Croatian Journal of Philosophy
Vol. XXIII, No. 69, 2023
https://doi.org/10.52685/cjp.23.69.5
Received: September 23, 2023
Accepted: November 14, 2023

The Syntax of Prominence

SAŠO ŽIVANOVIĆ
University of Ljubljana, Ljubljana, Slovenia
PETER LUDLOW
Status Holdings Pte. Ltd., Singapore

The standard view on discourse pronoun resolution is that determining the antecedents of discourse pronouns is typically a function of extra-linguistic reasoning. In contrast, Stojnić (2021) argues that pronoun resolution is a function of linguistic facts. In this article we offer what we take to be a friendly amendment to the technical aspects of Stojnić's proposal. Our point of departure will be with our idea that prominence is not determined by the position of the candidate antecedent within a stack, but rather by its position within standard syntactic tree structures, extended to include discourse-level trees. Our proposal leans on the notion of p-scope, a proof-theoretic accessibility relation among tree nodes which we develop in Ludlow and Živanović (2022), and the notion of closeness built on standard accounts of syntactic locality. The key idea is that a pronoun's antecedent resolves to its closest p-scoper; specifically, p-scope determines the potential antecedents, and the closeness relation orders these by prominence. Coherence relations, which we provisionally represent as syntactic heads, can be then seen as affecting accessibility and prominence indirectly, in virtue of their position in traditional LF tree structures.

Keywords: Pronoun resolution; discourse prominence; coherence relations; syntactic locality; logical form.

1. *Introduction*

The standard view on discourse pronoun resolution is that determining the antecedents of discourse pronouns is typically a function of extra-linguistic reasoning. In other words, when we encounter a pronoun that has an antecedent earlier in the discourse, often the determination of the proper antecedent cannot be determined by the syntax and

semantics of natural language alone but must be inferred with the help of pragmatics and common-sense reasoning.

In contrast to the standard view, the proposal in Stojnić (2021) argues that pronoun resolution is a function of linguistic facts, here taking "linguistic" to include logical forms that encode syntactic and semantic information. Of course, a proposal like this requires that things like deictic gestures inform the construction of the logical forms. But while gestures and other information may be involved in the construction of the logical forms, once constructed, only those logical forms are required to resolve the discourse antecedent.

Is it also possible to construe her proposal as suggesting that pronoun resolution could even be a function of syntax in the sense of computational data structures of some form? Maybe. Hereafter we will refer to this idea as involving *broad syntax* to distinguish it from an approach that relies on the syntactic structures in contemporary generative linguistics; we will refer to the latter approach as involving *narrow syntax*. We will also distinguish these two approaches when talking about logical forms that are posited. We will use "logical form" (lower case) when we speak of logical representations that involve broad syntax, and we will use "Logical Form" (upper case) or "LF" to talk about a level of syntactic representation in contemporary generative linguistics—narrow syntax, on our taxonomy.

The key element to the proposal in Stojnić (2021) is the idea that the antecedent of a pronoun will be determined by the *prominence* of elements in the discourse, and this is in turn determined by where those elements lie in a *stack*, a data structure populated during the interpretation of the logical form. Should we take this as being a broadly syntactic way of representing prominence? In appendix (A.1.1), a stack is characterized as "sequences of individuals from the model," which is certainly not syntactic, but it can be converted into a syntactic proposal if we replace individuals with representations of individuals. We note that on the latter understanding, such a stack and where something lies in that stack is a formal syntactic state (broadly syntactic). No semantic or pragmatic information need be accessed. Thus, just using a syntactic account of where things rest in the stack, one can give a broadly syntactic account of prominence, and thus of the resolution of discourse antecedents. We remain neutral on whether Stojnić offers a broadly syntactic approach or not.

In this article we offer what we take to be a friendly amendment to the proposal in Stojnić (2021), whether that proposal is ultimately taken to be broadly syntactic or not. Our proposal is sympathetic in that it offers a clearly linguistic account of prominence and discourse antecedence. Our point of departure will be with our idea that the prominent antecedent is not determined by the position of the candidate antecedent within a stack, but rather by its position within standard syntactic tree structures (extended to include discourse-level trees), and a notion of prominence that is, at its core, a syntactic relation holding between

nodes on those discourse trees. Another way to put our thesis is that we offer a "narrow syntax" account of prominence relations. That is to say, we will be representing prominence relations in the LF representations of contemporary generative linguistics.[1]

Our proposal will lean on the notion of p-scope which we develop in Ludlow and Živanović (2022). Because p-scope plays a critical role in the determination of what can be moved where within proof-theoretic derivations, this suggests that there is a strong relation between our notion of prominence and the fundamental elements of proof-theoretic semantics. In other words, prominence cannot be understood as an *ad hoc* device to resolve discourse antecedents, but it is a notion that is deeply wired into proof-theoretic accounts of inference in natural language, and thus into the most basic elements of the logical syntax of natural language. But because our proof theory is interwoven with the LF representations and the machinery of contemporary generative linguistics, we can also conclude that discourse level prominence is also deeply embedded in the LF representations and machinery of contemporary generative linguistics.

Part of our strategy is to develop and use a notion of closeness, building on standard accounts of syntactic locality. The key idea is that a pronoun's antecedent resolves to its closest p-scoper; specifically, p-scope determines the potential antecedents, and the closeness relation orders these by prominence. Coherence relations, which we will provisionally identify with syntactic heads, can be then seen as affecting accessibility and prominence indirectly, in virtue of their position in the Logical Form.

2. *The technical essentials of Stojnić's (2021) attention–coherence approach*

Stojnić (2021) persuasively argues that pronouns can be assigned uniform, unambiguous, context-independent meaning of resolving to the most prominent referent in the discourse, if prominence is correctly defined and tracked.

Following works on Centering Theory (Grosz and Sidner 1986; Grosz et al. 1995), Stojnić tracks prominence using a data structure she calls a stack. Her stack is a linear structure accessed via numerical indices.[2] In the dynamic semantics logical form below, the indefinite and the name push new individuals to the stack. Subjects interact with position 0 (the top of the stack), and direct objects with position 1 (one item below the top), so 'a woman' and 'Sue' contribute updates $\langle a0\rangle$

[1] Ultimately, the idea will be that just as some linguists have seen the reflex of narrow syntactic phenomenon at the micro level in morphology, we see it in the macro level in discourse structures.

[2] As Stojnić's (2021) stack allows random access, it might be better to call it an array. True stacks allow one only to push and pop the top element. We will keep to Stojnić's terminology for consistency with previous approaches.

and ⟨π1*s*⟩ to the logical form; the former introduces an unspecified individual (later restricted to be a woman), and the latter introduces 'Sue'. Pronouns refer to the topmost element of the stack (@) satisfying the condition inherent to the pronoun; @she therefore refers to the to the topmost element of the stack which is singular and of feminine gender, while @she^{x0} additionally requires this individual to be distinct from the individual in the subject position 0 (thereby implementing the effect of Principle B of Binding Theory). After the pronoun is resolved, the individual it refers to is reintroduced to the stack, to the position in accord with its grammatical function. For example, π0@… pushes the subject to the top of the stack.

(1) a. A woman met Sue. She greeted her.
b. ⟨*a*0⟩; [woman(x_0)]; ⟨π1*s*⟩; [met(x_0,x_1)]; ⟨π0@she⟩; ⟨π1@she^{x0}⟩; [greeted(x_0,x_1)]

(Note the absence of indices on the pronominal variable @. In standard approaches, a pronoun is indexed in syntax, and that index determines which variable the pronoun is mapped into, e.g. 'he$_{42}$' is mapped into x_{42}. In contrast, Stojnić's pronouns are index-free, and always translate to the same variable, @, which always refers to the same object: whatever is found at the top of the stack at the moment.)

Stojnić (2021) further argues that demonstrative gestures are fully linguistic items and ought to be represented in the logical form. The logical form reflex of the demonstrative gesture accompanying pronoun 'she' below is ⟨π0*b*⟩: the demonstratum is pushed to stack position 0 (0, because 'she' is the subject). Consequently, whatever used to be at the top of the stack before (below, the subject of the first sentence) is demoted in prominence. Once we interpret the pronoun itself (@she), it will thus resolve to "Betty" rather than "the woman who came in".

(2) a. A woman came in. She [pointing at a cat, Betty] sat down.
b. ⟨*a*0⟩; [woman(x_0)]; [came-in(x_0)]; ⟨π0*b*⟩; ⟨π0@she⟩; [sat-down(x_0)]

Finally, prominence ranking is also affected by coherence relations. In the example below, the pronoun resolution in the second sentence crucially depends on whether we understand the second eventuality to be a result of the first one, or its explanation.

(3) John was disappointed with Tim.
a. He$_j$ fired him$_t$. (Result)
b. He$_t$ disobeyed him$_j$. (Explanation)

Coherence relations are also argued by Stojnić to affect the stack, and thereby adjust the prominence ranking. For example, the coherence relation Result promotes the prominence of the subject of the preceding sentence (π0x_0), while Explanation promotes the object (π0x_1).[3]

(4) ⟨π0*j*⟩; ⟨π1*t*⟩; [was-disappointed-with(x_0,x_1)];

[3] The logical forms in (4) contain a simplification: the arguments of the coherence relations are the undefined e_0 and e_1, intended to represent the eventualities described by the two sentences of the discourse.

a. [Result(e_0,e_1)]; ⟨π0x_0⟩; ⟨π0@he⟩; ⟨π1@he^{x_0}⟩; [fired(x_0,x_1)];
b. [Explanation(e_0,e_1)]; ⟨π0x_1⟩; ⟨π0@he⟩; ⟨π1@he^{x_0}⟩;
[disobeyed(x_0,x_1)];

3. *Two ways to track prominence*

As we noted in the introduction, our goal is to offer a variation on Stojnić's (2021) pronoun resolution mechanism in which we represent prominence in terms of relations between nodes on syntactic trees rather than positions in a stack. We will go into details about our proposal in a later section, once we get acquainted with our Dynamic Deductive System and its central relation, p-scope. For now, we want to tackle a more basic question: How do these two approaches to data structures differ in how they track prominence?

Consider the following discourse.

(5) A: If Jane works for Harry, she will have a really good situation. His business seems well organized and well managed, and she would be paid six figures. On top of this, he's super chill, he seems generous and he listens to her ideas. On the other hand, if she works for Richard, her situation might not be as good. He has a reputation for being cheap, he isn't very supportive of his staff, he has shady friends, and he has anger management issues.
B: So obviously, the first option. But how can she convince him to hire her?

We are interested in the resolution of pronoun 'him' in B's utterance. Clearly, this pronoun has 'Harry' from the first option as its antecedent. How does this come about, given that at the end of A's utterance, 'he' has 'Richard' as its antecedent? Obviously, the trigger for the shift in interpretation is the phrase 'the first option' in B's utterance, but the question here is what mechanism underlies this shift?

On our view, the data structure utilized for tracking prominence will be a tree, but here let us think how this discourse might be modeled using stack data structure, using a linear structure with numeric access. Starting with the first part of A's utterance, the conditional pushes 'Harry' to the top of the stack, so it is 'Harry' which 'he' (construed broadly to include 'him' and 'his') has as its antecedent. Once 'Richard' is introduced in the second-option conditional, he is pushed to the top of the stack; 'Harry"s prominence is demoted. So, which entry in the stack refers to has 'Harry' as its antecedent within the second option? Well, this depends. Immediately after the conditional, probably the second entry, but as the second option unfolds, other items are pushed to the stack, so the exact position of 'Harry' at the end of the second option depends on the amount of material introduced by the second option. It might be 12, or 42, or whatever. To model the effect of 'the first option', one would thus need to keep track of the amount of

the material introduced by the second option.

Clearly, one needs to keep track of the amount of the material introduced by the first option as well, to know where it ends. For example, the conversation might continue by (6a) but not by (6b).

(6) a. A: Perhaps by telling him some of those ideas?
b. A: #Perhaps by discussing some of those issues with him?

The technical challenge facing the stack-based implementation is to ensure that 'those issues' in (6b) does not resolve to Richard's 'anger management issues'. One idea might be to delete the portion of the stack pertaining to the second option upon meeting 'the first option' in B's utterance, but this is not quite right, as the speakers may revisit the second option. One viable solution seems to be to mark the extent of the stack pertaining to the first option and to limit the potential pronoun antecedents to that option.

But notice that by keeping track of which material belongs to the first option, and which material belongs to the second one, we have started to reconstruct the hierarchical structure of the discourse. To cover all bases, the entire discourse will have to be structured hierarchically, because the number of fine distinctions we may make when referring to the previous discourse is unlimited (or, limited only by our memory). In short, we may end up supplementing our stack with a tree, or something very much like it. Our view is that if a hierarchical structure, something like a tree, is required to track prominence in cases like this, perhaps we should simply "embrace the tree" at the beginning when we select our data structure for this job.

Before we delve into the specifics of our pronoun resolution strategy, we want to address one further question regarding the representation of prominence in a data structure. Is it an independent structure, constructed for and dedicated to tracking prominence? Or can we appropriate an existing structure for this very same job? In other words, is there a kind of linguistic representation which we need anyway, and which contains all the information on prominence we might ever require? As we will see, there certainly already is such a data structure available if we are engaged in proof-theoretic semantics, a project wherein the central engine of logical inference involves operations on the syntactic forms of natural language constructions.

4. *Syntax and the preservation of information*

Assume that A's utterance in (5) receives a logical form largely consistent with the gloss outlined below,[4] and further assume that the first sentence of B's utterance somehow moves "the discourse marker" into

[4] Even the parts of the logical form in (7) that are present are much simplified. In particular, we do not wish to claim that the material implication is a faithful representation of a conditional. We also neglect the fact that information on Harry and Richard given by A is not conditional upon whether Jane works for them.

the consequent of the first conjunct, i.e. into the left triangle, so that whatever B utters will now be added into the part of the discourse related to the first option. Given that it is *h* (for 'Harry') and not *r* (for 'Richard') which explicitly occurs in the antecedent of the first option conditional, do we not now have a reason to assume that it is 'Harry' rather than 'Richard' which a masculine singular pronoun will resolve to? The idea is that whatever pronoun resolution mechanism we deploy—as long as it utilizes the information present in the logical form and utilizes it locally in some fashion—if 'he' resolved to 'Harry' in the first option while that option was the only option (i.e. before the second conditional was introduced and conjoined to the first one), then 'he' will again resolve 'Harry' once the first option is revisited.

(7)

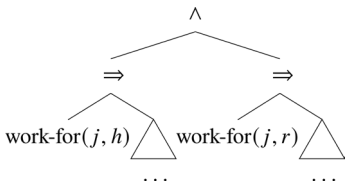

We will present our pronoun resolution mechanism in section 6. Here, we want do draw the reader's attention to the fact that certain information—notably, the positions of *h* and *r*—would be unavailable, or at least would not be transparently available, in a logical form of dynamic semantics, such as (1b). Those logical forms are a *sequence* of statements updating the context. They do not have hierarchical structure. At the end of *A* utterance, *A* and *B* are not left with a hierarchical logical form; they are left with a new context—a set of possible worlds. But the information about the previous discourse is not a part of the new context. For example, it is completely unrecoverable which option (the 'Harry' option, or the 'Richard' option) was discussed first. And given only the set of live possible worlds at some point in the discourse, it is also completely unrecoverable which antecedent noun phrase is the most prominent at that point—which is precisely why Stojnić (2021), being interested in prominence, observed that we need to track it independently, using the stack.

The general point is this: a formal representation carries more information than a model-theoretic representation of the same phenomenon. Moving from proof theory to model theory loses information. In particular, it loses syntactic information, like whether a conditional was rendered as $\phi \Rightarrow \psi$ or the logically equivalent $\neg\phi \vee \psi$.[5] Of course, the additional information offered by the proof-theoretic approach is not always pertinent to the phenomenon being investigated, so it often

[5] One can of course use additional model theory to model syntactic structures, but in doing so you are not doing model-theoretic semantics; you are using the resources of set theory to model syntactic information, which of course is certainly possible. But in doing so you are still doing syntax, even if syntax in disguise.

makes perfect sense to abstract away from it using model theory. However, every now and then, this information is crucial. It is crucial when we are trying to work out a pronoun resolution mechanism; witness Stojnić's (2021) stack. It was crucial for many authors trying to provide a compositional semantics of some phenomenon: we have seen them, again and again, enriching the concept of meaning with syntactic information, like indices or argument structure, in order to achieve compositionality; for a discussion, see Ludlow and Živanović (2022: §12.1.2). It is not for nothing that in the second edition of Meaning and Necessity (Carnap 1956), Carnap felt the need to supplement his possible worlds semantics with more fine-grained syntactic information to account for hyperintensional contexts like belief reports.

The informational robustness of syntactic forms compared to sets of possible worlds is one of the major reasons we advocate for a framework closer in spirit to the original DRT (Kamp 1981; Heim 1982) than to modern incarnations of dynamic semantics emanating from Groenendijk and Stokhof (1991). We choose to model the discourse by a Logical Form which grows as the discourse develops. In our opinion, there is no important difference in kind between the Logical Form of a sentence and the Logical Form of a discourse; the latter is simply a merger of many of the former. On this point, we echo a point made in Larson (1990). There, Larson argues the many standard linguistic relations track across sentential boundaries if we simply assume that discourse is represented by (growing) standard tree structures.

In our system, the ultimate result of asserting something is not shrinking of the set of live possible worlds; it is the integration of the Logical Form of the new assertion to the previous discourse Logical Form. (We say "integration" because the new material patently will not always be conjoined to the preceding discourse logical form at the root. The details of the options above illustrate that, as they must be joined into the consequent of a conditional.)

Our other major assumption is that the logical form of a sentence, as understood by philosophers, is nothing but the Logical Form of a sentence as understood by the practitioners of Generative Linguistics, including the Minimalist Program (Chomsky 1995, and many subsequent works). For reasons of space, we cannot justify this assumption in any detail in this paper, and refer the reader to Ludlow and Živanović (2022: §12). There, we argue that, despite the received view that the structure of natural language and predicate logic is completely different, they are in fact isomorphic.

Generative linguists, at least since Chomsky (1965) and his conception of Deep Structure, have observed that in language, what you see on the surface is not always what you get. The surface form is not always a good indicator of the underlying structure. In contemporary iterations of generative linguistics, the interpretable structure produced by syntax is the Logical Form (LF), which, unlike the surface form,

reflects the intended scope of quantifiers among other properties. It is the LF rather than the surface form of a sentence that we believe the logical form of philosophers is isomorphic to. This too is an idea dating back to the early days of generative linguistics (see Chomsky 1977), and which further dates back to the days of Deep Structure—see, for example, Harman (1970).

This brings us to the proposal in Ludlow and Živanović (2022). The guiding principle of that work is that much of the work that is currently carried out using model-theoretic semantics can be carried out in syntax—in some cases with superior results. For example, we typically use model-theoretic semantics to model the entailment relations between sentences of natural language, but we adopt a proof-theoretic approach to an account of such relations and carry out all such inferences in syntax.

This is certainly not the place to go into details about the execution of this very technical project, but we can give an informal gloss here, referring the reader to our book for details and a formal execution of the idea.

5. *P-scope*

P-scope is a central relation of the Dynamic Deductive System (DDS) developed in Ludlow and Živanović (2022). In this system, a proof can be seen as the evolution of a single logical formula, which we often envision and talk about as a tree, and which is changed upon every application of an inference rule. The starting point of a deduction of an argument is the conjunction of the premises, and the applications of inference rules step-wise transform this conjunction into the conclusion. In linguistic applications, like in this paper, the system is assumed to operate on the discourse Logical Form.

The inference rules of DDS are very simple operations guided by polarity. We work with a Boolean formal language, where the only sentential connectives are conjunction, disjunction and negation. In such a language, polarity is transparent: a constituent within the scope of an even/odd number of negations has positive/negative polarity. DDS deploys two inference rules which are sensitive only to polarity thus defined. Delete is a generalized Conjunction Elimination; it can eliminate a conjunct of positive polarity or a disjunct of negative polarity (8). Add is a generalized Disjunction Introduction; it can introduce a disjunct of positive polarity or a conjunct of negative polarity (9).[6]

[6] The reader familiar with *dictum de omni* and *dictum de nullo* will notice that our Delete and Add resemble these rules. In fact, they were inspired by them. Their positive polarity incarnations are instances of *dictum de omni*, while their negative polarity incarnations are instances of *dictum de nullo*.

DDS probably bears the greatest similarity to Peirce's *Beta System of Existential Graphs* (see e.g. Sánchez Valencia 1991), but it is also similar to Deep Inference systems (Brünnler and Tiu 2001; Guglielmi and Straßburger 2001), in particular

(8) a. Every black or gray dog is a friendly pet. $\forall x\colon \neg((Bx \vee Gx) \wedge Dx) \vee (Fx \wedge Px)$

b. ∴ Every black dog is pet. $\forall x\colon \neg(Bx \wedge Dx) \vee Px$

(9) a. Every dog is sleeping. $\forall x\colon \neg Dx \vee Sx$

b. ∴ Every black dog is eating or sleeping. $\forall x\colon \neg(Bx \wedge Dx) \vee (Ex \vee Sx)$

DDS has two further rules, Copy and Prune. Copy is an operation that takes something (the premise) from somewhere in the tree and copies it somewhere else (either overwriting the target, or joining the copy to it). Prune is a sort of anti-Copy, eliminating conflicts; in this paper, we ignore Prune, focusing on the simpler Copy. Clearly, not anything may be copied anywhere. We call the relation which governs what may be copied where *p-scope*.[7]

P-scope is a bit like the linguist's c-command, so let us remind ourselves how this relation is defined. Informally, a node in the tree *c-commands* another node iff it is possible to reach the latter from the former by first moving one node up, and then some (non-zero) number of nodes down the other branch. P-scope is very similar, only that we may move higher up than the first node we encounter. In principle, we may move any number of nodes up, but under certain conditions. At first, we hold a positive polarity pass, which allows us to move up into conjunctions. Moving into a negation reverses the polarity of the pass, and a negative polarity pass allows visiting disjunctions.

Let us provide a couple of simple examples of the p-scope powered Copy. In (10), we want to Copy *Dx* next to *Bx*. The initial positive polarity pass allows us to carry *Dx* into the conjunction above it, and the subsequent descent into *Bx* is condition-free, so *Dx* p-scopes over *Bx* and may thus be Copied next to it. In (11), the first node on the path from *Dx* to *Bx* is a negation, which reverses the initial positive polarity pass into the negative one. It is therefore possible to carry *Dx* through the disjunction and eventually reach *Bx*.

Calculus of Structures (Guglielmi 2007). The major difference is in treatment of negation. Deep Inference systems work on negation normal forms, making them suited for computer science applications; DDS pays great attention to polarity, making it suited for natural language applications.

[7] "P-scope" is short for "premise scope," and it governs application of more than just Copy. If α p-scopes over β, then it can function as a premise of *any* rule targeting β. Consider the following instance of Disamis from Figure 3. In modern terms, we are essentially applying Modus Ponens on $Dx \Rightarrow Ax$ and Dx to produce Ax. Crucially, this generalized version of Modus Ponens may be applied because $Dx \Rightarrow Ax$ p-scopes over Dx.

(i) Some dog is mortal. Every dog is an animal. ∴ Some animal is mortal.

(ii) a. $(\forall x\colon \neg Dx \vee Ax) \wedge (\exists x\colon Dx \wedge Mx)$

b. $(\forall x\colon \neg Dx \vee Ax) \wedge (\exists x\colon Ax \wedge Mx)$

(10) a. Some dog is black. $\exists x\colon \underline{Dx} \wedge \underset{\text{Copy}}{\underline{Bx}}$

b. ∴ Some dog is a black dog. $\exists x\colon Dx \wedge (\overline{Bx \wedge Dx})$

(11) a. Every dog is black. $\forall x\colon \neg\underline{Dx} \vee \underset{\text{Copy}}{\underline{Bx}}$

b. ∴ Every dog is a black dog. $\forall x\colon \neg Dx \wedge (\overline{Bx \wedge Dx})$

On the other hand, restrictions on the upward movement from the premise to the target ensure that the following invalid inferences cannot be deduced. In both examples below, the first node above the premise is a disjunction, which cannot be entered with the initial positive polarity pass.[8]

(12) a. Some cat or dog is black. $\exists x\colon (Cx \vee \underline{Dx}) \wedge \underset{*\text{Copy}}{\underline{Bx}}$

b. #. Some cat or dog is a black dog. $\exists x\colon (Cx \vee Dx) \wedge (\overline{Bx \wedge Dx})$

(13) a. It is raining, or we're on a trip. $\underline{R} \vee \underset{*\text{Copy}}{\underline{T}}$

b. #. It is raining, or it is raining and we're on a trip. $R \vee (\overline{R \wedge T})$

This concludes the brief outline of DDS and p-scope—we invite the reader to consult Ludlow and Živanović (2022: §6 and §7) for the complete story—but perhaps surprisingly, p-scope is useful outside DDS as well. In particular, we will see that it turns out to be a necessary condition for binding variables in a logical form.

Within the generative tradition, it is c-command which is usually assumed to govern variable binding; for example, the trace of a moved constituent is interpreted as a variable, and for this variable to be bound by the denotation of a quantified expression, that expression must c-command the trace at LF. It is well-known that once we move into the world of discourse anaphora, this assumption turns out to be incorrect. The failure to provide a principled extension of c-command has prompted many novel approaches to discourse anaphora, e.g. e-type anaphora (Evans 1977, 1980), DRT/FCS (Kamp 1981; Heim 1982), Dynamic Semantics (Groenendijk and Stokhof 1991) and their derivatives. We submit that p-scope is the correct and independently motivated extension to c-command required to deal with discourse anaphora. We outline our argument below and yet again invite the reader to consult Ludlow and Živanović (2022, §9–§11) for the full story.

The formal language of predicate logic we deploy in natural language logical forms is a bit non-standard: it does not deploy quantifier symbols. That is not to say that there is no quantification. There is, but

[8] The fact that (13a) entails (i) hints that under some conditions, it is the negation of the premise which may be copied. Indeed, this is the case when we can reach the target starting with a negative polarity pass; in Ludlow and Živanović (2022), we call the resulting p-scope *negative*. We disregard negative p-scope in this paper, as it has no bearing on pronoun resolution. We furthermore ignore the situations where the target is an ancestor or a descendant of the premise, and limit the discussion to what we call *relative* p-scope in the book. Consequently, whenever we write "p-scope" in this paper, we mean positive relative p-scope.

(i) It is raining, or it is not raining and we are on a trip. $R \vee (\neg R \wedge T)$

it arises exclusively via an interpretive rule we call Restricted Closure, and the crucial component of this rule is p-scope. For a quantifierless logical form to be interpretable via Restricted Closure, each variable x must be *restricted*, and it is restricted if the formula contains a *restrictor* for x—an occurrence of Px (for some monadic predicate P) which p-scopes over all other occurrences of x. If this condition is satisfied, Restricted Closure can then compute the location of the quantificational closure (the lowest node which contains all occurrences of x) and its type (existential/universal iff the polarity of the restrictor within the closure is positive/negative).

There is one further aspect of Restricted Closure important for our proposal. For us, linguist's LF and philosopher's logical form are one and the same thing, but linguist's LF is full of branching nodes with apparently no logical content. Alongside many semanticists working with event semantics (see e.g. Parsons 1990; Schein 2017), we assume that every LF branching node is interpreted as a conjunction—except when it is the locus of a universal closure yielded by Restricted Closure, when it is interpreted as a disjunction.[9]

Having no quantifier symbols in our logical forms makes it possible to see variable binding in linguistic rather than logical terms. In logic, variables are bound by quantifiers like $\forall x$; in linguistics, pronouns are bound by noun phrases. In our quantifierless logic, these noun phrases correspond to restrictors, so we can see the variable as being bound by its restrictor.

Now it should be becoming clear what our quantifierless logic and Restricted Closure have to do with pronoun resolution. One critical element of our analysis of prominence is that what is prominent for the resolution of a given pronoun depends in part on what antecedents are accessible from that pronoun, and in our system accessibility crucially relies on p-scope. As we will see in the next section, in our system resolving a pronoun amounts to deciding which variable it stands for in the logical form/LF, and this variable must be restricted for the resulting LF to be interpretable. Consequently, the antecedent of a pronoun must p-scope over the pronoun, or in other words, the potential antecedents of a pronoun are the noun phrases which p-scope over it. We develop this idea in more detail in the next section.

6. *Determining potential antecedents via p-scope*

Having familiarized ourselves with p-scope, we are now ready to deploy it in our reimplementation of Stojnić's (2021) pronoun resolution mechanism. In short, we will posit that a pronoun resolves to the closest referential expression p-scoping over it. To have this work, we will later modify Stojnić's assumption on *how* coherence relations and de-

[9] A variant of the system could additionally allow a branching node to be explicitly marked as a disjunction.

monstrative gestures affect the prominence ranking: not by modifying it explicitly, but by being integrated into different locations of the LF. Ultimately, we will suggest that it is the syntax which determines this location: if syntactician's LF and semanticist's logical form are one and the same, it is natural to assume that coherence relations are instances of functional heads, which come, by a core assumption of Minimalist Program, in a cross-linguistically fixed order. Note, however, that the mechanics of pronoun resolution does not hinge in any way on how the coherence relations obtain their location, only that they do.

Our first and simplest example, (14), will mainly serve to explain what resolving a pronoun actually amounts to in our approach. After the first sentence is uttered, the (much simplified) logical form of the discourse is Betty(x) $\wedge$ come-in(x). When the second sentence comes in, but before it is integrated into the discourse logical form, its logical form is sat-down(□), where the empty dotted box in the argument position is meant to indicate that the identity of the variable occurring there is as of yet unknown.[10] We assume that in this simple case, the new logical form is integrated into the discourse by a conjoining it to the root of the existing discourse logical form; we arrive at (14b), but with the dotted argument box still empty. It is now time to figure out the identity of the variable. There is a single candidate, x, and this candidate fulfills the requirement imposed by restrictedness: a monadic atomic formula containing this variable, i.e. Betty(x), p-scopes over sat-down(x).[11] (Starting at Betty(x) carrying a positive polarity pass, the p-scope easily passes through both conjunctions on the way to sat-down(x).)

(14) a.Betty came in. She sat down. b. c.

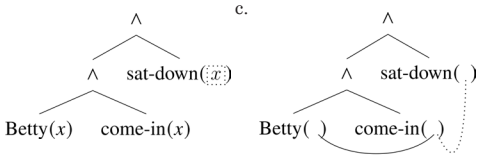

As a side note, we do not need to be realists about variable symbols. We could just as easily use Quine's (1981) *bonds*. With those, (14b) transforms into (14c). What we previously thought of as an unidentified variable is now simply an unlinked argument position, and pronoun

[10] In our system, variables are the only kind of an individual term, and thus the only possible argument of a predicate. We occasionally use individual constants, but they are only an abbreviation. For example, we would analyse names as predicates rather than individual constants.

[11] In generative syntax, an antecedent of a pronoun is usually understood to be a noun *phrase* (NP or DP), whereas our Restricted Closure requires a p-scoping *atomic* formula, which corresponds to the nominal head (N), or perhaps to the (nominal) root (√). The discrepancy is not as significant as it might seem at first. The Minimalist Program, for example, identifies heads and phrases anyway. We leave an explicit account of this detail of the isomorphism between the philosopher's logical form and the linguist's Logical Form for further research.

resolution is nothing but linking of a new argument position to some existing argument position (as indicated by the dotted arc).

Also remember that we are using the quantifierless format of predicate logic. When the first sentence is interpreted in isolation, existential closure over *x* applies at the root of its logical form. After integrating the logical form of the second sentence into the discourse logical form, that existential closure is automatically "lifted" to apply at the conjunction of both sentences.

Our next example involves two variables, *x* and *y*. It is easy to see both woman(*x*) and girl(*y*) p-scope over greeted(□,□). How do we decide which variable goes where, then? Assuming that the unidentified variables are resolved in argument order (we will provide a better motivation in a moment), this question can be rephrased thus: why does the first argument resolve to *x* rather than *y* (the second argument takes the leftover *y*). The clue comes from the syntactic structure of the first sentence, reflected in the logical form. Observing the tree representation of the formula in (15c), it is intuitively clear that woman(*x*) is closer to greeted(□,□) than girl(*y*) is. We will formalize this intuition in section 7. As far as the example below is concerned, it ultimately boils down to the syntactic fact that the subject 'a woman' is positioned higher (even graphically) than the object 'a girl'. In syntactician's parlance, the subject asymmetrically c-commands the object (for the reader unfamiliar with the concept, we will define it in section 7).

(15) a. A woman met a girl. She greeted her.

b. (woman(*x*) ∧ (girl(*y*) ∧ met(*x*,*y*))) ∧ greeted (□,□)

c.

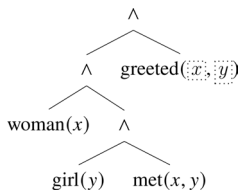

Syntactic height can also provide a better reason for identifying the first argument of greeted(□,□) first. We merely have to assume that the syntactically higher pronoun is resolved first. In (15a), the first and the second argument of 'greeted' are the subject and the object, respectively, and the subject is higher.[12]

The account in terms of syntactic height has a distinct advantage over the account in terms of grammatical roles (such as subject and ob-

[12] Alternatively, if you are worried about the fact that 'she' and 'her' do not contribute to the logical form in the same way as 'woman' and 'girl' do (there are no separate nodes for them), consider the event semantics decomposition of the second sentence, which mirrors the internal layered structure of the verb phrase in syntax: Agent(*x*,*e*) ∧ (greeting(*e*) ∧ Theme(*y*,*e*)). Again, *x* (occurring as an argument of theta role predicate Agent) winds up higher than *y* (occurring as an argument of theta role predicate Theme).

ject) where those roles are not represented by structural syntactic position. Being more general, it automatically applies to structural height differences (in a syntactic tree) that reflect grammatical phenomena. For example, it correctly predicts that topicalization and word order will affect pronoun resolution: "this preference for referents introduced by noun phrases in subject position is a grammatical feature of English, a reflection of the fact that English is a so-called subject-prominent language. This is not a feature that is universally shared across languages. Other languages, topic-prominent ones, grammaticize prominence with specialized morphemes, like topic markers; languages with flexible syntax utilize word order" (Stojnić 2021: 59).

Furthermore, this account helps explain Stojnić's (2021: 50) observation that "the presence of a deictic gesture is hard to override." We only need to add one, very natural assumption—that the pronoun and its accompanying deictic gesture form a morphosyntactic constituent. This guarantees that they are close enough in the logical form for the demonstrative gesture to always participate in and win the pronoun resolution race.[13] It participates because $y = b$, contributed by the deictic gesture of pointing to Betty, is conjoined to sat-down(…) and therefore p-scopes over it; and it wins because $y = b$ is certainly closer to sat-down(…) than woman(x) is—you cannot get any closer than the sister node![14]

(16) a. A woman came in. She [pointing at a cat, Betty] sat down.

b.

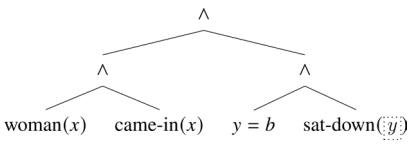

We will formalize the notion of closeness in the following section, where we will see that it can be defined deploying the standard syntactic re-

[13] Stojnić (2021) implicitly makes an analogous assumption. In general, logical forms such as (1b) would not yield the intended result without the update introduced by the deictic gesture (⟨π0*b*⟩) immediately preceding the stack lookup performed by the pronoun (@she). (In (1b), the full contribution of the pronoun is ⟨π0@she⟩, with π0 pushing the retrieved individual back to the stack to the subject position 0, but note that the stack lookup happens before this reintroduction, so the effects of the deictic gesture and the pronoun are adjacent after all.)

[14] In (16b), the simplifications we make for clarity actually get in the way of easily seeing who participates in the pronoun resolution race. To be restricted, whichever variable we put into sat-down(…) has to have a restrictor, which is, by definition, a monadic atomic formula p-scoping over it and containing an occurrence of the very same variable. woman(x) is clearly a candidate. However, came-in(x) is not: although it *looks* like an atomic formula, it is merely an abbreviation for event semantics decomposition came-in(e) ∧ Agent(x,e), and Agent(x,e) within this decomposition does not count, as it is not monadic. On the other hand, $y = b$ is a candidate, even if it does not look like a monadic formula. We should have really written something like $D_{\text{Betty}}(y)$, with D_{Betty} a one-off predicate created by the deictic gesture.

lation of (asymmetric) c-command, but before we do this, let us focus on the contribution of p-scope a bit more. In the examples above, all noun phrases occurring in the previous discourse p-scoped over the pronoun and were therefore all potential antecedents. This is not always the case. For example, when the new sentence is integrated into the discourse by a disjunction, an indefinite from the previous discourse cannot expand its scope to cover the pronoun, because p-scope cannot traverse the sequence of conjunction and disjunction (whatever polarity pass it holds at the start of the sequence).

(17) a. #Either there's a poltergeist in this house, or it is hiding very well.

b.

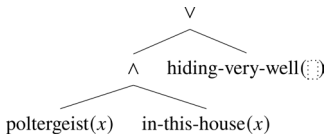

The situation is completely different in the well-known bathroom examples. Here, the negation contributed by the negative determiner 'no' intervenes between the conjunction and the disjunction, flipping the polarity of the pass just in time to allow for the p-scope to enter the disjunction at the root, thereby making it possible for poltergeist(*x*) to p-scope over hiding-very-well( ) containing the pronoun.

(18) a. Either there's no poltergeist in this house, or it is hiding very well.

b.

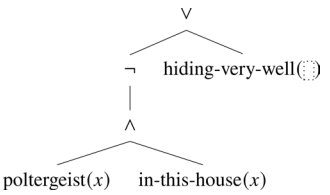

Finally, we see that even in the absence of this negation, i.e. with an indefinite noun phrase, the discourse can be felicitous if the antecedent can be found in the discourse preceding the disjunctive sentence. Crucially, 'my bag' wins over 'a poltergeist' despite the latter being closer (both linearly and geometrically) to the pronoun.

(19) a. I am looking for [my bag]$_{b}$. Either there's [a poltergeist]$_{p}$ in this house, or [it]$_{b,*p}$ is hiding very well.

b.

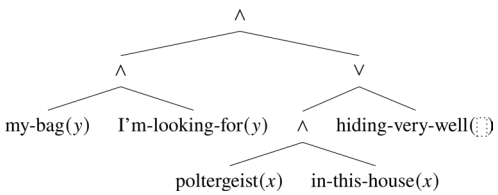

In Ludlow and Živanović (2022), we provide further similar examples of p-scope carving out the set of possible antecedents of a pronoun. Specifically, we deploy it in a novel approach to the donkey anaphora, showing both why the indefinite can receive the (universal) wide scope and why a negative determiner cannot receive it. In this paper, we will provide one more example of this, but we can only do that once we have introduced the role of coherence relations in our proposal. However, we turn to the definition of closeness first.

7. *A syntactic approach to prominence*

We have argued that p-scoping over the pronoun is a necessary condition for being its antecedent, but it is clearly not a sufficient condition, because many antecedent candidates might p-scope over a given pronoun. We have suggested that the closest candidate wins, but what precisely do we mean by "the closest"? Furthermore, we do not want to merely single out one candidate as the most prominent, but rather order all candidates by prominence, which allows one to consider the candidates in order of prominence until settling for the one which is suitable in the sense that it agrees with the pronoun in gender and number, and does not cause a Principle B violation—same as when selecting the suitable candidate from Stojnić's (2021) stack. The goal of this section is therefore to develop relation "closer to the given pronoun than" on the set of all the potential restrictors of the pronoun.

The notion of closeness, usually going under the name of *locality*, is a ubiquitous feature of generative syntax.[15] As we consider our project to be, at least broadly, a part of generative syntax, it is natural to check whether the syntactician's notion of locality can be applied to work out prominence. We will see that this is indeed the case.

The relation linearly ordering the set of nodes c-commanding a given node in terms of closeness to this node is c-command itself. Given α_1 and α_2 c-commanding π, α_1 is closer to π than α_2 is iff α_2 c-commands α_1. This works because it is always the case that one of α_i c-commands the other but not vice versa; in other words, one of α_i always *asymmetrically c-commands* the other.[16]

[15] The notion of locality, in one form or another, lies at the heart of many notions in generative syntax, notably Relativized Minimality (Rizzi 1990, 2001) and Linear Correspondence Axiom (LCA) (Kayne 1994). In fact, our definition of closeness and thus prominence will deploy LCA.

[16] In this section, we stick to the convention of marking the pronoun as π and the potential antecedents as α_i, where the index on α depends on the (intended) proximity to the pronoun: the smaller the index, the closer the potential antecedent. A dashed line stands for any number of branching nodes.

(20)

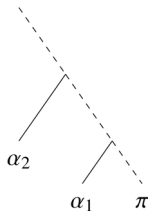

However, the closeness relation defined above will not do for our purposes, because it only orders the nodes c-commanding π, while we need to order the nodes p-scoping over it.[17] For example, α_1 and α_2 below (ignore β_1 and β_2 for now) might (depending on the content of the tree) both p-scope over π, but they cannot be ordered by asymmetric c-command, because neither of them c-commands the other. They are simply too deeply embedded to do so.[18]

(21)

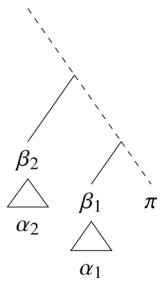

At this point, one might be tempted to rewrite the above definition of the closeness relation by substituting "p-scope" for "c-command"—so that if we have α_1 and α_2 p-scoping over π, α_1 is closer to π than α_2 is iff α_2 p-scopes over α_1—but this will not do, because α_1 and α_2 can easily p-scope over each other (for example, if all the branching nodes in (21) turn out to be conjunctions).

We therefore stick to (asymmetric) c-command, and proceed in the fashion almost identical to the formulation of Kayne's (1994) Linear Correspondence Axiom (LCA). The idea is to say that if we have phrases β_1 and β_2 c-commanding π, and β_2 (asymmetrically) c-commands β_1, as shown in (21) above, then not only is β_1 closer to π than β_2 is, any node dominated by β_1 is closer to π than any node dominated by β_2 is. We have already encountered such a situation in the demonstrative gesture example (16a), where $\alpha_1 = \beta_1 = [y = b]$, $\alpha_2 = [\text{woman}(x)]$ and $\beta_2 = [\text{woman}(x) \wedge \text{came-in}(x)]$, and our new definition of closeness correctly predicts the demonstrative to be closest node to the pronoun, and therefore the most prominent with respect to the pronoun.

[17] Actually, we will end up ordering all nodes not dominating or dominated by π.

[18] A triangle stands for an arbitrary phrase containing (but not equal to) the material indicated at its bottom.

Of course, we also need to order the nodes within each of β_1 and β_2. The situation is schematized in (22).[19] We have also already encountered such a situation in example (15), and that example indicates that we have to see the higher candidate, i.e., the one asymmetrically c-commanding the other, as closer to the pronoun.

(22)

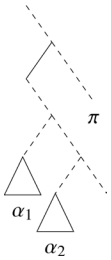

And if none of α_i c-commands the other? Well, we could continue recursively applying the LCA idea, but that would merely reimplement LCA. At this stage, the location of π becomes irrelevant, as we simply want to order the nodes by syntactic height, which is precisely what LCA does. We therefore choose to deploy LCA itself, and say that when α_1 and α_2 are both contained in the same node β c-commanding π, α_1 is closer to π than α_2 is iff LCA, applied to LF, linearizes α_1 before α_2.

Stating syntactic/hierarchical closeness in terms of linear precedence might sound strange, but it is not really, as we assume that linearization is indeed carried out by LCA and is therefore based exclusively on hierarchical information. Also note that we do not claim that closeness is based on linear order in the surface syntax. In the service of closeness, LCA is not applied to the state of the syntactic tree at Spell-Out; it is applied to LF.

That said, the deployment of LCA (and even of the LCA linearization idea) carries a major consequence. We cannot apply LCA to any tree structure and expect a linearized output. LCA requires a very specific form of the input tree. In particular, Kayne (1994) argues that it forces the natural language syntax into the X-bar format (23a), where the specifier asymmetrically c-commands the complement.[20] This is perfectly fine for our project, which we see as a part of generative grammar anyway, and which we see as operating on Logical Forms of natural language expressions. Furthermore, it is clearly impossible to develop a notion of closeness which only depends on the geometry of the tree

[19] Observe that for any two nodes α and π such that neither dominates the other, there is a unique node β which contains α and c-commands π. This is the node lying on the upward leg of the journey from α to π just below the top of the path. The schemas in (21) and (22) therefore cover all the possible configurations of α_1 and α_2.

[20] Note that in the X-bar format, the intermediate projection X′ is not a phrase and consequently cannot act as β_1 in (21). Without this standard syntactic assumption, one could not implement the linear ordering.

and works on any tree. For example, how could we say which of α and β below is closer to π *in absence of the linear order implied by the graphical representation*? We cannot, but as we have seen, we also do not have to. Our notion of closeness crucially requires that philosopher's logical form is the same as the linguist's Logical Form, and we consider this to be another argument in favor of proof-theoretic semantics based on LF.

(23) a.

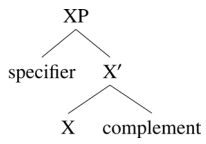

b.

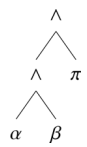

8. *Coherence Relations*

In Stojnić's (2021) account, coherence relations influence the prominence ranking directly. It is a part of the conventional meaning of Result that it pushes the subject of the first sentence to the top of the stack, and it is a part of the conventional meaning of (one incarnation of) Explanation that it does that to the object.

Our approach differs in that we believe, on independent grounds, that such coherence relations are reflected in the structural hierarchy of syntactic trees. Thus, questions about stack position are beside the point. All the relevant information is already encoded in the geometry of the tree.

Specifically, the idea is that each coherence relation integrates the new discourse material into a dedicated location in the syntactic tree. The new discourse material contains the pronouns undergoing resolution. These pronouns will therefore wind up in different locations for different coherence relations, and as it is the location which determines the set of p-scopers and their prominence order, the pronouns will be resolved differently for different coherence relations.

Let us illustrate this using one of Stojnić's canonical examples. For the result reading, we integrate the second sentence at the top of the first one, as shown in (25a). This places both potential antecedents, 'John' and 'Tim', into the same constituent c-commanding the result clause, i.e. the situation corresponds to schema (22) from the previous section. The higher potential antecedent, 'John', is therefore closer to the pronouns, and thus more prominent with respect to the pronouns. Consequently, the higher subject pronoun ('he') will resolve to the higher candidate ('John'), and the lower object pronoun ('him') will receive the leftovers ('Tim'). In the explanation reading, the situation is different. Here, we integrate the second sentence in the middle of the first one, between the subject and the object position. Consequently, we are in a situation schematized by (21). Here, the more prominent candidate is the one occurring in the lower position ('Tim' in vP); the higher pronoun ('he') therefore resolves to 'Tim', while the lower pronoun resolves to the remaining candidate, 'John'.

(24) John was disappointed with Tim.
a. (So) he$_j$ fired him$_i$. (Result)
b. (Because) he$_i$ disobeyed him$_j$. (Explanation)

(25) a.

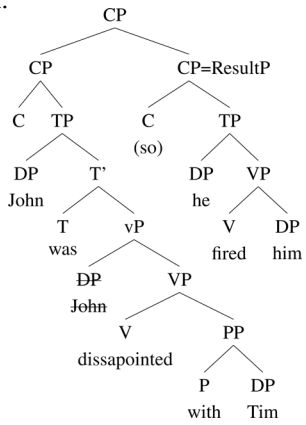

b.

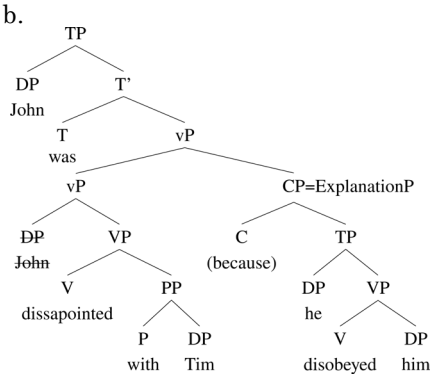

Given the integration position of the second sentence, our account yields the correct predictions. However, the real question is whether there is any independent evidence for this position. We believe that there is.

Haegeman (2012) distinguishes two broad classes of adverbial clauses, peripheral and central. Each class comes with its own set of syntactic properties. Peripheral adverbial clauses are discourse oriented and behave much like matrix clauses. For example, they may carry illocutionary force and contain speaker-related modal markers, and cannot occur in the scope of matrix negation. Conversely, central adverbial clauses are event oriented and behave unlike matrix clauses. They cannot carry illocutionary force or contain speaker-related modal markers, but may occur in the scope of matrix negation. Haegeman (2012) proposes that (as a first approximation) peripheral adverbial clauses are adjoined to the CP (complementizer phrase, found at the root of a sentence), while central adverbial clauses are adjoined to vP (the outer layer of a verb phrase) or TP (tense phrase), found in the middle of the extended verbal projection.

Each type of adverbial clause is introduced by a dedicated connective. However, a single connective typically introduces two kinds of adverbial clauses, one central and one peripheral. For example, 'so (that)' introduces both the central purpose clause and the peripheral result clause; (24a) is the instance of the latter. And 'because' introduces both the central event cause clause and the peripheral rationale clause; (24b) is the instance of the former. Below, we apply some of Haegeman's (2012) tests to our instances of result and explanation: (a) (directive) illocutionary force; (b) speaker-related modal marker 'probably'; (c) matrix negation. These tests clearly show that (24a) and (24b) are instances of a peripheral and central adverbial clause, respectively, thereby ad-

ditionally justifying the syntactic structures proposed in (25).

(26) a. You are disappointed with Tim, so fire him!
b. John was disappointed with Tim, so he probably fired him.
c. John was not happy with Tim, so he fired him (*but so …).

(27) a. *John is disappointed with you, because disobey him!
b. *John was disappointed with Tim, because he probably disobeyed him.
c. John was not disappointed with Tim because he disobeyed him (but because …).

The situation is somewhat trickier with Stojnić's example illustrating that Explanation may be either subject-based or object-based. Here, it is unclear whether the subject-based explanation should be considered an instance of a peripheral adverbial clause. It is perhaps better to assume that Explanation may be integrated in different central positions; remember that according to Haegeman (2012), central clauses may be adjoined either to vP or to TP. Our pronoun resolution mechanism yields correct results if we adjoin (28a) to TP, as shown below; (28b) should be adjoined to vP, resulting in a structure analogous to (25b). We leave the detailed syntactic investigation to further research.

(28) The city council denied the demonstrators a permit.
a. They feared violence.
b. They advocated violence.

(29) a.

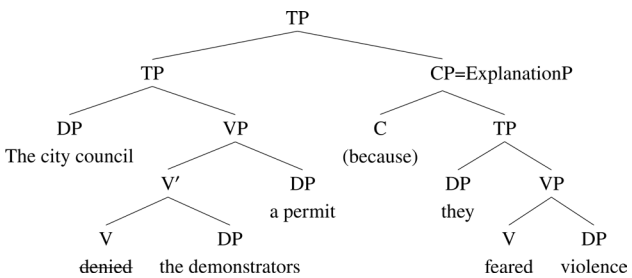

Another example discussed by Stojnić is (30). We are already familiar with our derivation of the result reading, what about Parallel, where 'him' resolves to the object? The strategy is the same as for Explanation. We get the correct prediction if we stick the new material between the subject and the object, but the real issue is to provide some independent evidence for such a move.

(30) Phil tickled Stanley,
a. (so) Liz poked him$_p$. (Result)
b. (while) Liz poked him$_s$. (Parallel)

The syntax of parallel sentences is quite involved, so we cannot go into much detail here. At the minimum, the syntactic structure involves a contrastive topic ('Phil' vs. 'Liz'), and probably a contrastive focus as

well ('tickled' vs. 'poked') (Kehler 2002; Hendriks 2004). The presence of these is made quite clear in Slovenian, a free word order language, where the preferred word order in the first clause of the translation of (30b) is as shown in (31), with contrastive topic beginning, and contrastive focus ending the sentence. On these grounds, something like the structure in (32) seems a likely representation of (30b).

(31) Filip je Stankota požgečkal, Liza ga je pa žoknila.
Phil$_{NOM}$ is Stanley$_{ACC}$ tickled, Liz him is $_{PARTICLE}$ poked.

(32) (parallel)

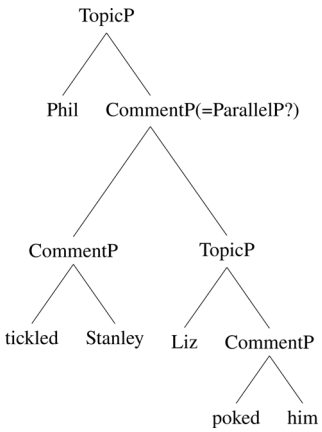

Whatever the details, one thing is certain: Parallel is not sensitive to subjects and objects *per se*. The story again revolves around syntactic height, and hinges on the fact that (contrastive) topics are positioned very high in the structure, within the split CP (Rizzi 1997, 2004). Look again at Slovenian, where it is easy to topicalize the object. In (33), 'Stanley', being topicalized, is out of the pronoun resolution game. The covert nominative subject pronoun pro therefore refers to 'Phil', the subject of the first sentence. Even English exhibits similar effects. In (34), we topicalize the time adverb, retaining both subject 'Phil' and object 'Stanley' in the comment, competing in the pronoun resolution. The higher one wins.

(33) Stankota je Filip požgečkal, Markota pa je pro$_f$ žoknil.
Stanley$_{ACC}$ is Phil$_{NOM}$ tickled, Marc$_{ACC}$ $_{PARTICLE}$ is pro poked.
'Phil ticked Stanley, and (in a parallel fashion) he poked Marc.'

(34) Today, Phil tickled Stanley. Tomorrow, he$_p$ will poke him$_s$.

In the examples above, all the noun phrases found in the first sentence were possible antecedents of the pronouns of the second sentence. However, our p-scope based system is not that permissive in general. Until now, we have only used examples where all syntactic branches were interpreted as conjunctions (and all quantification was consequently existential). We now turn to examples containing universals, which introduce a negation (and therefore a disjunction and universal quanti-

fication) and thereby prevent certain noun phrases p-scoping into the second sentence. In (35),[21] the coherence relation is the familiar Result, which integrates the new material at the top of the first sentence. The result is therefore conjoined above the (implicit) universal quantifier. Consequently, p-scope cannot pass through the sequence of disjunction and conjunction.

(35) #Every candidate walked to the stage. He sweated profusely.

a.

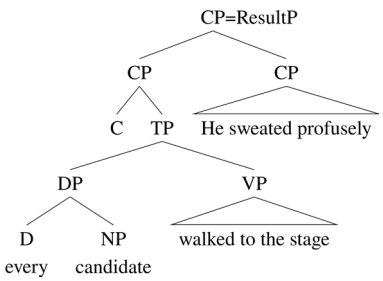

b.

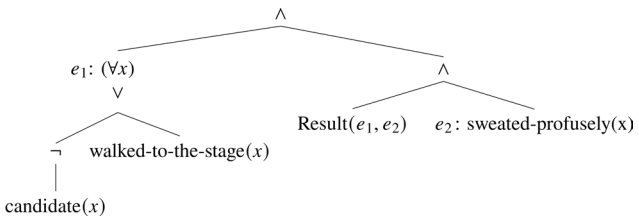

The above example is usually presented along an example of telescoping—a situation where binding out of a universal is exceptionally possible. Roberts (1989) suggests that these examples work because the sentences form a continuous narrative. In a coherence relation approach, this translates to the two sentences being related via Narration. It seems sensible to position Narration in the vicinity of tense; at the very least, narratives require that events unfold in temporal order (Kehler 2002; Wolf and Gibson 2006). Furthermore, in their feature-driven reimplementation of Quantifier Raising, Beghelli and Stowell (1997) argue that universals move into Dist(ributive)P. They locate this projection above NegP (if present), which is typically seen as occurring above TP (see e.g. Haegeman 1995). We thus arrive at an LF where the continuation of the narrative is positioned lower than the

[21] Typically, the failure of a universal from the first sentence to bind the pronoun from the second sentence is exemplified by having 'He was tall' as the second sentence. We avoid using this second sentence because there is no coherence relation linking it to the previous discourse. This is precisely the situation other authors need to discuss, but it will not work for us. We need an example where a coherence relation is imaginable, but the example is still ungrammatical. Then and only then can we conclude that the ungrammaticality is due to the failure to p-scope over the second sentence.

universal DP. Consequently, there is no fatal combination of a disjunction and conjunction on p-scope's path from candidate(x) to the second sentence so candidate(x) turns out to p-scope over the pronoun in the second sentence, allowing it to resolve to variable x.

(36) Every candidate walked to the stage. He shook dean's hand …

a.

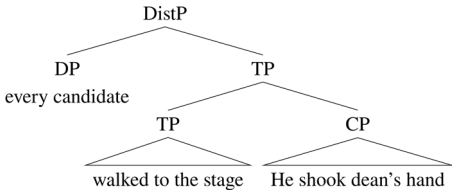

b.

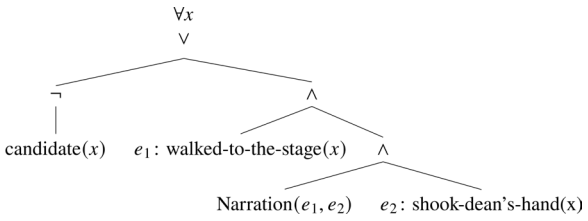

We intend to investigate all the intricacies of telescoping and related modal subordination through the prism of p-scope in a later paper.

Coherence relations started their life in pragmatics (Hobbs 1979) as non-linguistic entities. Stojnić (2021) argues that they are in fact linguistic items, represented in the logical form. We want to make the final step in identifying their nature and propose they are syntactic items, namely functional heads.

Indeed, upon a closer inspection, coherence relations seem intimately connected to the structure of a sentence. This is (explicitly or implicitly) acknowledged even in the pragmatics research. Kehler (2002) is an interdisciplinary work including both pragmatics and syntax, and within purely pragmatics research, it is a typical strategy (see e.g. Knott 1996; Wolf and Gibson 2006) to identify a coherence relation based on which explicit connective (or more generally, cue phrase) it deploys; in absence of an explicit connective, the identification relies on the judgment whether the discourse retains its meaning when a particular explicit connective is added. Furthermore, syntacticians study coherence relations as well, even if they do not call them that. Above, we have occasionally relied on Haegeman's (2012) study of adverbial clauses. Her typology of adverbial clauses (Haegeman 2012: 164) makes it quite clear that it is the same subject matter that is being studied (but from another, syntactic, perspective). We are certain that a comparison between the pragmatics literature on coherence relations, and syntactic cartographic studies (Cinque 1999, and subsequent works) would yield many matches.

In this section, we have integrated the new discourse material into the previous sentence using good old-fashioned adjunction (see e.g. Adger 2003). However, given the recent explosion of the functional structure of a sentence, powered by the cartographic studies, adjunction is slowly but surely becoming obsolete. This is why we are certain that, at the end of the day, coherence relations can be (formally) subsumed under the notion of functional heads, even if perhaps not in one-to-one fashion (witness the complicated situation with Parallel above).

The upshot of this is that in addition to having a way to represent prominence relations in LF syntax, we also have a promising tool for investigating hierarchical positions of phrases within linguistic structures, and more importantly it appears to be a tool that functions cross-linguistically. Whether the tool works for all structures in all languages is an open question for now.

9. *Conclusion*

We began this investigation with a very important observation from Stojnić (2021) that, contrary to standard assumptions, discourse anaphora can be resolved with linguistic resources alone, given basic assumptions about the nature of linguistic objects and how linguistic objects can be entered into the discourse representations. Stojnić offered that the relevant data structure for representing such information might take the form of stacks. We have argued that a perfectly acceptable alternative data structure would be the trees that are commonly deployed in linguistic theory.

We have also argued that there are already good reasons to believe that the relevant accessibility and prominence relations necessary to track discourse antecedents are compatible with current theories of generative grammar—indeed, they can be carried out once we deploy our independently motivated proof-theoretic notion of p-scope. Finally, we have argued that coherence relations can be seen as affecting accessibility and prominence in virtue of their representation in Logical Form—in particular by how they are reflected in the ordering of functional heads. In short, we believe that one can give a syntactic account of discourse prominence, and that the relevant prominence relations can be grounded in the LF representations of contemporary generative linguistics.

References

Adger, D. 2003. *Core Syntax: A Minimalist Approach*. Oxford and New York: Oxford University Press.

Beghelli, F. and Tim S. 1997. "Distributivity and Negation: the Syntax of *each* and *every*." In A. Szabolcsi (ed.). *Ways of Scope Taking. Vol. 65. Studies in Linguistics and Philosophy*. Dordrecht: Kluwer Academic Publishers 71–107.

Brünnler, K. and Alwen F. T. 2001. "A Local System for Classical Logic." In R. Nieuwenhuis and A. Voronkov (eds.). *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR). Vol. 2250. Lecture Notes in Computer Science*. Springer-Verlag, 347–361. DOI: 10.1007/3-540-45653-824. URL: http://cs.bath.ac.uk/ag/kai/lcl-lpar.pdf.

Carnap, R. 1956. *Meaning and Necessity: A study in semantics and modal logic*. Second edition. Chicago: University of Chicago Press.

Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge: The MIT Press.

Chomsky, N. 1977. *Essays on form and interpretation*. New York: North-Holland.

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

Cinque, G. 1999. *Adverbs and Functional Heads*. Oxford Studies in Comparative Syntax. Oxford and New York: Oxford University Press.

Evans, G. 1977. "Pronouns, Quantifiers and Relative Clauses." *Canadian Journal of Philosophy* 7 (3): 467–536.

Evans, G. 1980. "Pronouns." *Linguistic Inquiry* 11 (2): 337–362.

Groenendijk, J. and Stokhof, M. 1991. "Dynamic predicate logic." *Linguistics and Philosophy* 14: 39–100.

Grosz, B. J., Joshi, A. K. and Weinstein, S. 1995. *Centering: A Framework for Modelling the Local Coherence of Discourse*. Harvard Computer Science Group Technical Report TR-18–95. URL: http://nrs.harvard.edu/urn-3:HUL.InstRepos:25680326.

Grosz, B. J. and Sidner, C. L. 1986. "Attentions, intentions and the structure of discourse." *Computational Linguistics* 12 (3).

Guglielmi, A. 2007. "A System of Interaction and Structure." *ACM Transactions on Computational Logic* 8 (1:1): 1–64. DOI: 10.1145/1182613.1182614. URL: http://cs.bath.ac.uk/ag/p/SystIntStr.pdf.

Guglielmi, A. and Straßburger, L. 2001. "Non-commutativity and MELL in the Calculus of Structures." In L. Fribourg (ed.). *Computer Science Logic (CSL). Vol. 2142. Lecture Notes in Computer Science*. Springer-Verlag, 54–68. DOI: 10.1007/3-540-44802-0"¨5. URL: http://cs.bath.ac.uk/ag/p/NoncMELLCoS.pdf.

Haegeman, L. 1995. *The Syntax of Negation*. Cambridge University Press.

Haegeman, L. 2012. *Adverbial Clauses, Main Clause Phenomena, and the Composition of the Left Periphery*. Oxford University Press.

Harman, G. 1970. "Deep structure as logical form." *Synthese:* 275–297.

Heim, I. 1982. "The Semantics of Definite and Indefinite Noun Phrases." PhD thesis. University of Massachusetts, Amherst. 2011 Edition, typeset by Anders J. Schoubye and Ephraim Glick. URL: https://semanticsarchive.net/Archive/jA2YTJmN/HeimDissertationwithHyperlinks.pdf.

Hendriks, P. 2004. "Coherence relations, ellipsis and contrastive topics." *Journal of Semantics* 21 (2): 133–153.

Hobbs, J. R. 1979. "Coherence and Coreference." *Cognitive Science* 3: 67–90.

Kamp, H. 1981. "A theory of truth and semantic representation." In J. Groenendijk, T. Janssen and M. Stokhof (eds.). *Formal Methods in the Study of Language: Part 1*. Amsterdam: Mathematisch Centrum, 277–322.

Kayne, R. S. 1994. *The Antisymmetry of Syntax*. Vol. 25. Linguistic Inquiry Monographs. Cambridge: The MIT Press.

Kehler, A. 2002. *Coherence, Reference, and the Theory of Grammar*. CLSI Publications.

Knott, A. 1996. "A Data-Driven Methodology for Motivating a Set of Coherence Relations." PhD thesis. University of Edinburgh.

Larson, R. K. 1990. "Double objects revisited: Reply to Jackendoff." *Linguistic inquiry* 21 (4): 589–632.

Ludlow, P. and Živanović S. 2022. *Language, Form, and Logic: In Pursuit of Natural Logic's Holy Grail*. Oxford: Oxford University Press.

Parsons, T. 1990. *Events in the Semantics of English. Current Studies in Linguistics 19*. Cambridge, USA and London, UK: The MIT Press.

Quine, W. V. O. 1981. *Mathematical logic*. Harvard University Press.

Rizzi, L. 1990. *Relativized Minimality*. Cambridge: The MIT Press.

Rizzi, L. 1997. "The fine structure of the left periphery." In L. Haegeman (ed.). *Elements of Grammar*. Dordrecht: Kluwer, 281–337.

Rizzi, L. 2001. "Relativized Minimality Effects." In M. Baltin and C. Collins (eds.). *The Handbook of Contemporary Syntactic Theory*. Cambridge USA & Oxford UK: Blackwell Publishers, 89–110.

Rizzi, L. 2004. "Locality and Left Periphery." In A. Belleti (ed.). *Structures and Beyond: The Cartography of Syntactic Structures. Vol. 3. Oxford Studies in Comparative Syntax*. Oxford and New York: Oxford University Press, 223–251.

Roberts, C. 1989. "Modal subordination and pronominal anaphora in discourse." *Linguistics and Philosophy* 12 (6): 683–721.

Sánchez Valencia, V. M. 1991. "Studies on Natural Logic and Categorial Grammar." PhD thesis. University of Amsterdam.

Schein, B. 2017. '*And': Conjunction Reduction Redux*. Cambridge and London: The MIT Press.

Stojnić, U. 2021. *Context and coherence: The logic and grammar of prominence*. Oxford University Press.

Wolf, F. and Gibson, E. 2006. *Coherence in Natural Language: Data Structures and Applications*. The MIT Press.