN*
          AN   AN   AN
          | |
          HU   HU
          |
          EL

Building on a line of work on the PCC initiated by Anagnostopoulou (2003), they take the GCC to arise from an Agree operation that holds between a functional head and the arguments in its c-command domain. We will discuss the precise Agree mechanism that could be involved more fully in Section 4. But for now, we simply point out that, given the feature inventory in (17), it must enable cliticization of two arguments when the subject is more featurally specified than the object, e.g., 3.EL > 3.HU (12a), but not when it is less featurally specified than the object, e.g., 3.HU > 3.EL (12b). This mechanism can be extended, we will argue, to the LDB, though this requires elaborating how both pronouns and lexical DPs are represented grammatically, as well as a theory of Agree sufficiently articulated to interface with these representations. In doing this, our goal is to move away from construction-specific characterizations of both patterns and any direct reference to cliticization.

2.3. Toward a unified syntactic source

Importantly, we are not claiming that the LDB can simply be subsumed under the GCC. It holds for any DP subject and any pronominal object, regardless of animacy, including cases where the subject DP happens to denote an entity that is higher on the animacy hierarchy than the objective, e.g., (2c) above. Rather, we are proposing that the Agree mechanism underlying the GCC is also responsible for the LDB.

As in Anagnostopoulou's (2003) account of the PCC, we take this shared mechanism to include an important role for an Agree operation that holds between a functional head and the arguments that it c-commands. Importantly, we will argue that this Agree relation is subject to an intervention-based locality constraint like Attract Closest. Thus, for both LDB and the GCC, object cliticization is prohibited because the subject intervenes for probing and is closer than the object. Both patterns are, in other words, types of intervention effects, not so different in the end from superiority, under the familiar Agree-based account. The superficial differences amongst these patterns arise for independent reasons having to do with the specific properties of the goals involved, though the core computational mechanism underlying all of them is the same.

Initial evidence for this approach comes from two asymmetries that privilege the subject over the object in the LDB and the GCC. First, in both patterns, it is the properties of the higher argument that determine whether the lower argument can move. For the LDB, this is whether it is a lexical DP; for the GCC, the higher pronoun sets a threshold on a conceptual hierarchy of animacy which cannot be exceeded by an object clitic. Second, the repair mechanism for violations of both constraints is the same. When the GCC is violated, the subject pronoun must still move, with the object being realized as an independent pronoun (18a–b). Similarly, when the subject is a lexical DP, the object is also realized as an independent pronoun.

(18) a. Udi'in^4=**eb^4**$_1$ t$_1$ le^{2}ba'^{3}.
     bite.COMP=**3.AN**   3.HU
     'It (an animal) bit her/him.'   (FSR, SLZ1051-s, 5)
   b. *Udi'in^4=**ba'^{3}**$_2$ leb^{24} t$_2$.
      bite.COMP=**3.HU**   3.AN
     Intended: 'It (an animal) bit her/him.'   (FSR, SLZ1051, 13:45)

(19) Blenh^4 Xwa^{24}nh=a'^{4} **leb^{24}**.
     carry.COMP   Juana=DEF   **3.AN**
     'Juana carried it.'   (FSR, SLZ1051-s, 6)

In principle, it is possible that these two constraints could have different repairs, such as realizing the higher argument as a strong pronoun (for the GCC) or moving the higher argument. But they do not. Of course, to evaluate this second argument, the data points in (18–19) are not enough: a unified theory of both violations and their repairs is needed. We develop the details of such a theory in the rest of this paper.

The core mechanism underlying both the LDB and GCC is, as in Anagnostopoulou's theory of the PCC, a functional head which can Agree with both the subject and object, subject to Attract Closest. The LDB/GCC arises from how exactly this probe interacts with the arguments in its domain. In this respect, we advance two hypotheses:

1. **Extended person:**
   Lexical DPs and pronouns share a feature (δ) that is part of the structure of person; only pronouns are specified for additional person features.
2. **Probe activation:**
   A functional head must Agree with a goal, subject to Attract Closest. Additional Agree relations with other goals are possible, though not required, just in case they are not featurally distinct from the probe (in the requisite way).

The first ingredient is necessitated by our proposal that lexical DPs intervene for Agree with a pronoun. According to the logic of Attract Closest, pronouns and lexical DPs must be the "same" in some way that involves their featural content. In Section 3, we identify two kinds of features, those which bring lexical DPs and pronouns together and those which distinguish between them. In particular, we will argue that all nominals possess a person feature, which we call δ. At the same time, pronouns differ from lexical DPs in having additional person features. This places pronouns and lexical DPs in some sense on a "hierarchy," with pronouns exceeding lexical DPs because they are more featurally specified, by analogy to animacy.

To derive the LDB/GCC, the featural specification of lexical DPs and pronouns interact with the second ingredient in this theory. In both patterns, an object clitic cannot exceed a subject pronoun on some hierarchy, whether this is an animacy hierarchy or the admittedly less obvious "hierarchy" ordering lexical DPs and pronouns. In Section 4, we show how this can be derived from a *probe activation* model for Agree. We propose that a probe can only interact with both the subject and object if the former is more featurally specified than the latter. Importantly, this activation mechanism has nothing, in and of itself, to do with movement, thereby offering a unified account of both the LDB and GCC. The superficial difference between the two patterns

— namely that the subject still moves when the GCC is violated, while a lexical DP subject never moves — arises from the nature of pronominal cliticization. In both cases, the probe Agrees with the subject, though it only triggers movement of a pronoun in that position.

3. Extending the domain of person

By assimilating the LDB and GCC, our empirical claim is that the relation between a lexical DP subject and an object clitic is the same as the relationship between a subject clitic and object clitic in the ungrammatical alignments in (12b), (15b), and (16b). We are arguing, in other words, for a generalization across both patterns. Their grammatical source, which we will discuss in detail in Section 4, involves a functional head that Agrees with lexical DPs and with pronouns, subject to Attract Closest. This will account for the asymmetry inherent to both patterns: subjects intervene for objects, but not vice versa.

For both lexical DPs and pronouns to be eligible as goals for the same probe in the first place, though, they must count in some way as the "same." There must be some property that is shared by pronominal and non-pronominals alike, so that a clitic is sensitive to a higher lexical DP, but not to a higher PP or AP. At the same time, there must be some property that distinguishes them, since they are obviously not identical. Finally, these properties must be represented as features, rather than as categories, since this is the type of entity that is visible to the Agree mechanism.[4]

We propose that the category of person is implicated in both this sameness and this difference. In particular, we suggest that all nominals, including lexical DPs and pronouns, are specified for a person feature, which we call δ. Only personal pronouns have further person specifications, minimally including a π feature (Béjar 2003). These features stand in a structured relationship to each other: the presence of π asymmetrically requires the presence of δ.

(20) a. *Lexical DPs*   b. *Pronouns*
           δ δ
                                        |
                                        π
                                        ⋮

We are committed, minimally, to the existence of person specifications that both unify lexical DPs with pronouns and distinguish them from one another. While we adopt δ and π as features, it may also be possible, as we will discuss below, to conceive of them as syntactic categories. Whatever they are, they must stand in a structured relationship to each other and to the rest of the person domain, and they must exist in a currency that is visible to Agree. For presentational purposes, we have represented this structure in (20) using a feature geometry, but this is not the only way to encode the dependencies between these features, as we will discuss further below.

We should note that both lexical DPs and third person pronouns are sometimes taken to lack person features (e.g., Noyer 1997, Harley & Ritter 2002, Ackema & Neeleman 2018: 52–56). But to capture the LDB, personal pronouns and lexical DPs must be united by a positive

---

[4] We do not exclude the possibility that features are equivalent to categories; our point here is simply that the shared property cannot simply be that both lexical DPs and pronouns are DPs, since "DP" is not an entity that interacts with Agree. For more on the relationship between features and categories, see the discussion in Section 3.2 below.

property. Furthermore, the analytical moves which eliminate third person are only justified if the category of third person is homogenous. As Béjar (2003: 47–50) argues, though, many languages draw more fine-grained distinctions within third person. If "third person" is not a unified category and instead covers multiple sub-categories, a positive feature or set of features corresponding to "third person" is necessary to capture these distinctions.

Below, we motivate one account of δ and its relation to π and the rest of the person domain, which assumes that person features are privative (Harley & Ritter 2002, Bejar 2003) and denote one-place predicates that combine intersectively (e.g., Heim & Kratzer 1998, Sauerland 2006).[5] Under these assumptions, a person feature like Béjar's π is needed to distinguish pronouns from lexical DPs, a distinction that we relate to other oppositions within third person. If there are no categories independent of features, all nominals must share a common feature, which we name δ, the featural relative of D. We suggest that it is responsible for the shared ability of all nominals to individuate. While this property has been attributed to nominal roots (Baker 2003), we attribute it instead to nominal functional structure.

In sum, we are claiming that lexical DPs and pronouns share a person feature (δ) and that the relationship between this feature and other person features (e.g, π) must be stated in featural terms. While the second claim is substantive, the first claim is perhaps less controversial in light of bare phrase structure. Both pronouns and lexical DPs are frequently taken to belong to a single syntactic category (D), though see Cardinaletti & Starke (1999) and Déchaine & Wiltschko (2002) for different perspectives. If the syntactic identity of a constituent is constituted entirely by its featural content, then our claim that both pronouns and lexical DPs are specified for δ is equivalent, for all relevant purposes, with both classes "belonging to the same category."

3.1. The internal structure of person

In traditional grammars, the domain of person is typically conceived of as non-hierarchical, with the first, second, and third person categories arranged horizontally. But the cross-linguistic typology of pronominal inventories and the morphological processes making reference to person suggests that there is more structure to person (Ingram 1978, Noyer 1992: 145–175, Harley & Ritter 2002, building on earlier work). One commonly assumed structure, suggested by Benveniste (1956), builds on the discourse roles of speech situation participants. A PA(RTICIPANT) feature characterizes first and second person together, while another feature, SP(EAKER), characterizes just first person.[6] This means that the presence of SP implies the presence of PA, a coocurrence restriction that can be encoded in a feature geometry (Harley & Ritter 2002, Béjar 2003: 47–50).

(21) a. *1*   b. *2*   c. *3*
         PA   PA
         |
         SP

With just these two features, third person would be characterized by the absence of person features altogether, as in (21c), in line with the old idea that third person is simply the elsewhere category (Forscheimer 1953, Benveniste 1956, Zwicky 1977, and many others).

At the same time, many languages draw fine-grained distinctions within the third person, distinguishing between two (or more) sets of third person pronouns. In descriptive terms, languages can have logophoric and non-logophoric pronouns (Hagège 1974, Clements 1975), personal pronouns and d-pronouns (Wiltschko 1998, Sichel 2001, Patel-Grosz & Grosz 2017, Sichel & Wiltschko, to appear), proximate and obviative pronouns (Aissen 1997), or several pronouns expressing animacy distinctions, as we described for Sierra Zapotec in Section 2. These inventories require the existence of some person feature (or features), beyond just those in (21). Béjar (2003) proposes π as a privative feature corresponding to a cluster of properties characterizing the third person. Sichel & Wiltschko (to appear) further argue that once we incorporate π into the logic of markedness, it can be used to distinguish between classes of third person: when it is present, it underlies one class and when it is absent, it underlies another.

Following Béjar and Sichel & Wiltschko, we assume π is a universally available feature, whether or not a language makes use of its markedness potential. At the very least, we propose that, in many languages, π represents the boundary between pronouns and lexical DPs. Thus, instead of (21) above, we propose (22) for these languages, where lexical DPs lack π.

(22) a. *1*   b. *2*   c. *3*   d. *Lexical DP*
          π π π
          | |
          PA   PA
          |
          SP

We take π to have content related to the discourse roles represented in the local person paradigm. If the speaker and hearer together constitute the set of discourse subjects — those individuals participating in the speech situation — π would be associated with a more inclusive category. We will discuss what this might be below, and whether it can be tied directly to the notion of a discourse subject. It is important to emphasize that we understand π and its sub-categories SP and PA as *grammatical* features, whose presence distinguishes pronouns from lexical DPs.

At the same time, we take pronominal and non-pronominals to share a feature that is part of the structure of person, which we call δ, akin to Béjar & Kahnemuyipour's (2017) *d* feature (though see Section 3.3 for discussion of the differences). The presence of π asymmetrically entails the presence of δ, just like the presence of SP asymmetrically entails the presence of PA and π. Thus, all nominals share a person specification, though lexical DPs and pronouns are also distinguished from one another on the basis of person.

(23) a. *Pronoun*   b. *Lexical DP*
          δ δ
          |
          π
          ⋮

Importantly, to account for the intervention patterns, all personal pronouns must have the specification in (23a), whether they are clitic, weak, or strong pronouns. Cardinaletti & Starke (1999) propose that clitics contain a subset of the structure of weak pronouns, which in turn contain a subset of the structure of strong pronouns (see Wiltschko 1998 and Déchaine & Wiltschko 2002 for a related idea). But the featural content of pronouns is orthogonal to how much structure they contain. Whatever functional structure a clitic contains, it must have the person specification above, as must its weak and strong pronoun counterparts.

3.2. Features vs. categories

While we take δ and π to be features, there is another possibility: they could be syntactic categories that project as heads in the syntax.[7] As originally conceived, categories determine distribution through selection, while features are morphological, interacting with movement in a (morpho-)syntactic guise. The theory of bare phrase structure (Chomsky 1995) calls into question whether there is really a discrete boundary between these concepts. The choice between them thus might seem like a purely theory-internal matter, depending on how a particular theory views the relationship between categories and features. For us, the issue is more about the dependency between δ and π than about the ultimate substance of these entities, so that the substantive question is how the structure in the person domain is encoded: Are dependencies between person values represented featurally (e.g., feature geometries or other feature cooccurrence rules)? Or are they encoded in syntactic terms, via the selectional properties of heads, such that π is syntactically dominated by δ? The boundary between pronominals and non-pronominals turns out to be a particularly apt site for querying these basic questions, and we consider both possibilities below in the context of the LDB. In the end, however, we conclude there is nothing to be gained from encoding the structure of person in terms of syntactic hierarchy, and that the organization of person is more profitably stated in entirely featural terms.

If the elements of person were categories, they would be syntactic heads that project functional structure (e.g., Carstens 1991, Ritter 1993, Déchaine & Wiltschko 2002, Oxford 2017). The structure within the person domain would then be represented in terms of the hierarchical ordering of functional heads, as shown in (24): δP dominates πP in pronouns, while πP is not projected at all in lexical DPs. (More fine-grained person categories, such as PA and SP, could presumably also project as functional heads below πP.)
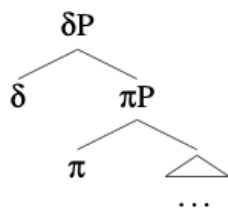
(24) a. *Pronoun*   b. *Lexical DP*

---
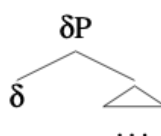
[7] It seems reasonable to assume that π and δ are of the same type because they interact. So either they are both features, or they are both categories.

Both pronominal and non-pronomimals would be δPs (essentially identical to DPs), which would contain other nominal structure, including but not necessarily limited to NP. With all nominals belonging to the same category, selection by external heads (e.g., V or P) could be stated simply in terms of selection for the highest category, which would (correctly) prohibit selection for person and other φ-features. Under this view, the co-occurrence restrictions within the person domain could be cashed out in terms of nominal-internal selection. The presence of δ is entailed by the presence of π because δ optionally selects for π.

Is there anything to be gained by taking the elements of person to be categories? It might be possible to capitalize on the additional internal structure in pronominals to derive the LDB without reference to any intervention-based notion of locality. In particular, if clitic pronouns correspond to the πP, they might be too deeply embedded to be directly extracted, assuming that the containing δP is a phase (for similar ideas, see Oxford 2017 and Preminger 2019). For a clitic pronoun to be subextracted, the phase would first need to be "opened up" by Agree (Rackowski & Richards 2005). The impossibility of object cliticization across a lexical DP might simply then be a product of phase impenetrability. If Agree with a higher lexical DP "uses up" the phase-opening potential of a probe, then the lower argument would remain "closed" to cliticization. Without saying anything more, however, this would mean that all subjects, including pronominal subjects, would prohibit object cliticization. Since this is not the case, an account of the LDB in terms of phase-based locality, which assumes the structures for pronouns and lexical DPs in (24), does not seem to be tenable.

It does not even help to give up on our initial assumption that all nominals belong to the same category, so that pronouns are πPs and only lexical DPs are δPs (with or without an internal πP). Besides the question of why δP and πP should interact this way, it is not clear why a δP should block movement of a lower πP, but a subject πP does not. One could appeal to the fact that when the higher argument is a pronoun it itself moves, thereby getting out of the way of the lower argument's movement trajectory. But then this would be an account based on position, not internal structure, and we already showed in Section 2 that any appeal to position, as in shape conservation approaches, does not cover the full empirical terrain. Instead, what we need is a theory in which both pronouns and lexical DPs are δPs, and at the same time, internal structure and classical phase-based locality do *not* play a role. By this, we do not mean to say that there can be no πP projection within δP, only that it plays no role in the calculation of why lexical DPs intervene for object cliticization.

If treating δ and π as projecting categories does not offer a solution for the LDB, then the alternative is to treat them as features, as we have been doing. On this approach, Agree interacts directly with these features in the familiar way. A functional head probes into its domain, subject to Attract Closest: it first finds the higher argument, and Agrees with it. If this goal has δ, what keeps it from finding the lower argument? Since a phase-based notion of locality is not relevant here, something else needs to be said. We will propose that the probe's value is fixed by the first goal it finds. This means that in order for a probe to find some feature on the lower argument, the higher argument would have to be specified for that same feature. In our case, in order to find a lower goal with π, the higher argument would have to be specified for π. Thus, a lexical DP subject blocks Agree with an object pronoun, but a pronominal subject does not. The details of this mechanism must be spelled out more fully, which we do in Section 4.

Before continuing, note that the two options for implementing δ and π are not necessarily entirely distinct. In bare phrase structure, there are no projecting category labels that are distinct

from the featural content of lexical items. This should include both lexical elements and functional elements, in particular pronouns. When the projecting element is functional, we might take it to be not the actual pronoun, but the kind of entity which encodes grammatical information: in other words, features like δ and π. This means that even if some hierarchical structure was implicated in the explanation of the LDB, the status of δ and π would still be featural, combined with the possibility that these features also project structure. Relatedly, within a nano-syntactic approach, there might be no difference between δ and π being "features" or "syntactic heads," since the elements which project are individual features.

If we accept these possibilities, the issue is not so much how we characterize δ and π, but rather what kind of locality is involved. Is it locality that deals in hierarchical depth, i.e., counts syntactic projections, as in classical phase-based locality, or locality that deals in features and relative distance from a probe, as in Attract Closest? Above, we suggested that classical structure-based locality does not work if pronominal and non-pronominal elements have the same external δP layer and clitics are further embedded within πP, as there is no natural way to draw the line between a non-pronominal DP which does intervene for a clitic, and a pronoun which does not.

However, even if the explanation for the LDB does not make reference to the amount of *internal* structure, i.e., a clitic within πP sub-extracting from δP, it might still make reference to outer structure, i.e., δP. On this scenario, what unites pronominals and non-pronominals is that both are δ-specified and both project δP. What distinguishes them is that only pronouns are π-specified, and the asymmetry of intervention follows from the interaction of π and δ with Agree as features, not structural layers. Again, our point here is not to deny that π projects, only that this projection, if it exists, plays no role in the calculation of locality for intervention by a lexical DP.

3.3. The semantics of person

While we have been discussing person in purely (morpho)syntactic terms, the elements in this domain are usually taken to have semantic content as well. Thus, features do not just characterize form classes, they also come with meanings that combine compositionally with the rest of a linguistic expression (Béjar 2003, Sauerland 2006, Kratzer 2009, among others). This raises another set of questions: Is the structure of person encoded in feature co-occurrence restrictions, perhaps imposed by a feature geometry, as in (22) above? Or can this organization be replaced entirely by structure inherent to the meanings of these features? While feature geometries enjoy widespread usage in the morphology and syntax literature, their theoretical status is not entirely clear. Harley & Ritter (2002) take them to be part of Universal Grammar, with their structure induced by looking at the crosslinguistic typology and acquisition of pronouns. They do not, however, specify a particular connection between geometries and the meanings of the features they contain. Béjar (2003) explicitly proposes that the structural relationships within a geometry derive directly from the semantics of the features involved, with the dominance relation corresponding to asymmetric semantic entailment. Our question about the semantics of person, therefore, is ultimately about the source of (23), and whether it can be grounded semantically.

It is important to keep in mind that, even if there is a semantic basis for how these features are structured, this would not necessarily affect how they interact with Agree. This mechanism has been proposed to interact with the structure encoded in feature geometries, often in ways that cannot obviously be translated into semantic terms. The LDB, for instance, is a fundamentally

syntactic phenomenon, which cannot obviously be boiled down to semantic notions without the mediation of (morpho)syntactic entities like features (or categories). Since we see no obvious way that the denotations of features, per se, could affect the intervention calculation for locality, it seems that π and δ must occur on nominals in the appropriate (morpho)syntactic currency.[8]

To give a plausible semantics for person which can accommodate its extension to π and δ, we start with the semantics of local person. A familiar semantics for SP and PA is given in (25a–b), assuming that these privative features denote one-place predicates that combine intersectively (e.g., Heim & Kratzer 1998, Sauerland 2006). Both make reference to discourse roles: either the speaker of the speech situation or another participant in that situation.

(25)   a. $[[\text{SP}]]^c = \lambda x : x$ is SPEAKER(c) . x
        b. $[[\text{PA}]]^c = \lambda x : x$ is a participant in the conversation of SPEAKER(c) . x

Sichel & Wiltschko (to appear) propose that the semantics of π can also be understood entirely in terms of discourse roles. They argue that π describes potential discourse participants: all those individuals who are possible discourse subjects, or in other words, potential interlocutors who can talk as well as be talked to.

(26) $[[\pi]]^c = \lambda x : x$ is a potential participant in a conversation . x

In Sichel & Wiltschko's system, all person features refer to the actual or possible ability of an individual to participate in a conversation (whether spoken or signed). Given this semantics, SP asymmetrically entails PA (their LOCAL), which in turn asymmetrically entails π (their PERSON), an arrangement that they call the Person Sphere, depicted in Figure 1. The Person Sphere stretches up to include discourse objects, or Discourse Referents (DR$_N$), the class of individuals that can be talked about. This inclusion relation between the two outer layers places the relationship between DR$_N$ and π on a continuum with the containment relationship between local and speaker, constituting a sphere that extends all the way from SP to non-pronominal DPs.
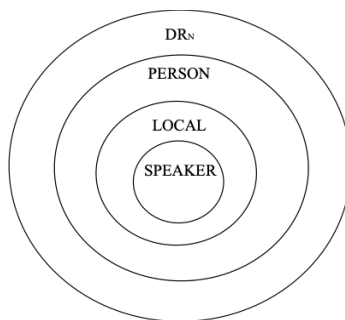
Figure 1: Person Sphere (Sichel & Wiltschko, to appear, p. 20)

---

[8] Harbour (2008:51–115, 2011, 2016:187–216) argues that grounding feature geometries in the semantics of the features involved makes the geometries themselves redundant. He proposes to eliminate the need for them by abandoning the assumption that features are first-order predicates that combine intersectively. This allows the traditional three-way person distinction to be captured by a simple two-feature system in which features are not subject to any co-occurrence restrictions. His approach can, in principle, be extended to deal with the finer-grained distinctions within person that we are concerned with, though we will not attempt to do so here.

Returning to the LDB in Sierra Zapotec, we are committed to all personal pronouns bearing π, since all pronominal clitics in object position are affected by a lexical DP subject in the same way. No object pronoun, regardless of its animacy, ever cliticizes when the subject is a lexical DP, even though their referents cannot plausibly be characterized as potential discourse subjects.

(27)   a.   *DP > 3.EL*
         *Blenh^4=e'$^2_1$ Xwanh^{24}=a'^{4} t$_1$.
         COMP.carry=3.EL   Juana=DEF
         Intended: 'Juana carried her/him.'   (FSR, SLZ1054, 10:00)

       b.   *DP > 3.HU*
         *Blenh^4=ba'$^3_1$ Xwanh^{24}=a'^{4} t$_1$.
         COMP.carry=3.HU   Juana=DEF
         Intended: 'Juana carried her/him.'   (FSR, SLZ1054, 10:45)

       c.   *DP > 3.AN*
         *Blenh^4=b$^4_1$ Xwanh^{24}=a'^{4} t$_1$.
         COMP.carry=3.AN   Juana=DEF
         Intended: 'Juana carried it.'   (FSR, SLZ1054, 12:00)

       d.   *DP > 3.IN*
         *Blenh=enh$_1$ Xwanh^{24}=a'^{4} t$_1$.
         COMP.carry=3.IN   Juana=DEF
         Intended: 'Juana carried it.'   (FSR, SLZ1054, 13:15)

How can we incorporate the behavior of animal and inanimate pronouns into the characterization of π as a potential discourse subject? One possibility would be to expand the denotation of π to include both discourse subjects (who talk) and discourse objects (who are talked about). While this would admit non-human referents, it would undercut the explanation for Sichel & Wiltschko's analysis of a pejorative effect that arises when pronominal demonstratives are used to refer to humans.

Another possibility is that both semantics are, in principle, available and languages choose which they encode. Béjar (2003) suggests that there may be a finely-articulated hierarchy of such semantic categories, and languages choose what the semantic range for π is. On this interpretation, the Person Sphere could be understood extensionally to define a range of semantic categories and the relations between them, to which linguistic entities such as SP, PA, π, or δ would be mapped. While the semantic categories, along with the relations of asymmetrical entailment encoded in the Person Sphere would be universal, languages could vary in terms of where they place their π: immediately above discourse subjects, in which case π would be restricted to potential interlocutors, i.e., humans, or slightly higher, in which case it may also include non-potential interlocutors, as in the Zapotec animate and inanimate pronominial

categories.[9] Regardless, it is clear that a substantive discourse-role semantics for π is possible that creates the entailment relations depicted in Figure 1.

The semantics of δ is less clear at first glance. We are proposing that all nominals bear this feature, so its meaning must be general enough to encompass both pronouns and lexical DPs. Sichel & Wiltschko extend their Person Sphere to include an outer layer (DR_N), which is associated with the broadest class of referential nominals, subsuming both pronouns and referential lexical DPs. Semantically, they identify this as the class of discourse subjects and discourse objects, or all those individuals that can be referred to by a nominal. But in Sierra Zapotec, it is not just referential expressions (definite description, possessive descriptions, proper names) that prohibit object cliticization, but all lexical DPs, including quantificational ones.

(28) a. Ts-ja^2-se'e^2-naw^{23} to^{4}to^4 bi^{2}dao'^{1} **leb^{24}**.
      CONT-AND-PL-follow   every   child   **3.AN**
      'Every boy chased the dog.'   (RM, GZYZ083, 1:40)

    b. *Ts-ja^2-se'e^2-naw^2=**b^4**$_1$ to^{4}to^4 bi^{2}dao'^{1} t$_1$.
       CONT-AND-PL-follow=**3.AN**   every   child
                                     (RM, GZYZ083, 1:45)

So the δ feature cannot be tied to reference too directly, as for instance Harley & Ritter's (2002) REFERRING EXPRESSION or Béjar & Kahnemuyipour's (2017) *d* features are. A semantic property is required that characterizes both referential and quantificational nominals, and importantly does not hold of other linguistics expressions (e.g., adjectives or verbs). The extensional interpretation of the Person Sphere can incorporate the behavior of quantificational DPs consistent with a modified semantics for δ, which we now turn to.

Baker (2003:94–189) identifies individuation as the property that is unique to nominals (following Geach 1962, Gupta 1980, and Larson & Segal 1983). Whereas all lexical categories have "criteria of application" that make it possible to decide whether an individual is a member of that category, nominals have an additional "criterion of identity." This makes it possible to individuate entities and to determine whether particular individuals are the same sort of thing or not. This individuation is a precondition for reference tracking as well as for quantification, explaining why only nominal elements can directly form referential or quantificational expressions. (For adjectives or verbs to serve these functions, additional morphology is usually required.)

We adopt the characterization of nominals as linguistic individuators, and assume that the use of a nominal implies individuation. However we see no compelling reason to associate this

---

[9] One way in which π might be distinguished from δ in a language like Sierra Zapotec involves the concrete-abstract distinction. An inanimate pronoun does not, for instance, seem to be able to refer to propositions:

(i) Q: E^1 nhezd^4=u'^{4} wxe^2 ye^{2}lha'a^2 bi^{14}zanh^4=a'^{4}?
       Q   STAT.know=2SG tomorrow   POT.arrive   sister=1SG
       'Do you know that my sister is arriving tomorrow?'

   A: Nhezd^{14}=a'(*=nh)^4.
      STAT.know=1SG=3.IN
      'I know (it).'   (RM, GZYZ096, 2:00)

This follows if π only describes concrete entities (see Sichel and Wiltschko, to appear for discussion).

property with the category N, as Baker does, since it is also associated with pronouns, which contain no N, and with nominalizations, which similarly do not contain an N, under the hypothesis that they are mixed nominal-verbal projections (Borsley & Kornfilt 2000 and others). We associate the individuation property with δ instead, which by hypothesis is found inside all nominals (see also Ghomeshi & Massam 2020 for a different argument that person is responsible for individuation). A linguistically individuated entity corresponds to the notion of discourse referent (entities that we can talk about). This is the outer layer of the Person Sphere reserved for discourse objects. If we understand the outer layer to refer to types of entities, rather than to types of expressions, the characterization of nominals in terms of individuation can also cover the intervention potential of quantificational nominals.

To be clear, in advancing a semantic characterization for π and δ, we are not suggesting that this is opposed to the formal representation for person that we proposed above. To the extent that the LDB is a syntactic phenomenon and to the extent that the denotations of features themselves do not factor into the locality calculation for Agree, these features must appear on nominals in the appropriate morphosyntactic currency. With these featural representations in place, we can now turn to the properties of the Agree mechanism that derives the LDB as well as the GCC.

4. Overprobing and its source

In the preceding section, we explored the idea that lexical DPs and pronouns share featural content, a person feature (δ) that characterizes all nominals. This creates a featural asymmetry across nominal classes: pronouns are more specified than lexical DPs, since they bear other person features. This specficiational asymmetry enables a uniform understanding of the locality effect involving lexical DPs, which we state explicitly again in (29), and the GCC, which is repeated in (30) below.

(29)   *Lexical DP Blocking (LDB)*
        An object pronoun cannot cliticize when the subject is a lexical DP.

(30)   *Gender–Case Constraint (GCC)*
        An object clitic pronoun cannot exceed a subject clitic pronoun on the animacy hierarchy, i.e., EL > HU > AN > IN.

In both patterns, cliticization of an object pronoun is prohibited when it is more featurally specified than the higher argument.

(31)   *
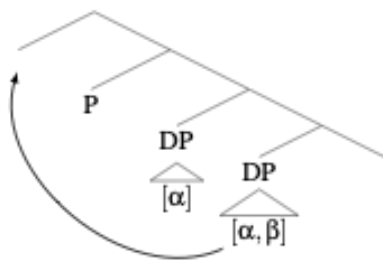
In the LDB, object cliticzation is prohibited when the subject is a lexical DP, which by hypothesis only has a subset of the features of a pronoun (just δ). And for the GCC, the animacy hierarchy is transparently related to the feature structures in (22) above: pronouns higher on the hierarchy have more features than those lower on the hierarchy.

This correspondence can be translated into a grammatical mechanism by building on the idea that cliticization is, in some way, the reflex of agreement (Borer 1984, Suñer 1988, Sportiche 1993, Anagnostopoulou 2003: 249–320, among others). If agreement is produced by an Agree operation, then Agree would be a precondition for pronominal cliticization and a failure to Agree would result in a failure to cliticize. Building on these assumptions, Anagnostopoulou (2003, 2005) proposes that constraints like the GCC can be understood in terms of how a single functional head interacts with multiple arguments in its domain via Agree. She hypothesizes that whether a probe can Agree with and move an object pronoun depends on the Agree relation it enters into with the higher argument. There is more than one possible implementation of Anagnostopoulou's basic idea, though it is stated in neutral terms in (32): the availability of what we call *overprobing*, a probe's entering into an Agree relation with a lower goal, depends on the features it finds when it Agrees with a higher goal.

(32)   *Overprobing*
        A probe P can Agree with a goal G when there is a closer goal G′ only if P Agrees with G′ in some feature(s) F.

What these features are depends on what grammatical mechanism gives rise to overprobing, as well as auxiliary assumptions such as the position of the probe relative to the goals. Whatever its ultimate grammatical source, though, it is clear how overprobing might underlie both the LDB and the GCC. If Agree is a necessary precondition for pronominal cliticization, then the possibility of object clicization depends on the possibility of Agree with the object, which in turn depends on the Agree relation between the probe and the higher argument.

Since Anagnostopoulou's work, the literature has witnessed a growing number of proposals for the grammatical source of overprobing (e.g., Nevins 2007, 2011, Deal 2020, Coon & Keine to appear, Foley & Toosarvandani, to appear), motivated by its superficial incompatibility with the classical Minimalist theory of agreement and movement. In Chomsky's (2000, 2001) original proposal, Agree only takes place when a probe bears some uninterpretable features, which must be matched against interpretable feature(s) on a goal. Once the closest suitable goal with matching features has been found, subject to Attract Closest, the uninterpretable feature(s) on the probe are deleted. If the matching goal is eligible for movement, it is Merged in a local configuration to the probe. Once its uninterpretable feature(s) are deleted, the probe cannot interact with any other goals. This permits only a single Agree relation between the probe and the closest goal and the possibility of movement for just this one goal. So, to enable overprobing, the classical theory must either be abandoned or revised.

Anagnastopoulou's (2005) strategy is to throw it, more or less, out. She instead adopts Multiple Agree (Hiraiwa 2001), which is not subject to Attract Closest, specifically in order to derive the Weak PCC. This constraint prohibits object cliticization if structurally higher argument is higher on a certain person hierarchy (see Bonet 1991:180 for Catalan, Perlmutter 1971:62–63, Pancheva & Zubizarreta 2017:19 for Spanish, Bianchi 2006:2028 for Italian, and Nevins 2011:963 for Kashmiri).

(33)   *Weak PCC*
       An object clitic pronoun cannot exceed a structurally higher clitic pronoun on a
       hierarchy of person, i.e., 1/2 > 3.

As Anagnostopoulou points out, the Weak PCC requires overprobing because it is inherently
relational. (An object pronoun is only prohibited from cliticizing if its person exceeds that of a
higher pronoun on the person hierarchy.) Multiple Agree enables this overprobing because it is
not subject to locality, allowing simultaneous interaction with all the goals in a probe's domain.
At the same time, Multiple Agree must be constrained: overprobing is only allowed, according to
(32), when a higher goal has certain features. Building on Anagnastopoulou's work, Nevins
(2007, 2011) proposes a condition called *Contiguous Agree*, which directly blocks Multiple
Agree when a lower goal is less featurally specified than a higher goal.

(34)   *Contiguous Agree (cf. Nevins 2007:291)*
       For a probe P relativized to a feature F with a goal G that bears F, there can be no G′
       such that: (i) P c-commands G′ and G′ c-commands G, and (ii) G′ does not bear F.

A constraint like Contiguous Agree is clearly sufficient to rule out ungrammatical configurations
like (31). But as Coon & Keine (to appear) and Foley & Toosarvandani (to appear) argue, a
constraint like this simply encodes an empirical generalization in a grammatical principle.

More recently, there have been attempts to revise the classical theory to enable overprobing
in a more explanatory way. In principle, this could happen in one of two ways. First, Agree could
be redefined to force a single probe to Agree with every goal in its domain, albeit still subject to
locality. Deal (2015, 2020) offers one proposal along these lines, systematically reconfiguring the
Agree operation to allow for "insatiable probes" that interact in a cyclic fashion with all
accessible goals in their domain. While this permits overprobing, it must be constrained in some
way so that the probe does not Agree with just any goal. Deal proposes this happens through
dynamic update of a probe's specifications in the course of the derivation. However, as we will
show below, while this can derive the GCC, it cannot derive the LDB.

Second, the Agree operation could be maintained essentially in its classical form, but the
structure of probes could be elaborated to enable a single probe to interact with more than one
goal. Coon & Keine (to appear) advance a theory of this kind, adding two kinds of structure to
probes: (i) a probe's requirements are *sequenced* into "subprobes" that can be satisfied through
separate Agree relations (Taraldsen 1995, Anagnostopoulou 2003, Béjar & Rezac 2003, among
others) and (ii) each subprobe is *articulated*, so it can be checked by more than one goal (Béjar
2003, Béjar & Rezac 2009). When a probe Agrees with more than one goal on a single subprobe,
this leads to what Coon & Keine call "feature gluttony." While this configuration is not
inherently problematic, they take it to lead to a problem for pronominal clitization. Two
pronouns cannot both move to the probe if they Agree on the same subprobe. But again, we will
argue that this only derives the GCC. Feature gluttony cannot be responsible for the LDB, which
involves lexical DPs that do not cliticize.

The LDB, it turns out, provides a crucial wedge into the grammatical mechanism that
underlies overprobing. As we will argue, assimilating the LDB to the GCC provides a new
perspective on what a restrictive theory of Agree — one that can allow for overprobing —

should look like. In particular, we suggest a *probe activation* model for Agree that incorporates two hypotheses in order to derive the LDB/GCC: (i) the probe Agrees with the highest goal first because it c-commands all goals and is subject to intervention-based locality (i.e., Attract Closest), and (ii) the probe is "activated" through its Agree relation with the highest goal, so that it can interact with subsequent matching goals, even if it need not do so. The first hypothesis distinguishes our account from Deal's insatiable probes, while the second hypothesis distinguishes it from Coon & Keine's feature gluttony. Thus, on the conception of Agree that we advance here, a probe must enter into an Agree relation only with the highest goal. Subsequent relations with lower goals, while not entirely free, are not absolutely required.

4.1. The probe activation model

The model of Agree that we propose to account for the LDB/GCC in Sierra Zapotec, as well as for parallel patterns like the Weak PCC, does not have a built-in requirement for movement or cliticization. We only assume that Agree is a precondition for movement. Therefore, if something like a clitic needs to move, it will need to Agree. Nor do we assume any particular sequencing, beyond Attract Closest and the timing imposed by hierarchical structure. Since the probe in T is higher than both the subject and object, a subject is targeted by Agree before an object, simply because it is closer. Finally, we do not assume that anything special needs to be invoked to allow or to force multiple Agree relations.

    The probe activation model incorporates two hypotheses, which are stated more precisely as follows:

(35)   *Probe activation*
        (i) **Locality:**
            The probe P c-commands all goals and Agrees subject to intervention-based locality (i.e., Attract Closest). Thus, it must Agree with the highest goal G in its domain.
        (ii) **Activation:**
            After Agreeing with G, the probe is "activated" and is able to Agree with another goal G', though it is not required to, just in case G' is not more featurally specified than the probe (i.e., $G' \subseteq P$).

We begin with the second hypothesis. Our central claim is that, from the perspective of the Agree mechanism, only the first round of Agree is obligatory. Anything beyond that is strictly optional, and depends solely on the needs of the goal(s). For a probe that has already Agreed with the subject, if an object also needs to Agree, this may be granted subject to two further constraints. The first one is trivial, while the second one captures the specificational asymmetry inherent to the LDB/GCC. The central situation in which an object will need to Agree is if it is a clitic; as a clitic, it is required to move, presumably becuase of a prosodic requirement. To meet this requirement, it must satisfy a syntactic requirement, i.e., Agree.[10] As a result, there is no "feature

---

[10] We are building here on Foley & Toosarvandani's (to appear) account of the GCC. They locate the source of both grammatical and ungrammatical combinations in the mechanism underlying cliticization. They propose that a pronoun can cliticize just in case the condition in (i) is satisfied.

gluttony" in our model. Quite the opposite, in fact: beyond the obligatory first round with the subject, Agree may further extend itself to a clitic in need if it can, subject to clause (ii) in (35).

Nor under this model is the probe "insatiable." The only obligatory Agree relation is with the closest goal. All subsequent relations, when they are possible, are optional: the syntax does not mandate one way or another whether they take place. At the same time, this does not mean that we expect surface optionality in whatever morphological realization these lower Agree relations might have. While the syntax is indifferent to whether these subsequent Agree operations take place, the morphology and phonology may not be. Post-syntactic considerations, including but not limited to expressiveness and economy, may in fact favor derivations in which Agree takes place. So while the operation itself, when applied to lower goals, is not obligatory, a unique surface output is still possible.[11]

So, for the highest goal, the obligatoriness of Agree is attributed entirely to the syntax. But for lower goals, any "obligatoriness" that we see is due to non-syntactic factors.[12] Given this somewhat more opaque mapping from syntax to surface form, what empirical support is there that Agree with lower goals is optional? In Sierra Zapotec, the evidence is indirect, but clear.

(i) *Condition on Pronominal Cliticization*
    For a functional head H that has been valued (i.e., VALUE(H) ≠ Ø), a clitic pronoun P can (internal) Merge with H iff, for the set of relevant features F on P, F ⊆ VALUE(H).

While the motivation for the additional cliticization is much the same, as are the featural preconditions that must be met for this additional stage of cliticization to take place, we attribute the mechanism explicitly to Agree.

[11] Evidence that there is indeed an Agree relation with the lower goal may come from complementizer agreement in Nez Perce. Deal (2015: 4) reports that the complementizer can agree with either the subject, the object, or both, depending on their person and syntactic position. When there is only one local person argument, the complementizer agrees with it. But when the subject is first person, the complementizer agrees both with it and the second person object (i). In the inverse configuration, it only agrees with the second person subject (ii).

(i) *1 > 2*
    ke-**m**-**ex**   kaa   pro   cewcew-téetu   pro
    C-**2**-**1**   then   1SG   telephone-TAM   2SG
    'when I call you'
(ii) *2 > 1*
    ke-**m**   kaa   pro   cewcew-téetu   pro
    C-**2**   then   2SG   telephone-TAM   1SG
    'when you call me'   (Deal 2015: 6)

In (i), when the subject is more featurally specified than the object, the probe Agrees with both goals, with the morphology transparently registering both relations. The activation model accounts for the absence of both markers in (ii): the higher argument is less specified than the lower argument, and hence Agree with the object is impossible. There is the question of why Agree with the object in (i) is obligatory. (If it indeed is: Deal provides no evidence one way or another.) We speculate that this has to do with why both Agree relations can be exponed in the first place. We might imagine that there are morphological constraints that preference exponing as much information as possible (Foley 2020: 137–197).

[12] There is certainly an obligatory component here, but it is the obligatoriness of using a clitic whenever possible, and limiting the independent pronoun to cases in which a clitic would not work, such as those given in (35ii). This we attribute to an economy principle which favors local licensing, i.e., internal Merge over "case licensing," whenever possible. Or to any other economy principle that would favor a clitic and could replace Cardinaletti & Starke's (1999) Minimize Structure, which is no longer an option if all pronouns are specified for δ, and hence are possibly δPs.

Since there is no palpable morphological effect produced by Agree with the object, nothing would be morphologically amiss if it did not occur. But the optionality of object Agree can be detected in two related ways: first, by comparing a given object pronoun, say the third person human pronoun, in two different alignments, and second, in the nature of the strategy for overcoming violations of the LDB/GCC.

For the first type of evidence, we consider in more detail the conditions under which an additional round of Agree can occur, i.e., (35ii). And when it can, is the Agree operation obliged to enter into a relation with the lower argument or not?

(36) a. *3.EL > 3.HU*
         Blenh^4=e'$^2_1$=**ba'$^3_2$**   t$_1$ t$_2$.
         bite.COMP=3.EL=**3.HU**
         'S/he carried her/him.'   (FSR, SLZ1012, 15:16)
      b. *3.AN > 3.HU*
         *Udi'in^4=eb$^4_1$=**ba'$^3_2$**   t$_1$ t$_2$
          bite.COMP=3.AN=**3.HU**
         Intended: 'It bit her/him.'   (FSR, SLZ1012, 19:45)

In the configuration in (36a), the object is less specified than the subject. Agree with the object is sanctioned, and cliticization can occur. In (36b), by contrast, Agree is impossible, ruled out by (35ii) because the object is more specified than the subject. Since the object cannot Agree, it cannot cliticize, and it must be realized as an independent pronoun, as in (37b) below. Now, is Agree obligatory when it can occur? The answer is no. Comparing the object pronoun in (36a) to its counterpart in (36b) shows that an object human pronoun can Agree, but it can also fail to Agree. The ungrammaticality of (36b) is to be attributed to the non-moved status of the clitic, and not to the Agree mechanism. This is not affected by the absence of Agree with the object. This can be seen when the clitic is replaced by an independent pronoun. A pronominal object not conforming to (35ii) is realized as an independent pronoun, both when the GCC is violated (37a) and the LDB is violated (38a). By hypothesis, Agree does not apply in the grammatical realizations with an independent pronoun, in (37b) and (38b).

(37) *3.AN > 3.HU*
      a. *Udi'in^4=eb$^4_1$=**ba'$^3_2$**   t$_1$ t$_2$
         bite.COMP=3.AN=**3.HU**
         Intended: 'It bit her/him.'   (FSR, SLZ1012, 19:45)
      b. Udi'in^4=eb$^4_1$ t$_1$ **le^{2}ba'^{3}**.
         bite.COMP=3.AN   **3.HU**
         'It bit her/him.'   (FSR, SLZ1051-s, 5)
(38) *DP > pro*
      a.*Blenh^4=**eb$^4_1$**   Xwanh^{24}=a'^{4} t$_1$.
         carry.COMP=**3.AN**   Juana=DEF
         Intended: 'Juana carried it (an animal).'   (FSR, SLZ1051, 7:30)
      b. Blenh^4 Xwa^{24}nh=a'^{4} **leb^{24}**.
         carry.COMP   Juana=DEF   **3.AN**
         'Juana carried it.'   (FSR, SLZ1051-s, 6)

Why does an object independent pronoun, with the very same features as its clitic counterpart, not run afoul of the activation constraint in (35ii), in both GCC and LDB contexts? One thing is clear: since it is not a clitic, it need not move. We can attribute this directly to the Agree mechanism, because movement requires Agree. Since independent pronouns need not move, they also need not Agree, and activation is not invoked. On this account, when a pronoun undergoes cliticization, this is not triggered by the probe, but by the clitic, which needs to move. Before we say more about how Agree derives both compliance with and violations of the LDB/GCC, we provide independent evidence for the role of Agree in distinguishing between clitics, which need to Agree and independent pronouns, which do not. This, in turn, supports our view that Agree beyond the first round is optional.

The absence of Agree with an object independent pronoun is somewhat more transparent when it is local person. This is because, typically, local independent pronouns must be doubled by a clitic in Sierra Zapotec (Sichel & Toosarvandani 2020), similar to French (Kayne 2000). This is true across grammatical positions, including for subjects (39a), possessors (39b), and prepositional objects (39c).[13]

(39) a. (*Nhu^{42}de'^{4} ye^{1}ga'an^4?* 'Who is going to stay?')
          Ye^{1}ga'an^4***(=u'^{4})**   lhe'^{2}.
          stay.POT**=2SG**   2SG
          'YOU are going to stay.'   (FSR, SLZ1061, 5:45)
      b. (*Nhu^{42} xna'^{4} benh^{4}bia^4 Xwa^{2}nh=a'^{4}?* 'Whose mother did Juana meet?')
          Benh^{4}bia^4=ba'^{3} xna'^{4}***(=o'^{4})**   lhe'^{2}.
          meet.COMP=3.HU   mother**=2SG**   2SG
          'She met YOUR mother.'   (RM, GZYZ101, 6:50)
      c. (*Nhu^{42} tse^4 dzekd^4 Xwa^{2}nh=a'^{4}?* 'Who does Juana love?')
          Dzekd^4=ba'^{3} {tse^4, *tsi^2**=a'^{4}**}   nada'^{24}.
          love.CONT=3.HU   {of,   of**=1SG**}   1SG
          'She loves ME.'   (RM, GZYZ101, 22:17)

Similarly, when the LDB is violated, only an independent pronoun appears (40a), with doubling being ungrammatical (40b).[14]

---

[13] A local person pronoun cannot cliticize in object position due to a constraint independent of the LDB/GCC. The Strong PCC excludes all local person clitics in object position, regardless of what the subject is. This does not follow directly from probe activation, as 1 > 2 is ruled out just as 2 > 1 is. There are two ways of deriving the Strong PCC that are compatible with our account. First, local person pronouns could be subject to an independent licensing constraint like the Person Licensing Condition (Béjar & Rezac 2003). Second, the probe could be permitted to copy not just features, but subtrees in a feature geometry. As in Foley & Toosarvandani (to appear), this can derive the Strong PCC, alongside the GCC (and hence also the LDB).

[14] In this respect, Sierra Zapotec differs from Icelandic, where it is not possible to repair a PCC violation by suspending agreement with the offending local person pronoun. Agreement with the lower nominative local pronoun is just as ungrammatical as a default third person form (Schütze 1997:117, 122).

(i) a. Þú$_1$ varst   gefinn   t$_1$ honum.
          you   were.2SG   given   him.DAT
          'You were given to him.'
      b. *Honum$_1$ var/varst   gefinn   t$_1$ Þú.

(40) a. Dza⁴la⁴lle'⁴ *Xwa²nh=a'⁴* **lhe'²**.
forget.CONT Juana=DEF **2SG**
'Juana forgot you.' (FSR, SLZ1061, 11:45)

b. *Dza⁴la⁴ll=**u'⁴**₂ *Xwa²nh=a'⁴* **lhe'²**.
forget.CONT=**2SG** Juana=DEF **2SG**
'Juana forgot you.' (FSR, SLZ1061, 12:50)

Despite there never being doubling of local person pronouns in object position for comparison, we propose that the absence of an accompanying clitic in LDB/GCC contexts is a sign that the probe is not activated. Clitic doubling is required in subject position, because here Agree is obligatory, and this seems to be the case for Agree in the syntactic domains associated with the possessor in (39b) and with the prepositional object in (39c) as well. Considering the absence of clitic doubling in the LDB/GCC-cases in this broader context suggests that Agree is not invoked. In other words, the reason for the grammaticality of an independent pronoun in all these cases is that, unlike a clitic, it does not rely on Agree. If so, then we have shown that when the probe is present, the second round of Agree is optional.

Independent evidence for this conclusion is provided by other environments in which clitic doubling of a local person pronoun is not mandated. Below we demonstrate an alternation, in subject position, where doubling is usually required, between an obligatorily clitic doubled pronoun and amn independent pronoun which cannot be doubled. And we conclude, following Sichel (2001, 2002) and Sichel & Toosarvandani (2020), that the requirement for doubling of an independent pronoun, as in (39a), is imposed by the probe (not by any properties of the independent pronoun itself). By extension, absence of doubling is due to absence, or inactivity, of the relevant probe. In other words: no doubling, no Agree. We examine two contexts where the requirement for doubling is suspended: non-verbal predications and fragment answers.[15]

In non-verbal predications, cliticization is impossible, and only an independent pronoun appears, whether for local person (41a) or third person (41b). These are the same forms that emerge when the LDB/GCC is violated.

(41) a. (*Bi⁴² llin⁴ dzonh²³u'³?* 'What do you do?')
Be²ne'² skwel²⁴(*=a'⁴) **ne⁴da'⁴**.
person school=1SG **1SG**
'I am a teacher.' (FSR, SLZ1061, 14:45)

b. (*Bi⁴² llin⁴ dzonh²³ Xwanh²a'⁴?* 'What does Juana do?')
Be²ne'² skwel²⁴(*=ba'³) **le²ba'³**.
person school=3.HU **3.HU**
'She is a teacher.' (FSR, SLZ1061, 15:45)

---

him.DAT was.3SG/2SG given you
Intended: 'He was given to you.'

[15] Another context in which doubling of DP is suspended is coordination, but this is due to general constraints on movement, rather than to any requirements or restrictions that apply specifically to clitics. See Sichel & Toosarvandani (2020) for discussion.

We take the probe to be either absent or inactive in this context, because it is associated with the inflectional domain and this higher functional structure is missing or defective in non-verbal predication. Agree thus cannot apply, and because cliticization depends on Agree, a clitic is impossible. But since the probe is missing, the absence of Agree does not disqualify the derivation. Thus, this is another context in which Agree does not occur, though for a different reason than in contexts where the LDB/GCC is violated. Here it is the subject that does not Agree and is repaired by an independent pronoun, simply because of properties of the external syntactic environment in non-verbal predications.

A related context is which doubling of a subject independent pronoun can be grammatically suspended is in fragment answers.

(42)   *(Nhu^{42} ye^{2}yej^2? 'Who is going to go?')*
        Lhe'^{2}.
        2SG
        'You.'   (RM, GZYZ052, 1:02:58)

Here, too, it may be that non-doubling is tolerated because, as in non-verbal predications, the probe is missing. Fragment answers plausibly involve no more structure than what meets the eye (Stainton 1998, Ginzburg & Sag 2000, Jacobson 2013).[16] The fact that the same repair — an independent pronoun — surfaces in fragments, non-verbal predications and in the context of LDB/GCC violations supports the claim that the independent pronoun is good because Agree is not invoked. And this, in turn, supports our claim that the violation incurred by the clitic in LDB/GCC contexts is due to a failure of Agree. It is difficult to see how surface-oriented approaches, like the shape conservation approaches we surveyed in Section 1, could account for these shared patterns of cliticization and doubling across construction types.

Moving on now, if we are correct, it is possible, but not necessary for the probe to Agree with the lower argument. This characterization raises two questions: (i) Why is Agree with an object possible at all? (ii) Why is it possible only within the window specified in (35ii)? Starting with the first question, we can keep the answer short and simple: this is just how it is. Agree is, in general, unconstrained and it can iterate *if it has to*, where the necessity has an external source, such as with object clitics in Sierra Zapotec. Often the effects of iterated Agree will go unnoticed, when the morphological signature is inscribed exclusively by the subject. Here also lies the source of the specification asymmetry in (35ii): Agree with the object in Sierra Zapotec is a free rider, and it has to remain free: it cannot affect the signature left by the subject. This is why the object must be less specified. If it were more specified, something would need to be added, and that is impossible. As long as Agree with the object can fit into the footprint created by the subject, it is permitted for the object to initiate a second round of Agree.

Why are subsequent rounds of Agree unable to change what happened before? It might be tempting to think of this in purely morphological terms: there is a morphological signature to Agree with the subject, which cannot be altered. This, however, would lead us to expect that the probe can only Agree with both subject and object if they are featurally *identical*. But this is not

---

[16] Alternatively, the entire structure is present, Agree is triggered by the probe, and everything except the fragment is subsequently deleted at PF (Giannakidou 2000, Merchant 2004, Weir 2014), in which case, absence of doubling in fragment answers is not directly relevant for the argument that absence of doubling signals absence of Agree; there could be more than one reason for the absence of overt doubling.

the case: the object may be less specified than the subject. The specification asymmetry gives us a closer look into the physiology of Agree and its inner workings. Beyond its morphological outcome, Agree must involve abstract features of the sort we have been referring to all along: privative and structured (in some way). The object must make do with whichever features the subject determined; Agree with the object cannot *add* anything. We might imagine these features mechanically, as levers that get pulled: if Agree with the subject pulls the HU lever and the AN lever, Agree with the object cannot pull the EL lever. We can then think of overprobing as resulting from "activation": the subject activates the levers that it needs, and this then primes the probe for subsequent rounds. Alternatively, we could think of this in terms of "deactivation:" whatever levers were not activated by Agree with the subject atrophy and become deactivated. Either way, the result is the same. In the cases that we are looking at, Agree with the subject is primary, as dictated by Attract Closest; it is obligatory, and it sets the parameters for future rounds of Agree. Beyond that, the object can get a free ride just as long as it can fit into the footprint left by the subject.

We can now compare the probe activation model to the two alternatives for enabling overprobing that we introduced above. As we will see, neither insatiable probes nor feature gluttony allows for a unified understanding of the LDB/GCC as arising from the mechanism of Agree.

4.2. Insatiable probes

Deal (2015, 2020) systematically overhauls the Agree operation to allow for overprobing, while still taking the operation to apply cyclically, subject to Attract Closest. Probes are specified separately for their *interaction* conditions (what they can match and be valued by) and their *satisfaction* conditions (what stops them from interacting further). For the Weak PCC, the probe is taken to be "insatiable": it has no satisfaction condition and so it is capable of interacting with any goal in its domain. The restrictions on cliticization follow from the possible Agree relations, which are dictated by the probe's interaction conditions and how these are dynamically updated in the course of the derivation.

The Weak PCC, as a reminder, prohibits a direct object from cliticizing in a ditransitive when it exceeds the indirect object on the person hierarchy: 1, 2 > 3; see (33) above. For a grammatical combination like (43), where the indirect object is second person and the direct object is first person, the probe can Agree with both arguments. Since Agree is a necessary and sufficient condition for cliticization, both pronouns can subsequently move.

(43)   *2 > 1*

29
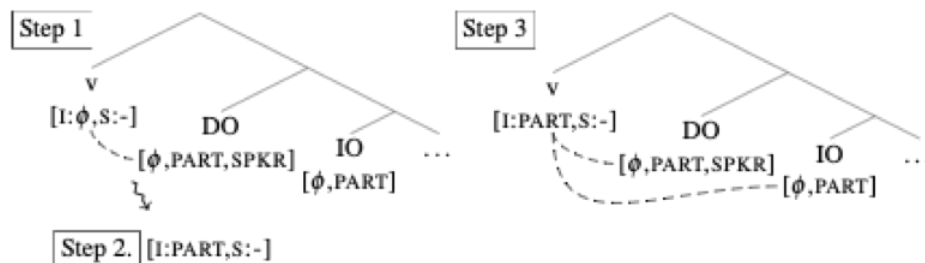
Deal's innovation is what happens after the probe Agrees with the first argument it finds. In ditransitives, it is assumed that this is the direct object, since the direct object is structurally higher than the indirect object.[17] To start, the probe's interaction condition (abbreviated as I) is specified as φ, so it can Agree with an argument of any person. But this probe is also specified with a dynamic interaction feature [PART]$^\uparrow$: once it finds a goal with [PART], its interaction condition is updated to [PART]. Since the probe has no satisfaction condition, it can continue to probe, though from this point on it can only interact with arguments which are specified for this feature, i.e. local person pronouns. In (43), it Agrees with both arguments, as desired.

In the ungrammatical combination in (44), where the indirect object is third person and the direct object is first person, the probe is not able to Agree with the third person indirect object, since this argument lacks the required [PART] feature.

(44)   *3 > 1
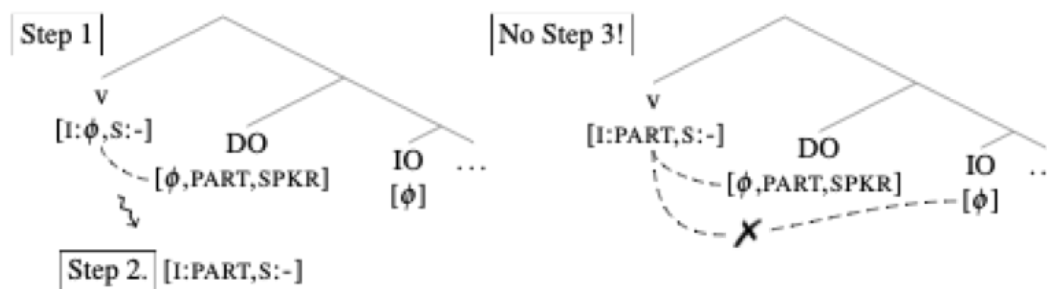
This account, like the one introduced in Section 4.1, derives a pattern of cliticization like the Weak PCC from restrictions on Agree. It is also similarly based on a specificational asymmetry, though in the opposite direction: whereas under the probe activation model, the generalization across LDB/GCC is that the higher argument cannot be less specified than the lower one, the generalization underlying Deal's account is that the higher argument (really, the argument that is Agreed with first) cannot be *more* specified than the lower one. Otherwise, the probe's interaction feature would be set to a feature that the lower argument lacks. As we will see below, this is what allows us to tease apart the insatiable probe account from a probe activation account , using the LDB.

Deal's account can be extended to the GCC in Sierra Zapotec, *but only if the probe is located between the two arguments.* Since the probe needs to find the less specified argument first, Deal must assume that direct objects asymmetrically c-command indirect objects if the probe is located above both arguments. However, this is not possible with the GCC, which restricts subject–object combinations, and subjects unambiguously asymmetrically c-command objects.[18] For the logic of interaction condition update to go through, the probe would have to be located between the subject and object. One possibility, which Deal mentions, is that it is hosted by v,

---

[17] It is unclear whether this is a general claim or not, as in many languages the indirect object is the structurally superior argument. See Harley & Miyagawa (2017) for a recent survey of this variation.

[18] In Sierra Zapotec, the subject is rigidly ordered before the object. See Adler et al. (2018) for an account of how basic word order is derived in the language, as well as binding data showing that the subject asymmetrically c-commands the object.

looking first into its complement before then looking into its specifier via cyclic expansion (Béjar & Rezac 2009). This shifts the order of interaction, such that the probe interacts first with the lower argument, i.e., the object in interactions like the LDB/GCC.

Consider, now, a GCC-compliant third person combination like (45), where the subject is elder and the object is human. The probe Agrees first with the object. If it possesses dynamic interaction features for all three animacy categories ([EL]$^\uparrow$, [HU]$^\uparrow$, and [AN]$^\uparrow$), its interaction condition will be updated, and it will then be able to Agree with the subject in its specifier, which is specified for all the features that the object has.
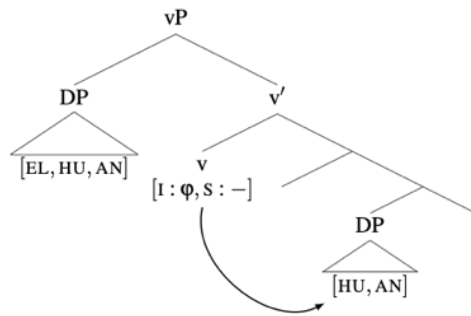
(45)   *3.EL > 3.HU*
       Step 1:   Step 3:
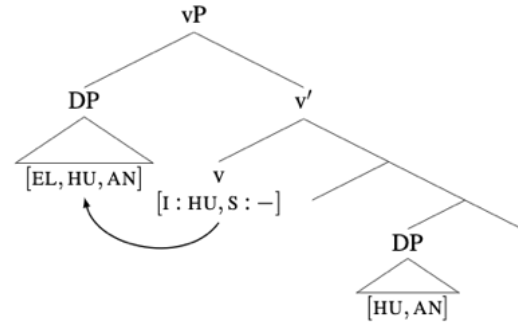
       Step 2: [I : HU, S : −]

The combination of these assumptions correctly predicts that the combination in (46), where the subject is human and the object is elder is ungrammatical. After the probe's interaction condition is updated by Agreeing with the object, it cannot interact with the subject, which lacks [EL].

(46)   *3.HU > 3.EL*
       Step 1:
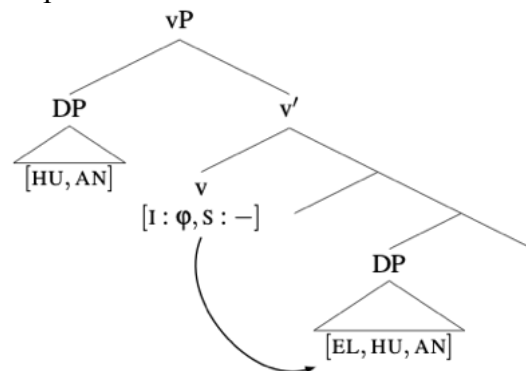
       Step 2: [I : EL, S : −]

To preserve Deal's generalization about the specificational asymmetry and the order of interaction, it is crucial that the probe be located between the two arguments and interact with the lower argument first.

This order of interaction cannot be applied, however, to the LDB, when we hold all other assumptions constant. This is because lexical DP subjects do not move. While we have claimed that the subject DP creates an intervention effect, in Deal's system this should be beneficial. In a configuration with an object clitic specified for [EL], the probe Agrees first with the object. As argued in Section 3, lexical DPs are not specified for [EL], so no Agree relation is established with the subject. Since the subject DP need not move, no problem should arise.

(47)　　*DP > 3.*EL*
　　　　Step 1:
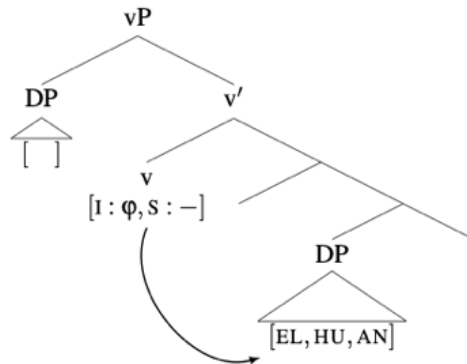
　　　　Step 2: [I : EL, S : −]

This order of interaction should also bear no ill consequences for the object clitic: since the object is interacted with first, the subject should not intervene, and the object should Agree and move. Yet, this configuration is ungrammatical.

The account also incorrectly predicts ungrammaticality for the inverse situation, where the probe first interacts with a nominal that does not need to move, i.e., a lexical DP object, and then subsequently interacts with a subject clitic. Deal considers this type of configuration in Tseltal, where only indirect objects cliticize: direct objects are not realized as weak pronouns. Nonetheless, in Tseltal and other Mayan languages, a local person direct object is ungrammatical when the indirect object is a local person clitic pronoun.

(48)　*Lah　　　y-aʔ-**bat**　　　　　　　　joʔon-eʔ.
　　　　PFV　　　3.ERG-give-**APPL.2.ABS**　　　**1SG-CL**
　　　　Intended: 'She gave you me.'　　　　　　　　(Shklavosky 2012: 445)

Since the probe Agrees first with the direct object, it is not able to reach the indirect object, which must Agree in order to cliticize. In the subject-object counterpart in Sierra Zapotec, however, this configuration is perfectly grammatical: see (2b) above. Thus, reversing the order of interaction in order to account for the GCC, first object and then subject, leads to false predictions for the LDB in both directions: pro > DP is predicted to be ungrammatical and DP > pronoun is expected to be grammatical.
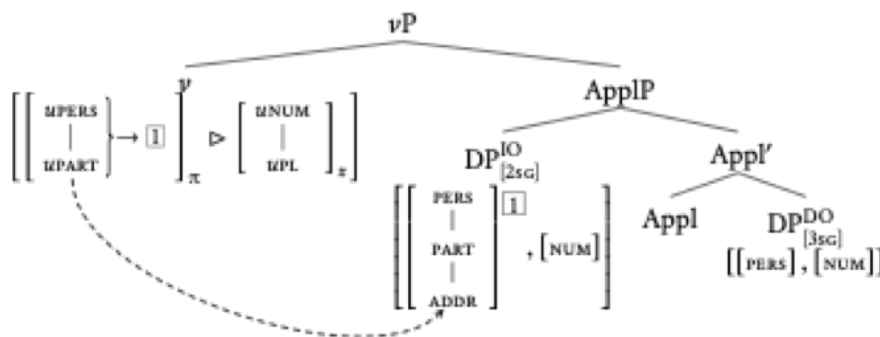
32

Thus, we conclude that this combination of assumptions is untenable. The probe must (i) c-command all goals and (ii) be subject to Attract Closest, so that it Agrees with the highest goal first. This implies that the generalization regarding the specificational asymmetry required by the LDB/GCC, as well as Weak PCC, is as stated above: the higher argument cannot be less specified than the lower argument. And this, in turn, has the consequences for the inner workings of Agree that we described in Section 4.1.

4.3. Feature gluttony

Coon & Keine (to appear) take a different approach to the problem of overprobing. They assume that a probe c-commands all of its goals and is subject to Attract Closest. And, they assume that the Agree operation does not have separate interaction and satisfaction conditions. However, to enable interactions beyond the first goal, they take the probe to have a richer structure in two ways. First, they assume that the probe's requirements are *sequenced*: it is divided into a sequence of "subprobes", searching for person, number, and gender features, in that order (Taraldsen 1995, Anagnastopoulou 2003, Béjar & Rezac 2003, among others). Second, within each of these φ-domains, the probe is *articulated*, in Béjar & Rezac's (2009) sense (see also Béjar 2003). After Agreeing with one goal, it can continue to Agree with other goals, subject to locality, as long as it still has unvalued features. Coon & Keine call this configuration, where one of the probe's subprobes has Agreed with more than one goal, "feature gluttony." This condition is not inherently problematic — it is a licit output of the Agree operation — though it may create problems for other syntactic operations or for the morphology. In particular, Coon & Keine take the ungrammatical combinations of clitic pronouns in the Weak PCC to arise when the probe has entered into a state of feature gluttony, which violates certain conditions on cliticization.

We begin with a grammatical combination, where the higher argument is local person and the lower argument third person. The probe Agrees, first, in person features with the highest goal it can find. The probe's person requirement is completely satisfied (49a), and it can Agree again, this time in number (49b). Both arguments Agree, and thus can cliticize.
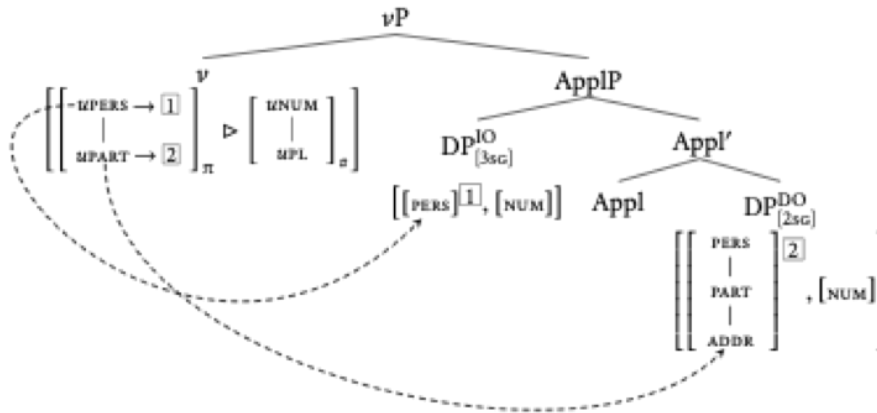
(49)   *2 > 3*
    a.

b.

vP
D_IO=v   ApplP

$\left[\left[\begin{matrix} u\text{PERS} \\ | \\ u\text{PART} \end{matrix}\right] \rightarrow \boxed{1}\right]_{\pi} \triangleright \left[\begin{matrix} u\text{NUM} \rightarrow \boxed{2} \\ | \\ u\text{PL} \end{matrix}\right]_{\#}$

DP$^{IO}_{[2SG]}$ Appl'

$\left[\left[\begin{matrix} \text{PERS} \\ | \\ \text{PART} \\ | \\ \text{ADDR} \end{matrix}\right] \boxed{1}, [\text{NUM}]\right]$ Appl   DP$^{DO}_{[3SG]}$

$[[\text{PERS}], [\text{NUM}]\boxed{2}]$
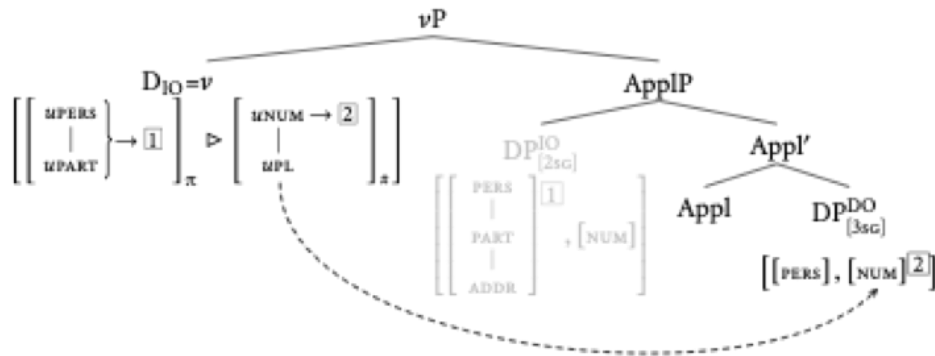
In the ungrammatical configuration in (50), the probe's person features are not fully checked by the highest goal. Since these are articulated, the probe must try again to check the rest of them. It does so by Agreeing with the lower argument, again in person.

(50) *3 > 2

vP
v   ApplP

$\left[\left[\begin{matrix} -u\text{PERS} \rightarrow \boxed{1} \\ | \\ u\text{PART} \rightarrow \boxed{2} \end{matrix}\right]_{\pi} \triangleright \left[\begin{matrix} u\text{NUM} \\ | \\ u\text{PL} \end{matrix}\right]_{\#}\right]$

DP$^{IO}_{[3SG]}$ Appl'

$[[\text{PERS}]\boxed{1}, [\text{NUM}]]$ Appl   DP$^{DO}_{[2SG]}$

$\left[\left[\begin{matrix} \text{PERS} \\ | \\ \text{PART} \\ | \\ \text{ADDR} \end{matrix}\right]\boxed{2}, [\text{NUM}]\right]$

The probe is now in a state of feature gluttony: a single subprobe, uPERS, has been checked by more than one goal. While not inherently problematic, this configuration runs afoul of the mechanism underlying cliticization. Coon & Keine propose that if a probe Agrees with a pronoun, it must internal (Merge) with it, as in (51) below. So, both pronouns in (50) must cliticize.

(51)   If a segment of a clitic-doubling probe on a head H has agreed with a DP, this DP
       must cliticize onto H.   (Coon & Keine, to appear, p. 21)

But they also propose that this Merge operation can only apply *once* for each subprobe: that is, once for person, once for number, and so on.[19] Since the two pronouns in (50) have checked the same person feature of the probe, they cannot both cliticize, and the derivation is ruled out by the constraint in (51). While the contour of the restriction follows from multiple iterations per

---

[19] They suggest that this comes from general computational considerations. Both pronouns cannot cliticize simultaneously since Merge is binary. Nor can they cliticize sequentially since every step in the derivation must be well-formed: after only one pronoun has moved, the constraint in (50) is still violated.

34

subprobe (because it is articulated), the constraint in (51), in essence, shifts the explanatory burden away from Agree to the theory of cliticization and the structure building component.

In this framework, the GCC receives a straightforward account, directly parallel to the Weak PCC. With the probe c-commanding both arguments, feature gluttony results when the subject has fewer gender features than the object, and cliticization of both pronouns is impossible. But it cannot be extended to the LDB, for the simple reason that one of the goals is not a pronoun. Consider a configuration where cliticization of a (general) third person object is impossible because the subject is a lexical DP:

(51)   *DP > pro
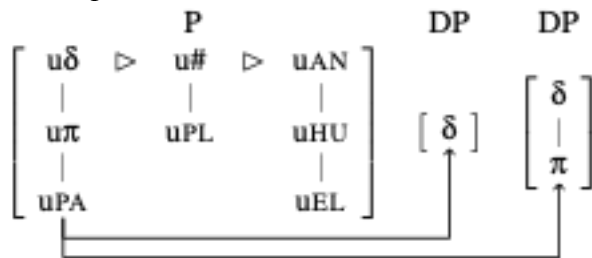
The probe Agrees with both arguments in person, since the lexical DP does not check all of its features. The result is feature gluttony. However, unlike with the PCC or GCC, there is no problem satisfying the condition in (51): only one of the goals is eligible to cliticize, and so there should be no problem with the object pronoun moving.[20]

Thus, feature gluttony and the articulated probes that underlie it cannot be at the root of the LDB. While probes may indeed be sequenced and articulated, and feature gluttony may indeed result in ill-formed derivations in other cases as Coon & Keine propose, these considerations do not underlie the locality effects we have been examining here, in particular the LDB.

4.4. Parameters of cross-linguistic variation

This account of the LDB/GCC and its repair has led us to a view of Agree in which it is both maximally free — it can iterate freely in an unrestricted fashion — and maximally rigid — subsequent iterations cannot modify the outcome of the initial relation. It is improbable that these properties of Agree hold exclusively in Sierra Zapotec. Why, then, do we not always see the effects of the LDB, nor always the effects of the GCC or PCC? We conclude this section with some brief comments on the dimensions along which languages might vary in these respects under our account.

First, the pronoun inventory of a language might be more restricted than in Sierra Zapotec. While we have not explicitly addressed what underlies the differences amongst pronoun types, it is clear that some languages only have non-clitic pronouns (e.g., weak or strong). If a language lacked clitic pronouns altogether, we might expect it to lack a LDB, GCC, or PCC. While clitic (and possibly weak) pronouns must Agree, other classes of pronouns need not. (This distinction

---

[20] There is a question here about what drives cliticization. The condition in (51) suggests that it is motivated by attributes of the probe (something like an EPP), but such a requirement would not be satisfied in every derivation, e.g., a sentence with only lexical DP arguments.

was the foundation of our account of the pronominal repair in Serra Zapotec.) Simply lacking clitic (or weak) pronouns is thus one possible locus of crosslinguistic variation.

Even if a language did have clitic pronouns, they could still fail to exhibit some of these intervention effects. There is significant variation in the PCC across languages: some languages adhere to the person hierarchy more completely than others (e.g., the Weak vs. Strong PCC, as well as the Ultrastrong PCC; see Pancheva & Zubizaretta 2017 for a survey). Similarly, there is variation across Sierra Zapotec dialects in the degree to which they adhere to the animacy hierarchy (Foley & Toosarvandani, to appear). It is fairly standard to attribute this kind of variation to the specification of the probe, following Anagnostopoulou (2005) and Nevins (2007, 2011). In particular, Nevins proposes that, in a language that enforces a person (or animacy) hierarchy more fully, the probe is specified so that it looks for more features. Insensitivity to a hierarchy corresponds to lack of feature specification on the probe. If a probe does not need *any* person or animacy features, it would not Agree with its goals in these features, and thus would not have a PCC or GCC. The probe could even lack a specification for δ, in which case it would lack the LDB, too.

It does not seem improbable that a language could exhibit the LDB for some arguments and not others. We have not considered ditransitives to any significant extent, but we might expect that a subject would intervene for object cliticization, while an indirect object would not. Here, multiple paths forward are available. The indirect object might be embedded under additional functional material, and hence invisible to Agree with the probe (Preminger 2014: 137–140). Whether this functional material is present and is sufficient to enable the probe to skip past the DP inside it are factors that are plausibly subject to variation, so we would not necessarily expect all indirect objects to behave alike.

Beyond these considerations, there may be other loci of variation that have consequences for whether a language exhibits the intervention effects we have been concerned with. In outlining these parameters, our intention has been to demonstrate that the probe activation model does not entail any uniformity in whether languages exhibit the GCC, PCC, or LDB. There are at least these three independent properties (and likely more), whose interaction with the Agree mechanism could plausibly yield the variation that we know to exist.

5. Conclusions and open questions

We have proposed that the LDB is akin to constraints like the GCC and Weak PCC, though with a twist: a lexical DP intervenes for cliticization, even though it itself need not move. We have suggested that this supports the view that locality constraints are not inscribed directly in the movement component, and that Attract Closest should be understood, instead, as a locality condition on Agree. This is consistent with an approach to movement that reduces it to internal Merge, in other words to an operation that is nothing "special."

Our examination of the LDB has led to a number of other conclusions, of which (i) and (iv) are new:

(i)   The person domain stretches to include lexical DPs. There is a person feature δ, which is maximally underspecified, shared by both pronominals and non-pronominals; in some languages, the presence of this feature makes lexical DPs interveners for cliticization.

(ii)   Specificational asymmetries of the sort observed in LDB, as well as in the GCC and PCC, support an approach to feature specification in which traditional categories such as person and animacy are internally arranged vertically, rather than horizontally. This is the representational basis for the specificational asymmetries at the core of these constraints.

(iii) Agree is associated with a single clausal probe (i.e., there is no Agr_S, Agr_O, or their equivalents).

(iv)   Only the first round of Agree with the closest goal, typically the subject, is obligatory.

The last point, if correct, raises a host of issues which we can only gesture towards here. The fact that subsequent rounds of Agree are grammatically optional does not entail that there is optionality in well-formed outputs, relative to a given derivational starting point. As we have shown, in a non-LDB compliant derivation, the realization of the object as an independent pronoun is obligatory. What the optionality of subsequent rounds implies, then, is that any appearance of obligatoriness must be attributed to some other mechanism, which may be syntactic or not.

In Sierra Zapotec, specifically, we have seen obligatory movement and Agree by the clitic, which we have attributed to the preference for a clitic over an independent pronoun whenever possible, a requirement that may have a morphophonological source, rather than a syntactic one. Our analysis of nominals, in which pronominal and non-pronominals alike possess δ, makes a structural principle such as Cardinaletti & Starke's (1999) Minimize Structure unlikely. Their economy principle would require that clitic pronouns contain less nominal functional structure than a lexical DP, even though they also possess δ. While this is not inconsistent with the letter of our approach, it leaves very little substance to its spirit. It is more likely, given the distribution of δ that we propose, that all nominals possess the same external structure, so there is no sense in which pronouns are smaller than other nominals. But Minimize Structure is not the only way to account for the difference between clitics and independent pronouns. It may be possible to do this within a realizational model such as Distributed Morphology by having the notion of a minimal pronoun correspond to clitic pronouns (Kratzer 2009, Rooryck & Wyngaerd 2011, Safir 2014), with an optimizing calculus (Burzio 1998, Bresnan 2001), some other global calculation (Rezac 2011), or other economy and expressivity constraints on morphological exponence (Foley 2020). We leave this choice, and the nature of the clitic-independent pronoun alternation, open to further exploration.

Finally, if we are correct to attribute the LDB to syntax proper, our analysis may also apply to those subcases of Holmberg's Generalization in which a dative argument blocks object shift of an accusative pronoun in Scandinavian (Holmberg 1999). This effect is essentially the same as the LDB, except that the intervening argument is a dative. But that difference may be due to independent factors. In Scandinavian, the subject (in specTP) is presumably higher than it is in Sierra Zapotec (which has verb-initial word order). If the probe is in the T area generally, it will be above the subject in Sierra Zapotec, but below it in Swedish. This could explain why subjects are only interveners in the former. We leave a more complete account of this particular case of crosslinguistic variation for the future.

References

Ackema, Peter & Ad Neeleman. 2018. Features of person: From the inventory of persons to their morphological realization. Cambridge, MA: MIT Press.

Aissen, Judith. 1997. On the syntax of obviation. Language 73: 705–750.

Adler, Jeff, Steven Foley, Jedd Pizarro-Guevara, Kelsey Sasaki & Maziar Toosarvandani. 2018. The derivation of verb initiality in Santiago Laxopa Zapotec. A reasonable way to proceed: Essays in honor of Jim McCloskey. Jason Merchant, Line Mikkelsen, Deniz Rudin, and Kelsey Sasaki, eds. 31–49.

Anagnostopoulou, Elena. 2003. The syntax of ditransitives: Evidence from clitics. Berlin: Mouton de Gruyter.

Anagnostopoulou, Elena. 2004. Holmberg's Generalization and locality. Generative Linguistics in the Old World (GLOW) 27.

Anagnostopoulou, Elena. 2005. Strong and weak person restrictions: A feature checking analysis. In Clitic and affix combinations: Theoretical perspectives, eds. Lorie Heggie and Francisco Francisco Ordóñez, 199–235. Amsterdam: John Benjamins.

Anagnostopoulou, Elena. 2006. Clitic doubling. In The Blackwell companion to syntax, eds. Martin Everaert and H. C. van Riemsdijk, 519–581. Oxford: Blackwell.

Avelino Becerra, Heriberto. 2004. Topics in Yalálag Zapotec, with particular reference to its phonetic structure. Ph.D. Dissertation, University of California, Los Angeles. Béjar, Susana. 2003. Phi-syntax: A theory of agreement. Ph.D. Dissertation, University of Toronto.

Baker, Mark. 2003. Lexical categories: Verbs, nouns and adjectives. Cambridge University Press.

Béjar, Susana. 2003. Phi-syntax: A theory of agreement. Doctoral dissertation, University of Toronto.

Béjar, Susana & Arsalan Kahnemuyipour. 2017. Non-canonical agreement in copular clauses. Journal of Linguistics. 53: 463499.

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

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

Benveniste, Emile. 1956. La nature des pronoms. In For Roman Jakobson: Essays on the occasion of his sixtieth birthday, 11 October 1956, ed. Morris Halle, 34–37. The Hague: Mouton.

Bianchi, Valentina. 2006. On the syntax of personal arguments. Lingua 116.

Bobaljik, Jonathan. 2008. Where's phi? Agreement as a post-syntactic operation. Phi theory: Phi features across modules and interfaces, ed. by Daniel Harbour, David Adger, and Susana Béjar, 295–328. Oxford University Press.

Bonet, Eulàlia. 1991. Morphology after syntax: Pronominal clitics in Romance. Ph.D. Dissertation, Massachusetts Institute of Technology. 81

Borer, Hagit. 1984. Restrictive relatives in Modern Hebrew. Natural Language and Linguistic Theory 2:219–260.

Borsley, Robert D. & Jaklin Kornfilt. 2000. Mixed extended projections. The nature and function of syntactic categories, vol. 32 Syntax and Semantics, ed. by R. D. Borsley, 101–132. Academic Press.

Bresnan, Joan. 2001. The emergence of the unmarked pronoun. Optimality theoretic syntax, ed. by Géraldine Legendre, Sten Vikner and Jane Grimshaw, 113–142. MIT Press.

Burzio, Luigi. 1998. Anaphora and soft constraints. Is the best good enough? Optimality and competition in syntax, ed. by Pilar Barbosa, 93–114, MIT Press.

Cardinaletti, Anna and Michal Starke. 1999. The typology of structural deficiency: A case study of the three classes of pronouns. In Clitics in the languages of Europe, ed. H. C. van Riemsdijk, 145–233. Berlin: Mouton de Gruyter.

Carstens, Vicki. 1991. The morphology and syntax of Determiner Phrases in Kiswahili. Doctoral dissertation, UCLA.

Chomsky, Noam. 1995. Bare phrase structure. In Evolution and revolution in linguistic theory, eds. Héctor Campos and Paula Kempchinsky, 57–109. Washington: Georgetown University Press.

Chomsky, Noam. 2000. Minimalist inquiries: The framework. In Step by step: Essays on Minimalist syntax in honor of Howard Lasnik, eds. Roger Martin, David Michaels, and Juan Uriagereka, 89–156. Cambridge, MA: MIT Press.

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

Clements, George N. 1975. The logophoric pronoun in Ewe: Its role in discourse. *Journal of West African Languages* 10: 141–177.

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

Cooper, Robin. 1979. The interpretation of pronouns. In Selections from the third Groningen round table, eds. Frank Heny and Helmut S. Schnelle, volume 10, 61–92. San Diego: Brill.

Deal, Amy Rose. 2015. Interaction and satisfaction in φ-agreement. North East Linguistic Society (NELS) 45:1–14.

Deal, Amy Rose. 2020. Interaction, satisfaction, and the PCC. North East Linguistic Society (NELS) 50.

Déchaine, Rose Marie & Martina Wiltschko. 2002. Decomposing pronouns. Linguistic Inquiry 33.3: 409-442.

Dowty, David and Pauline Jacobson. 1988. Agreement as a semantic phenomenon. Eastern States Conference on Linguistics (ESCOL) 95–108.

Foley, Steven. 2020. Case, agreement, and sentence processing in Georgian. PhD dissertation, University of California, Santa Cruz.

Foley, Steven and Maziar Toosarvandani. To appear. Extending the Person-Case Constraint to gender: Agreement, locality, and the syntax of pronouns. Linguistic Inquiry.

Forscheimer, Paul. 1953. The category of person in language. Walter de Gruyter.

Fox, Danny and David Pesetsky. 2005. Cyclic linearization of syntactic structure. Theoretical Linguistics 31: 1–45.

Geach, Peter. 1962. Reference and generality: An examination of some medieval and modern theories. Ithaca, NY: Cornell University Press.

Ghomeshi, Jila and Diane Massame. 2020. Number is different in nominal and pronominal phrases. Linguistic Inquiry 51:597–610.

Giannakidou, Anastasia. 2000. Negative… concord? Natural Language and Linguistic Theory 18.3:457–523.

Ginzburg, Jonathan and Ivan Sag. 2000. Interrogative investigations. Stanford CA: CSLI

publications.

Gupta, Anil. 1980. The logic of common nouns: An investigation in quantified modal logic. Yale University Press.

Hagège, Claude. 1974. Les pronoms logophoriques. *Bulletin de la Société de Linguistique de Paris* 69: 287–310.

Harbour, Daniel. 2008. Morphosemantic number: From Kiowa noun classes to UG number features. Dordrecht: Springer.

Harbour, Daniel. 2016. Impossible persons. Cambridge, MA: MIT Press.

Harley, Heidi & Shigeru Miyagawa. 2017. The syntax of ditransitives. O*xford Research Encyclopedia of Linguistics.* DOI:10.1093/acrefore/9780199384655.013.186.

Harley, Heidi and Elizabeth Ritter. 2002. Person and number in pronouns: A feature-geometric analysis. Language 78:482–526. 83

Heim, Irene and Angelika Kratzer. 1998. Semantics in Generative Grammar. Oxford: Blackwell.

Hiraiwa, Ken. 2000. On nominative-genitive conversion. MIT Working Papers in Linguistics 39:66–124.

Hiraiwa, Ken. 2001. Multiple Agree and the defective intervention constraint in Japanese. MIT Working Papers in Linguistics 40:67–80.

Holmberg, Anders. 1986. Word order and syntactic features in the Scandinavian languages. PhD Dissertation, University of Stockholm.

Holmberg, Anders. 1999. Remarks on Holmberg's generalization. Studia Linguistica 53.1:1–39.

Ingram, David. 1978. Typology and universals of personal pronouns. Universals in Human Language. 3:213–248.

Jacobson, Pauline. 2013. The short answer: Implications for direct compositionality (and vice versa). Ms. Brown University.

Kayne, Richard. 2000. A note on clitic doubling in French. Parameters and Universals. Oxford University Press.

Kitahara, Hisatsugu. 1997. Elementary operations and optimal derivations. Cambridge MA: MIT Press.

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

Larson, Richard & Gabriel M. A. Segal. 1995. Knowledge of meaning: An introduction to semantic theory. Cambridge MA: MIT Press.

López, Filemón and Ronaldo Newberg. 2005. La conjugación del verbo zapoteco, zapoteco de Yalálag. Mexico City: Instituto Lingüístico de Verano.

Merchant, Jason. 2004. Fragments and ellipsis. Linguistics and philosophy 27.6:661–738.

Müller, Gereon. 2000. Shape conservation and remnant movement. Proceedings of NELS 30, ed. by M. Hirotani, A. Coetzee, N. Hall, J.-Y. Kim, 525–539, Amherst, UMass: GLSA.

Murasugi, Kumiko. 1992. Crossing and nested paths: NP movement in accusative and ergative languages. Doctoral dissertation, MIT>

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.

Noyer, Rolf. 1992. Features, positions, and affixes in autonomous morphological structure. Ph.D. Dissertation, Massachusetts Institute of Technology.

Noyer, Rolf. 1997. Features, positions and affixes in autonomous morphological structure. New York: Garland, revised version of 1992 MIT dissertation.

Ochi, Masao. 1998. Move or Attract? Proceedings of the Sixteenth West coast conference on formal linguistics, ed. by Emily Curtis, James Lyle, and Gabriel Webster, 319–333. Stanford CA: CSLI publications.

Oxford, Will. 2017. Size matters: Deriving person hierarchies from nominal structure. Ms., University of Manitoba. https://ling.auf.net/lingbuzz/004534.

Pancheva, Roumyana and Maria Luisa Zubizarreta. 2017. The Person Case Constraint: The syntactic encoding of perspective. Natural Language and Linguistic Theory .

Patel-Grosz, Pritty & Patrick Grosz. 2017. Revisiting pronominal typology. Linguistic Inquiry 48.2:259–297.

Perlmutter, David. 1971. Deep and surface structure constraints in syntax. New York: Holt, Reinhart, and Winston.

Pesetsky, David. 2001. Phrasal movement and its kin. Cambridge MA: MIT Press.

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

Preminger, Omer. 2019. What the PCC tells us about "abstract" agreement, head movement, and locality. Glossa 4:13.

Rackowski, Andrea & Norvin Richards. 2005. Phase edge and extraction: A Tagalog case study. Linguistic Inquiry 36.4:565–599.

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

Richards, Norvin. 1997. What moves where when in which languages? Doctoral dissertation MIT.

Ritter, Elizabeth. 1993. Where's gender? Linguistic Inquiry 24: 795–803.

Rizzi, Luigi. 1990. Relativized minimality. Cambridge, MA: MIT Press.

Rizzi, Luigi. 2004. Locality and left periphery. Structures and beyond: The cartography of syntactic structures 3, ed. by Adriana Beletti, 223–251. Oxford University Press.

Rooryck, Johan & Guido Vanden Wyngaerd. 2011. Dissolving Binding Theory. Oxford University Press.

Safir, Ken. 2014. One true anaphor. Linguistic Inquiry 45:91–124.

Sauerland, Uli. 2006. On the semantic markedness of phi-features. In Phi theory: Phi-features across interfaces and modules, eds. Daniel Harbour, David Adger, and Susan Béjar, 57–82. Oxford: Oxford University Press.

Schütze, Carson. 1997. INFL in child and adult language: Agreement, case and licensing. Doctoral dissertation, MIT.

Sells, Peter. 2001. Structure, alignment and optimality in Swedish. Stanford CA: CSLI publications.

Shklovksy, Kirill. 2012. Person–Case effects in Tseltal. *The Linguistic Review* 29(3): 439–490.

Sichel, Ivy. 2001. Studies in the syntax of pronouns and features. Doctoral dissertation, CUNY Graduate Center.

Sichel, Ivy. 2002. Pronoun movement and attraction. Proceedings of the Israeli Association of Theoretical Linguistics (ITAL) 18.

Sichel, Ivy & Maziar Toosarvandani. to appear. Pronouns and attraction in Sierra Zapotec. Syntax and Semantics at Santa Cruz (SASC) 4, ed. by Andrew Hedding and Morwenna Hoeks.

Sichel, Ivy & Martina Wiltschko. to appear. The logic of person markedness. Language.

Sportiche, Dominique. 1993. Clitic constructions. In Phrase structure and the lexicon, eds. Johan Rooryck and Laurie Zaring, 213–276. Dordrecht: Springer.

Stainton, Robert, J. 1998. Quantifier phrases, meaningfulness "in isolation", and ellipsis. Linguistics and philosophy 21:311–340.

Suñer, Margarita. 1988. The role of agreement in clitic-doubled constructions. Natural Language and Linguistic Theory 6:391–434.

Taraldsen, Knut Tarald. 1995. On agreement and nominative objects in Icelandic. In Studies in comparative german syntax, eds. Hubert Haider, Susan Olsen, and Sten Vikner, 307–327. Dordrecht: Kluwer.

Wechsler, Stephen. 2004. Number as person. Empirical issues in syntax and semantics 5, ed. by Olivier Bonami and Patricia Cabredo Hofherr, 255–274.

Weir, Andrew. 2014. Fragments and clausal ellipsis. Doctoral dissertation, UMass Amherst.

Williams, Edwin. 2003. Representation Theory. Cambridge MA: MIT Press.

Wiltschko, Martina. 1998. On the syntax and semantics of (relative) pronouns and determiners. The journal of comparative Germanic linguistics 2:143–181.

Zwicky, Arnold. 1977. On clitics. Bloomington IN: Indiana University Linguistics Club.

isichel@ucsc.edu
mtoosarv@ucsc.edu