CROATIAN
JOURNAL
OF PHILOSOPHY

Vol. XXIII · No. 69 · 2023

Articles

Introduction

The COVID-19 pandemic stopped us for two years in organizing our longstanding course Philosophy of Language and Linguistics which started in 2005 and is always held at the Interuniversity Center (IUC) in Dubrovnik. We continued in September 2022. Selected papers from all our conferences are customarily printed in the Croatian Journal of Philosophy.

In 2022, the course was primarily dedicated to the discussion of the book by Una Stojnić Context and Coherence *(Oxford University Press 2021). Part of the course included the discussion of Fabrizio Cariani's book,* The Modal Future. A Theory of Future-Directed Thought and Talk *(Cambridge University Press 2021) The discussion of Cariani's book was scheduled for 2020, the year the course was not held, so we decided to have his book discussed in 2022. The first five papers in this volume are on Stojnić's book while a jointed paper by Cariani and Glanzberg on Cariani's work follows.*

Una Stojnić gives a valuable précis of her book Context and Coherence. *The book develops and defends a thoroughly linguistic account of the meta-semantics of context-sensitivity: the interpretation of context-sensitive expressions is fully determined by linguistic rules, discourse conventions. If this is right, the dominant, extra-linguistic account must be rejected. This précis also outlines some other key themes in* Context and Coherence. *The papers that follow are different answers given and questions posed to Stojnić's provocative claim.*

Peter Pagin in his contribution "Linguistic conventions or open-ended reasoning: Some questions for Una Stojnić" is mainly concerned with Stojnić's strong claim that linguistic phenomena related to prominence and coherence, in particular the interpretation of pronouns, are governed by linguistic conventions and are not pragmatic in nature. Pagin's presented views are opposite to Stojnić's. Pagin also questions Stojnić's view that coherence relation has priority over the interpretation of pronouns. Magdalena Kaufmann's paper "From coherence relations to the grammar of pronouns and tense" argues against Stojnić's strong claim that the content of linguistic utterances is determined by the rules of natural language grammar more than it is generally assumed. Kaufmann also takes a close look at the empirical evidence from English and Serbian that Stojnić offers in support of her position on coherence. Kaufmann adds examples from German and Japanese. She argues that there is no compelling evidence for the assumption that coherence relations directly determine the resolution of pronouns. Inten-

tionalism is the view that a demonstrative refers to something partly in virtue of the speaker intending it to refer to that thing. The article "Intentionalism and the natural interpretation of discourses" by Aleksandru Radulescu is a critical assessment of Stojnić's contrary claim that the natural interpretation of demonstratives is that they refer to the objects but not by speakers' intentions. Radulescu further argues that many phenomena presented by Stojnić can be explained from an intentionalist point of view. Sašo Živanović and Petar Ludlow in their contribution "The Syntax of Prominence" offer what they label "a friendly amendment" to the proposal in Stojnić. The notion of prominence at its core, they argue, is a syntactic relation holding between nodes on the discourse trees. Michael Devitt in his article "Incoherent meanings" argues against the radical view that coherence relations determine the reference of context-sensitive language. His starting point is that the theoretical interest in language comes from an interest in thoughts and their communication. A person can have any thought at all, however incoherent. Thus, a thought's meaning and reference are independent of its coherence and coherence has no place in the theory of meaning or reference. He concludes that the error in Stojnić's approach exemplifies the widespread confusion of the metaphysics of meaning with the epistemology of interpretation.

One of the driving themes of Fabrizio Cariani's book The Modal Future *concerns the interplay of tense and modality in powering future reference. Cariani's book opens by contrasting a 'symmetric' paradigm in which languages have three tenses (past, present and future) with an alternative on which past and present are the 'just' tenses. Building on prior work in semantics Cariani argues that the devices languages recruit to power future-directed discourse are modals. Fabrizio Cariani and Michael Glanzberg in the joined article "What is tense, anyway?" explain that an implicit corollary of the above thesis is that because expressions like will are modals, they cannot also be tenses. The article ends with a question: Does identifying modal features in will, or any other future expression, entail that it's not a tense? In this paper, the authors argue that the answer to this question is in an important sense indeterminate. There are multiple conceptions of tense which yield diverging answers to the question whether tense and modality are compatible—thus illuminating the relationship between tense and modality in a different way.*

DUNJA JUTRONIĆ

Croatian Journal of Philosophy
Vol. XXIII, No. 69, 2023
https://doi.org/10.52685/cjp.23.69.1
Received: November 15, 2023
Accepted: November 30, 2023

Précis for Context and Coherence: The Logic and Grammar of Prominence

UNA STOJNIĆ
Princeton University, Princeton, USA

This précis outlines some of the key themes in Context and Coherence. *At the core of* Context and Coherence *is the meta-semantic question: what determines the meaning of context-sensitive language and how do we interpret it as effortlessly as we do? What we can express with language is obviously constrained by grammar, but it also seems to depend on various non-linguistic features of an utterance situation, for example, pointing gestures. Accordingly, it is nearly universally assumed that grammar underspecifies content: the interpretation of context-sensitive language depends in part on extra-linguistic features of the utterance situation. Contra this dominant tradition, the book develops and defends a thoroughly linguistic account: context-sensitivity resolution is entirely a matter of grammar, which is much more subtle and pervasive than has typically been noticed. In interpreting context-sensitive language as effortlessly as we do, we draw on our knowledge of these subtle, but pervasive, linguistic cues—what I call discourse conventions. If this is right, the dominant, extra-linguistic account must be rejected. It not only mischaracterizes the linguistic conventions affecting context-sensitivity resolution, but its widespread, and often implicit, endorsement leads to philosophically radical conclusions. The recent arguments for non-truth-conditional and non-classical semantics for modal discourse provide just one illustration of this point. But appeals to context are quite common within a wide range of debates across different subfields of philosophy, and they typically assume the extra-linguistic model of context-sensitivity resolution. If the account of context-sensitivity developed in* Context and Coherence *is on the right track, such arguments have to be reconsidered.*

Keywords: Context; content; discourse coherence; semantics/pragmatics interface; logical form.

At the core of *Context and Coherence* is the question of the meta-semantics of context-sensitivity: what determines the meaning of context-sensitive language and how do we interpret it as effortlessly as we do?

Suppose I want to convey to you that you forgot your keys at the desk in front of me. I could say to you, in this situation, pointing at the keys, "You forgot these." Being a competent English speaker, you will have understood that I said you forgot your keys. If you take me to be sincere and reliable you might come to believe this, and this might impact your action: you might grab your keys. What allows us to coordinate our thoughts and actions through language in this way? Intuitively, what facilitates such exchanges is the fact that my thought has content, which represents the world a certain way—as such that in it you forgot your keys, here and now—which my utterance expresses, and you as a competent speaker can understand it to express, and which you can further come to believe and act upon, if you take me to be sincere and reliable.

But an utterance of the sentence, "You forgot these," in principle, can express indefinitely many contents. For instance, it would express something quite different if I were talking to someone other than you or pointing at something else, e.g., the stack of books on my table; or if instead you uttered it talking to me and pointing at something different still. Such context-sensitivity is stunningly pervasive in natural languages. It is indeed hard to find an utterance that is not in some way context-sensitive.[1] Yet, even though context-sensitive utterances can express indefinitely many different contents on different occasions of use, we still interpret context-sensitive language effortlessly, on the fly. The pervasiveness of context-sensitivity in natural languages does not hinder our capacity to coordinate thoughts through linguistic communication. How is this possible? What determines the meaning of

[1] This is not to deny that there is controversy over which expressions are context-sensitive, and how context-sensitivity is to be modeled and resolved. Some theorists maintain that the list of context-sensitive expressions is small, containing perhaps only pure indexicals such as 'I,' 'you,' 'here' and 'now' and demonstrative terms such as 'he,' 'she,' and 'this' and 'that' (viz. Cappelen and Lepore 2005). Others hold that nearly all expressions are massively context-sensitive (Travis 1989). But most theorists lie somewhere in between these extremes. For instance, it is common to posit context-sensitivity in analyses of nominal and adverbial quantification, tense, aspect, mood, modality, conditionals, relational expressions, gradable adjectives, predicates of personal taste, attitude and knowledge ascriptions, among many other types of expressions. And while there might be disagreements over whether any such particular expression is context-sensitive, and how any such context-sensitivity is realized and resolved, it is safe to say that most authors agree context-sensitivity is pervasive in natural languages. For a sample of the debates over whether and how modals are context sensitive, see, e.g., Egan, Hawthorne, and Weatherson (2005), Yalcin (2007), von Fintel and Gillies (2008, 2009), Kolodny and MacFarlane (2010), Dowell (2011); for those over predicates of personal taste, see, e.g., Cappelen and Hawthorne (2009), Egan (2010), MacFarlane (2014), *inter alia*; for knowledge ascriptions, see, e.g., DeRose (1995, 2009), Lewis (1996), Cohen (1998), Hawthorne (2004), Stanley (2005), Schaffer and Szabó (2013), Moss (2023).

context-sensitive utterances on an occasion of use, and what cognitive and linguistic resources allow us to interpret them so effortlessly?

Part of the answer, of course, must be in the meaning of the words, e.g., 'you,' 'forgot,' and 'these,' in English, and how they are put together; we draw on knowledge of grammar of our shared language in interpreting one another's speech. But while what we can communicate with language is obviously constrained by grammar, it also seems to depend on various non-linguistic features of an utterance situation, for example, which gestures accompany my utterance and whether I was speaking literally or figuratively. Accordingly, most theorists endorse the common-sense view that grammar underspecifies content: what 'that' picks out depends not just on its linguistic meaning, but also on extra-linguistic features of the utterance situation—what the speaker intends and/or what's salient in the speech situation. Audiences must exploit whatever epistemic cues a speaker and her situation provide to discern the speaker's intentions. It is thus nearly universally accepted that interpretation largely relies on general reasoning about communicative situations and intentions.

Context and Coherence urges a departure from this tradition. It argues that context-sensitivity resolution is entirely a matter of grammar, which is much more subtle and pervasive than has typically been assumed. In interpreting context-sensitive language as effortlessly as we do, we draw on our knowledge of these subtle, but pervasive, linguistic cues—what I call *discourse conventions.*

If this is right, the nearly universally accepted view that context-sensitivity resolution is mediated by extra-linguistic factors—speaker intentions and/or other extra-linguistic contextual cues—is misguided. This in turn has far-reaching philosophical consequences. Appeals to context and context-sensitivity have played an important role in philosophical theorizing about the foundational issues in philosophy of language—e.g., over the nature of linguistic meaning and its relation to speech and attitude content—as well as a wide range of debates in other subfields of philosophy, where philosophers frequently appeal to context-sensitivity in analyses of philosophically interesting expressions (e.g., 'know,' 'believe,' 'ought,' 'good,' 'true,' counterfactuals) in order to draw interesting conclusions about the underlying phenomena that these expressions denote or otherwise help elucidate (e.g., knowledge, belief, obligation, goodness, truth, causation). Yet, invariably, the model of context-sensitivity resolution these theorists assume in their arguments is one following the dominant tradition, whereby extra-linguistic parameters—speaker intentions and salient worldly cues—combine to determine the overall most plausible interpretation. The account I defend shows that in constructing and assessing such arguments, philosophers will have to think of context-sensitivity quite differently than has been customary.

1. *Discourse conventions*

A context can be thought of as an abstract representation of the features of an utterance situation required to interpret context-sensitive language, at least including (but not limited to) the speaker, the addressee, and the location, time, and world of the utterance (Kaplan 1989a, 1989b; Lewis 1980). The linguistic meaning of a context-sensitive expression—its *character* (to use Kaplan's familiar jargon)—selects its semantic content as a function of a particular parameter of the context in which it was uttered. For instance, because the character of 'I' requires that its referent be the speaker, when I utter (1), it means that Una Stojnić is a philosopher; 'I' simply selects me—the speaker—as its referent.

1. I am a philosopher.

What if instead I utter (2a)–(2b)?

2.

a. She is a philosopher.
b. That is made of glass.

The character of 'she' constrains its referent to be third-person, singular, and female; but this does not suffice to determine its referent: there can be more than one candidate female referent in a given situation, so something has to single out a particular one: e.g., a pointing gesture, or her perceptual salience in our speech situation. Moreover, the target referent nee not even be present in the speech situation; viz.:

3. Mary is away, attending a conference. She's a philosopher.

Similarly for the demonstrative 'that' in (2b).

This difference between 'I,' on the one hand, and 'she' and 'that,' on the other, motivates a theoretical distinction between the so-called "pure indexicals," like 'I', the character of which alone determines the referent given a context, and "true demonstratives," like 'she' and 'that,' the character of which is incomplete, and requires extra-linguistic supplementation to fix the interpretation (Kaplan 1989a, 1989b).

Pure indexicals are generally assumed to be few and far between. Most context-sensitive expressions are thought to be like true demonstratives, in that they require extra-linguistic supplementation. And thus, the resolution of context-sensitivity by and large depends on extra-linguistic resources.[2]

The idea that context-sensitivity resolution requires extra-linguistic supplementation appears obvious: after all, does this not just follow from the fact that the meaning of a demonstrative can vary with seemingly non-linguistic features of an utterance situation such as pointing

[2] See, e.g., Grice (1957, 1975), Schiffer (1972, 1981, 2005), Wettstein (1984), Kaplan (1989a, 1989b), Neale (1990, 2004), Reimer (1992), Stanley and Szabó (2000), Glanzberg (2007), King (2014a, 2014b), Dowell (2011), Lewis (2020), *inter multa alia*. The extra-linguistic model is often implicitly assumed, even when not explicitly endorsed.

gestures? Such extra-linguistic cues are messy and multifarious: there is no in-principle limit to the elements of world-knowledge, and the information about the speaker and speech situation, that one can factor into what is salient in a given situation. Further, one has to weigh these resources against one another in determining the overall most plausible interpretation.

Context-sensitivity resolution would not be any less dependent on extra-linguistic supplementation if we thought of contextual salience as a parameter of context, and built it into the linguistic character of, say, 'she,' that what it picks out, in addition to being third person, singular, female, must also be contextually salient. For, this would only mask the theoretically important distinction: real-world salience in a speech situation would still be a product of complex extra-linguistic parameters, which might pull in different directions, and which must be weighed holistically against one another, to determine the overall most plausible interpretation.

By contrast, in *Context and Coherence,* I argue that the extra-linguistic model is genuinely misguided. Context-sensitivity in general operates on a model of pure indexicals: context-sensitive expressions automatically select their content from the context as a matter of their character. Most context-sensitive expressions appear to behave as true demonstratives because their character is sensitive not to real-world salience of a particular interpretation, but to the linguistically determined prominence. So, for instance, the English demonstrative pronoun 'she' picks out a third person, singular, female referent that is prominent in the linguistic context. However, building the sensitivity to linguistic prominence into the character does not merely mask the dependence on extra-linguistic factors, as would building in the sensitivity to worldly salience. For unlike worldly salience, prominence, in the relevant sense, is fully linguistically governed. It is determined by a set of linguistic rules—discourse conventions—which are triggered as a matter of linguistic meaning of particular expressions within the discourse itself. These items induce changes in the context, marking certain interpretation as prominent at a particular point in discourse, demoting others.

Since prominence is dictated by linguistic contributions of the expressions that are a part of the discourse itself, the context must keep track of prominence as it evolves with the unfolding discourse, word-by-word. To capture this, I model context as a dynamically evolving conversational scoreboard (Lewis 1986). It still provides the abstract representation of information needed for the resolution of context-sensitive expressions, including, at least, the speaker, world, time, and location of the utterance, but also the prominence ranking of candidate interpretations that dynamically evolves as the discourse progresses. Since this prominence ranking is exclusively governed by linguistically contributed updates—i.e., by discourse conventions—the inter-

pretation of context-sensitivity is fully linguistically determined.[3] A context-sensitive expression simply selects what discourse conventions determine is the most prominent element of the ranking that satisfies the constraints of the character: e.g., 'she' picks out the currently most prominent third person, singular female. The interpretation is settled by linguistic rules through and through: by the linguistic character and discourse conventions.

But in what sense is prominence determined by discourse conventions if, as we have seen, it can vary with e.g., pointing gestures? Drawing on work in Stojnić et al. (2013, 2017, 2020) and Stojnić (2017, 2019), I argue that many of the parameters affecting context-sensitivity resolution have been either missed or mistaken for extra-linguistic cues, but are in fact grammaticized in language. So, for instance, the interpretation of demonstrative gestures varies with a particular form of a gesture. Distinct forms receive distinct meanings, and there is significant cross-linguistic variation in both the range of gestures recognized as demonstrative, and in the interpretation of specific forms of gestures (Kendon 1988, 2004; Wilkins 2003). Such variation and arbitrariness in form-to-meaning mapping is a hallmark of linguistic conventionality.[4]

Similarly, discourse conventions that are triggered by discourse relations that signal how individual utterances connect into a coherent whole are often mischaracterized as byproduct of holistic reasoning drawing on general world knowledge. To illustrate what is at stake consider (4) (Hobbs 1979):

4. John took a train from Paris to Istanbul.
   a. He has family there.
   b. He likes spinach.

(4a) is natural, and its second sentence is readily understood as providing an explanation for the event described in the first: John took a train from Paris to Istanbul *because* he has family there. (4b), by contrast, sounds off. The audience is left wondering how the train trip explains the preference for spinach. This observation is captured within Discourse Coherence Theory by positing an implicit organization of a discourse, a network of discourse coherence relations that hold among individual utterances, and signal how they are connected into a coherent discourse (Hobbs 1979; Kehler 2002; Asher and Lascarides 2003). In (4a) the coherence relation Explanation signals the explanatory connection between the two sentences, which is why it receives its natural interpretation. In (4b), one still expects the same relation, Explanation, but fails to confirm it: this is why one is left wondering, what is it about spinach that explains the train trip in question.

Establishing coherence in discourse affects the resolution of con-

[3] Thus, my understanding of the evolution of the conversational scoreboard is closer to that of Lepore and Stone (2015), than Lewis (1986), for Lewis allows that non-linguistic factors can update the scoreboard.

[4] The relevant notion of convention is that of Lewis (1969).

text-sensitive items such as demonstrative pronouns (see, e.g., Kehler 2002; Wolf, Gibson and Desmet 2006; Kehler et al. 2008; Kaiser 2009, and reference therein); viz. (5):

5. Phil tickled Stanley. Liz poked him. (Smyth 1994)

Speakers tend to understand an utterance of (5) out of the blue in one of two ways: either Liz's action is a result of Phil's (perhaps she is acting in disapproval), or Liz's action is described as similar to Phil's. In one case, the discourse is organized around an event-result relation, and in the other, around a parallel, resemblance one. Crucially, in the former case, 'him' is understood to refer to Phil, and in the latter, to Stanley. One might be tempted to understand these effects as mere by-products of holistic reasoning: after all, if we are comparing what Phil and Liz did to Stanley, it only makes sense that the pronoun is resolved to Stanley; similarly, if we are describing how Liz reacted in response to Phil tickling Stanley, it makes sense that the pronoun refers to Phil.

But while this understanding might be natural—and, indeed, is one that coherence theorists tend to endorse—I argue it is mistaken (Stojnić et al. 2013, 2017, 2020). The effects of discourse relations on the resolution of demonstrative pronouns are not a mere byproduct of general pragmatic reasoning about the epistemic cues that guide interpretation. They are grammatically encoded and are one among many such discourse conventions that, together, fully settle the interpretation on an occasion of use, without an appeal to extra-linguistic cues.

The conventionality of these effects is nicely illustrated by examples like the following one, from Kehler (2002):

6. Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships her.

(6) is generally judged infelicitous; it is understood as if the speaker has made an error referring to Reagan with a female-gendered pronoun. Yet, this is surprising if coherence relations merely pragmatically, but defeasibly, favor a particular resolution of the pronoun. For, 'Thatcher' is a perfectly available antecedent, and Thatcher is generally known to be admired by Bush. Moreover, if we resolved the pronoun to Thatcher, we would get a plausible, relevant interpretation, and one that is charitable to the speaker, not charging them with a mistake. So, if we were holistically searching for the overall most plausible interpretation, this one should win out. Yet, Kehler reports that his subjects judge (6) infelicitous instead.

The infelicity judgement, however, is perfectly expected if coherence relations force a particular resolution of the pronoun as a matter of an underlying convention. (6) is organized by the coherence relation Parallel, comparing Thatcher's and Bush's respective attitudes. Parallel requires that a pronoun in the object position be resolved to an antecedent introduced in the object position; so, 'her' must pick out Reagan. But since the pronoun is feminine, this results in gender mismatch.

Consequently, (6) is infelicitous.

Context and Coherence argues that such discourse conventions are pervasive, and that they affect the resolution of context-sensitivity quite generally, not just for demonstrative pronouns. In short: context-sensitivity in general is linguistically resolved.

2. *Discourse conventions, dynamic meaning and truth-conditions*

On the account developed in *Context and Coherence* then, linguistic items are associated with a layer of content that encapsulates their effect on context—how they change the conversational scoreboard by updating the prominence ranking. Such context-changing aspect of meaning is naturally thought of within the framework of dynamic semantics, as a *dynamic context-change potential*, realized as a relation between input and output prominence ranking assignment functions, reflecting the change in ranking an utterance brings about.[5] But while dynamic semantics is sometimes presented as in tension with traditional truth-conditional accounts of meaning, the model still allows us to capture the ordinary truth-conditional content.

To illustrate, consider the distinction between the following:

7.

a. Mary came in. She sat down.
b. Mary came in. She [pointing at Betty, the cat] sat down.

Simplifying somewhat, the following are the key effects (7a) has on context: the subject NP, 'Mary,' updates the prominence ranking, making Mary the top-ranked element, and requiring that she came in. The second sentence continues the narrative about Mary, maintaining Mary as the top-ranked element. The pronoun 'she' selects the top-ranked element that is third person, singular, and female—which, given the effect of the first sentence, is Mary. Further, (7b) requires that this individual sat down. The whole discourse is true just in case Mary came in and sat down, the intuitively correct truth-condition.

The first sentence in (7b) has exactly the same effect as that in (7a). The second, however, features a pointing gesture accompanying the pronoun. As a matter of its linguistic contribution, the gesture promotes the individual pointed at—here, Betty the cat—as the new top-ranked element. The pronoun again selects the top-ranked third person, singular female referent. But given the effect of pointing, this referent is now Betty, the cat. It is further required that this referent sat down. The whole discourse is thus true just in case Mary came in and Betty sat down.

Not only do we capture the differences in truth-conditions between

[5] The dynamic approach to semantics is due to Kamp (1981) and Heim (1982). Dekker (2011) offers an accessible overview. For more details see ch. 3 of *Context and Coherence*.

(7a) and (7b), but we also capture the differences in entailment patterns they give rise to as a matter of the differences in the underlying logical forms of the two discourses. (7a) but not (7b) entails that Mary came in and sat down. The difference is underwritten by a difference in the logical form, because the form of (7b), but not (7a), features a pointing gesture, which, we have seen, must be linguistically represented.

These considerations extend beyond pronouns. Quite generally, we can think of expressions as carrying two layers of content—the dynamic context-change potential, describing how an expression updates the context in which it occurs, and its representational, truth-conditional contribution—what it contributes to the truth-conditions of an utterance it occurs in. The two are interrelated—the context-change potential models the effects of discourse conventions on the context, which in turn affect the interpretation of subsequent context-sensitive items; and the interpretation a context-sensitive item receives, in turn, affects how this item itself updates the context downstream. In this way, linguistic rules fully determine truth-conditional content. As the context change potential updates the context, it builds the underlying truth-conditional content. Discourses can thus be thought of as recipes for building truth-conditional content, expression-by-expression.

3. *The dangers of the extra-linguistic model assumption: An illustration*

We can now illustrate with a concrete example how the implicit assumption of the extra-linguistic account of context-sensitivity resolution can lead to radical philosophical consequences. In recent literature a growing number of theorists have argued that modal discourse fails to express representational, truth-conditional content. Some of the key arguments are fueled by the data that seems to suggest that context cannot determine the propositional, representational content of many modal utterances; for instance, consider (8a)–(8c) (Yalcin 2007):[6]

8.

a. If it's not raining and it might be raining, I'm misinformed about the weather.
b. If it's not raining and for all I know it is raining, I'm misinformed about the weather.
c. If it's not raining and the body of information *i* doesn't rule rain out, then this body of information *i* lacks some information about the weather.

[6] (8a)–(8b) are Yalcin's original examples. (8c) generalizes his point: the contrast remains whichever body of information the context might supply.

On the standard account, modals are quantifiers over contextually restricted domains of possibility (Kratzer 1977, 1981; Kripke 1980). So, "It might be raining" means that the (contextually determined) body of information—which is typically assumed to be or include that of the speaker—is compatible with rain. But if so, there should be no difference in truth-conditions between (8a) and (8b); yet, (8a) sounds incoherent, while (8b) sounds perfectly felicitous. Indeed, whatever body of information i the context delivers as the domain for the modal, the corresponding utterance featuring explicit reference to i still remains coherent, as illustrated in (8c), while (8a) remains incoherent. Yalcin (2007) takes this type of data to show that context cannot determine the representational, truth-conditional content for utterances like (8a): there is no coherent representational content of this sort; so, modals do not express propositional, representational content.

This type of data has fueled a departure from representational accounts of meaning: the idea—dominant since at least Frege (1892)—that (declarative) utterances express content that represents the world a certain way. The departure led to now increasingly more prominent accounts—various forms expressivism, relativism, and certain types of dynamic accounts of meaning.[7] These accounts maintain that modal assertions, e.g., "It might be raining," only contribute a dynamic effect on context, which reflects the interlocutors' mutual acceptance of a non-representational attitude—roughly, being in a state of mind that does not rule out that it is raining. This dynamic meaning is not reducible to representational content, nor can such content be recovered from it.

This departure from the representational tradition is further fueled by the apparent counterexamples to certain classical patterns of inference, which seems to arise in the presence of modal vocabulary. The following apparent counterexample to *modus tollens* from Yalcin (2012a) is a case in point:

9. Take an urn with 100 marbles. 10 of them are big and blue, 30 big and red, 50 small and blue, and 10 are small and red. One marble is randomly selected and hidden (you do not know which). Given this setup, (9a) and (9b) are licensed; yet (9c) does not follow.
    a. If the marble is big, then it is likely red.
    b. The marble is not likely red.
    c. So, the marble is not big.

This type of data seems to further support the departure from the representational paradigm, since the standard implementations of expressivist, relativist and dynamic accounts give rise to a non-classical logic that invalidates the relevant patterns. Many philosophers have thus concluded that we must endorse non-representational accounts

[7] See, e.g., Veltman (1985), Gillies (2004, 2010), Swanson (2006); von Fintel and Gillies (2008, 2009), Yalcin (2007, 2011, 2012a, 2012b); Kolodny and MacFarlane (2011), Willer (2013, 2014), Bledin (2015), Charlow (2015), Starr (2016), Moss (2015), *inter alia.*

of meaning for modal discourse and recognize a deep incompatibility between classical logic and natural language.

In *Context and Coherence,* drawing on Stojnić (2017, 2019), I argue that this reaction is misguided. At the core of the arguments drawing on these data is the implicit reliance on the dominant yet flawed extra-linguistic conception of context-sensitivity resolution. Once we appreciate the import of discourse conventions that have been missed, these arguments dissipate.

Like pronouns, modal expressions are, I argue, prominence-sensitive. A modal searches for the prominent possibility that serves as a restrictor on its domain of quantification (Stone 1997, 1999). Crucially, the prominence of the relevant restrictor possibility is determined by discourse conventions, such as the prominence-resetting updates triggered by coherence relations. Here is a sketch of how these conventions operate in (8a)–(8c). The first conjunct in the antecedent of (8a) introduces a hypothetical non-raining scenario. The second, "It might be raining," elaborates on this scenario. The coherence relation Elaboration between the two conjuncts has an effect on prominence: it makes the proposition elaborated on—the one introduced by the first conjunct and comprising the epistemically accessible worlds in which it is not raining—prominent. 'Might' in the second conjunct selects the most prominent possibility as its restrictor; consequently, it selects this proposition, and is thus understood as quantifying over the epistemically accessible worlds in which it is not raining. But as a result, the antecedent as a whole delivers the proposition that it is not raining, and that within the set of epistemically accessible worlds *in which it is not raining*, there is at least one raining world. This, of course, leads to a contradiction, and hence the infelicity.

The antecedents of (8b) and (8c) similarly feature a conjunction, in which the first conjunct introduces a hypothetical non-raining possibility, and the second one further elaborates on it. The Elaboration relation between the conjuncts still makes this possibility prominent. However, since there is no modal expression in the second conjunct to select this possibility as the restrictor, we get a perfectly consistent reading: it is not raining and the speaker's information/the contextually relevant body of information i does not rule out raining. This explains the contrast between the two examples.

So, it is not that the context cannot fix the representational, truth-conditional meaning for modal constructions like those in (8a); rather, the discourse conventions do fix such meaning, but it is one that is inconsistent. The mistake was to implicitly assume the standard extra-linguistic model of context-sensitivity resolution, so that general epistemic cues work together to determine the overall most plausible interpretation. For, if the effect of Elaboration on prominence were a mere defeasible byproduct of such holistic reasoning about the available cues, the interpretations like those in (8b) and (8c) would be per-

fectly possible for (8a). Indeed, not only should they be possible, but they should be favored, for considerations of charity, relevance, and plausibility would all point in their direction. So, assuming the general extra-linguistic factors conspire to determine the overall most plausible interpretation, it would indeed be mysterious why we get the contrast between (8a) and (8b)–(8c). If instead, it is a part of the linguistic contribution of Elaboration to promote the possibility elaborated on, the contrast is predicted. (8a) is infelicitous because its antecedent receives an inconsistent truth-condition as a matter of grammar.

Similar considerations apply to putative counterexamples to classical patterns of inference, like Yalcin's counterexample to *modus tollens*. The consequent of the conditional in (9a) elaborates on the possibility introduced by its antecedent (the one corresponding to the set of epistemically accessible worlds in which the marble is big). The Elaboration relation again promotes this possibility. The modal 'likely' in the consequent selects this possibility—the currently most prominent one—as the restrictor for its domain. The consequent thus receives the intuitively correct, restricted, reading—the marble is likely red, given that it is big. (9b), in turn, stands in Contrast relation to (9a). The two utterances are understood as contributing contrasting information relative to the body of information available discourse initially, which describes the situation concerning the urn: they contras the likelihood of the marble being red, given some or no assumption about its size. Contrast makes this initial body of information prominent, and the modal 'likely' in (9b) selects it as its restrictor. Thus, (9b) conveys that the marble is not likely red given this overall body of knowledge (so, given no particular assumption about its size). But this means that (9b) does not contradict the consequent of (9a), for the two occurrences of 'likely' are interpreted differently. So, (9a)–(9c) is not an instance of *modus tollens*, and hence, not a counterexample to it.

The lesson is that once we properly capture the effects of discourse conventions, the seaming counterexamples to classical patterns of inference disappear. With linguistic contributions of discourse coherence relations properly reflected in the logical form, we see that (9a)–(9c) is not an instance of *modus tollens,* nor is it associated with a valid form. More generally, a semantic account that adequately tracks the contribution of discourse conventions provably preserves classical logic (Stojnić 2017).

This is but one illustration of how a model of context-sensitivity resolution that fails to account for discourse conventions can lead to radical philosophical conclusions. If the arguments in *Context and Coherence* are on the right track, the pessimism about the representational, truth-conditional accounts of meaning is unwarranted, as is the embrace of the failures of classical validities. These reactions rest on an overly simplistic account of content, context, and content-context interaction, which presupposes the dominant but faulty extra-linguistic

model of context-sensitivity resolution. While it is true that, as expressivists, relativists, and dynamic semanticists urge, an important aspect of modal meaning concerns the dynamic effect modals have on context, these theorists have mischaracterized this dynamic aspect of meaning. Properly characterized, the dynamic meaning is a reflex of discourse conventions, and it, I show, fully determines the representational, truth-conditional content. Contrary to a widespread assumption, the dynamic aspect of meaning does not exhaust the contribution of modal discourse, nor does modal discourse fail to express truth-conditional, representational content; the truth-conditional, representational content is, instead, fully determined by the dynamic meaning contributed by discourse conventions. And properly characterized, the underlying semantics provably preservers classical (modal) logic.

This result also shapes how we should think about the relation between natural language and logic. Context-sensitivity has long presented a challenge for the proper treatment of validity in natural language discourse. The traditional strategy, rooted in Kaplan (1989a, 1989b), teaches that in assessing the validity of an argument expressed in a natural language like English, the context must be fixed. This is to avoid utterances like 'he [pointing at Tom] is happy; therefore he [pointing at a different person, Bill] is happy' qualifying as counterexamples to classical patterns of inference (here: *φ; therefore φ*). But if the account of context-sensitivity developed in *Context and Coherence* is on the right track, imposing a ban on context-shifting is impossible. Natural language arguments are structured discourses, which trigger discourse conventions that affect the context which determines the meanings of context-sensitive items those very discourses harbor. Moreover, discourse conventions are not isolated to the contribution of individual sentences but are also encoded in the linguistically specified discourse coherence relations between them.

Research in dynamic semantics has long stressed the importance of the dynamics of context-change for capturing intuitions about validity, a lesson that is also adopted by expressivists and relativists. But these semantic accounts still only characterize dynamic meaning as of individual sentences and represent arguments as relations between sets of sentences and sentences, the premises and conclusion. This is a mistake: an adequate account requires individuating argument patterns as structured discourses, the structure of which determines the content expressed by the premises and conclusion.

4. *Conclusion*

This précis outlines some of the key themes in *Context and Coherence*. The book develops and defends a thoroughly linguistic account of the meta-semantics of context-sensitivity: the interpretation of context-sensitive expressions is fully determined by linguistic rules, discourse conventions. We interpret context-sensitive language as effortlessly as

we do by employing our linguistic competence with these conventions.

If this is right, the dominant, extra-linguistic account must be rejected. It is not only faulty, missing or mischaracterizing the linguistic conventions affecting context-sensitivity resolution, but its widespread, and often implicit, endorsement leads to philosophically radical conclusions. The recent arguments for non-truth-conditional and non-classical semantics for modal discourse provide just one illustration of this point. But appeals to context are quite common within a wide range of debates across different subfields of philosophy, and they typically assume the extra-linguistic model of context-sensitivity resolution. If the account of context-sensitivity developed in *Context and Coherence* is on the right track, such arguments might have to be reconsidered.

References

Asher, N. and Lascarides, A. 2003. *Logics of Conversation*. Cambridge University Press.

Bledin, J. 2015. "Modus Ponens Defended." *Journal of Philosophy* 112: 57–83.

Cappelen, H. and Hawthorne, J. 2009. *Relativism and Monadic Truth*. Oxford: Oxford University Press.

Cappelen H. and Lepore E. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Wiley-Blackwell.

Charlow, N. 2015. "Prospects for an Expressivist Theory of Meaning." *Philosophers' Imprint* 15: 1–43.

Cohen, S. 1998. "Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery." *Australasian Journal of Philosophy* 76: 289–306.

DeRose, K. 1995. "Solving the Skeptical Problem." *Philosophical Review* 104: 1–52.

DeRose, K. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context*. Volume 1. Oxford: Oxford University Press.

Dekker, P. 2011. "Dynamic Semantics." In K. von Heusinger, C. Maienborn and P. Portner (eds.). *Semantics: An International Handbook of Natural Language Meaning*. Volume 1. Berlin: Mouton De Gruyter, 923–945.

Dorr, C. and Hawthorne, J. 2014. "Embedding Epistemic Modals." *Mind* 122: 867–913.

Dowell, J. L. 2011. "A Flexible Contextualist Account of Epistemic Modals." *Philosophers' Imprint* 11: 1–25.

Egan, A. 2010. "Disputing About Taste." In R. Feldman and T. A. Warfield (eds.). *Disagreement*. Oxford: Oxford University Press, 247–286.

Egan, A., Hawthorne, J. and Weatherson, B. 2005. "Epistemic Modals in Context." In G. Preyer and G. Peter (eds.). *Contextualism in Philosophy*. Oxford: Oxford University Press, 131–169.

von Fintel, K. and Gillies, A. S. 2008. "CIA Leaks." *Philosophical Review* 117: 77–98.

von Fintel, K. and Gillies, A. S. 2009. "'Might' Made Right." In A. Egan and B. Weatherson (eds.). *Epistemic Modality*. New York: Oxford University Press.

Frege, G. 1892. "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik* 100: 25–50.

Gillies, A. S. 2004. "Epistemic Conditionals and Conditional Epistemics." *Noûs* 38: 585–616.

Gillies, A. S. 2010. "Iffiness." *Semantics and Pragmatics* 3: 1–42.

Glanzberg, M. 2007. "Context, Content, and Relativism." *Philosophical Studies* 136: 1–29.

Grice, H. P. 1957. "Meaning." *The Philosophical Review* 66 (3): 377–388.

Grice, H. P. 1975. "Logic and Conversation." In P. Cole and J. L. Morgan (eds.). *Syntax and Semantics 3: Speech Acts*. New York: Academic Press, 41–58.

Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis. University of Massachusetts.

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

Hobbs, J. R. 1990. *Literature and Cognition (CSLI Lecture Notes 21).* Chicago: University of Chicago Press.

Kaiser, E. 2009. "Effects of Anaphoric Dependencies and Semantic Representations on Pronoun Interpretation." In *Anaphora Processing and Applications*. Heidelberg: Springer, 121–130.

Kamp, H. 1981. "A theory of truth and semantic representation." In J. A. G. Groenendijk, T. M. V. Janssen, and M. B. J. Stokhof (eds). *Formal methods in the Study of Language*, Mathematical Centre Tracts 135. Amsterdam: Mathematisch Centrum, 277–322.

Kaplan, D. 1989a. "Afterthoughts." In J. Almog, J. Perry, and H. Wettstein (eds.). *Themes From Kaplan*. Oxford: Oxford University Press, 565–614.

Kaplan, D. 1989b. "Demonstratives." In J. Almog, J. Perry, and H. Wettstein (eds.). *Themes From Kaplan*. Oxford: Oxford University Press, 481–563.

Kehler, A. 2002. *Coherence, Reference and the Theory of Grammar*. Stanford: CSLI Publications.

Kehler, A., Kertz, L., Rohde, H. and Elman, J. 2008. "Coherence and Coreference Revisited." *Journal of Semantics* 25: 1–44.

King, J. C. 2014a. "The Metasemantics of Contextual Sensitivity." In A. Burgess and B. Sherman (eds.). *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press, 97–118.

King, J. C. 2014b. "Speaker Intentions in Context." *Noûs* 48: 219–237.

Kendon, A. 1988. *Sign Languages of Aboriginal Australia: Cultural, Semiotic and Communicative Perspectives*. Cambridge: Cambridge University Press.

Kendon, A. 2004. *Gesture: Visible Action as Utterance*. Cambridge: Cambridge University Press.

Kolodny, N. and MacFarlane, J. 2010. "Ifs and Oughts." Journal of Philosophy 107: 115–143.

Kratzer, A. 1977. "What Must and Can Must and Can Mean." *Linguistics and Philosophy* 1: 337–355.

Kratzer, A. 1981. "Notional Category of Modality." In H.-J. Einkmeyer and H. Rieser (eds.). *Words, Worlds, and Contexts: New Approaches in Word Semantics*. Berlin: de Gruyter, 38–74.

Kripke, S. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.

Lepore, E. and Stone, M. 2015. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford: Oxford University Press.

Lewis, D. K. 1969. *Convention: A Philosophical Study*. Cambridge: Harvard University Press.

Lewis, D. K. 1980. "Index, Context and Content." In S. Kanger and S. Öhman (eds.). *Philosophy and Grammar*. Amsterdam: D. Reidel, 79–100.

Lewis, D. K. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339–359.

Lewis, D. K. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–67.

Lewis, K. 2020. "The Speaker Authority Problem for Context-sensitivity (or: You Can't Always Mean What You Want." *Erkenntnis* 85 (6): 1527–1555

MacFarlane, J. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.

Moss, S. 2018. *Probabilistic Knowledge*. Oxford: Oxford University Press.

Moss, S. 2023. "Knowledge and Legal Proof." *Oxford Studies in Epistemology* 7: 176– 213.

Neale, S. 1990. *Descriptions*. Cambridge: MIT Press.

Neale, S. 2004. "This, That and the Other." In M. Reimer and A. Bezuidenhout (eds.). *Descriptions and Beyond*. New York: Oxford University Press, 68–182.

Reimer, M. 1992. "Three Views of Demonstrative Reference." *Synthese* 93: 373–402.

Schaffer, J. and Szabó, Z. G. 2013. "Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions." *Philosophical Studies* 2: 1–53.

Schiffer, S. 1972. *Meaning*. Oxford: Clarendon Press.

Schiffer, S. 1981. "Indexicals and the Theory of Reference." *Synthese* 49: 43–100.

Schiffer, S. 2005. "Russell's Theory of Definite Descriptions." *Mind* 114: 1135–1183.

Smyth, R. 1994. "Grammatical Determinants of Ambiguous Pronoun Resolution." *Journal of Psycholinguistic Research* 23: 197–229.

Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Stanley, J and Szabó, Z. G. 2000. "On Quantifier Domain Restriction." *Mind and Language* 15: 219–261.

Starr, W. B. 2016. "Dynamic Expressivism about Deontic Modality." In N. Charlotte and M. Chrisman (eds.). *Deontic Modality*. Oxford: Oxford University Press, 355–394.

Stojnić, U. 2017. "One's *Modus Ponens*…: Modality, Coherence and Logic." *Philosophy and Phenomenological Research* 95: 167–214.

Stojnić, U. 2019. "Content in a Dynamic Context." *Noûs* 52: 394–432.

Stojnić, U., Stone, M. and Lepore, E. 2013. "Deixis (Even Without Pointing)." *Philosophical Perspectives* 27: 502–525.

Stojnić, U., Stone, M. and Lepore, E. 2017. "Discourse and Logical Form: Pronouns, Attention and Coherence." *Linguistics and Philosophy* 40: 519–547.

Stojnić, U., Stone, M. and Lepore, E. 2020. "Pointing Things Out: In Defense of Attention and Coherence." *Linguistics and Philosophy* 43: 139–148.

Stone, M. 1997. "The Anaphoric Parallel Between Modality and Tense." *IRCS Report 97*–06. Philadelphia: University of Pennsylvania.

Stone, M. 1999. "Reference to Possible Worlds." *Technical Report 49*. New Brunswick: Rutgers University, RuCCS.

Travis C. 1989. *The Uses of Sense: Wittgenstein's Philosophy of Language*. New York: Oxford University Press.

Veltman, F. 1985. *Logics for Conditionals*. Ph.D. thesis. University of Amsterdam.

Wettstein, H. K. 1984. "How to Bridge the Gap Between Meaning and Reference." *Synthese* 58: 63–84.

Wilkins, D. 2003. "Why Pointing With the Index Finger is Not a Universal (in Sociocultural and Semiotic Terms)." In S. Kita (ed.). *Pointing: Where Language, Culture and Cognition Meet*. Abingdon: Lawrence Erlbaum Associates.

Willer, M. 2013. "Dynamics of Epistemic Modality." *Philosophical Review* 122: 45–92.

Willer, M. 2014. "Dynamic Thoughts on Ifs and Oughts." *Philosophers' Imprint* 14: 1–30.

Wolf, F., Gibson, E. T. and Desmet, T. 2004. "Discourse Coherence and Pronoun Resolution." *Language and Cognitive Processes* 19: 665–675.

Yalcin, S. 2007. "Epistemic Modals." *Mind* 116: 983–1026.

Yalcin, S. 2011. "Nonfactualism about Epistemic Modality." In B. Weatherson and A. Egan (eds.). *Epistemic Modality*. New York: Oxford University Press, 295–332.

Yalcin, S. 2012a. "A Counterexample to Modus Tollens." *Journal of Philosophical Logic* 41: 1001–1024.

Yalcin, Seth. 2012b. "Bayesian Expressivism." *Proceedings of the Aristotelian Society* 112: 123–160.

Croatian Journal of Philosophy
Vol. XXIII, No. 69, 2023
https://doi.org/10.52685/cjp.23.69.2
Received: October 11, 2023
Accepted: November 20, 2023

Linguistic Conventions or Open-Ended Reasoning: Some Questions for Una Stojnić

PETER PAGIN*
Stockholm University, Stockholm, Sweden

This short paper has the character of a critical notice of Una Stojnić's book Context and Coherence: The Logic and Grammar of Prominence *(Stojnić 2021). It is mainly concerned with Stojnić's strong claim that linguistic phenomena related to prominence and coherence, in particular the interpretation of pronouns, are governed by linguistic conventions and are not pragmatic in nature. On these matters, my views are opposite to Stojnić's.*

Keywords: Coherence; convention; pragmatics; prominence; pronouns.

1. *Introduction*

This short paper has the character of a critical notice of Una Stojnić's (2021) book *Context and Coherence: The Logic and Grammar of Prominence*. It is mainly concerned with Stojnić's strong claim that linguistic phenomena related to prominence and coherence, in particular the interpretation of pronouns, are governed by linguistic *conventions* and are *not* pragmatic in nature.

Before moving into the discussion of these matters, I would like to give some brief indications of the contents of this brilliant book.

— The book contains a new formalism that can represent the semantic role of pronouns without the use of arbitrary indices. This is achieved by means of a stack algorithm that ranks individuals according to contextual salience / prominence.

— This formalism is also able to represent the changes of context that

takes place as sentences unfold, not just between sentences. This allows for temporally more fine-grained contexts than standard conceptions allow.
— The formalism implements a theory of shifts of prominence ranking that determines the interpretation of pronouns and is a driver of context change.
— The book further contains a formalism that integrates *coherence* relations in the representation of discourse.
— It contains a theory of the effects of coherence relations on the resolution of pronoun anaphora.
— It contains an application of this theory to the interpretation of *epistemic modals* as well as of *modal subordination*.
— It contains another application of this theory to handle alleged counterexamples to *modus ponens* and *modus tollens* without giving up the idea that *propositions* are the semantic values of sentences.[1]

In what follows, I shall focus on a theme in the book that runs through the various accounts of pronouns and modals, the appeal to *linguistic rules / conventions* and the opposition to pragmatic reasoning. Stojnić (2021: 8–10) calls such reasoning *open-ended*, *defeasible*, *holistic*, *abductive*.

I think all these labels are to the point. The first two emphasize the provisional nature of pragmatic considerations: they have no definitive end-point but can typically be strengthened or undermined by further considerations.

The third one, *holistic*, is also to the point because knowledge of a situation as a whole as well as general world knowledge often play a role in pragmatic interpretation.

The fourth one, *abductive*, is again to the point because it focuses on the explanatory aspects of pragmatic reasoning. Typically, the interpreter tries to come up with the interpretation that best satisfies certain conditions that derive from the standing meaning of a sentence used, often in combination with general principles of interpretation. I shall occasionally refer to such a process of interpretation as *constraint satisfaction*.

In this paper, I shall argue that constraint satisfaction plays a more central role, and linguistic convention a less central role, for some of the theories in the book. This concerns both prominence ranking and coherence relations.

2. *Pronouns*

The main idea in the book, concerning pronouns, is that they refer to the individual of the right ϕ features (gender, person, number) that is *top-ranked*, i.e. at the center of attention:

[1] In this respect, Stojnić is on the same side of the propositionalist fence as Kathrin Glüer and myself, in our work on switcher semantics (e.g. Glüer and Pagin 2005, 2008, 2012, 2022), where we also have a goal of preserving the classical proposition as the meaning in context of a sentence.

> At any given point in a discourse, the context provides a ranking by prominence of candidate interpretations for a pronoun, tracking what's most prominent—that is, at the center of attention. The prominence ranking changes and updates, as the discourse unfolds, as a function of the meaning of linguistic items the discourse harbors, word by word. In this way, the resolution of a pronoun requires no extra-linguistic supplementation. It is linguistically determined, through and through: by its standing linguistic meaning, and the *linguistically* set up context. (Stojnić 2021: 40)

The application of this general idea provides a definition that gives the linguistic meaning of particular pronouns of English, exemplified by 'she':

> Definition 4.1: The standing linguistic meaning of 'she' Interpreted relative to an assignment *g* 'she' denotes *g*'s highest ranked entity that is singular, feminine, and disjoint from the speaker and addressee of the utterance, and that yields an interpretation where the occurrence of the pronoun is free in its governing category. (Stojnić 2021: 56)

This theory accounts for the interpretation of a discourse like

(1) A woman came in. She sat down. (Stojnić 2021: 33)

The idea is that the indefinite noun phrase 'a woman' introduces an indefinite woman and places her, i.e. the implied witness to the existential first sentence of (1), at the center of attention. Technically, this means placing it at the top of the stack of values to the relevant assignment function. The pronoun 'she' in the following sentence will be assigned as value by the assignment function the female individual that is highest ranking on the stack. That is, in this case, the presumed witness to 'a woman'. In Stojnić's formalism, this is represented as follows.

(1′) ⟨α⟩;[*woman*(@)];[*came.in*(@)];[*sit.down*(@)]

The first element in this sequence, '⟨α⟩', is a dynamic existential quantifier. It changes the context by placing a new (witness) individual at the top of the stack. The semicolon represents context update. '@' is the formal variable that refers to the individual at the top of the stack. Hence, the second element, '[woman(@)]', predicates of the individual at the top of the stack that it is a woman. The following two updates predicate of the same individual that it came in and that it sat down.

Example (1) shows how an indefinite can introduce a new individual that is pushed to the top of the stack. Another linguistic means of doing that is by means of a demonstrative. Stojnić also has the alternative example

(2) A woman came in. She [pointing at a cat, Betty] sat down. (2021: 44)

Here, the pointing gesture induces another shift in attention and places the demonstrated individual at the top of the stack, demoting the indefinite woman to second place. Formally, this is handled by means of a demonstrative update operator which combines with a name of the individual pushed. As a formal representation of we then get

(2′) ⟨α⟩[*woman*(@)];[*came.in*(@)];[π*b*];[*sit.down*(@)]. (Stojnić 2021: 45)

Here '*b*' names the cat Betty. There is then a demonstrative update represented by '*πb*', after which the last clause, stating that that the top-ranked individual sat down, now refers to Betty the cat, not the indefinite woman, since Betty is now top-ranked (and female).

These are the basic elements of the formal account and they work well for the cases they handle. The problem is that there are other cases which are not handled well by this account. These are cases where the individual referred to by a pronoun do not have prominence *before* the utterance but acquires prominence *after* the utterance. This happens in cases where a referential but non-demonstrative pronoun occurs discourse initially, that is without linguistic antecedent. We shall look at a few examples.

(3) X: I will leave him.

The utterance, as said by X to a hearer Y, can be easily understood even if the referent has not been mentioned earlier in the conversation, nor been in a salient set of individuals that *have* been mentioned. We can easily sketch a scenario where the communication nevertheless easily succeeds. There are typically few people that a particular speaker X can potentially leave, fewer still for which there is common knowledge between X and hearer Y that the X has this relation to, and typically only one that would merit the information. In particular, only one that would be the obvious referent when referred to by a pronoun without antecedent.

The referent of 'him' is the highest-ranking male (distinct from the hearer) that has these properties. It need not have been the highest-ranking male *before* the utterance but it will have become the highest-ranking male *after* the utterance.

A second example:

(4) X: How was the conference? Y: She did it again.

In the case of (4), there can easily be common knowledge between X and Y who are the potential female referents in the domain related to the contextually salient conference, and also who is the most salient referent in this sub-domain with a record of repeating a noteworthy pattern of behavior or achievement.

That referent need not have been the highest-ranking female before the utterance but will have become the highest-ranking female after the utterance. In this case, the potential for being raised to prominence is enhanced by the choice of topic—a particular conference and common knowledge between X and Y about who attended—but the referent need still not have been the highest-ranking female before the utterance.

Even in the case of

(5) X: He is back.

there can easily be common knowledge between X and hearer Y of who is the most salient male in the category of having been *absent* and possibly having *returned* from that absence. That person need not have been the highest-ranking male before the utterance but will have become so after the utterance.

It seems therefore that there are many possible cases, and probably also many actual cases, where pronoun references are successfully communicated in a way that the prominence ranking theory cannot account for.

What would a successful account look like? One possibility would be to *complicate* the Stojnić attention theory into one that allows for a tree-like prominence hierarchy, with rank relations between individuals relative to different categories, properties, or relations.

It would be a very complex theory, with a prominence hierarchy for each relevant property; who is the highest-ranking male with property F, who with property $F \& G$, etc. Note that a linear ranking will not suffice: a person X may be more salient than a person Y with respect to property F, while Y still be more salient than X in relation to a property G. The theory would require, for explanatory power, a highly definite prominence structure. It would plausibly be a tree structure with no properties, or perhaps only Φ properties at the top. Such a theory is not impossible, but it is highly implausible that it would be a *linguistic* structure. Rather, the more natural idea would be that such a prominence or salience structure would be a complex feature of the non-linguistic context, and that the semantics can draw on the salience features of the context in the semantics for pronouns. Going in this direction would therefore be to move away from the linguistic convention position that Stojnić occupies.

Although this theory would not appeal to linguistic convention, it would still be a view close to the semantics-pragmatics interface. Just as the *context of use*, in the sense of Kaplan (1989), delivers values to automatic indexicals, like 'I', and to demonstratives, like 'that cat', so on the *salience tree* conception, the context would simply *provide* values to referential discourse initial pronouns.

Further out along the pragmatic road we find the position that Stojnić explicitly distances herself from. On an optional account in this position, what is going on is exactly the kind open-ended abductive reasoning that Stojnić in the book says is not normally taking place. Such an account would say that the hearer looks for the most salient individual commonly known to satisfy the conditions imposed in the exchange. The speaker, on the other hand, has an intended individual in mind, and implicitly takes the property ascribed to the individual to be identifying.

Since it is a free search on the hearer's part, however, there is in principle the possibility that the hearer fixates on an individual that satisfies the property constraint to a lower degree than the individual intended by the speaker but to a sufficiently high degree for the hearer

to stop searching. There is also the converse possibility that the speaker intends an individual that satisfies the constraint to a lower degree, but does not at the moment keep the other, more suitable referent, in mind.

Nevertheless, in many cases, communication succeeds without noticeable effort. In some cases, the intended referent might already have been salient, but in many cases, prominence changes as a result of the speaker's utterance, leading to an interpretation update, something akin to what Davidson (1986: 10) characterize as the *passing theory*, as opposed to the *prior theory*, of interpretation. This interpretation is *post hoc*, requiring a change of assignment of prominence to individuals. One version of such an account has it that this process, with an intended satisfier on the part of the speaker, and a search for a satisfier on the part of the hearer, always takes place. The cases where the intended and found referent was already prominent is just a special case, even if common, where the search is immediately successful.

The opposite view of the phenomenon is that of deeming the *post hoc* prominence cases *abnormal*, and outside the conventions of language. This is basically Stojnić's stance.

Concerning a somewhat different case of discourse initial use of pronouns ('She is happy'), Stojnić writes:

> Though admittedly, even in such cases it might sometimes be possible for the audience to eventually "figure out" what the speaker had in mind, the process by which that occurs is markedly different from the seamless interpretation of pronouns on the fly that we see in normal circumstances. The potentially open-ended reasoning about what the speaker wanted to convey kicks in precisely after one is faced with the infelicity of the utterance. The utterance cannot be properly interpreted on its own, so some kind of repair is needed in order to help guide understanding. My account would simply maintain that a part of the linguistic material in the utterance is missing, and the reasoning is about which material one would have to posit to arrive at an utterance that has a plausible interpretation in this case. (Stojnić 2021: 49–50)

Stojnić's picture is that occurrences of such referential but non-demonstrative, discourse initial pronouns are simply infelicitous, and that interpretation involves some kind of repair. If this view is wholly descriptive, i.e. without any normative verdict on appropriateness, it should be a matter of ease of interpretation, something which to some extent can be measured, e.g. by EEG studies of event-related potentials.

No doubt there are cases where uses of pronouns *are* infelicitous. Some speakers are prone to egocentric speech, in the sense of not taking the perspective of the hearer into account when using indexicals or ambiguous expressions. In such cases, the hearer can easily be at a loss of trying to figure out what the speaker is trying to say.

But there are also frequent examples of such discourse-initial pronouns where little or no additional effort of interpretation is noticed. The speaker's intended referent pops up immediately in the mind of

the hearer. That such cases are infelicitous is subject to debate. What is hardly subject to debate is that such cases occur and are not overly rare. They seem to be within the range of ordinary language use. As such, we should try to understand what makes them successful. Moreover, there are also cases, as we shall see next, of pronouns *with* discourse antecedents where there *are* ambiguities and difficulties of interpretation. The view that we have a clear separation between the convention-governed felicitous use of pronouns on the one and the aberrant and infelicitous pragmatics-needed use on the other, is at least controversial.

3. Coherence

Another main tenet of the book is that *coherence* relations (*discourse* relations), i.e. relations between propositions expressed in a sentence or discourse, determine the resolution of anaphora. Jerry Hobbs was one of the leading pioneers in discourse relations theory, and his theory is succinctly presented in Hobbs 1985. Andrew Kehler (2002) later followed and developed Hobbs's theory. After a suggestion by Hobbs, Kehler uses the categories of *connections between ideas* of David Hume (1748) as his basic categories of discourse relations: *Resemblance*, *Cause-Effect*, and *Contiguity*.

Kehler's *Resemblance* relations are *Parallel*, *Contrast*, *Exemplification*, *Generalization*, *Exception*, and *Elaboration*. As an example, *Parallel* is exemplified by (2002: 16):

(6) Dick Gephardt organized rallies for Gore, and Tom Daschle distributed pamphlets for him.

Here there is a relation of *doing something in support of*, a relation which subsumes both *organize rallies for* and *distribute pamphlets for*. So, the relation of *doing something in support of* holds both between Gephardt and Gore and between Daschle and Gore. In addition, Gephardt and Daschle have the shared property of being high-ranking democratic politicians.

This background of ideas in the theory coherence relations is employed by Stojnić as providing tools for the interpretation of pronouns. A central example of this phenomenon is the following (Stojnić 2021: 61):

(7) John was disappointed with Tim.
  a. He fired him.
  b. He disobeyed him.

How do we resolve anaphora in these examples? The suggestion is that in (7a), the coherence relation is that of *Result*: John fired Tim as a *result* of being disappointed with him. The anaphora resolution that follows from this coherence relation is that 'he' is resolved to John and 'him' to Tim.

In (7b), the coherence relation is that of *Explanation*: that Tim had disobeyed John *explains* why John fired Tim. The anaphora resolution

that follows from this coherence relation is that 'he' is resolved to Tim and 'him' to John. As this example shows, our interpretation how a discourse hangs together can be closely related to the interpretation of context sensitive elements in the discourse.

Prima facie, this appeal to discourse relations does not sit well with the emphasis on linguistic conventions over pragmatics, as coherence relations have typically been taken to belong to pragmatics. However, it is also a central tenet of the book (esp. 68–71) that coherence relations are instantiated in discourse precisely as a matter of linguistic convention.

Stojnić herself stresses that hers is a minority view. On this view, speakers do not infer the obtaining of coherence relations as a result of open-ended abductive reasoning. It is a feature of language itself. Central to the view is that in the determination of meaning of a sentence or discourse, the obtaining of coherence relations are established *first*. The resolution of anaphora follows. The view comes out clearly in her discussion of example (8) (Stojnić 2021: 64):

(8) Phil tickled Stanley. Liz poked him.

This discourse can be understood as exemplifying the *Result* relation: Liz poked *Phil*, and this action was prompted by the action described in the first sentence. 'him' is then resolved to Phil. It can also be understood as exemplifying the *Parallel* relation: Liz's action is *similar* to Phil's. 'him' is then resolved to Stanley. Concerning this case, Stojnić says:

> Note that general reasoning can still have a role to play, but this role, again, is not one of assigning content to the form, but rather one of disambiguating which form has been uttered to begin with. So, for instance, [(8)] is ambiguous between a form containing Result and one featuring Parallel. Some general reasoning might be invoked in disambiguating between these, much as it might be involved in figuring out whether a speaker means a financial institution or a river bank, with a use of 'bank,' or which quantifier scope is intended with a use of 'Every boy kissed a girl' or who they named when they uttered 'Betty,' or, as we have seen before, in disambiguating a particular form of a gesture. To interpret, a hearer must first settle disambiguations. (2021: 70)

One might think that disambiguation itself is a pragmatic operation but this is not so for Stojnić:

> But *disambiguation* is pre-semantic, in Kaplan's sense: it involves the interpretive work needed to settle *the linguistic form* of an utterance, not to assign content to the form (Kaplan 1989a). Disambiguation is distinct from semantic interpretation: it is only once the form is disambiguated that it can be semantically interpreted. Semantics determines what an expression means, but not which expression was uttered. And, while general pragmatic reasoning about the speaker's intentions and available epistemic cues plays no role in semantic interpretation, it *can* play a role in guiding the audience to recognize *which form* of a pointing gesture was performed. But this is the role they can play in the disambiguation of any ambiguity, for example, in

> the disambiguation of a use of the word 'bank,' or a name, 'Betty.' And, as with other ambiguities, conventions governing demonstrative actions constrain possible disambiguations. A flat hand shape with the palm up, fingers toward the audience, allows for a certain range of interpretations, but not others; similarly, for an extended index finger, or the word 'bank,' or the name 'Betty'. (Stojnić 2021: 55)

Of course, part of the general view of the book is that coherence relations belong to the *form* of an utterance. In the formalism, clauses like 'Explanation(x_0,x_1)' occur in the representation of logical form. Thus, from this perspective, determination of coherence relations in a discourse is a determination of logical form and this, in turn, is strictly speaking syntactic disambiguation.

However, this view of settling coherence strikes me as implausible: the interpretation of a discourse that involves settling coherence relations typically involves a hypothesis about the *propositions* that the coherence relations relate. In interpreting, we do not first settle on *Parallel* and then resolve the anaphoric relation. Coherence relates propositions, not propositional functions. Rather, we compare the package of *Parallel+him*$_{Stanley}$ with the package *Result+him*$_{Phil}$.

This means that determining coherence already includes the resolution of anaphora. It is not clear what would even be the *basis* for settling the coherence relation prior to considering the anaphoric relation. In some respects, it seems Stojnić agrees with this. The quoted passage ending with 'disambiguation' (2021: 70–1) continues:

> This may involve assessing the plausibility of possible coherence relations that could be operative in a given context. It may involve evaluating whether a particular disambiguation of coherence relations delivers a plausible resolution of pronouns.

This remark strikes me absolutely spot on. And it is not unique to coherence. We have similar phenomena in relation to implicature. Consider

(9) A: Are you coming to Martha's party on Saturday? B: My mother will be visiting.

We try to find a suitable overall interpretation of the answer, and we get this by interpreting B's answer as expressing the proposition that B's mother will be visiting *on Saturday*. This supports the relevance of the answer, via the implicature that B cannot come to the party.

This exemplifies what Stephen Levinson (2000: 186) has called "Grice's Circle": disambiguation and other determinations of *what is said* may depend on processes that "look indistinguishable" from implicature.

Nevertheless, Stojnić is adamant that settling the coherence relation has priority over the interpretation of pronouns. The quoted passage continues:

> In short, it can serve in recognizing which form of the available ones that grammar delivers was uttered, but not in determining which meaning a

> particular form takes on. Once a coherence relation is established, pronoun resolution is determined by grammar, not by general reasoning. And, as I have been arguing throughout, any extra-linguistic parameters such interpretive reasoning may invoke do not serve to *determine* meaning. (Stojnić 2021: 71)

Why this view? What supports it? Some of the motivation seems to come from a particularly striking example of coherence relations and anaphora, to be considered next.

4. *Parallel*

Stojnić uses an example from Andrew Kehler (2002: 159) which she takes to show that the attention-shifting operations prompted by coherence relations are grammatically encoded.

> They privilege linguistic ones, over the broader constraints of background knowledge and rational inference that they might potentially consider. (2021: 68)

The example is:

(10) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships her.

The point of the example is that (Kehler's) informants judge (10) *infelicitous* despite the availability of a gender-matching antecedent to the pronoun. It is judged infelicitous since the pronoun is gender incongruous with its expected antecedent, 'Ronald Reagan'. And the idea is that 'Ronald Reagan' is the expected antecedent because the sentence exemplifies *Parallel*. Hence, the argument goes, the coherence relation seems to be established *before* the anaphora resolution, and even trumps the grammatical incongruence. And this is so *despite* the fact that general reasoning can produce a reading that satisfies the congruence requirement.

This seems to speak in favor the pre-semantic status of establishing coherence. However, there is reason to suspect that the effect is not wholly, and not even predominantly, due to the *Parallel* relation. The adverb 'absolutely' functions as an intensifier in (10). It induces a certain expected stress contour:

(10′) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely WORSHIPS her.

The stress peak on 'worships' has the further effect that the pronoun 'her' is *deaccentuated*. Deaccentuation is often represented by underlining:

(10″) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely WORSHIPS her.

Deaccentuation is associated with avoidance of either of two features (Shapiro and Anttila 2021). One is semantic, the other phonological. It is, on the one hand, associated with avoidance of stress on expressions

with given, entailed, coreferential, or contextually accessible meanings. It is, on the other hand, also associated with avoidance of stress on the second of two segmentally identical strings.

Deaccentuation of an expression that is coreferential with an expression already given is coupled with contrastive stress on an immediately preceding expression. An example would be exactly

(11) Margaret Thatcher admires Ronald Reagan, and George W. Bush WORSHIPS him.

without the adverb 'absolutely'. The explanation for this is phonological/semantic. In (10"), we have the same stress contour of the VP 'worships' + pronoun as in (11), but the source is different. The similarity in stress contour leads to a conflation: the hearer/reader expects a similarity of meaning between the pronoun and its antecedent. That is, the deaccentuation itself *indicates* that there is a similarity, correctly in (11), incorrectly in (10"). Shapiro and Anttila say:

> What is fundamentally a phonological alternation has acquired a semiotic function. Deaccentuation is a signal that invites the hearer to establish a similarity between two strings. (Shapiro and Anttila 2021: 8)

I do not claim that the apparent instantiation of *Parallel* has nothing at all to do with the reported infelicity reactions to (10). The availability of an alternative explanation does not by itself license that conclusion. As Stojnić says (2021: 68), the *Parallel* instantiation requires that the antecedent of the pronoun in object position has an antecedent in object position. 'him' should corefer with 'Ronald Reagan' to instantiate it. Thus, if *Parallel* is *perceived* as instantiated, the expected antecedent should be 'Ronald Reagan'. But the availability of an alternative explanation that relies on phonology + semantics instead does undercut some of the explanatory force of the appeal to *Parallel*. Also, as indicated in the preceding section, why should the reader settle on *Parallel* in the first place, before the reference of the pronoun is determined, as there are then not yet two propositions to relate.

Furthermore, if we remove 'absolutely' from (10), we get

(12) Margaret Thatcher admires Ronald Reagan, and George W. Bush worships her.

(12) is much more felicitous than (10), despite instantiating *Parallel* as much as (10) itself. A stress on the pronoun is here to be expected.[2] Stojnić herself (2021: 68n) does acknowledge that stress does play a role: a stress on 'her' in *itself* would make if felicitous:

(10*) Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships HER.

Finally, however, Stojnić defends the conventionality of coherence relations by claiming that although (10*) is felicitous, it cannot exemplify

[2] During the Q&A after the talk in Dubrovnik, Stojnić denied that there is much difference in felicity between and (12) and (10).

Parallel. But this is far from obvious. As far as I can see, we could have a *different Parallel* instantiation:

(13) Conservative leaders tend to like conservative leaders. Margaret Thatcher admires Ronald Reagan, and George W. Bush absolutely worships HER.

(10*) occurs as the second sentence of (13), and is naturally read as instantiating the relation expressed in its first sentence, in two parallel examples. What is needed is just *some* available basis for *Parallel*.[3]

In conclusion, I deny that the Kehler example, (10), shows that coherence relations are established before the interpretation of pronouns, and that it shows that they can trump certain pronoun interpretations. The case for the conventionality of coherence relations remains to be made.

5. *Conventionality*

I have tried to make a case for the prevalence of constraint satisfaction in interpretation, both when it comes to the prominence ranking of pronouns and when it comes to the settling of coherence relations. My stance is that when constraint satisfaction is employed, it does determine everything that gets settled by means of it. Thus, when constraint satisfaction is employed in selecting a pair of a coherence relation and a pronoun resolution, they do get determined together, in a general pragmatic way.

By contrast, Stojnić's strategy is to separate core conventional parts of interpretation—that are strictly linguistic—from the general pragmatic ones. In the case of pronoun prominence, on her view, there are simply two distinct kinds of interpretation; the core linguistic one, based on prominence ranking, and the open-ended pragmatic one, which only kicks in when the first one fails.

In the coherence case, her view is that coherence relations *may* be established by open-ended reasoning, which can even involve considering the resulting resolution of anaphora, but coherence is nonetheless determined *first*, and anaphora determined as a consequence.

Stojnić tends to contrast open-ended reasoning with the demands of linguistic conventions: there can be a conflict with what convention *requires* and open-ended reasoning *allows*. This is taken to be especially exemplified with (10).

But an opposing view, that I myself have put forward in work on coherence (Pagin 2014, 2017, 2019), is that some pragmatic forces are in conflict with others. In particular, I have claimed that the demands

[3] During the Q&A Stojnić objected by saying that the coherence relation instantiated here is rather that of *Exemplification*: both conjuncts of the second sentence express propositions that exemplify the proposition expressed by the first sentence. This is correct, but the relation between the conjuncts themselves is still that of *Parallel*.

of *coherence* conflict with the demands of *charity* (maximizing truth potential): requiring an interpretation where propositions hang together makes it more difficult for them to be true. Many cases of pragmatic enrichment testify to this, for instance the following example from Robin Carston (2002: 71):

(14)

a. He handed her the key and she opened the door.
b. He handed her the key and she opened the door [*with the key that he had handed her*].

A normal and typical interpretation of (14a) associates with it a content that is more completely articulated in (14b), which includes additional linguistic material in brackets. Clearly, (14a) itself is true in more situations than (14b). Hence, the enrichment runs counter to charity but makes the two conjuncts cohere better. Linguistic convention seems not to be involved.

We can also get a similar example where charity trumps coherence:[4]

(15)

a. He handed her the key and she wiped the table.
b. He handed her the key and she wiped the table [*with the key that he had handed her*].

Since, by world knowledge, we can reject the possibility that one wipes a table with a key, the proposed enrichment is in conflict with the demands of *charity*, and indeed strikes one as odd. It will not easily be made. Thus, we do not need to appeal to *convention* to see that pragmatic considerations can be overridden. Pragmatic considerations can be in conflict with and override each other.

I do not claim to have refuted Stojnić's position. It seems perfectly viable. One can suspect that the two approaches are empirically equivalent, when it comes to the results of interpretation. One may hold that the more pragmatic view is simpler in respect of being more uniform when it comes to prominence, and in respect of saving a step in the case of coherence and anaphora. One may then wonder whether there are other considerations of interpretation that actually require or at least support the linguistic convention view. I leave this as a question for future research.

References

Carston, R. 2002. *Thoughts and Utterances*. Oxford: Blackwell.

Davidson, D. 1986. "A Nice Derangement of Epitaphs." In E. LePore (ed.). *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. London: Basil Blackwell.

Glüer, K. and Pagin, P. 2006. "Proper Names and Relational Modality." *Linguistics & Philosophy* 29: 507–35.

[4] Thanks to Fabrizio Cariani for requesting such an example.

——. 2008. "Relational Modality." *Journal of Logic, Language and Information* 17: 307–22.

——. 2012. "General Terms and Relational Modality." *Nous* 46: 159–99.

——. 2023. *Switcher Semantics. From Propositions to Ingredient Sense*. Under contract with Oxford University Press.

Hobbs, J. R. 1985. *On the Coherence and Structure of Discourse*. Tech. rep. CSLI-85-37. Stanford University: CSLI.

Hume, D. 1748. *An Inquiry Concerning Human Understanding*. 1955 edition. New York: Liberal Arts Press.

Kaplan, D. 1989. "Demonstratives." In J. Almog, J. Perry and H. Wettstein (eds.). *Themes from Kaplan*. Oxford: Oxford University Press, 481–563.

Kehler, A. 2002. *Coherence, Reference, and the Theory of Grammar*. Stanford University: CSLI Publications.

Levinson, S. C. 2000. *Presumptive Meanings. The Theory of Generalized Conversational Implicature*. Cambridge: MIT Press.

Pagin, P. 2014. "Pragmatic Enrichment as Coherence Raising." *Philosophical Studies* 168: 59–100.

——. 2017. "Radical Interpretation and Pragmatic Enrichment." *Argumenta* 3: 87–107.

——. 2019. "Enrichment, Coherence, and Quantifier Properties." *Journal of Pragmatics* 154: 92–102.

Shapiro, N. T. and Anttila, A. 2021. "On the Phonology and Semantics of Deaccentuation." *Proceedings of the Annual Meetings on Phonology* 8.

Stojnić, U. 2021. *Context and Coherence: The Logic and Grammar of Prominence*. Oxford: Oxford University Press.

Croatian Journal of Philosophy
Vol. XXIII, No. 69, 2023
https://doi.org/10.52685/cjp.23.69.3
Received: September 20, 2023
Accepted: November 13, 2023

From Coherence Relations to the Grammar of Pronouns and Tense

MAGDALENA KAUFMANN*
University of Connecticut, Storrs, USA

Stojnić (2021) argues that the content of linguistic utterances is determined by the rules of natural language grammar more stringently than what is generally assumed. She proposes specifically that coherence relations are encoded by the linguistic structures and determine what individuals count as most prominent, thereby serving as the referents of free ("demonstrative") pronouns. In this paper, I take a close look at the empirical evidence from English and Serbian that she offers in support of this position. Considering these data points in connection with additional linguistic data (also from German and Japanese), I argue that there is no compelling evidence for the assumption that coherence relations directly determine the resolution of pronouns. Instead, grammatical restrictions imposed by different types of pronouns and tenses have a larger impact on the meaning conventionally expressed by complex utterances than what is generally assumed in the literature on coherence relations.

Keywords: Pronouns; tense; coherence relations; information structure; discourse relations.

1. *Introduction*

In her 2021 book *Context and Coherence*, Una Stojnić develops and defends the claim that grammar determines the content expressed by

* This research was supported in part by the author's NSF grant #2116972 "Research on conditional and modal language". For discussion of data and theories, I am indebted to Ivana Jovović, Stefan Kaufmann, Teruyuki Mizuno, Rick Nouwen, Tyler Poisson, Philippe Schlenker, Julian Schlöder, Walter Shaw, and the participants of *Philosophy of Linguistics and Language*, Interuniversity Center Dubrovnik, September 5–9, 2022.

complex natural language expressions to a significantly higher degree than what is standardly assumed:

> [L]anguage–grammar–itself is far more expressive and pervasive than has been assumed; the resolution of context-sensitivity is entirely a matter of linguistic convention. (Stojnić 2021: 5)

In developing an account along these lines, she assigns a uniform, non-ambiguous interpretation to pronouns and to modals, which offers strikingly elegant accounts of quandaries for even a dynamic truth-conditional theory or apparent failures of otherwise intuitive inference patterns. Stojnić (2021) argues, moreover, that crucial linguistic conventions that determine content for the items in question...

> have gone unnoticed, because their principal domains are entire discourses, not just their constituent words and sentences. While it is not controversial that the way sentences are constructed depends on conventions of syntax and semantics which specify the rules by which individual expressions combine, I similarly argue for rules–*discourse conventions*–that specify how individual sentences combine to form a discourse. (Stojnić 2021: 5)

It is this latter point that I take issue with: I do not think that Stojnić (2021) establishes convincingly that coherence relations directly address the resolution of pronouns, and I would like to defend the more traditional picture of the interplay between pronominal resolution and coherence relation resolution in parallel, a position Stojnić explicitly rejects.

Reconciling the points with Stojnić's (2021) main claim as reproduced initially, I would like to argue that it is rather the grammar of pronouns and tenses that is richer than what has been assumed in some parts of the literature. Restrictions imposed by the grammar of these elements will thus reduce the choice of otherwise possible (that is, salient enough) referents for a given pronoun; concerning the remaining options, the choice has to be made between pairings of compatible coherence relations with the content resulting from the alternative pronominal resolutions.

2. *A unified interpretation for pronouns*

At first glance, Stojnić's take on pronouns may sound like a standard credo in linguistic semantics:

> The meaning of a pronoun is simple, uniform, and unambiguous; as a first pass, a pronoun selects the most prominent candidate interpretation—what is "at the center of the attention" at the point in discourse at which it occurs. (2021: 40)

Upon closer inspection, this is probably better considered a standard desideratum that most accounts tacitly stop short of delivering on.

With the advent of dynamic semantics (specifically, *file change semantics*, Heim 1988; *Discourse Representation Theory* (DRT, Kamp and Reyle 1993; *Dynamic Predicate Logic*, Groenendijk and Stokhof 1991), it has become standard to assume that specific linguistic phe-

nomena require us to take into account that context evolves between utterances and also over the course of a single utterance. These developments affect the interpretation of subsequent linguistic material, a phenomenon that is captured by reconceptualizing linguistic meanings as *context change potentials*: functions from input contexts to output contexts, and Stojnić is firmly rooted in that tradition.

Dynamic notions of meaning have been put to use fruitfully when explaining coreference across sentences (e.g. (1b)) and covariation effects within (e.g. (1a)), standardly indicated by coindexation:

(1) a. If a boss$_i$ has employees$_j$ she$_i$ should treat them$_j$ well.
b. A woman$_i$ met a girl$_j$. She$_{i/j}$ greeted her$_{j/i}$.

In capturing such phenomena, dynamic theories also seek to explain that the availability of an antecedent or binder is constrained by the linguistic context; negation, for instance, blocks phenomena as illustrated in (1):[1]

(2) a. Mary does not have a violin. #She will (not) bring *it* to the party.
b. #If Mary does not have a violin, she will not bring *it* to the party.

Despite their impressive traction for explaining these phenomena, when it comes to a treatment of the pronouns themselves, mainstream versions of dynamic semantics tacitly accept infinite ambiguity. Pronouns are translated as variables to be assigned referents by assignments. Treating contexts as sets of assignments or world-assignment pairs[2] provides a way of modeling limited information as associated with an intended referent (Heim 1982, Kamp and Reyle 1993, Groenendijk et al. 1996). For instance, the discourse in (1b) can be integrated successfully even when the interlocutors are unable to identify more precisely the two individuals the speaker has in mind.

(3) a. A phonetic form [hi:] realizes the lexical element he_i for some $i \in$ N.
b. For any $i \in$ N, he_i is translated as x_i.
c. For any $i \in$ N, the interpretation of x_i with respect to a variable assignment g is $g(i)$ (an element of the domain of individuals provided by the model of interpretation).

If we assume moreover that assignments are undefined unless a suitable referent has been stored at an index, then the treatment sketched in (3) provides significant restrictions on what indexations could be available at any given point. Pronouns themselves, however, are still associated with infinitely many possible translations.

A uniform interpretation as 'the most prominent male/female' would be preferable but does not intuitively match the behavior of the pronominal elements. For instance, in the individual domain, Stojnić (2021) shows that in (4a) (in the absence of contrastive intonation),

[1] Dynamic theories also capture the similarity between possible pronominal resolution patterns and presupposition projection (van der Sandt 1992).

[2] Contexts can also be treated as sets of sets of assignments to model plurality, (see van den Berg 1996; Brasoveanu 2006).

he has to refer to John even at the expense of infelicity. In contrast, the definite description *the contextually salient male* is understood straightforwardly as referring to a different individual in (4b).

(4) a. #John came to the party, and he did not come.
b. John came to the party, and the contextually salient male did not come.

An attempt at analyzing the pronoun as picking out the contextually most salient male thus seems doomed. This reasoning, however, is fallacious. By developing a formal account, Stojnić shows that an implementation of this idea need not behave like the object language definite descriptions we can use to describe it, i.e., English phrases like *the contextually salient female*/*the contextually salient male*. Technically, for her, interpretation proceeds with respect to a sequence of possible referents (the *attentional state*), which Stojnić calls a stack, emphasizing the special status of the top-most element.[3] Pronouns are translated as strings beginning with '@', and these expressions are interpreted as picking out the top-most element of the stack that meets the requirement imposed by their lexical content (e.g. being male for *he*, or being female for *she*: these pronouns are translated as *@he* and *@she*, respectively). Specific assumptions about how the stack is affected when updated with a simple predication as expressed by the first conjunct in (4a) ensure that at the point where the pronoun *he* is encountered, the top-most element of the sequence that meets the requirement of being male is John. Specifically, the subject referent gets stored in the top-most position by default, followed by the object referent for a two-place predicate. Pronouns with suitable features are therefore expected to refer to the subject by default. And indeed, subject orientation seems a reasonable default for English pronouns.[4]

[3] Other positions are accessible directly via indexation, making the formal object behave more like a sequence proper.

[4] It is worth noting that the use of stacks to provide an index-free (and thereby semantically unambiguous) treatment of pronouns is not unprecedented. Motivations and data coverage differ across the approaches that come to mind. For instance, van Eijck (2001) develops a stack-based dynamic semantics for pronouns to address the issue of *destructive assignments*: in dynamic accounts, indefinites are taken to introduce new referents. If coreference and binding is handled through indices (variables), a separate requirement has to ensure that the process uses a fresh variable (and does not therefore overwrite the information associated with a variable that is already in use). Nouwen (2003) observes that placing indices in the structure struggles to account for certain patterns of plural coreference in discourse as exemplified in (i) (his (5.8)): despite the difference in interpretation, the plural pronoun *them* has the same antecedent *exactly two papers*.

(i) Three students each wrote exactly two papers. They each sent them to L&P.
a. 'each student sends all the written paper'
b. 'each student sends just their own papers'

Schlenker (2005) develops a stack-based semantic version of classical binding theory. Modal Centering Theory as mentioned in Stojnić (2021) develops related accounts specifically to capture phenomena in the modal and temporal domain (Bittner 2011; Murray 2014). In contrast, sign languages might provide evidence for the existence

Stojnić (2021) does not work out an account for definite descriptions like *the contextually salient male*. As a first pass, these could be analyzed similarly: they would then pick out the top ranked referent in the attentional state that meets the descriptive content (of being a contextually salient male).[5] At first glance, this seems to bring back the issue we set out to avoid—the definite description should behave the same as the pronoun, and the contrast in (4) would remain unaccounted for. This, however, fails to take into account that the technical notion of being listed at the top of the stack thanks to grammatical mechanisms is crucially different from the notion of being the individual that counts as contextually most salient. For instance, there might be a particular individual nobody dares to name but which dominates a conversation about other individuals. Or else, the discourse goal (the question under discussion) may consist in characterizing a particular individual (maybe regarding their chances in an upcoming election), making that person conversationally most salient in this respect. But in doing so, the individual sentences may ascribe properties to their aides, thus associating these individuals with the top position of the stack temporarily. In short, the predicate *conversationally (most) salient* need not be interpreted as true only of the individual at the top of the stack with respect to which interpretation proceeds. The difference between the English phrase *conversationally most prominent* and the technical notion of being at the top of the stack would probably become most obvious when extending Stojnić's (2021) account to quantifier bound pronouns, as in *Every student thinks that he is smart*. A natural way of doing this while retaining the uniform semantics of the pronoun would be to 'borrow' the top-position to run over all the values in the domain, which would of course not lead to each individual becoming conversationally most prominent at least for a split second (see also Stojnić 2021: 42, for a suggestion along these lines). Beyond this, Stojnić's interpretation with respect to an attentional state offers the possibility to directly compare the behavior of pronouns and definite descriptions. As *R(eferential)-expressions* (Binding Theory, Chomsky 1981), definite descriptions are generally expected not to appear in the scope of a co-indexed (binding or coreferential) expression. It is therefore not implausible to assume that they come with a restriction or at least a bias against picking out referents that are salient enough to be accessible for pronominal reference.[6] Systematic exceptions, on which definite descriptions behave

of indices in natural languages (Lillo-Martin and Klima 1990; Schlenker 2018). A comparison of these different stack-theoretic accounts (as well as variable-free approaches more in general, Jacobson 2012) and an evaluation of the arguments from sign linguistics have to be left to future research.

[5] An analysis along these lines is provided by von Heusinger (2004), who models the changing referential prominence of various referents of a NP as choice functions and identifies *he* with *the male* for this purpose.

[6] A treatment along these lines is given in Schlenker's (2005)–also stack-based–semantic implementation of the Binding Principles. Definite descriptions and proper

like anaphoric pronouns, result, however, when the definite description is deaccented as in (5b). In this case, *the shed* has to be interpreted as anaphoric to the cottage, which is thus additionally classified as being merely a shed. In contrast, the accented occurrence in (5a) cannot be interpreted as anaphoric to John's old cottage: with pitch accent on the noun phrase, a new referent is introduced (Umbach 2002, her (1); *bridge anaphora*):

(5) (John has an old cottage.)
a. Last summer he reconstructed the SHED.
b. Last summer he RECONSTRUCTED the shed.

Related effects with definite descriptions in the place of anaphoric pronouns occur in newspaper jargon, where, as Riester (2009) points out, they serve to activate contextually available information that has not yet been discussed (or which is introduced by accommodation):

(6) Gerhard lives in Munich. The father or triplets is 42 years old.

The interaction with prosody is reminiscent of an effect Stojnić observes for pronouns: (4a) can become felicitous when the pronoun is stressed (see discussion in Sect. 4.1).

Ultimately, a full assessment of Stojnić's predictions for the contrast in (4) awaits an application of definite descriptions into the framework. The system seems flexible enough, however, to capture both the similarities and the differences between pronouns and definite descriptions like *the contextually salient male*.

Building on the idea that modals come with contextually supplied domain arguments, Stojnić can reduce the contrast in (7) to a similar problem (Stojnić 2021: 100, her (62) vs. (63)):

(7) a. #If it is not raining and it might be raining, then I am uninformed about the weather.
b. If it is not raining and the body of information *i* does not rule raining out, then I am uninformed about the weather.

Crucially, (7a) 'is not incoherent because there is no body of information that the context can select that yields a plausible interpretation', Stojnić (2021: 121) attests; 'rather, it is incoherent because the context determines the body of information that delivers an inconsistent interpretation' (Stojnić 2021: 121). Findings like (7a) have been adduced as evidence for non-propositionalism about epistemic modals. Stojnić's response is decisive: 'the problem is not, [...], in the idea that these expressions express truth-conditional content; the problem is in the underlying assumption of how a context operates to determine these truth-conditions'. (2021: 7)

As shown for the individual domain, Stojnić's specific implementation of 'most prominent/contextually salient' as the top-ranked one

names are subject to *Avoid Redundancy*, which blocks their use for referents already stored in the sequence of individuals that have been introduced in the linguistic discourse.

of all referents tracked in the on-going conversation avoids the issue elegantly. As a sort of corollary to her main thesis about the power of linguistic convention in the determination of content, we should thus also embrace a word of caution: "Beware of metalinguistic naivete!". Linguistic expressions come with various layers of static and dynamic meaning; faithfully paraphrasing formal accounts in English, our informal metalanguage, is tricky and cannot be taken to refute these accounts.

3. *Coherence relations and pronominal resolution*

Standard dynamic theories focus on what referents are available to be picked up by pronouns at any given point of the conversation. To allow for this, different syntactic objects are standardly differentiated by coindexation (or suitable choices of variables when translating to a formal representation language, Nouwen 2003 for discussion). For (1b) the two salient options that are felicitous out of the blue are given in (8), where indefinites are taken to modify assignments by storing individuals that have the noun phrase property (being a woman and being a girl, respectively) under the indices they carry.

(8) a. A_1 woman met a_2 girl. She$_1$ greeted her$_2$.
b. A_1 woman met a_2 girl. She$_1$ greeted her$_2$.

Stojnić (2021) argues that it is a genuinely grammatical phenomenon that in such cases only one resolution is available, and that content is therefore determined by grammar beyond what is usually assumed. Specifically, the burden is placed on coherence relations. Operative at the level of entire texts, these determine what referent pronouns like *that* or *she* pick out on any given occasion of use. The effect is illustrated with the classical example in (9) (modified from Smyth 1994: see also Hobbs 1979; Kehler 2002; Asher and Lascarides 2003 for discussion of this and similar examples):

(9) Phil tickled Stanley. Liz poked him.

The second sentence in the sequence can be understood as specifying a result of the event in the first (perhaps a show of disapproval), or as an event in parallel to what is described by the first. The resolution of the pronoun *him* stands and falls with the resolution of the coherence relation to Result or Parallel: we understand *him* as referring to Phil when we take the poking to be a result of the tickling, and to Stanley when we conceive of the tickling and the poking as two parallel events.

> [M]ost extant theories treat it as a pragmatic default. Standard coherence theoretic accounts interpret this correlation as evidence of an inferential relationship between a speaker's intention in organizing the discourse and her referential intentions. I argue that this is a a mistake: there is a tighter connection between discourse coherence and pronoun resolution, one underscored by linguistic convention. [...] I argue that discourse relations that connect and organize utterances are a part of the grammar of a language,

> and that they govern the resolution of context-sensitivity as a matter of grammar, too. (Stojnić 2021: 6)

Once a discourse relation is inferred, it determines the resolution of the pronouns by linguistic convention. Emphasizing the difference to the standard picture, Stojnić writes:

> There is reason to think that it's a mistake to treat (3) [our (9)] as harboring *two separate* ambiguities, or *two separate* underspecified elements that must be resolved in turn—one involving discourse coherence, another concerning pronoun resolution. The examples suggest that, once a coherence relation is established, a certain pronoun resolution is automatically set up. [...] there's good reason to conclude that pronoun resolution is *settled* by whichever coherence relations organize a discourse. (Stojnić 2021: 65)

As much as I agree with Stojnić that grammar has a significantly more decisive say in pronominal resolution than what is standardly assumed, I do think that this is a move in the wrong direction: there is good reason to think that the standard account has it right. While discourse relations are known to interact with, and to constrain, pronominal resolution, Stojnić's arguments that coherence relations asymmetrically determine pronominal resolution do not strike me as convincing. I will ultimately argue that the grammar of pronouns and the grammar of tense have a larger role to play than much of the standard literature on pronominal resolution acknowledges. Once the standard picture is updated to reflect this,[7] the data discussed in Stojnić (2021) can be integrated into the standard picture very naturally.

To explore this, like Stojnić, I will assume that, when interpreting utterances in context, language users integrate them into the given linguistic discourse that can be represented in Stojnić's translation language or also in the one of *Segmented Discourse Representation Theory* (SDRT, Asher and Lascarides 2003). Predicates representing coherence relations are a crucial building block in this (according to the SDRT convention, they appear typeset in small caps in the following).

Stojnić emphasizes that the construction of a coherent discourse is subject to constraints similar to what we observe at the sentence level; consider for instance well-known Principles A and B of binding theory (Chomsky 1981). While the reflexive pronoun *himself* in (10a) has to be coreferential with a c-commanding expression in the relevant binding domain (roughly, the smallest domain containing also a subject), the personal pronoun *him* in (10b) cannot be coreferent with another expression in this domain. At the level of interpretation, we obtain the effect of obligatory coreference and disjoint reference as illustrated in (10).[8]

[7] Note that especially DRT already includes very fine-grained studies of temporality.

[8] See Heim (1993) for the need to make the constraint sensitive to presupposed coreference. Consider her example in (i):

(i) Zelda must be the author. She praises her to the sky.

When coreference is not taken for granted but is established in the very sentence, the reflexive pronoun is felicitously replaced by the regular personal pronoun.

(10) a. Phil is shaving himself.
b. Phil is shaving him.

Similarly to this sentence internal constraint on the interpretation of object pronouns as depending on the referent of the subject, specific coherence relations can indeed go hand in hand with constraints on the resolution of personal pronouns. It is, however, far from clear that this requires the coherence relations to directly modify the ranking that underlies pronominal resolution. For instance, the coherence relation PARALLEL enforces a particular interpretation of the pronoun. On the SDRT account this follows, because two sentential structures can be related by PARALLEL only if their arguments match (Asher 1993).[9] Resolutions to individuals other than Stanley fail to instantiate a structure that could count as parallel to the one assigned to the first sentence *Phil poked Stanley*, and the two utterance units can thus not be connected with the relation PARALLEL. Therefore, PARALLEL imposes a constraint on the resolution of the pronoun also on the standard picture, but this constraint is indirect: there is no need to associate PARALLEL with an effect on the stack representing the referents that have been introduced in the discourse; the requirement it imposes (that the predication expressed by two units be parallel in the part of the predicate and each of the arguments) can be met only under one specific resolution of the pronoun and it can therefore not co-occur with the other option.

Other discourse relations, like RESULT or EXPLANATION lack even an indirect formal effect of that sort: they can be inferred as long as the contents of the two sentences under whatever resolution of the pronouns can plausibly be conceived of as standing in the relevant relation; restrictions are imposed, however, on the temporal order of the eventualities involved. For RESULT, the first discourse unit has to describe an eventuality that causes and hence precedes an eventuality described by the second (Asher and Lascarides 2003: 155). For EXPLANATION, the eventuality described by the first discourse unit may not precede the eventuality described by the second. If the second describes an event, this has to strictly precede the eventuality described by the first discourse unit (Asher and Lascarides 2003: 159).

In contrast, Stojnić (2021) argues for a direct connection between coherence relations and pronominal resolutions. She maintains that, in the following discourse, they stand and fall together: if the second sentence is understood as describing a result, its subject pronoun *he* is resolved as referring to John (the matrix subject); if it is understood as an explanation, the subject pronoun *he* is resolved as referring to Tim (the matrix object). To derive this, a representation of the discourse that relates the translations of the two individual sentences by EXPLANATION comes with an operator that reorders the attentional state by

[9] For instance, building on Kehler (2002), Altshuler and Truswell (2022) require that for two discourse units to be related by parallel, they have to share a common theme and all elements that differ have to be similar.

promoting the object over the subject. The intuition seems to be that 'something about the object' explains John's disappointment. In contrast, RESULT maintains the default order with subject prominence.

(11) John was disappointed with Tim.
a. He fired him. RESULT
b. He disobeyed him. EXPLANATION

In contrast, the standard picture (as for instance SDRT) allows for all possible combinations that are compatible with the requirements of the individual coherence relations (at least four in this case), the most coherent one of which will be selected by language users (following the principle of *Maximize Discourse Coherence* (MDC) as laid out in Asher and Lascarides 2003, their section 5.10). Crucially, this selection operates on the different options of pairing coherence relations with pronominal resolutions, rather than letting a coherence relation directly determine the pronominal resolution.[10] In short, according to SDRT (which represents the standard theory in this respect), interpreting a discourse thus involves choosing between all possible combinations of coherence relations and pronominal resolutions. According to Stojnić (2021), interpreting a discourse relation involves choosing between various possible coherence relations (characterized in parts by their impact on pronominal resolutions).

In the following, I will turn to two pieces of empirical evidence that Stojnić offers in support of the idea that coherence relations themselves determine pronominal resolution. I will argue that they rather support the standard theory as reflected in the SDRT framework. In section 4.3 I discuss a third empirical phenomenon which, despite at first glance supporting Stojnić's conceptualization, ultimately also supports independence of coherence relations and pronominal resolution.

4. *Disentangling coherence relations and pronominal resolution*

4.1 *Crosslinguistic variation in coherence relations?*

While processes relying on general purpose reasoning are standardly expected to be invariant across languages, effects tied to linguistic conventions are known to vary across languages, making data along these lines important arguments in favor of the one or the other. Aiming to support the assumption that coherence relations are part of grammar rather than of general all-purpose reasoning, Stojnić (2021) argues that

[10] The discussion in this section simplifies the SDRT picture by ignoring another structural restriction imposed on the resolution of pronouns by the distinction of coordinating and subordinating coherence relations. This difference impacts what sites are, at any point in the discourse, available for the attachment of a subsequent sentence (*Right Frontier Constraint*) that is, with what previous sentence it can be connected. The attachment site then constrains further what referents are available to resolve pronouns to at any given point in the discourse.

the impact of coherence relations on pronominal resolution varies even across Indo-European languages. While English shows the determination by coherence relation discussed above, Serbian does not. Serbian can realize subject pronouns overtly or covertly. Independently of the coherence relation inferred between the two sentences, in subject position, Serbian overt pronouns pick up the object of a previous sentence, covert pronouns pick up the subject of the previous sentence (Stojnić 2021: 69, (48)).[11]

(12) Džon je bio razočaran Timom.
John-NOM is-PRS-3MS be-PPA-3MS disappointed-ADJ-3MS Tim-INS
'John was disappointed with Tim.'
a. Otpustio ga je.
Fired-PPA-3MS him is-PRS-3MS
'He fired him.' (= 'John fired Tim.')
b. On ga je otpustio.
He him is-PRS-3MS fired-PPA-3MS
'He fired him.' (= 'Tim fired John.')

Obviously, Serbian coherence relations cannot determine prominence ranking (contrary to what Stojnić has argued for English); if they did, they would overwrite the grammatical information the pronouns need to access. Stojnić (2021) considers this difference between English and Serbian evidence for the linguistic/conventional nature of CRs.

Upon closer inspection, however, it seems more natural to locate the variation in the pronouns. Across languages, different pronominal elements are well-known to differ in interpretative possibilities: *weak pronouns* (including covert pronouns and clitic pronouns) appear, for instance, as bound variables or anaphoric to linguistically salient elements, whereas *strong pronouns* pick up other referents and can typically not be bound (see Patel-Grosz and Grosz 2017 for overview and references to earlier work). For English, Stojnić herself shows a difference in behavior depending on the prosodic contour realized on the pronoun: with unmarked intonation, (13) cannot be read as referring to Mary, no matter how much Mary's jumping up and down and yelling loudly may make her the most salient female in the conversational context and the center of our attention. As pointed out in Stojnić (2021) fn. 14, stress on the pronoun in combination with a pointing gesture will make Mary the referent (note that pointing at Mary is still infelicitous if the pronoun is not stressed). Stress on the pronoun in the absence of a pointing gesture may still retain Betty as a referent if (13) is used to correct a previous utterance about some other female sitting down.

(13) Betty came in, and {she/ SHE} sat down.

[11] Serbian clitics have to appear in second position, in (12), they therefore either follow the verbal participle or the overt subject pronoun (which is not a clitic). The resulting change in word order is orthogonal to her point.

While these interpretative changes go hand in hand with the phonological shape of the pronoun as well as changes in the overall discourse setting (like previous commitments with respect to which individual sat down), they are perfectly compatible with keeping the discourse relation constant (the two conjuncts in (13) are naturally related by narration, for instance).

For German, relevant contrasts are discussed as obtaining between regular personal pronouns and *d(emonstrative)*-pronouns[12] (Patel-Grosz and Grosz 2017, their (55)):

(14) a. Hans$_1$ wollte mit Paul$_2$ joggen, aber er$_{1,2}$ war krank.
Hans wanted with Paul jog but he was sick
b. Hans$_1$ wollte mit Paul$_2$ joggen, aber der$_2$ war krank.
Hans wanted with Paul jog but DEM was sick
'Hans wanted to go jogging with Paul, but he was sick.'
(adapted from Bosch et al. 2003)

Here, too, without a change in coherence relations (in both cases, the second coordinate stands in contrast to the first, as cued by *but*, and specifies an obstacle to the realization of the desire), we find a difference in pronominal resolution: the personal pronoun can pick up either Hans or Peter, while the d-pronoun can only be resolved to the non-subject participant (Peter). Hinterwimmer (2015) states this as an information-structural constraint: d-pronouns cannot refer to aboutness topics, which is the discourse status typically assigned to referential, unstressed preverbal subjects. The findings from Serbian are best understood along these lines, as well. Jovović (2022, 2023) explores the occurrence restrictions of overt and covert subject pronouns in Serbian. In non-subject positions, where covert pronouns are unavailable, the contrast is replicated by strong pronouns in contrast to clitics. Like the German d-pronouns, Serbian overt pronouns (or more generally, strong pronouns) cannot be anaphoric to a topic.[13]

The relevant linguistic conventions and crosslinguistic differences are thus better located in the grammar of the pronouns than in the grammar of coherence relations. A move along these lines reflects independent findings that different types of pronominals (covert pronouns, reflexives, phonologically reduced pronouns, full pronouns, stressed pronouns) differ in terms of resolution possibilities (Montalbetti 1984;

[12] Stojnić (2021) labels all free personal pronouns demonstrative pronouns; in the linguistic literature, the term is sometimes restricted to specific types of pronouns that are more closely related to pronouns like *this* or *that* (e.g. in German the pronoun series that is homophonous to the definite determiner).

[13] In Jovović's rendering, strong pronouns are acceptable in three environments: (i) when the pronoun's antecedent bears new information focus, (ii) when the pronoun itself is constrastively focused, or (iii) in the presence of a co-sentential focalized adverb. Jovović mostly aims to show that acceptability contrasts that were previously considered to involve syntactic violations (Despić 2011) are better analyzed in terms of the information structural requirements of different pronominal types. She does not develop an analysis of the patterns described.

Patel-Grosz and Grosz 2017; see also the binding principles as exemplified partly in (10)).

4.2 *Different explanations in English*

In (11), repeated here for convenience, we see a clear effect towards a preference of pronoun resolution depending on result, which keeps the default prominence of the subject, and explanation, which promotes the object.

(11) John was disappointed with Tim.
  a. He fired him. RESULT
  b. He disobeyed him. EXPLANATION

However, Stojnić (2021) points out that explanation need not always promote the object. It is therefore proposed that explanation comes in two variants, depending on what argument is promoted to be the most salient one.

(15) The city denied the demonstrators a permit.
  a. They feared violence. $\text{EXPLANATION}_{\text{Subject-based}}$
  b. They advocated violence. $\text{EXPLANATION}_{\text{Object-based}}$

Stojnić adduces contrasts with overtly encoded causal relations as independent evidence for this distinction: similarly to the two different explanation-relations, lexical verbs induce an implicit bias towards the causally implicated agent.[14]

(16) a. Sue frightened Mary because she was boisterous.
  'because Sue is boisterous' (*she* = Sue/subject)
  b. Sue feared Mary because she was boisterous.
  'because Mary is boisterous' (*she* = Mary/object)

Upon closer inspection, however, we find that the preference for pronominal resolution is determined not so much by the higher verb and the agent it causally implicates, but rather by the content of the propositions expressed by the embedded clause (with the pronoun resolved to either participant):

(17) a. Sue frightened Mary because she found her boisterous.
  'because Mary finds Sue boisterous' (*she* = Mary/object)
  b. Sue feared Mary because she found her was boisterous.
  'because Sue finds Mary boisterous' (*she* = Sue/subject)

The *because*-clauses with the evaluative predicate *find* are naturally understood as an explanation for the emotional state that is ascribed to a participant in the matrix clause. The pronoun *she* is thus resolved to the emotionally impacted participant, not the causally implicated one. We thus find the oppositive preferences for the subject pronoun *she* of the embedded clause even though the matrix clause as well as the

[14] In (16), I replace Stojnić's (2021) original choice *scary* with *boisterous* to at least reduce the bias from the lexical predicate which is applied to the pronoun in question.

coherence relation (overtly cued by *because*) have remained the same. This suggests that contrary to our first intuitions about (11), explanation does not determine whether the subject or the object is the most salient referent.

More in general, when playing with the contexts also for these sequences, it is possible to overwrite the initial preferences also for these specific examples while keeping the coherence relations constant. For (11a), imagine that Tim is a very uncertain person and, as a boss, does not take it well if he notices that someone is disappointed with him. John lets it transpire that he is disappointed with his boss, therefore Tim fires him. In this case, result is paired with a resolution of the second sentence's subject to the first sentence's object (contrary to the original intuition for the example). (11b) is harder to understand in a way such that an eventuality of John disobeying Tim is the reason for John's disappointment with Tim. Even if we imagine that John is disappointed with bosses that are weak enough to tolerate disobedience, this situation would be expressed preferably with an overt indication that the disobedience preceded the disappointment as in (18):

(18) John was disappointed with Tim. He had disobeyed him.
EXPLANATION, ok: *he* = John; ok: *he* = Tim

If the pronoun is anchored to the object, speakers apparently do not require this indication of temporal precedence. One speaker, however pointed out that they strongly prefer (18) to (11b) even on the reading discussed in Stojnić (2021), where the subject pronoun of the second sentence is resolved to Tim, the object of the first sentence. These last observations suggest that the initially dispreferred resolution of the pronoun cannot be paired up with a dispreferred resolution of the temporal ordering. While this phenomenon requires further investigation, it puts on the map tenses as yet another player in the game of how to resolve coherence relations and pronominal reference. All in all, the data discussed in this section strongly suggest that, even in English, coherence relations do not asymmetrically determine pronominal resolution.

4.3 *Might Japanese be coherence dominant?*

Stojnić (2021) aims to show that languages differ in the grammatical rules associated with particular coherence relations (specifically their impact on pronoun resolution). She uses this as an argument for the position that coherence relations form part of natural language grammar. In Sect. 4.1, we have seen that the evidence from Serbian is not entirely compelling in this regard. It might be interesting to investigate Japanese as a language which appears to display strong connections between coherence relations and grammatical markers, for instance in the inventory of conditional connectives (Arita 2004; Takubo 2020).

In English, sequences of eventive sentences in simple past tense can be related by EXPLANATION (as in (19b)), which requires the second event

(the pushing) to precede the first (the falling).[15] This contrasts with a discourse related by NARRATION or RESULT, in which the events are presented in the order of occurrence.

(19) Max fell.
a. John helped him up. NARRATION/RESULT
b. John pushed him. EXPLANATION

In Japanese, however, Explanation appears to require overt marking by the copula construction with *no da*, roughly 'it is...' (Kaufmann 2020: 416).

(20) Makkusu ga taore-ta.
Max NOM fall-PST
'Max fell.'
a. Zyon ga osi-ta.
John NOM push-PST
'[Then] John pushed him.' not: EXPLANATION
b. Zyon ga osi-ta no da.
John NOM push-PST NMLZ COP.NPST
'[That's because] John pushed him.' ok: EXPLANATION

Upon closer inspection, however, the issue seems to be with the temporal order of events, and not about a requirement to overtly mark a coherence relation: inserting a perfect marker makes the explanation reading available even without the nominalizer, cf. (21).

(21) Zyon ga osi-tei-ta (no da).
John NOM push-PERF-PST (NMLZ COP.NPST)
'[That's because] John pushed him.' ok: EXPLANATION

In this case, the crosslinguistic difference is probably better placed on the behavior of the tenses than on the behavior of the coherence relations. As evidenced above, a sequence of sentences in the simple past with eventive predicates is naturally interpreted in terms of temporal progression, (Kamp 1979; Kamp and Reyle 1993) but can be interpreted as going back in time to the cause of the event described by the first sentence. The temporal order presumed stands and falls with the choice of coherence relation: RESULT (and NARRATION) require one ordering, EXPLANATION requires the other. Things are different, however, in Japanese. We can make sense of the Japanese data if we assume that event progression is strict, that is, sequences of sentences in the simple past tense shift forward the time at which the events are located (the *reference time*, Reichenbach 1947). This is schematized in (22):

(22) For Japanese:
$Clause_1$-Past. $Clause_2$-Past.
entails: event-time($Clause_1$) < event-time($Clause_2$)

[15] As pointed out at the end of Sect. 4.2, not all speakers of English seem to be equally permissible in this respect.

As explanations cannot temporally follow their explanandum (Asher and Lascarides 2003 for restrictions on the parameters associated with discourse relations), relating these clauses by EXPLANATION is inconsistent with the very interpretation of EXPLANATION.

In contrast, the construction with nominalizer and copula construction (*no da* 'it is (the case) that') but also with the perfect *tei* describe the $Clause_2$-event as concluded at utterance time. The utterance time follows the time at which the event in the first clause is anchored, which renders the sequences in (20b) and (21) felicitous. The apparent effect of coherence relations is thus better understood as deriving from the grammar of tense.[16]

5. *Conclusions*

Stojnić (2021) maintains that the conventionally encoded meaning of complex natural language expressions determines the content conveyed to a much higher degree than what is standardly assumed. Over the course of a conversation, individuals and bodies of information talked about, are tracked in the attentional state, a sequence that records them ordered by prominence as determined (up to a presumedly small remnant of ambiguity) by the linguistic context. Coherence relations as inferred between subsequent sentences are part of the grammar and can directly impact how referents are ordered in the attentional state. Pronouns (and the parameters for the interpretation of modal verbs and adverbs) receive a unform treatment as picking out the highest ranked referents meeting their lexical requirements (as imposed for instance by gender features).

I have argued that Stojnić makes a valid point: data shown in the

[16] While Japanese main clauses thus fail to display a clear effect of coherence relations, the effect might still be real in conditionals. Japanese has a large inventory of conditional connectives, each subject to their own constraints on tenses, temporal order,... (Takubo 2020)

(i) *Mary-ga {ku-ru nara / ki-tara / ku-reba / kuru to / ki-tewa}, John-mo ku-ru.*
Mary-NOM come-NPST NARA / come-TARA / come-BA / come-NPST TO / come-TEWA John-also come-NPST.
'If Mary comes, John also comes.'

Only *nara* is fully felicitous for *indicative backtrackers* as in (ii), that is, conditionals where the antecedent specifies currently existing evidence that the event described in the consequent has taken place.

(ii) Ima sinku-ni koohii magu-ga { ar-u nara / ?ar-eba / ?at-tara}, John-wa
now sink-at coffee mug-nom { be-NPAST nara / be-COND / be-TARA} John-TOP
kesa kokoni ki-ta (hazu-da)
this.morning here come-PAST (must)
'If there is a coffee mug in the sink (right now), John was here this morning.'

In this case, inserting the perfect marker *-tei-* would not resolve the conflict, which suggests that it is not merely an issue of temporal order. Further research is required to fully understand the interaction between Japanese conditional markers and coherence relations.

book and novel data I have added support the finding that pronominal resolution is severely constrained by grammatical phenomena. However, contrary to what is argued in Stojnić (2021), there is no convincing evidence for the assumption that coherence relations directly determine pronominal resolution. The overlooked rules of grammar that Stojnić locates in the coherence relations are better located in the complex pronominal systems and temporal forms that we find in natural languages. Where the grammar of pronouns and tense allows for multiple resolutions, general all-purposes reasoning resolves combinations of coherence relations and pronominal resolutions to determine the most plausible contender if one exists. If several combinations are equally coherent according to independent criteria (as developed for instance in Asher and Lascarides 2003), we are facing an instance of ambiguity. Where, however, one combination is favored by criteria for discourse coherence, it remains to be determined if it is then also conventionally encoded. This is, for instance, the position advocated almost in passing by Asher and Lascarides (2003: 236):

> After all, degree of discourse coherence is a partial rather than a total order, and there could be more than one maximally coherent. update. In such a case, MDC won't pick a unique updated logical form from among the candidates, and this amounts to semantic ambiguity.

The quoted passage suggests that SDRT is committed to the assumption that, up to ambiguity as reflected in equally coherent segmented discourse representations, the propositional content of a declarative utterance[17] is a fact of the matter independent of the speaker's intentions and determined solely by the linguistic material uttered in a given conversational setting. Thanks to the non-monotonic nature of the logic underlying the construction of segmented discourse representations (*glue logic*), contextual knowledge plays a systematic role absent from Stojnić's framework, leaving room for extra-linguistic factors. The two approaches thus differ on how much role is granted to extralinguistic factors (a difference to be explored in more detail), but they agree that the content conveyed is settled objectively (and independently of speaker intentions). A main difference between Stojnić (2021) and SDRT remains in any case that the grammatical principle determining content for SDRT is not the grammar of coherence relations, but an overarching principle of Maximize Coherence.

To what extent we do take content to be determined to this high degree to begin with has to await further investigation. Potentially relevant data might come from considerations of how speakers react to moves in which their referential intentions were misrepresented: as far as clearly semantically encoded content is concerned, they can easily resort to 'but I said that...'. For combinations of pronominal resolu-

[17] Same for its equivalent in interrogatives, e.g. a set of propositions (Hamblin 1973; Karttunen 1977).

tions and coherence relations, it might be more natural to stick to 'but I meant that...'. To the best of my knowledge, data along these lines still remain to be investigated systematically.

References

Altshuler, D. and R. Truswell. 2022. *Coordination and the Syntax – Discourse Interface*. Oxford University Press.

Arita, S. 2004. *(Fu)kanzenjiseisetsu to nihongojoukenbun. ((In)complete tensedness and Japanese conditionals.)*. PhD thesis. Kyoto University.

Asher, N. 1993. *Reference to Abstract Objects in Discourse*. Dordrecht, Boston, and London: Kluwer.

Asher, N. and A. Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.

van den Berg, M. 1996. *Some aspects of the Internal Structure of Discourse. The Dynamics of Nominal Anaphora*. PhD thesis. University of Amsterdam.

Bittner, M. 2011. "Time and modality without tenses or modals." In R. Musan and M. Rathert (eds.). *Tense Across Languages*. Tübingen: Niemeyer.

Bosch, P., Rozario, T., and Zhao, Y. 2003. "Demonstrative pronouns and personal pronouns: German *der* vs. *er*." In *Proceedings of the EACL 2003 Workshop on the Computational Treatment of Anaphora*. Association for Computational Linguistics, East Stroudsburg, 61–68.

Brasoveanu, A. 2006. *Structured Nominal and Modal Reference*. PhD thesis. Rutgers University.

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

Despić, M. 2011. *Syntax in the Absence of Determiner Phrase*. PhD thesis. University of Connecticut.

van Eijck, J. 2001. "Incremental dynamics." *Journal of Logic, Language, and Information* 10 (3): 319–351.

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

Groenendijk, J., Stokhof, M., and Veltman, F. 1996. "Coreference and modality." In S. Lappin (ed.). *Handbook of Contemporary Semantic Theory*. Oxford: Blackwell.

Hamblin, C. L. 1973. "Questions in Montague English." *Foundations of Language* 10: 41–53.

Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis. University of Massachusetts, Amherst.

Heim, I. 1988. "On the projection problem for presuppositions." *Proceedings of the West Coast Conference on Formal Linguistics* 2: 114–125. Reprinted in Davis (Ed., 1991).

Heim, I. 1993. "Anaphora and semantic interpretation: A reinterpretation of Reinhart's approach." In O. Percus and U. Sauerland (eds.). *The Interpretive Tract*. MIT Working Papers in Linguistics 25, 205–246.

von Heusinger, K. 2004. "Choice functions and the anaphoric semantics of definite NPs." *Research on Language and Computation* 2: 309–329.

Hinterwimmer, S. 2015. "A unified account of the properties of German demonstrative pronouns." In P. Grosz, P. Patel-Grosz, and I. Yanovich (eds.). *NELS 40: Semantics Workshop on Pronouns*. University of Massachusetts, Graduate Linguistic Student Association, Amherst, 61–107.

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

Jacobson, P. 2012. "Direct compositionality." In W. Hinzen, E. Machery and M. Werning (eds.). *Oxford Handbook of Compositionality*. Oxford University Press.

Jovović, I. 2022. "Condition B and other conditions on pronominal licensing in Serbo-Croatian." *Linguistic Inquiry:* 1–20.

Jovović, I. 2023. *Pronouns at the interface of syntax and discourse structure*. PhD thesis. University of Connecticut.

Kamp, H. 1979. "Events, instants, and temporal reference." In R. Bäuerle, U. Egli, and A. von Stechow (eds.). *Semantics from Different Points of View*. Berlin: Springer.

Kamp, H. and U. Reyle. 1993. *From Discourse to Logic*. Dordrecht: Kluwer.

Karttunen, L. 1977. "Syntax and semantics of questions." *Linguistics and Philosophy* 1 (1): 3–44.

Kaufmann, S. 2020. "Formal treatments of tense and aspect." In W. Jacobsen and Y. Takubo (eds.). *The Handbook of Japanese Semantics and Pragmatics*. Berlin, Boston: Moulton de Gruyter, 371–422.

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

Lillo-Martin, D. and Klima, E. S. 1990. "Pointing out differences: ASL pronouns in syntactic theory." In S. D. Fischer and P. Siple (eds.). *Theoretical Issues in Sign Language Research, Volume 1: Linguistics*. Chicago: The University of Chicago Press, 191–210.

Montalbetti, M. 1984. *After binding: On the interpretation of pronouns*. PhD thesis. MIT.

Murray, S. E. 2014. "Varieties of update." *Semantics and Pragmatics* 7 (2): 1–53.

Nouwen, R. 2003. *Plural Pronominal Anaphora in Context: Dynamic Aspects of Quantification*. PhD thesis. University of Utrecht, The Netherlands.

Patel-Grosz, P. and Grosz, P. G. 2017. "Revisiting pronominal typology." *Linguistic Inquiry* 48 (2): 259–297.

Reichenbach, H. 1947. *Elements of Symbolic Logic*. New York: Free Press.

Riester, A. 2009. "Partial accommodation and activation in definites." In *Proceedings of the 18th International Congress of Linguists*. Seoul: Linguistic Society of Korea, 134 –152.

van der Sandt, R. 1992. "Presupposition projection as anaphora resolution." *Journal of Semantics* 9: 333–377.

Schlenker, P. 2005. "The lazy Frenchman's approach to the subjunctive." In T. Geerts, I. van Ginneken and H. Jacobs (eds.). *Romance languages and linguistic theory 2003: Selected papers from 'Going Romance' 2003, Nijmegen, 20–22 November*. Amsterdam: John Benjamins, 269–310.

Schlenker, P. 2018. "Visible meaning: Sign language and the foundations of semantics." *Theoretical Linguistics* 44 (3–4): 123–208.

Smyth, R. 1994. "Grammatical determinants of ambiguous pronoun resolution." *Journal of Psycholinguistic Research* 23: 197–229.

Takubo, Y. 2020. "Conditionals in Japanese." In W. Jacobsen and Y. Takubo (eds.). *The Handbook of Japanese Semantics and Pragmatics*. Moulton de Gruyter.

Umbach, C. 2002. "(de)accenting definite descriptions." *Theoretical Linguistics* 27: 251–280.

Croatian Journal of Philosophy
Vol. XXIII, No. 69, 2023
https://doi.org/10.52685/cjp.23.69.4
Received: October 12, 2023
Accepted: November 27, 2023

Intentionalism and the Natural Interpretation of Discourses

ALEXANDRU RADULESCU
University of Missouri, Columbia, USA

Intentionalism is the view that a demonstrative refers to something partly in virtue of the speaker intending it to refer to that thing. In recent work, Una Stojnić has argued that the natural interpretation of demonstratives in some discourses is that they do not refer to the objects intended by the speaker, and instead refer to other things. In this paper, I defend intentionalism against this charge. In particular, I argue that the data presented by Stojnić can be explained from an intentionalist point of view. The explanations take two forms: either the audience's reaction to the discourse does not concern reference, or the natural interpretation is wrong. This latter claim has been defended by Stojnić in other work as applied to word identification and is neutral between intentionalism and Stojnić's objectivism. It is also very plausible. But it takes away the import of the argument from natural interpretation, at least in the form discussed here.

Keywords: Demonstratives; reference; intentionalism; objectivism.

1. *The issue, briefly*

Intentionalism about the referential mechanism of demonstratives is the claim that the speaker's intentions play a role in determining the referents of demonstratives when they use them. It is fair to say that intentionalism has been the dominant family of views in the metasemantics of demonstratives, with most of the discussion made up of disagreements between different kinds of intentionalism: about the nature of these intentions, about the extent of their role, about factors

other than intentions, if there are any, etc.[1] One main reason why intentionalism has become the default view is that it seems difficult to come up with alternatives: what else could guide a demonstrative to its referent? We use demonstratives in the company of pointing gestures, but also on their own. We use demonstratives when one particular object is salient to everyone, but also when there is no such object. We use demonstratives, and care about the audience figuring out the referents, but we know that they can get it wrong, and we do not always defer to their opinions. By contrast, it seems reasonable to say that speakers always have intentions when they use demonstratives. So we have no obvious competitors for offering such explanations, or a natural candidate.

In recent work, Stojnić has presented a radically anti-intentionalist view, or what she calls an "objectivist" view (Stojnić 2021a: 4).[2] This view takes issue with both of the alleged advantages of intentionalism: it claims that other features of the discourse are more natural candidates for determining the referents of demonstratives, and that intentionalism often gets things wrong anyway.

Stojnić's defense and development of the objectivist view contain many interesting, interwoven claims, and they should be judged as a whole. My goal in this paper is more modest: I only aim to defend intentionalism from a particular objection relating to intuitions about natural interpretations. Furthermore, I think that objectivism can survive without this argument, and would be better off without it, so the paper could also count as a way to improve objectivism.

Here is Stojnić's argument, in a nutshell: plausible, natural, seemingly unavoidable interpretations of various discourses assign to demonstratives in those discourses certain referents that were not intended by the speakers. It follows that the speaker's intentions cannot play any role in determining these referents.

I will offer two types of rejoinders on behalf of intentionalism. First, I go through the details of these cases, and I argue that they can be interpreted in an intentionalist-friendly way. Then, I argue that if we change the cases slightly, the intuitions go away, although according to objectivism they should not. Finally, and more broadly, I argue that metasemantic theories must allow the audience to get things wrong. So figuring out the natural interpretation of a discourse is not the end of the story; the question that should be answered is whether the natural interpretation is the correct interpretation. If the intentionalist can always explain away the intuitions that Stojnić appeals to, they cannot be used to show that intentionalism is false.

[1] Much of the modern literature can be traced back to Kaplan (1989); for a more recent overview, see Braun (2017).

[2] Stojnić's view has been presented in two papers and a book (Stojnić et al. 2013, Stojnić et al. 2017, and Stojnić 2021a, respectively.)

Here is the plan: in §2, I present a case against intentionalism discussed in Stojnić et al. (2013), and I argue that the basic intuitions can be explained in an intentionalist-friendly way. In §3, I address a later discussion of a similar case from Stojnić (2021a). This longer presentation contains three arguments against intentionalism, based on variations of the original case, and I discuss them in separate subsections. In §4, I present my own variation on the case, which generates no objectivist intuitions, although by objectivist standards, it should. I take this to show that the earlier cases do not work against intentionalism either. Finally, in §5, I argue that any plausible metasemantic theory must allow audiences to misinterpret any feature of an utterance, a requirement that both intentionalism and objectivism can accept. If this requirement is accepted, the arguments addressed in this paper cannot help us decide between intentionalism and objectivism.

2. *An earlier discussion of Stuck Arms*

In this section, I discuss two cases offered by Stojnić, where the natural interpretation of a discourse seems to go against intentionalism. For the purposes of this paper, I am interested in two features. In this section, I challenge the robustness of those intuitions, by pushing back against those intuitions. In the next section, I offer a dilemma, and argue that one horn should be avoided by any metasemantic account, and the other horn makes the current argument irrelevant.

I begin with a case discussed by Stojnić et al. (2017):

> *Stuck Arm*: Consider a speaker who intends to refer to Ann, but her hand becomes suddenly stuck, and so, she accidentally points at Sue, while uttering, "She is happy." It would be odd to say she intended to refer to Sue, or indeed, anything in the general direction of her pointing gesture. Quite clearly, though, intuitively, it is Sue, not Ann, who is the referent of 'she'. After all, the audience can follow up her utterance with "So, you are saying Sue is happy" and can challenge her with "That is false. Sue is not happy at all." (Note, the audience could not felicitously ask, "So, are you saying that Ann is happy?" or follow up with "That's false; Ann is not happy," or "True! Ann is happy.") The speaker cannot felicitously deny she said Sue is happy (or claim she said Ann is happy). (Stojnić et al. 2013: 508)

The claim is that our intuitions would have it that, contrary to the speaker's intentions, Sue is the referent of the demonstrative.

The notion of *felicity* does a lot of work here. We will come back to the intuitions themselves; for now, I will grant that there is something infelicitous with the mentioned reactions. We have two kinds of reasons offered in support of the authors' claim. First, we have appeals to intuitions about what *the audience* can felicitously say next, and what it cannot. Second, we have appeals to intuitions about what *the speaker* can felicitously say. I will take them in turn.

What can the audience felicitously say? If they take the speaker to have said something about Sue, it would be strange of them to ask something about Ann, and it would be quite appropriate to continue

the conversation as if it was about Sue. This only reinforces the claim that the audience interpreted the demonstrative in a particular way. But misinterpretation is always possible. So we have no reason yet to think that intuitions about what that audience would naturally say tell us anything about the correctness of that interpretation.

What can the speaker go on to say? Note that we are imagining that the audience has already made manifest that they take it that the speaker was talking about Sue. So the speaker is in the following situation: she intended to refer to Ann and was taken to have referred to Sue. If you prefer, we can move to a formal mode of speaking: the speaker intended for the demonstrative to refer to Ann, but the audience takes it that the demonstrative refers to Sue, and, presumably, that this is what the speaker had intended. The speaker can let things go, and go on as if the audience's interpretation was correct. Since the audience thinks that she intended Ann, not Sue, as the referent, by continuing as if the intended referent was Sue, the speaker is misleading about her original intentions. Stojnić et al. can insist at this point that the intuitions they are interested in are not about what the speaker had intended, but about the referent of the demonstrative. But is it all that easy to separate the two? If the audience later finds out that the speaker was deceitful about her intentions, will they be content to stick to their intuitions about the referent of the demonstrative? I am not so sure. And I really do mean it: I am not sure that they would care to make the distinction, or that they would have much clarity on the issue. Claims about the referent of a demonstrative are theoretical. Intuitions, like the law, are highly pragmatic: people care about what they care about (often, speaker intentions), and they will care about other things only if they have to (e.g. truth value, what is said, etc.). Appeal to intuitions about how a simple conversation would naturally continue will not decide intra-semantics disagreements.

My responses depend on what we think the audience and the speaker would do next, given the setup of Stuck Arm. In particular, one driving feature has been the fact that the audience were not in a good position to figure out what the speaker meant. This leaves open the question whether that latter fact should be changed, in order to provide a better case for the objectivist. Fortunately, Stojnić has done exactly that in later work, so we turn to that discussion.

3. *A later discussion of Stuck Arms*

§2 was focused on the earlier discussion of Stuck Arm, as presented in Stojnić et al. (2013). The case gets discussed again, with some details left out, but, more importantly, others added. I provide the full quote below. In order to organize the subsequent discussion, I added some parenthetical numbers and divided the text into more paragraphs than the original.

Stuck Arm 2 :

[1] Suppose I want to say that Mary is my best friend, but due to some accident, perhaps a muscle spasm, or confusion, I point at Sue while saying 'She is my best friend.' While you might realize that a mistake of sorts happened, because, say, you might have good reasons to think that Mary is the one I in fact wanted to talk about, it is essential to the case that a mistake happened: I accidentally said something I did not mean, just as I would have said something I did not mean to say had I uttered a wrong word.

[2] To make the case more dramatic, suppose that I wanted to insult Ann, but accidentally pointed at Sue, while uttering 'You are a jerk.' It seems that in that case I would have to apologize to Sue. I could not simply say that I did not say she was a jerk, or that I said that Ann was. Of course, I could say that I meant that, but the fact that I'm apologizing and making this excuse is precisely explained by the fact that I did not say what I meant.

[3] Notice that the predicament is analogous to the one I would be in if I were to accidentally utter a wrong name. Suppose I say 'Sue is a jerk,' accidentally uttering 'Sue' instead of 'Ann.' You might, in this case, if you have enough evidence, conclude that I meant to say something about Ann, rather than Sue. (Perhaps you know I dislike Ann but not Sue.) However, this does not make 'Sue' mean Ann on this occasion. Similarly, realizing that I made a pointing error does not make my pointing any less an instance of pointing at the person I actually pointed at, Sue, and does not make Sue any less the referent of the accompanying occurrence of 'she.'

[4] Note that had a pointing gesture merely served as a kind of a (defeasible) cue indicating an underlying intention to single out a particular individual (or, alternatively, had the pointing gesture itself had a context-sensitive interpretation), one would expect that if the audience had enough evidence to figure out who the speaker had in mind, the fact that the speaker obviously did not point at whomever she actually had in mind, would not result in infelicity. As always the audience would just work out the overall most plausible interpretation given their evidence; from this standpoint, the speaker would not have made an error. But this would be a wrong prediction: the speaker clearly would have made an error in such a case.[3]

As I see it, there are four sections in these two paragraphs, which contain three related strands that make up the later version of Stojnić's argu-

[3] Stojnić (2021a: 50–51). I am focusing on Stuck Arm cases because they are discussed in detail, but talk of natural interpretations, or natural ways to understand a discourse, or related notions, is abundant in the book. See pages 11, 42, 49, 64, 68–70, 74, 75, 80, 86, 114, 122, 123, 130, 136, 149, 150, and 180.

ment from natural interpretation. I will take the three strands in turn.

3.1 *The speaker's alleged mistake ([1] and [4])*

The first claim about this modified case, as I understand it, comes from [1] and [4]: the speaker made a mistake, and that fact is recognized by both the audience and the speaker, even though the speaker's intentions are known to all involved.

In [1], Stojnić characterizes the mistake as one about what was said: the speaker said something other than what she intended. But this is just the claim that the intentionalist denies, so it cannot be the end of the story. Other ways of grouping the arguments are possible, but I think that the most direct reason for the claim in [1] is given by [4]: if both semantics and the audience's interest were limited to the speaker's intentions, there would be no intuition of infelicity, or of the speaker making a mistake. Since, by hypothesis, everyone knows all that there is to know about the speaker's intentions, it follows that the mistake must be at the semantic level. In other words, it must be that what was said is different from what was intended.[4]

Intentionalists can choose between two responses. First, they can say that the mistake happened in the pointing: the speaker did not point at the person they intended to point at. Stojnić describes the situation as one in which the audience "realize[s] that a mistake of sorts happened", but presumably cannot tell exactly which mistake and why. That could be explained either by the speaker being confused about whom they are pointing at (e.g. by confusing Ann for Mary), or by some error in the gesture itself (which is what in fact happened). If confusion is unlikely in that situation, it would be reasonable of the audience to suppose, or even to figure out, that the pointing had somehow gone awry. In that situation, although the audience have good evidence that the speaker intended to say something about Mary, there is the seemingly contrary evidence coming from the pointing. That leaves what is said to be fully determined by the speaker's intentions, and allows that a mistake did indeed happen.

Second, intentionalists may prefer to insist that no mistake happened. The speaker did everything she could to refer to Mary, and to make her intentions clear to the audience. Her arm got stuck, but if that was not within her control, nor something she could have predicted, she made no mistake. When the audience finds out what happened, it would be strange for them to insist that the speaker made a mistake. They may say that it was difficult to figure things out, and that might be correct. But if we describe their reaction this way, there is little room for them to also insist that the demonstrative referred to Ann.

[4] I assume that "what I meant", which is the expression used in the text, amounts to the same thing as "what I intended."

3.2 *The speaker's apology ([2])*

Stojnić then asks us to imagine that the speaker uttered the sentence "You're a jerk", while intending to point at Ann, but accidentally pointing at Sue. The claim is that the speaker would need to apologize to Sue, even if she makes it clear that her intention was to point at and insult Ann.

I do not think that the switch to "you" is innocent here. I have argued elsewhere that "you" is a pure indexical, not a true demonstrative (Radulescu 2018). In short, the reason is that "you" picks out the addressee, where the addressee may well be fixed by the speaker's intentions but is not fixed in order to give "you" its referent. By contrast, according to intentionalism, a demonstrative gets its referent because the speaker intended that referent for that particular demonstrative. Utterances often have addressees, and if they contain a "you", it will pick out whoever the addressee is. If they contain no pronouns, they may still have an addressee, which is why it makes sense to ask who the addressee is, concerning an utterance of "2+2=4", while it makes no sense to ask what the referent of "that" is with respect to that utterance. So, if Stojnić's case works, it is at best an objection against intentionalism about fixing the addressee, not about the referential mechanism of the second person pronoun.

But we can set that aside. The intentionalist, whether of the reference-determining kind, or of the addressee-determining kind, might respond that you can insult someone unintentionally, without saying anything about them. The idea here would be that insults happen out in the open, and the fact that an utterance would naturally be interpreted as an insult may well be enough for the addressee to be insulted by it. I do not find this response very convincing, because it should be possible for the audience to say later "I thought I was being insulted, but it turns out I wasn't". I find a different intentionalist response more plausible. Suppose that the speaker's apology is appropriate, or even required. What is she apologizing for? Stojnić's interpretation is that she is apologizing for insulting the addressee, by having said something insulting about her. The intentionalist may prefer a different option: that the speaker is apologizing for the addressee's reasonably feeling insulted by her utterance. If I make someone feel bad, and if it is reasonable for them to think I did it on purpose, the right thing for me to do is to repair the unintentional damage. It is right for me to apologize for certain things I did unintentionally. What needs to be resisted is the claim that I'm apologizing for an actual insult. So, the apology is focused on the effects of the utterance, and the content of the utterance is addressed separately. When the speaker says "I didn't mean you!", they are focusing on their own part in the whole situation, and they are trying to repair the person's interpretation of the utterance. This is an intentionalism-friendly analysis of the case that makes sense of the original intuitions, while retaining a non-objectivist metasemantics for demonstratives.

3.3 *Demonstratives and proper names ([3])*

A central challenge to intentionalist responses is that they overgeneralize. I take this to be the point of part [3] of the quotation, where Stojnić points out that if we substitute a name for a pronoun in cases like this, we go back to the dispute between Donnellan and Kripke about the possibility of using a name to refer to a person who does not have that name.[5]

One option for intentionalists is to accept the comparison and take it in the opposite direction to Stojnić. They can just bite the bullet, and claim, against Kripke, that in certain circumstances, the referent of "Sue" is Ann, namely whenever the speaker so intends.[6]

The other option, which I prefer, is to reject the comparison. People have names; when you use a name, it seems reasonable to say that the name cannot refer to someone who does not bear that name. People do not have demonstratives; when you use a demonstrative, it can refer to anyone, or indeed, for some demonstratives, to any*thing* (modulo gender, proximity, and such other coarse-grained restrictions). And names have something like that feature too: in some sense, many people bear the name "Ann", so when we say that it refers to Ann, a particular Ann, we need to explain why it refers to her, rather than any other Ann. Famously, Kripke (1972: 7–8) sets aside that issue, but others have taken it up.[7] And taking it up might require an appeal to speaker intentions, as Kaplan argued. That is a separate intentionalism, and I am not planning on defending it here. My point is only that to the extent that the comparison between names and demonstratives works, it can be appropriated by the intentionalist.

Stojnić is not saying that people bear demonstratives, of course. Her claim is that a demonstrative just refers to whatever it refers to, no matter what the speaker intends, and names do the same. That one can refer to one person rather than an intended other person with a demonstrative just as well as with a name. And, of course, the intentionalist denies this. My point here is only that claims about how an utterance is naturally interpreted will not help decide the dispute.

4. *A different Stuck Arm*

Thus far, I have been granting Stojnić's intuition that something goes wrong with the speaker's attempt to use the demonstrative to refer to Ann while the arm was pointing at Sue. In response, I argued that the feeling of infelicity can be explained at a level that tells us nothing about the referent of the demonstrative. In this section, I present a

[5] See Donnellan (1966, 1978), Kripke (1977).

[6] See Pepp (2019) and Capuano (2020) for recent defenses of such Donnellan-inspired views.

[7] See Strawson (1950), Evans (1982), Kaplan (1990), Matushanski (2008), and Cumming (2014), for example.

modification of the original case, where it seems to me that the intuition of infelicity goes away. I then argue that reflecting on this case puts pressure on the earlier intuitions as well.

Transparently Stuck Arm: The speaker intends to refer to Ann, and her hand gets stuck in Sue's direction, while she utters "She is happy". The audience know that the speaker has poor control over her arm movements, and do not take the hand movement as good evidence of the speaker's intentions. Furthermore, the audience know that the speaker intends to say something about Ann. Quite clearly, it would be infelicitous for the audience to follow up with "So, you're saying that Sue is happy", and cannot follow up with "That is false. Sue is not happy at all". The speaker herself can also, if challenged by the audience, felicitously deny that she said anything about Sue.[8]

In the earlier version of Stuck Arm, the audience were not in any position to recover the speaker's intentions. In the later version, they are in a position to know something about it, but they are ignorant about the arm control issue. In this last version, they know everything there is to know. These differences should not matter, according to the objectivist. After all, the referent it determined by the arm's pointing, so the things that the audience can go on to say should just track those facts. Yet this is not what we find by making these subtle adjustments to the initial scenario. The more the audience knows about the situation, the less inclined they will be to interpret it as containing an utterance about Sue. This is just what the intentionalist would say: interpretation can fail, but when it succeeds, it correctly tracks the referent of the demonstrative.

Of course, Stojnić can respond that in Transparently Stuck Arm the audience are attending to speaker reference, not to semantic reference. I have said nothing that would argue otherwise. But at this stage we have a disagreement between the two views about how to categorize certain cases and certain intuitions, not an argument against one of the views. And that suffices for the purposes of this paper.

5. *A dilemma*

Are we at the stage where neither view has an advantage in this arena? It may look like any argument based on intuitions about interpretation can be spelled out so that it comes out favorable to either side. But let us zoom out a bit. The broader issue concerns the use of intuitions about the natural interpretation of a discourse by an audience who may or may not know everything that we, the observers, know.

So, I propose a dilemma: either the audience can get things wrong, or they cannot. At one end, we could have the view that the audience can get nothing wrong; so long as their interpretation is the reasonable

[8] For more on cases where the audience is in a position to figure out the speaker's intentions, or where the speaker falsely assumes that that is the case, see Radulescu (2019).

one, or the best one, given their epistemic situation, that is the correct interpretation. At the other end, we have the view that the audience can get everything wrong, at all levels of interpretation: force, implicatures, reference, word identification, etc., even when that is the best interpretation, or even the only reasonable interpretation.[9]

In other work, Stojnić has defended the view that which word is uttered depends on certain facts about the process by which the utterer comes to produce certain sounds, facts that are determined neither solely by the speaker's intentions, nor solely by the audience's interpretation (Stojnić 2021b). The claim is that the speaker may utter a word they did not intend to utter, and the audience may mistakenly, but reasonably, take the speaker to have uttered a different word than what actually got uttered. So she agrees that some natural interpretations are incorrect.

Must the objectivist hold this? Not necessarily. But there is at least one level of interpretation where objectivism is just too implausible. We have been going along with Stojnić's description of the case as one where the speaker pointed to the wrong person. But did she? We were only told that the arm got stuck. Was the speaker in the process of pointing, and the arm only went halfway? Or was it a muscle spasm that ended up looking like a pointing? If it was a spasm, it was not a pointing, since pointings are intentional actions, whereas spasms are not.[10] In which case, that arm movement should not pick out anything, since there is no convention that mere positions of arms have the power to refer. The audience may naturally interpret the arm's movement as a pointing; if it is not, they would be wrong. So at least at the level of identifying whether an arm movement constitutes a pointing, objectivism strays too far from our conception of which things in the world have representational powers.

So, at some point, natural interpretations must be allowed to go wrong. The question is where. Must reasonable, natural interpretations be correct at the level of reference? Must they be correct, at least in cases where there is no competing, equally reasonable interpretation? What about force, or implicature, and so on?

I submit that more basic than the debate between intentionalism and objectivism is the intuition that the audience can always get things wrong, whether it be the identity of the words uttered, the force of the utterance, the referents of terms, etc. And they can get it wrong even when that is the only plausible interpretation, given what they know. This intuition is strictly weaker than either view, since it is compatible with both. After all, both intentionalists and objectivists agree that the referent of the demonstrative is fixed by facts that are independent of the state of mind of the audience, whether that be the actual audience

[9] An even more extreme view, that the reasonable interpretation is never correct, can be set aside as too implausible to discuss.

[10] I do not focus on this aspect of the case here; but see Pavese and Radulescu (2023).

or an idealized version thereof.[11] Any plausible metasemantic account needs to allow audiences to get it wrong. Otherwise, if some ants in the sand spell out "help!", we would need to count that as them asking for help, or at least as uttering the word "help". Nothing can be a sign merely in virtue of looking like a sign. Similarly, nothing can *refer* merely in virtue of being naturally taken to refer. And, finally, nothing can refer *to a particular thing* merely in virtue of naturally being taken to refer to that thing.

If this is correct, the argument from Stuck Arm, in all its versions, fails. Even if we agree that the natural interpretation of such cases is the one posited by Stojnić, a claim that I have offered some reasons to doubt, it cannot be used to decide between intentionalism and objectivism.

6. *Conclusion*

When we discuss cases, we need to start somewhere. We start with our own intuitions, and sometimes we ask others too, in a more or less organized manner. The clearest reactions to cases happen when something has gone wrong. The difficulty is diagnosing exactly what went wrong, and why. Couple that with the fact that the way we describe the case can make a significant difference in the audience's judgment, and we have a complicated situation.

Appeals to natural ways to interpret a discourse are a good place to start. In this paper, I have argued that intentionalists have good responses to these initial judgments. They can reject them, either as a matter of reporting intuitions, or by deflecting the charge, and saying that the intuitions of infelicity are about something else. Finally, there must be room for the audience to get things wrong. And if there is, pointing out the natural interpretation of a discourse does not suffice to distinguish between objectivism and intentionalism.

I find intentionalist explanations of the various cases we have been discussing to be more plausible than objectivist ones. The reader may disagree. Success can only be measured by how well the overall theories deal with all manner of data. I look forward to seeing how things turn out.

References

Braun, D. 2017. "Indexicals." In E. N. Zalta, (ed). *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2017 edition.

Capuano, A. 2020. "In defense of Donnellan on proper names." *Erkenntnis* 85: 1289–1312.

[11] There are philosophers who advance audience-centered views. I cannot argue against that view here, but for a good overview, and for what I take to prove that such views do not work, see Nowak and Michaelson (2021).

Cumming, S. 2014. "Discourse content." In A. Burgess and B. Sherman (eds.). *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press, 214–230.

Donnellan, K. 1966. "Reference and definite descriptions." *Philosophical Review* 75: 281–304.

Donnellan, K. 1978. "Speaker reference, descriptions, and anaphora." In P. Cole, (ed.) *Syntax and Semantics 9: Pragmatics*. New York: Academic Press, 47–68.

Evans, G. 1982. *Varieties of Reference*. Oxford: Clarendon Press.

Kaplan, D. 1989. "Demonstratives." In J. Almog, J. Perry and H. Wettstein (eds.). *Themes from Kaplan*. Oxford: Oxford University Press, 481–563.

Kaplan, D. 1990. "Words." *Proceedings of the Aristotelian Society* 64: 93–119.

Kripke, S. 1972. *Naming and Necessity*. Cambridge: Harvard University Press.

Kripke, S. 1977. "Speaker's reference and semantic reference." In P. French, T. Uehling and H. Wettstein, (eds.). *Midwest Studies in Philosophy, Vol. 2*. Minneapolis: University of Minnesota Press, 255–276.

Matushanski, O. 2008. "On the linguistic complexity of proper names." *Linguistics and Philosophy* 31: 573–627.

Nowak, E. and Michaelson, E. 2021. "Who's your ideal listener?" *Australasian Journal of Philosophy* 99 (2): 257–270.

Pavese, C. and Radulescu, A. 2023. "Intentionalism out of control." Manuscript.

Pepp, J. 2019. "What determines the reference of names? What determines the objects of thought." *Erkenntnis* 84: 741–759.

Radulescu, A. 2018. "The difference between indexicals and demonstratives." *Synthese* 195: 3173–3196.

Radulescu, A. 2019. "A defence of intentionalism about demonstratives." *Australasian Journal of Philosophy* 97 (4): 775–791.

Stojnić, U. 2021a. *Context and Coherence*. Oxford University Press.

Stojnić, U. 2021b. "Just words: Intentions, tolerance and lexical selection." *Philosophy and Phenomenological Research* 105: 3–17.

Stojnić, U., Stone, M. and Lepore, E. 2013. "Deixis (even without pointing)." *Philosophical Perspectives* 27 (1): 502–525.

Stojnić, U., Stone, M. and Lepore, E. 2017. "Discourse and logical form: Pronouns, attention, and coherence." *Linguistics and Philosophy* 40: 519–547.

Strawson, P. F. 1950. "On referring." *Mind* 59: 320–344.

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)

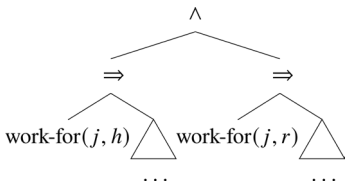

We will present our pronoun resolution mechanism in section 6. Here, we want do draw the reader's attention to the fact that certain information—notably, the positions of *h* and *r*—would be unavailable, or at least would not be transparently available, in a logical form of dynamic semantics, such as (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.

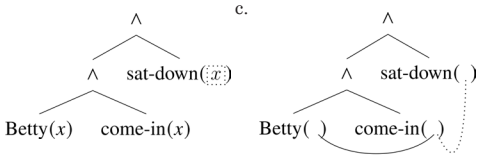

As a side note, we do not need to be realists about variable symbols. We could just as easily use Quine's (1981) *bonds*. With those, (14b) transforms into (14c). What we previously thought of as an unidentified variable is now simply an unlinked argument position, and pronoun

[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.

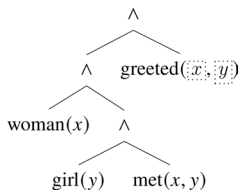

Syntactic height can also provide a better reason for identifying the first argument of greeted(□,□) first. We merely have to assume that the syntactically higher pronoun is resolved first. In (15a), the first and the second argument of 'greeted' are the subject and the object, respectively, and the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

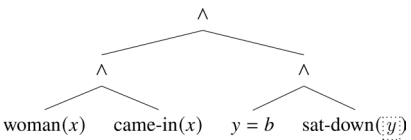

We will formalize the notion of closeness in the following section, where we will see that it can be defined deploying the standard syntactic re-

[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.

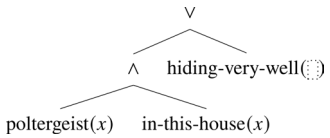

The situation is completely different in the well-known bathroom examples. Here, the negation contributed by the negative determiner 'no' intervenes between the conjunction and the disjunction, flipping the polarity of the pass just in time to allow for the p-scope to enter the disjunction at the root, 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.

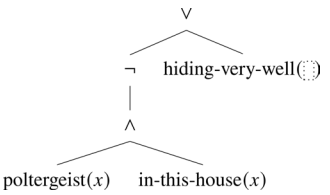

Finally, we see that even in the absence of this negation, i.e. with an indefinite noun phrase, the discourse can be felicitous if the antecedent can be found in the discourse preceding the disjunctive sentence. Crucially, 'my bag' wins over 'a poltergeist' despite the latter being closer (both linearly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

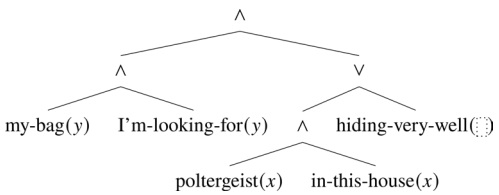

In Ludlow and Živanović (2022), we provide further similar examples of p-scope carving out the set of possible antecedents of a pronoun. Specifically, we deploy it in a novel approach to the donkey anaphora, showing both why the indefinite can receive the (universal) wide scope and why a negative determiner 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)

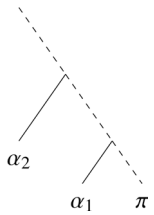

However, the closeness relation defined above will not do for our purposes, because it only orders the nodes c-commanding π, while we need to order the nodes p-scoping over it.[17] For example, α_1 and α_2 below (ignore β_1 and β_2 for now) might (depending on the content of the tree) both p-scope over 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)

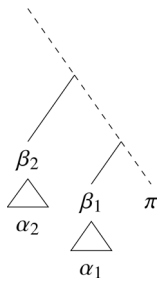

At this point, one might be tempted to rewrite the above definition of the closeness relation by substituting "p-scope" for "c-command"—so that if we have α_1 and α_2 p-scoping over π, α_1 is closer to π than α_2 is iff α_2 p-scopes over α_1—but this will not do, because α_1 and α_2 can easily p-scope 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)

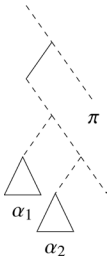

And if none of α_i c-commands the other? Well, we could continue recursively applying the LCA idea, but that would merely reimplement LCA. At this stage, the location of π becomes irrelevant, as we simply want to order the nodes by syntactic height, which is precisely what LCA does. We therefore choose 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. 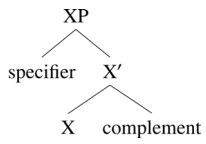

b. 

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.

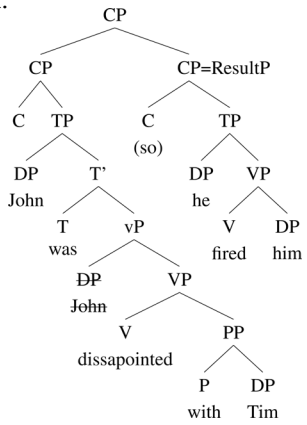

b.

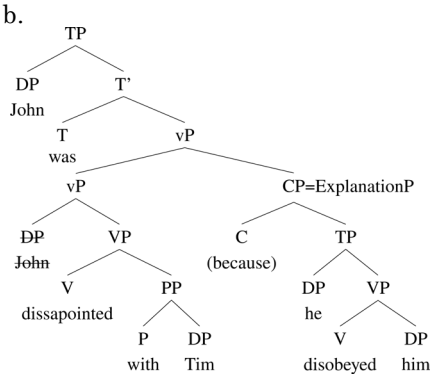

Given the integration position of the second sentence, our account yields the correct predictions. However, the real question is whether there is any independent evidence for this position. We believe that there is.

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.

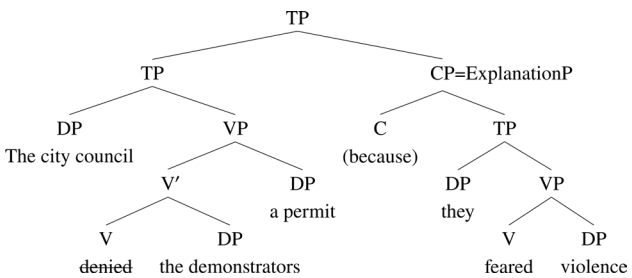

Another example discussed by Stojnić is (30). We are already familiar with our derivation of the result reading, what about Parallel, where 'him' resolves to the object? The strategy is the same as for Explanation. We get the correct prediction if we stick the new material between the subject and the 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$_{\text{NOM}}$ is Stanley$_{\text{ACC}}$ tickled, Liz him is PARTICLE poked.

(32) (parallel)

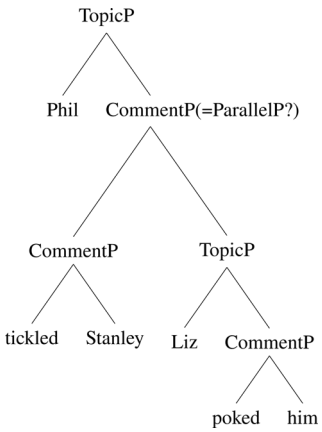

Whatever the details, one thing is certain: Parallel is not sensitive to subjects and objects *per se*. The story again revolves around syntactic height, and hinges on the fact that (contrastive) topics are positioned very high in the structure, within the split CP (Rizzi 1997, 2004). Look again at Slovenian, 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$_{\text{ACC}}$ is Phil$_{\text{NOM}}$ tickled, Marc$_{\text{ACC}}$ 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.

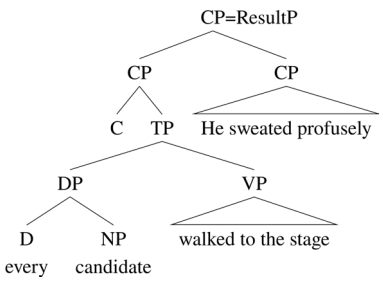

b.

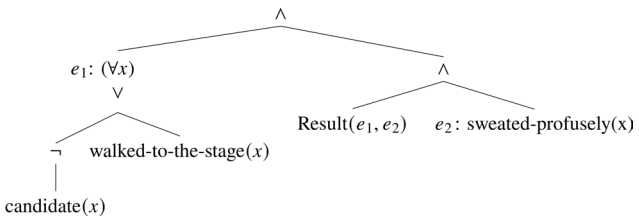

The above example is usually presented along an example of telescoping—a situation where binding out of a universal is exceptionally possible. Roberts (1989) suggests that these examples work because the sentences form a continuous narrative. In a coherence relation approach, this translates to the two 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.

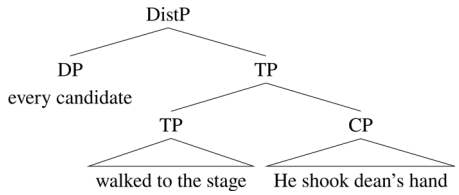

b.

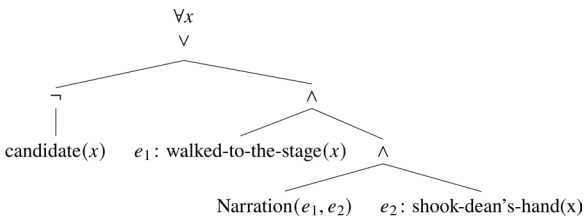

We intend to investigate all the intricacies of telescoping and related modal subordination through the prism of p-scope in a later paper.

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.

Croatian Journal of Philosophy
Vol. XXIII, No. 69, 2023
https://doi.org/10.52685/cjp.23.69.6
Received: October 12, 2023
Accepted: November 27, 2023

Incoherent Meanings

MICHAEL DEVITT
City University of New York, New York, USA

Stojnić holds the radical view that coherence relations determine the reference of context-sensitive language. I argue against this from the theoretical perspective presented in Overlooking Conventions *(2021). Theoretical interest in language comes from an interest in thoughts and their communication. A language is a system of symbols, constituted by a set of governing rules, used (inter alia) to communicate the meanings (contents) of thoughts. Thought meanings, hence speaker meanings, are explanatorily prior to semantic meanings. So, we start our consideration of the theoretical place of coherence by considering the bearing of coherence on thought meanings. The paper argues that a person can have any thought at all, however incoherent. So, a thought's meaning is independent of its coherence. Any thought can be expressed in an utterance. The semantic meaning of any utterance governed by the linguistic rules will be the meaning of the thought it expresses. So, the utterance's meaning is independent of its coherence. The paper concludes that coherence has no place in the theory of meaning or reference. Nonetheless, it has a place in the theory of communication. I suspect that the error exemplifies the widespread confusion of the metaphysics of meaning with the epistemology of interpretation.*

Keywords: Coherence; context sensitivity; reference; thought meaning/content; speaker meaning; semantic meaning; communication.

1. *Introduction*

What place does *coherence* have in theorizing about language? In her engaging book, *Context and Coherence: The Logic and Grammar of Prominence* (2021), and in several related articles coauthored with Matthew Stone and Ernie Lepore (2013, 2017, 2020), Una Stojnić takes "mechanisms of discourse coherence" (Stojnić 2021: 5) to be constitutive

of reference. She argues that the meaning of context-sensitive language is not "partially determined by non-linguistic features of utterance situation", as traditionally thought, but rather "is determined entirely by grammar—by rules of language that have largely been missed" (Stojnić 2021: vii). The missed rules are ones of discourse coherence. Coherence plays its constitutive role because "[s]uccessive contributions to the discourse must be linked into a coherent whole by a recognizable flow of interpretive relationships" (Stojnić 2021: 61).

In "Demonstratives, Context-Sensitivity, and Coherence" (forthcoming), I argue against this radical view as it applies to demonstrations, demonstratives, and the indexical 'I'. I find Stojnić's theories of reference to be seriously incomplete, failing to meet the demands on any such theory. Furthermore, I argued that, so far as Stojnić's theories of these terms go, the theories are false. My argument appeals to perception-based theories of demonstratives, a part of the tradition going back at least to Husserl that Stojnić strangely overlooks. I use my own perception-based theory as an example (Devitt 1974, 1981, 2022).

That forthcoming paper ends with some brief remarks about coherence in general. I claim that though coherence has a place in a theory of *understanding* and *communication*, it has no place in a theory of *meaning*. My aim in the present paper is to provide an argument for those brief final remarks.

My view of the theoretical place of coherence arises from the perspective on language and communication presented in *Overlooking Conventions: The Trouble with Linguistic Pragmatism* (Devitt 2021; also, 2013a), as I shall now demonstrate, drawing on that work.

2. *Human thoughts*

Why are we interested in language in the first place? Our theoretical interest in language comes from our theoretical interest in thoughts and their expression, usually in communication.

It is a piece of folk wisdom that people have thoughts, which is to say that they have beliefs, desires, and other such "propositional attitudes", mental states with intentional contents or meanings. So, the folk are "intentional realists". I think that we have a very good reason for supposing that we do indeed have thoughts (Devitt 2006a: 125–127). We need to ascribe them to people for at least two reasons: to explain people's behaviors and to explain the way they use others as a guide to a largely external reality.

Consider the explanation of behavior first. We observe Mark putting on a raincoat and picking up an umbrella before leaving a room. Why is he doing that? Central to our explanation is that Mark believes that it is raining. Such "intentional" explanations of "intentional" behavior are familiar and central parts of ordinary life, of history, of economics, and of the social sciences in general. They all ascribe thoughts.

Ascribing beliefs serves another remarkably valuable purpose. If a

person believes that the world is such and such, and the person is reliable, then we have good reason to believe that the world is such and such. Thus, ascribing to Mark the belief that it is raining not only helps to explain his behavior but also gives us evidence about the weather. We have a wide range of interests in learning about the world. The direct way to serve these interests is to examine the world. The indirect way is to use reliable indicators. Sometimes these indicators are "natural" ones like tree rings. Sometimes they are artifacts like thermometers. Very often they are the beliefs of others. Some belief ascriptions serve our *theoretical* interest in explanation. Many, however, are like ascriptions of desires, hopes, and so on in serving interests that are not really theoretical at all. We have the most immediate *practical* interest in finding out quite humdrum facts about the world to satisfy our needs for food, shelter, a mate, and so on. So, it helps to know what is on sale at the supermarket, where there is a hotel, who is available, and so on. Ascribing beliefs is a very good way of finding out about anything at all.

This practice of ascribing thoughts is generally *successful* at serving these two purposes. Day in and day out we explain people's behaviors with these ascriptions. Almost everything we know about the world—what we learn at mother's knee, in classrooms, and from books—we get from ascribing beliefs to people and assessing them for reliability. If there really were not any thoughts, this success would be very hard to explain. We clearly have a great theoretical interest in the details of this process of explaining behavior and learning from each other.

It is a familiar piece of folk psychology that, without any involvement of language, we can sometimes use our insight into other minds and knowledge of the world to figure out what a person thinks. Thus, we came to our view that Mark believes it is raining from observing his rain-avoidance behavior. And he might deliberately communicate his belief to us, without using language, by pointing upwards meaningfully as he puts on the raincoat.

3. *Animal communication*

We have a similar theoretical interest in the inner states of other organisms and their communication. We posit these states to explain behavior and to explain how one organism can communicate "information" to another. There is much debate in cognitive ethology and comparative psychology about these matters. There is no presumption that an organism's learning from another must involve a language. At one extreme, chemical detectors may sometimes do the job. At the other extreme, the idea is seriously entertained that this learning should sometimes be explained by attributing to an organism something like human insight into other minds.

So, we do not always have to posit languages to explain this learning. Still, scientists often do. What are they thus positing? What is a language? It is a system of representations or symbols that is consti-

tuted by a set of governing rules, and that a group of organisms use to communicate with each other. Most such languages are not very interesting because they simply communicate information about the animal's own current state; for example, that the animal is hungry, or wants a mate. The interesting ones are the ones known as "referential", ones that convey information about the environment. The honey bee provides a famous, and very surprising, example. The bee uses a "waggle dance" to communicate the direction and distance of a food source. Gunnison's prairie dogs provide another example: they have a system of "barks" that convey information about which sort of predator is threatening and about the characteristics of a particular predator of that sort. Clearly, the whereabouts of food is a pressing concern for the bee, the presence and nature of a predator, for the prairie dog. A bee that has returned from a food source has reliable information about the former, a prairie dog that has observed a predator, the latter. Their languages enable them to communicate this valuable information.[1]

The rules of the bee's language are very likely entirely innate. The rules of the prairie dog's language seem to be partly learned and, perhaps we should say, "conventional": its alarm calls vary a bit from colony to colony; and when an experimenter used a plywood model to simulate a new sort of predator, the prairie dogs introduced a new call (Slobodchikoff 2002). In any case, whether a language used to communicate information is innate or conventional, we have a powerful theoretical interest in that language and its rules. Serious scientists work to discover the natures of the symbols in these representational systems, to discover their *meanings*.

Karl von Frisch is a notable example. He won a Nobel Prize for his discoveries about the bee's dance. I shall simplify by ignoring what he discovered about how the dance conveys the distance of the food source, attending only to what it conveys about direction. Von Frisch found the following remarkable rule:

> To convey the direction of a food source, the bee varies the angle the waggling run makes with an imaginary line running straight up and down...If you draw a line connecting the beehive and the food source, and another line connecting the hive and the spot on the horizon just beneath the sun, the angle formed by the two lines is the same as the angle of the waggling run to the imaginary vertical line. (Frank 1997: 82)

In hypothesizing that a certain behavior in members of a species involves a symbol that represents something in their language, we are supposing that the behavior was produced *because*, in some sense, it involves that symbol representing something in their language; and it is *because* of what the symbol represents in their language that other members of the species respond to the behavior as they do. So, it is

[1] And it is worth noting that sometimes we are confident that an animal has a language because we have *taught* it one; think of some dolphins and primates that have been taught surprisingly complex languages.

because of what it represents that the symbol plays its striking role in the life of an organism.

Evidence for such hypotheses is to be found, of course, in regularities in behavior. Thus, von Frisch's hypothesis was offered as an explanation of his many painstaking observations of the bee's behavior. But is it the *best* explanation? For some time, it was not obvious that it was. A rival hypothesis was that a bee heading off in the direction of the food source was not responding to information communicated by a bee's dance but rather was following an odor trail left by other bees. But this rival did not stand up to ingenious experiments. The consensus now is that the best explanation of the bee's behavior is indeed that the bee is using the language described by von Frisch (Dyer 2002; Riley et al. 2005; Vladusich et al. 2006).[2]

4. *Human language*

Return to humans. It is a truism that they have languages which they use to communicate their thoughts: as the folk say, "language expresses thought". This idea seems irresistible once one has accepted intentional realism, accepted that humans have thoughts (2006a: 127–8). As Fodor, Bever, and Garrett say, "there is much to be said for the old-fashioned view that speech expresses thought, and very little to be said against it" (1974: 375). So, just as the bees and the prairie dogs have representational systems used to communicate the contents of inner states to each other, so do we.[3] The evidence for this in our behavior seems overwhelming.[4]

Consider again our example of Mark and the ascription to him of the belief that it is raining. Suppose that the people present ascribe this belief on the basis of his production of the sound, /It is raining/. According to the rules of English, this sound means that it is raining at the location in question. If the people assume that Mark is being literal and straightforward, they will take that meaning to be the meaning (content) that the speaker intentionally communicates, his "speaker

[2] For more on this issue see Devitt (2006b: 585–6) responding to Smith (2006: 440–1).

[3] Just as the non-referential languages of animals (sec. 3) have other functions that do not utilize *representational* properties, so has ours: we greet ("Hi"), cheer ("Bravo"), abuse ("Bastard"), and curse ("Shit"). My focus is on the representational properties.

[4] Chomskians have a different view. They see a human language as an internal state not a system of external symbols that represent the world. I argued (2003, 2006a: 17–41) that this is deeply misguided. This led to some always lively and sometimes nasty exchanges: Collins 2006, Matthews 2006, Rattan 2006, Rey 2006, and Smith 2006, responded to in Devitt 2006b; Collins 2008a,b and Rey 2008, responded to in Devitt 2008a,b,c, 2013c; Antony 2008 and Pietroski 2008, responded to in Devitt 2008c; Longworth 2009 and Slezak 2009, responded to in Devitt 2009; Ludlow 2009, responded to in Devitt 2013d; Collins 2020 and Rey 2020, responded to in Devitt 2020; Rey and Collins 2023, responded to in Devitt 2023.

meaning". As a result, they have evidence of his thoughts. Taking him to be sincere in his expression, they conclude that he has a belief with that meaning (content), the belief that it is raining in that location. In this way, language is an extraordinarily effective way of making the thoughts of others accessible to us, thoughts that otherwise would be largely inaccessible; and of making our thoughts accessible to others, often in the hope of changing their thoughts and hence their behavior. Even though, as we noted, the thoughts of others are sometimes accessible to us without language, they mostly are not.

Just as we have a powerful theoretical interest in the languages of bees and prairie dogs, we have one in human languages and their rules: we need to know about the natures of the representations used to communicate in these systems.[5]

The rules of human languages arise largely from conventions. Indeed, it is a truism that symbols in a language (like English) have their meanings by convention. As David Lewis points out at the beginning of his classic, *Convention*, it is a "platitude that language is ruled by convention" (1969: 1). Still, I say only that the rules of human languages are "largely" conventional. The qualification is necessary for two reasons. First, if Chomsky is right then quite a lot of syntactic structure is innate. I think he probably is right (Devitt 2006a: Ch. 12). Second, the language of each human, her idiolect, is to some extent, mostly small, idiosyncratic (like Mrs. Malaprop's). So, the rules of her language are largely conventional but probably partly innate and partly her own work. Whatever the origin of a rule in her language that governs a certain linguistic form, it *is* a rule in virtue of her disposition to associate that form, in language production and understanding, with a certain aspect of thought content (2021: 75–77).

So, conventions should loom very large in our view of human language. On some occasions linguistic conventions are established by some influential people in a community stipulating that a certain form has a certain meaning and the community concurring. However, following Paul Grice (1989) and Stephen Schiffer (1972), I think that the conventional use of a linguistic form in a community—a sound, an inscription, etc.—typically come from the form's regular use in utterances to convey a certain part of thoughts, a certain concept or structure; it comes from the regular use of that form to "speaker mean" that content or structure. This regular use in utterances leads, somehow or other, to that form having that meaning conventionally in the language of that community. That meaning has become the literal semantic meaning of the form in the community's language. Crucially, *thought meanings, hence speaker meanings, are explanatorily prior to semantic meanings*.[6]

[5] Some philosophers and linguists, impressed by the great difference between a human language and the representational systems used by other animals, resist calling those systems "languages". I can see no theoretical point to this resistance. In any case, the point is merely verbal.

[6] In support of this crucial Gricean idea, see (Devitt 2021: Ch. 5).

Consider the English word 'train', for example. According to the OED, this word had several uses prior to the nineteenth century. Then came the railways and the word got a new conventional meaning referring to railway trains. How? We note first that this new meaning is conceptually related to old ones referring to a sequence of persons or things. The word is polysemous. It got its new conventional meaning from people using it in successful communications to speaker mean railway trains. The communications were successful, of course, because this speaker meaning traded on old conventional meanings of 'train'. The success led to the *regular* use of 'train' to speaker mean railway trains. In time this regularity led to the new conventional semantic meaning.

5. *Coherence and meanings*

In light of the priority of thought meaning, we should start our consideration of the theoretical place of coherence by considering the bearing of coherence on thoughts.

One thought coheres with its predecessor if the two are linked in some appropriately *rational* way. Here are some truisms. (a) Coherence comes in degrees, from highly rational thinking all the way down to mere "association of ideas". (b) People differ in the coherence of their thinking. (c) The coherence of a person's thinking varies from time to time; it tends to get worse after a few drinks.

Now consider any thought that a person, Fiona, might have; for example, one she would express, "He likes spinach", with John in mind. Label the meaning (content) of this thought '*M1*'. Would any former thought that Fiona might have make it (metaphysically) impossible for Fiona to think a thought meaning *M1*? In particular, would the failure of an *M1*-thought to *cohere* with some immediately preceding thought prevent Fiona from thinking an *M1*-thought? Suppose, for example, that Fiona has a thought that means *M2* and that she would express, "John took the train from Paris to Istanbul". Would the failure of a thought meaning *M1* to cohere with one meaning *M2* make it (metaphysically) impossible for Fiona to follow her *M2*-thought with an *M1*-thought? I take it as obvious that the answer to all these questions is a resounding "No". Thus, Fiona might "associate" her *M1*-thought with her *M2*-thought because the *M2*-thought reminded her immediately of a previous encounter with John in which he rhapsodized about spinach. In sum, *a person can have any thought at all, however badly its meaning coheres with its predecessor*. Indeed, its degree of coherence with its predecessor is a *function* of their meanings. So, crucially, *a thought's meaning, and hence reference, are independent of its coherence*.

Turn now to language. Fiona may express any thoughts she has, however incoherent they are; people do crazy things. Her language will typically include rules for literally expressing any such thought (rules that may demand "saturation" in context, as the expression of Fiona's

M1-thought does). The semantic meaning of the resulting utterance will *be* the meaning of the thought that the utterance expresses. That is a consequence of the utterance *being* the literal expression of the thought according to the rules of the language (sec. 4). (Of course, a person may express the thought non-literally, resulting in an utterance that has a speaker meaning that differs from any semantic meaning it may have.) Thus, Fiona may express her series of thoughts in the following discourse, which is one of Stojnić's examples:

(40) John took the train from Paris to Istanbul. He likes spinach. (2021: 62)

This discourse consists of an utterance meaning *M2* followed by one meaning *M1*, those being the meanings of the thoughts expressed. We can conclude that *since the meaning and reference of a thought are independent of its coherence, so too are the meaning and reference of the utterance expressing that thought. So, coherence has no place in a theory of meaning or reference for language as well as for thought.*

The point here is that the meaning, *M1*, of "He likes spinach" is *not* so constituted that this sentence cannot be uttered after "John took the train from Paris to Istanbul" meaning *M2*, or indeed after *any* sentence meaning *anything*. This is not to say, of course, that it would be felicitous for Fiona to utter (40), nor that an audience would find (40) easy to understand. But *utterances that are infelicitous, even incomprehensible, can nonetheless be perfectly meaningful expressions of thoughts.*

6. *Coherence and understanding*

This introduces the next point. Coherence is very relevant to a hearer's *process of understanding* an utterance, to *successful* communication. That understanding involves using multiple clues to figure out, given the context, which meaning of an ambiguous term is likely, what saturations are likely to have occurred, what the utterance might mean non-literally, and so on. The likelihood of any interpretation being correct depends on whether it implies an *appropriate* degree of coherence in the speaker's thinking in the context. So, a hearer should interpret David Lewis' utterances so that they come out highly coherent, even after a drink or two; and a hearer should have much lower expectations of Donald Trump's utterances. So, coherence has a place in the theory of communication.

Consider (40). Stojnić claims "that the requirement that a discourse must be coherent is strikingly evident in the interpretive effort (40) elicits. Given apparently unrelated facts about John in (40), we search for a connection" (Stojnić 2021: 62). This is right about the interpretive effort, but that effort is not evidence that discourse must be coherent. It is evidence of the role of coherence in linguistic understanding.

7. *Diagnosis?*

Where has Stojnić gone wrong? In earlier works (Devitt 2013b, 2021: Ch. 7), I have identified a widespread flaw in the work of linguistic pragmatists/contextualists, the confusion of the *metaphysics of meaning* with the *epistemology of interpretation*. I wonder if the same confusion explains Stojnić's view that coherence relations are constitutive of meanings.[7]

Consider the "meaning-properties" of utterances in as broad a sense as you like, covering semantic meanings and speaker meanings, including conversational implicatures and the like. What *constitutes* an utterance having one of those properties is one thing, how a hearer *discovers* the property, another. The utterance's having the property is constituted by what the *speaker* does, by the conventions she participates in, the objects she has in mind, or the thoughts she intentionally expresses.[8] That is where we look for the "metaphysics of meaning". And what needs emphasizing is that none of these meaning-properties is constituted *in any way at all* by what the *hearer* does in trying to *interpret* what is said or meant.[9] The hearer's problem is an epistemic one of understanding an utterance. Grice (1989) made very clear that something like his "Cooperative Principle" and its associated maxims must play a role in the hearer's decision about what the speaker implicated but did not say. Later, pragmatists have demonstrated that something like that principle—perhaps the "Principle of Relevance" (Sperber and Wilson 1995)—must play a role also in the hearer's interpretive decision about what is said. Some such principle, along with contextual clues, will guide her in figuring out what conventions the speaker is using (including what language or dialect the speaker is using), what objects the speaker has in mind, and so on. And Stojnić has demonstrated, with examples like (40), the role that coherence plays in understanding. Any of these processes that the hearer uses to interpret an utterance might indeed provide *evidence* about an utterance's meaning-property but they do not *constitute* it. The hearer might do everything right, acting in accord with all appropriate communicative principles, and still get the wrong interpretation: she might *mis*understand.

[7] The common concern with "Grice's Circle" (Devitt 2021: 125–6.) is a sure sign of the confusion. The appearance of a problem here arises from equivocation between the constitutive and epistemic senses of 'determine' (2021: 125–6).

[8] But note two things. (1) All these meaning properties of utterances determined by the speaker are themselves ultimately constituted by the contents of thoughts. (2) The conventions that the speaker participates in are not of course constituted solely by her. They are constituted by the interdependent linguistic dispositions of the speech community that she is a member of (Devitt 2021: 79–80).

[9] This speaker-centered view of meaning flies in the face of Davidsonian "interpretationism" (1984). I have argued against this interpretationism elsewhere (1981: 115–18; 1997: 186–99; see also Simchen 2017). It rests on an unacceptable behaviorism: "Meaning is entirely determined by observable behavior, even readily observable behavior" (Davidson 1990: 314).

If Stojnić was confusing the metaphysics of meaning with the epistemology of interpretation that would explain her view that coherence relations are constitutive of meanings. But there is no persuasive independent evidence that she does confuse them. There is, however, a hint. Throughout the book, in discussing the likes of demonstratives, Stojnić talks of the "resolution" of context-sensitive reference (e.g. 2021: 4–5). Yet, given our concern with what *constitutes* the reference, it would be more appropriate to talk of the "fixing" of context-sensitive reference. For, talk of "resolution" is quite likely to misdirect us to how hearers *figure out* reference. So, I wonder if Stojnić's talk is a small sign that she has been misdirected.

Even if Stojnić has been misdirected, this is not to say that her considered opinion is that hearers' epistemic processes constitute meanings. But, as I emphasize in discussing linguistic pragmatism (Devitt 2021: 127, 132), the problem is not the *considered opinion* of theorists but rather their *theoretical practice* of taking meanings to be constituted by those epistemic processes.

8. *Conclusion*

Una Stojnić holds the radical view that coherence relations determine the reference of context-sensitive language. I have argued against this from the theoretical perspective presented in *Overlooking Conventions* (2021). Theoretical interest in language comes from an interest in thoughts and their communication. A language is a system of symbols, constituted by a set of governing rules, used (*inter alia*) to communicate the meanings (contents) of thoughts. Thought meanings, hence speaker meanings, are explanatorily prior to semantic meanings.

So, we start our consideration of the theoretical place of coherence by considering the bearing of coherence on thought meanings. I have argued that a person can have any thought at all, however incoherent. So, *a thought's meaning and reference are independent of its coherence*. Indeed, its coherence is a *function* of its meaning. Any thought can be expressed in an utterance. The semantic meaning of any utterance governed by the linguistic rules will be the meaning of the thought it expresses. So, *the utterance's meaning and reference are independent of its coherence*. I conclude that *coherence has no place in the theory of meaning or reference*. Nonetheless, it has a place in the theory of communication. I suspect that the error exemplifies the widespread confusion of the metaphysics of meaning with the epistemology of interpretation.[10]

[10] My thanks to Andrea Bianchi and Dunja Jutronić for comments on a draft.

References

Antony, L. 2008. "Meta-Linguistics: Methodology and Ontology in Devitt's *Ignorance of Language*." *Australasian Journal of Philosophy* 86: 643–56. URL: http://devitt.commons.gc.cuny.edu/spirited-exchanges/.

Collins, J. 2006. "Between a Rock and a Hard Place: A Dialogue on the Philosophy and Methodology of Generative Linguistics." *Croatian Journal of Philosophy* 6: 469–503.

___. 2008a. "Knowledge of Language Redux." *Croatian Journal of Philosophy* 8: 3–43.

___. 2008b. "A Note on Conventions and Unvoiced Syntax." *Croatian Journal of Philosophy* 8: 241–7.

___. 2020. "Invariance as the Mark of the Psychological Reality of Language." In A. Bianchi (ed.). *Language and Reality from a Naturalistic Perspective: Themes from Michael Devitt*. Cham: Springer, 7–44.

Davidson, D. 1984. *Inquiry into Truth and Interpretation*. Oxford: Clarendon Press.

___. 1990. "The Structure and Content of Truth." *Journal of Philosophy* 87: 279–328.

Devitt, M. 1974. "Singular Terms." *Journal of Philosophy* 71: 183–205.

___. 1981. *Designation*. New York: Columbia University Press.

___. 1997. *Realism and Truth*. 2nd edn. with a new afterword. 1st edn. 1984. Princeton: Princeton University Press.

___. 2003. "Linguistics is not Psychology." In A. Barber (ed.). *Epistemology of Language*. Oxford: Oxford University Press, 107–39.

___. 2006a. *Ignorance of Language*. Oxford: Clarendon Press.

___. 2006b. "Defending *Ignorance of Language*: Responses to the Dubrovnik Papers." *Croatian Journal of Philosophy* 6: 571–606.

___. 2008a. "Explanation and Reality in Linguistics." *Croatian Journal of Philosophy* 8: 203–31.

___. 2008b. "A Response to Collins' Note on Conventions and Unvoiced Syntax." *Croatian Journal of Philosophy* 8: 249–55.

___. 2008c. "Methodology in the Philosophy of Linguistics." *Australasian Journal of Philosophy* 86: 671–684.

___. 2009. "Psychological Conception, Psychological Reality." *Croatian Journal of Philosophy* 9: 35–44.

___. 2013a. "What Makes a Property 'Semantic'?" In A. Capone, F. Lo Piparo and M. Carapezza (eds.). *Perspectives on Pragmatics and Philosophy*. Cham: Springer, 87–112.

___. 2013b. "Three Methodological Flaws of Linguistic Pragmatism." In C. Penco and F. Domaneschi (eds.). *What is Said and What is Not: The Semantics/Pragmatics Interface*. Stanford: CSLI Publications, 285–300.

___. 2013c. "The 'Linguistic Conception' of Grammars." *Filozofia Nauki* 21: 5–14.

___. 2013d. "Responding to a Hatchet Job: Ludlow's Review of *Ignorance of Language*." *Revista Discusiones Filosóficas* 14: 307–12. URL: http://devitt.commons.gc.cuny.edu/spirited-exchanges/.

___. 2020. "Stirring the Possum: Responses to the Bianchi Papers." In A. Bianchi (ed.). *Language and Reality from a Naturalistic Perspective: Themes from Michael Devitt*. Cham: Springer, 371–455.

___. 2021. *Overlooking Conventions: The Trouble with Linguistic Pragmatism*. Cham: Springer.

___. 2022. "The Irrelevance of Intentions to Refer: Demonstratives and Demonstrations." *Philosophical Studies* 179: 995–1004. DOI: https://doi.org/10.1007/s11098-021-01682-5.

___. 2023. "Linguistic Luck; A Response to Rey and Collins." In A. Fairweather and C. Montemayor (eds.). *Linguistic Luck: Safeguards and Threats to Linguistic Communication*. Oxford: Oxford University Press, 124–58.

___. Forthcoming. "Demonstratives, Context-Sensitivity, and Coherence." *Linguistics and Philosophy*.

Dyer, F. C. 2002. "The Biology of the Dance Language." *Annual Review of Entomology* 47: 917–49.

Fodor, J. A., Bever, T. G. and Garrett, M. F. 1974. *The Psychology of Language: An Introduction to Psycholinguistics and Generative Grammar*. New York: McGraw-Hill Book Company.

Frank, A. 1997. "Quantum Bees." *Discover*: 80–7.

Grice, P. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.

Levinson, S. C. 2000. *Presumptive Meaning*. Cambridge: MIT Press.

Lewis, D. K. 1969. *Convention: A Philosophical Study*. Cambridge: Harvard University Press.

Longworth, G. 2009. "Ignorance of Linguistics." *Croatian Journal of Philosophy* 9: 21–34.

Ludlow, P. 2009. Review of Devitt 2006a. *Philosophical Review* 118: 393–402. URL: http://devitt.commons.gc.cuny.edu/spirited-exchanges/.

Matthews, R. J. 2006. "Could Competent Speakers Really be Ignorant of their Language?" *Croatian Journal of Philosophy* 6: 457–67.

Pietroski, P. 2008. "Think of the Children." *Australasian Journal of Philosophy* 86: 657–9. URL: http://devitt.commons.gc.cuny.edu/spirited-exchanges/.

Rattan, G. 2006. "The Knowledge in Language." *Croatian Journal of Philosophy* 6: 505–21.

Rey, G. 2006. "Conventions, Intuitions and Linguistic Inexistents: A Reply to Devitt." *Croatian Journal of Philosophy* 6: 549–69.

___. 2008. "In Defense of Folieism: Replies to Critics." *Croatian Journal of Philosophy* 8: 177–202.

___. 2020. **"**Explanation First! The Priority of Scientific Over 'Commonsense' Metaphysics." In A. Bianchi (ed.). *Language and Reality from a Naturalistic Perspective: Themes from Michael Devitt*. Cham: Springer, 299–327.

___ and John Collins. 2023. "Laws and Luck in Language: Problems with Devitt's Conventional, Common-Sense Linguistics." In A. Fairweather and C. Montemayor (eds.). *Linguistic Luck: Safeguards and Threats to Linguistic Communication*. Oxford: Oxford University Press, 88–123.

Riley, J. R., Greggers, U., Smith, A. D., Reynolds, D. R. and Menzel, R. 2005. "The Flight of Honey Bees Recruited by the Waggle Dance." *Nature* 435: 205–7.

Schiffer, S. 1972. *Meaning*. Oxford: Clarendon Press.

Simchen, O. 2017. *Semantics, Metasemantics, Aboutness*. Oxford: Oxford University Press.

Slezak, P. 2009. "Linguistic Explanation and 'Psychological Reality'." *Croatian Journal of Philosophy* 9: 3–20.

Slobodchikoff, C. N. 2002. "Cognition and Communication in Prairie Dogs." In M. Bekoff, C. Allen and G. M. Burchardt (eds.). *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*. Cambridge: MIT Press, 257–64.

Smith, B. C. 2006. "Why We Still Need Knowledge of Language." *Croatian Journal of Philosophy* 6: 431–56.

Sperber, D. and Wilson, D. 1995. *Relevance: Communication and Cognition*. 2nd edn. 1st edn. 1986. Oxford: Blackwell Publishers.

Stojnić, U. 2021. *Context and Coherence: The Logic and Grammar of Prominence*. Oxford: Oxford University Press.

___, Stone, M. and Lepore, E. 2013. "Deixis (Even without Pointing)." *Philosophical Perspectives, 27, Philosophy of Language*: 502–25.

___, ___ and ___. 2017. "Discourse and Logical Form: Pronouns, Attention and Coherence." *Linguistics and Philosophy* 40: 519–47. DOI: https://doi.org/10.1007/s10988-017-9207-x.

___, ___ and ___. 2020. **"**Pointing Things Out: In Defense of Attention and Coherence." *Linguistics and Philosophy* 43: 139–148. DOI: https://doi.org/10.1007/s10988-019-09271-w.

Vladusich, T., Hemmi, J. M. and Zeil, J. 2006. "Honeybee Odometry and Scent Guidance." *The Journal of Experimental Biology* 209: 1367–75.

Croatian Journal of Philosophy
Vol. XXIII, No. 69, 2023
https://doi.org/10.52685/cjp.23.69.7
Received: September 12, 2023
Accepted: October 31, 2023

What is a Tense, Anyway?

FABRIZIO CARIANI
University of Maryland, College Park, USA
MICHAEL GLANZBERG*
Rutgers University, New Brunswick, USA

We study three different conceptions of tense emerging from semantics, syntax and morphology, respectively. We investigate how they bear on the question of the relationship between tense and modality as they emerge in Cariani's The Modal Future *(2021).*

Keywords: Tense; modals; selection; semantics; syntax.

1. *Introduction*

One of the driving themes of Cariani's *The Modal Future* (2021, henceforth TMF) concerns the interplay of tense and modality in powering future reference. Building on prior work in semantics (Enç 1996; Condoravdi 2002; Kaufmann 2005; Copley 2009; Klecha 2014; Cariani and Santorio 2018, a.o.), Cariani argues that the devices languages recruit to power future-directed discourse are modals. In TMF's framing, an implicit corollary of this thesis is that because expressions like *will* are modals, they cannot also be tenses. Indeed, the book opens by contrasting a 'symmetric' paradigm in which languages have three tenses (past, present and future) with an alternative on which past and present are the 'just' tenses.

Does identifying modal features in *will*, or any other future expression, entail that it's not a tense? Answering this question in turn requires a grasp—preliminary as it might be—on the category of tense.

* This paper is the result of Cariani and Glanzberg collaborating on some themes in a commentary Glanzberg delivered on Cariani (2021) at the 2022 Philosophy of Language and Linguistics conference in Dubrovnik. The author of the book is treated here as a third person by both authors of this piece.

In this paper, we argue that the answer to this question is in an important sense indeterminate. There are multiple conceptions of tense which yield diverging answers to the question whether tense and modality are compatible—thus illuminating the relationship between tense and modality in a different way. We are aware that the territory can be carved in a much finer grained way than we are going to attempt here. The present paper stands as a public record on a series of ongoing conversation we hope to enrich and develop in future work.

2. *The semantic account of tense*

We begin our discussion by looking at the semantics of tense, and whether it can be used as the basis for a characterization of the category of tense itself. For the most part, we restrict our discussion to absolute, unembedded tenses. Though that leaves out some interesting subtleties, it is enough to illustrate what might be semantically distinctive about tense.[1]

2.1 *Two families of theories of tense*

The semantics literature offers up two families of theories about the meanings of tenses (Ogihara 2007). According to one, tenses are quantifiers over times—or perhaps quantifiers over intervals (Ogihara 1996; Kusumoto 1999, 2005). According to the other, tenses are pronoun-like, in that they make reference to times (or intervals) (Partee 1973; Heim 1994; Abusch 1997; Kratzer 1998). For illustration purposes, we will sketch a pronominal analysis. It is hard to say if either of these two approaches is more standard, for reasons we will return to below; but the pronominal approach is widely adopted, and a good representative of current work in the semantics of tense.

Pronominal analyses start with the observation from Partee (1973) that tenses pattern with pronouns in having deictic uses as in (1-a), anaphoric uses as in (1-b), and bound uses as in (1-c):

(1) a. Steve didn't turn the stove off.
b. Sheila had a party last Friday and Sam got drunk.
c. Whenever John came in, Sue left.

[1] Embedded tenses can display different semantic properties than their unembedded counterparts, and there is interesting cross-linguistic variation in how they do so. See Ogihara and Sharvit (2012) for a good survey of these issues. Absolute or 'simple' tenses provide one time, be it present or in the past. Following Reichenbach (1947) and then more recently Klein (1994), it has been observed that some tense constructions require more information, such as an utterance time or an event time. With these resources, more fine-grained distinctions among tenses can be made. The simple present roughly places the time given by the tense at the utterance time, while the simple past places the tense time before the utterance time. But, for instance, the perfect places the event time before the tense time. These relative tenses often overlap with aspect. For more on aspect, a good starting place is Smith (1997). For more on the perfect, see Portner (2011).

Taking this analogy seriously leads to a treatment of tenses as time pronouns.

This turns out to be doable and elegant. Like any other pronoun, a tense picks up its semantic value—a time interval—from an assignment function, which can reflect context in deictic uses and the effect of a quantifier in bound uses. We also need to ensure they have the right temporal properties. Unbound present tense should generally pick out the time of utterance, and past tense should pick out times in the past. Following Heim (1994), it is common to see these properties as presupposed. The semantic value of a tense is a time, but it presupposes the location of the time with respect to now. Formally, let *c* be the context, let *g* be the assignment function, and let be the relation that two temporal points bear to each other when they are near enough to each other. Let be the utterance time of a context:

(2) a. $[\![\text{PAST}_i]\!]^{g,c}$ is defined only if $g(i) < t_c$; if defined, $[\![\text{PAST}_i]\!]^{g,c} = g(i)$

b. $[\![\text{PRES}_i]\!]^{g,c}$ is defined only if $g(i) \approx t_c$; if defined, $[\![\text{PRES}_i]\!]^{g,c} = g(i)$

Composition of tense with verb phrases is now relatively easy. We need to assume verbs and other predicates have a temporal argument position.[2] Once we combine a verb with its ordinary arguments, there is still a temporal position to be filled. Syntactically, the verb combines with its ordinary arguments at a position called vP, and tense is a position above that, called T. Ignoring modals (and aspect), we have a structure like:

(3)

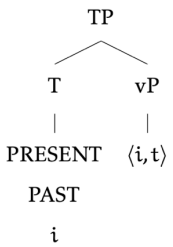

Here, *i* is the name of the type of times (or intervals).

We will comment more extensively on the labels TP, T, and vP in the next section, as we look at the syntactic account of tense. For now, all we are assuming is that syntax provides us with a tense phrase, headed (in the case of English, at least) by one of two simple tenses, and taking as complement a verb phrase. Function-argument composition suffices for this case.[3]

[2] See Enç (1986), Heim (1994) and Abusch (1997) for discussion of this idea.

[3] Embedded tenses, which show up in attitude contexts, make all of this more complicated. But this is a good illustration about how a compositional semantics with tenses works.

The other prominent idea about tenses is that they are quantifiers over times. This can be implemented in a standard generalized quantifier framework, as found in many semantics textbooks (e.g. Heim and Kratzer 1998). Tenses are quantifiers that take as input predicates of times, of type . We then have:

(4) a. $[\![\text{PAST}_i]\!]^{g,c} = \lambda p_{\langle i,t\rangle}.(\exists t' < t_c (p(t') = 1))$

b. $[\![\text{PRES}_i]\!]^{g,c} = \lambda p_{\langle i,t\rangle}.(p(t_c) = 1)$ (semantically vacuous)

Composition is also not difficult. If we assume, as we did above, that the vP is of type $\langle i,t\rangle$, then it can compose with a quantifier directly. As we have mentioned already, there are complications about embedded tenses, and there are questions about how a quantifier winds up in a T position. But again, we have a relatively clear beginning of a semantic analysis of tense.[4]

These two approaches give different semantics to tenses, but they are surprisingly hard to tease apart empirically. Though the pronominal approach is designed to explain the Partee analogy with pronouns, so can the quantificational theory. The key ingredient in this explanation is the idea that the quantifier must be contextually restricted, and the general observation that quantificational restrictors can quite generally be involved in deictic and anaphoric uses.

The pronominal theory makes the temporal content of tenses presupposed, while quantificational theories make it asserted. This seems like a substantial difference, but again, it is hard to spot in the data. The reason is that the presuppositional status of temporal information is itself a delicate issue. On a pronominal theory, temporal information is treated like other features on pronouns, such as gender and number (called 'phi-features'). The content of these features is like presupposition, in that it is backgrounded, but it is not at all clear that it projects just like presupposition (e.g. Kratzer 1998; Heim 2008; Sudo 2012). Here is one illustration:

(5) John thinks it is 10:00.

Suppose this is uttered at 11:00. A standard account of how presuppositions project out of attitudes (Heim 1992) predicts that this presupposes that John thinks 10:00 is 11:00, or at least that 10:00 ≈ 11:00. But that is not right. Because the way tense projects is delicate, it is not easy to find clear examples decisively refuting either analysis.

[4] Philosophers might have been expecting a semantics of tense along the lines of tense logic (e.g. Prior 1957, 1967; van Benthem 1983). Indeed, some early work in semantics (e.g. Montague 1970) used such an analysis. Subsequent work has shown it not to be promising. As such work has focused on embedded tenses, we will not discuss it in detail. See among places Richard (1981), Enç (1986), King (2003), Kusumoto (2005), Glanzberg (2011), and Glanzberg and King (2020). For a general comparison of quantifier versus operator theories, see Cresswell (1990).

We might think that quantifier scope would distinguish the two theories. When quantifiers scope, their behavior looks different from what we get with presuppositions. But simple tenses do not really show much of any quantifier scope. They do not scope with negation, for instance:

(6) a. John cried.
b. John did not cry.

Both require there to be a time in the past when John did/did not cry. As we will see below, this does not reveal much. Syntactically, tense is already in a position that limits scope. So, this is compatible with a pronominal analysis, but also with a quantificational analysis that puts tenses in a syntactic position that limits scope.

Ogihara (1995) offers one argument in favor of a quantifier view:

(7) a. Did you see Mary?
b. I saw her, but I don't remember exactly when.

Ogihara claims the second sentence gets a purely existential reading, without any anaphora or binding. If so, it suggests we can sometimes get a purely quantificational reading of past tense.

But we doubt this is conclusive. What we need is a purely existential reading, along the lines we see with:

(8) John ate.

This has a reading (the most natural one) where John at something or another, and it is unconstrained what (perhaps beyond it being normal food). Ogihara's example seems to us not so unrestricted. It would be sufficient to have a contextually provided and fairly large time interval. If so, a pronominal theory can explain it.

2.2 *Semantics of tense vs. semantics of* will

There is no doubt much more to be said about the proper semantics for tense. What matters for our purposes is what happens if we adopt either approach as a necessary ingredient in the category of tense.

If we assume that all tenses must have the pronominal semantics in (2) or the quantificational semantics in (4), the questions we led with get to have straightforward answers. To start, under this assumption, tense and modality are naturally understood to be incompatible categories. Zooming in on English, under this assumption it is hard to escape the conclusion that *will* is not a tense (and similarly for other predictive expressions of English). Much of the argument in chapter 3 of TMF and the semantics literature it references is an argument to the effect that *will* is not well understood as meaning the same as:

(9) $[\![\text{FUT}_i]\!]^{g,c}$ is defined only if $g(i) > t_c$; if defined, $[\![\text{FUT}_i]\!]^{g,c} = g(i)$

That is to say, it is not well understood semantically as simply being the mirror image of past tense. Indeed, under the assumption it is not clear that English has a future tense.

T MF—following and expanding on Cariani and Santorio (2018) and Klecha (2014)—advances four arguments in defense of a modal theory of *will*. In rough summary, these are:

- *The argument from common morphology* (§3.2): *will* shares morphology with would; would is a modal, so is will.
- *The argument from present-directed uses (§*3.3*)*: *will* seems to have present directed uses that appear to have a vaguely modal flavor (as in *the president will be in his* office *by now).*
- *The argument from modal subordination (§*3.4*)*: *will* goes in for modal subordination, which is something that modals do. Cariani highlights this as the centerpiece of the overall argument.[5]
- *The argument from the acquaintance inference (§*3.5*)*: *will* appears to obviate the acquaintance inference, a property which it generally shares with other modals, and that distinguishes it from past tense.

Let us assume that these arguments collectively work to support a theory according to which *will* is given a modal semantics. The low-effort option for a modal semantics is to assimilate *will* to universal modals in a Kratzer-style semantics (Kratzer 2012). Let us notate the modal base $f(\cdot)$ and the ordering source $os(\cdot)$.[6] Furthermore, we package all the domain formation mechanics of Kratzer's semantics into a single domain-construction function, notated as domain (f,os,w) (see, e.g. von Fintel and Heim 2011, for more details). We then have:

(10) $[\![\text{will}]\!]^{g,c} = \lambda\mathbf{p}\lambda w.\forall w' \in \text{domain}(\text{f},\text{os},w'),\mathbf{p}(w')$

Cariani (2021) contrasts this with the selectional account, which in its simplest form looks like this: let *sel* be a function that inputs a set of worlds and a world, and outputs a 'selected' world in the modal base. Suppose further that *sel* is subject to two constraints, that we state as part of the entry in (11):

(11) a. centering : for all w and modal base f, if $w \in \text{f}(w)$, $\text{sel}(\text{f}(w),w) = w$

b. success : for all w and f, if $\text{f}(w) \neq \emptyset$, $\text{sel}(\text{f}(w),w) \in \text{f}(w)$.

c. $[\![\text{will}]\!]^{g,c} = \lambda\mathbf{p}\lambda w.\mathbf{p}(\text{sel}(\text{f}(w),w))$

Both the universal and the selectional entries above are in need of refinement. In particular, neither reflects the temporal orientation of *will* (this matter is discussed in chapter 7 of TMF). But however we decide to expand on them, it is clear that the end result is not going to match either the pronominal or the quantificational theory.

[5] The argument comes in for some interesting criticism in Boylan's (2023) review of the book and for some expansion in Cariani (forthcoming).

[6] Standard notation *f* or ordering sources is $g(\cdot)$, but we have already recruited '*g*' for the assignment function.

Overall, our conclusion is that tense has a range of specific semantic properties, and the future *will* differs from tense in important ways. But if anything, this strengthens our confidence that we cannot read the nature of tense off its semantics. The empirical situation does not nail down what semantic type a tense must have; and the typing of tense semantically does not constrain the semantics of the future.

If we did assume that the semantics of tense—be it the referential or the quantificational variety—is a guide to the nature of tense, the cross-linguistic picture would also become significantly more puzzling. Some languages, including e.g. Romance languages, have dedicated morphology for future reference. Under the semantic conception of tense, this morphology only gets to count as *tense* morphology if it turns out that its correct semantics is as in (9). Not only does this seem to not be guaranteed *a-priori*, but insofar as the arguments for a modal semantics carry over to these other languages we cannot consistently assume that (i) the future morphology in Romance languages is a type of tense, (ii) that tenses are associated with a particular kind of lexical entry, and (iii) that the future in these languages is like the English future in demanding a modal semantics. Throughout Chapter 3 of TMF Cariani suggests that at least some of the arguments for a modal analysis of *will* do carry over to the Italian language.

If it turns out that the right semantics for simple tense in languages like English is referential, and the right semantics for the future *will* is modal, then we have a clear difference. But our discussion here has shown that even with simple past and present, we do not (so far) have a clear-cut semantic category. And of course, we also do not have to accept the background assumption that the category of tense is homogeneous in its semantic behavior. In the rest of this paper, we consider two more ways of conceptualizing tense that do not have this implication. When it comes to semantics, we doubt that there is really a goal of providing a definition of tense; rather, the goal is to provide semantic analyses of the various puzzling semantic properties of tense. Embedded tense has provided a rich diet of such puzzles, so there is much work to be done.

3. *The syntactic account of tense*

A glance at the syntax literature shows a special place for a functional category of T for Tense.[7] So, one answer to the question of the nature of tense is that it is what occupies a special syntactic position.

The basic idea is that clauses, the main units we utter and otherwise use, come in layers. It is not easy to put this idea in an entirely theory-neutral way, so we will make use of a common tradition in gen-

[7] This can be found in many contemporary syntax textbooks, such as Adger (2003) that we rely on heavily, as well as Carnie (2021). The main idea can be found in Chomsky (1986), and important work of Edmonds (1980) and Pollock (1989).

erative linguistics. We lean on the Chomskian project as it grows out of the 'Principles and Parameters' tradition (e.g. Chomsky 1986) and evolves into the 'Minimalist' tradition (e.g. Chomsky 1995). In this kind of framework, one important layer that occurs fairly high in a syntactic tree is Tense Phrase or TP. Tenses are heads of TPs.

We begin unpacking this idea by discussing two important layers. The first is now known as the vP layer. This is where basic descriptions of events occur, and it typically involves verbs, whose main job is to describe events and states. But to do so, verbs need to add the participants in the event. The verb *to give*, for instance, describes events of giving. But making a clause requires specifying who is doing the giving (the agent of the event), what is being given (the theme of the event), and a recipient (the 'goal' of the event). Thus, a verb needs to combine with its arguments: an intransitive verb requires one argument, a transitive two, and a ditransitive three. A verb can also combine with adjuncts that further specify the participants in the event. Some verbs, like *cut* take an instrument.[8] Syntactically, there is a place where a verb merges with its arguments and any appropriate adjuncts—a predicate meets its arguments and together they describe something (Glanzberg 2011). In current theories, this layer is called vP. Languages seem to have many types of predicates: some are formed by combining nouns and adjectives with other materials (e.g. copulas). But there is a special place for verbs in building clauses, that is captured by a vP analysis.

A vP is not a sentence. It is not really the kind of thing a speaker may utter, and it is not a full clause semantically or syntactically. Semantically, a vP describes an event and its participants, but it is neither temporally nor aspectually determinate. It does not locate the event in time, nor does it tell us if the event is completed or still happening. Syntactically, it leaves out all the inflectional elements that language requires for a clause.

Inflectional elements, like auxiliaries in English, live high in the syntactic tree, as has been clear since Chomsky (1957). Current theories indicate there is a very high layer of TP above vP, where temporal information is added. In most theories of the sort we are considering, T is the point where you get a fully inflected clause—the sort of thing we can normally assert, for instance. So, a TP is good candidate for being the first place where we get a 'sentence'. A sentence, in this theory, is headed by T. Also, in many theories, subjects of sentences get special treatment and occupy the syntactic position of the 'specifier' of TP. TP is the layer where subjects appear where they are supposed to.[9]

At this point of our description, we have identified two clausal layers. Next, we observe that they come in a distinct structural order:

[8] The status of these as arguments versus adjuncts is actually somewhat controversial, but we do not need to take a stand on this issue here. See among places Larson (1988) and Bhatt and Pancheva (2017).

[9] Any of the syntax textbooks we mentioned will explain this, but see also classic work of Stowell (1981) and McCloskey (1997).

the TP is higher in the clause than the vP. Evidence for this structure comes from a number of sources, including observations about word order. Here are some textbook examples, from Adger (2003). First, modal auxiliaries, including *will*, occupy a position outside of vP. We see this from the grammatical impossibility of certain inversions that would put them there:

(12) a. * Gilgamesh seek will/must/may Ishtar.
b. What Gilgamesh will/must/may do is [seek Ishtar].

The same holds for the auxiliary *do* and its inflected forms *does* and *did*:

(13) a. Enkidu did free animals.
b. * Enkidu free did animals.

So far, we have a rough division into two layers, one of which hosts inflectional elements. It is also telling that these inflectional elements have a close relation to tense. In this position, *will* and *do* appear inflected for tense, and indicate temporal information. So, at the very least, we can conclude with Chomsky that a very high inflectional layer is where we expect to find tense and related elements.

Why single out tense, TP, as a distinct layer and high among inflections. Why make T the head of a sentence?[10] Here matters get more delicate. One reason to think the TP layer is higher than the position of modals comes from the way modals—including *will* and *would*—inflect for tense. It suggests that a tense applies to a lower common modal, often labeled *woll*. When *woll* combines with present tense it spells out as *will*. When it combines with past tense it spells out as *would*.[11] Evidence that this is inflection for tense comes from the way it patterns with tense in embedded contexts:

(14) a. I thought she was happy.
b. * I thought she is happy.

(15) a. I thought she would go.
b. * I thought she will go.

Of course, the markings of tense over modality are also apparent in romance languages and other languages in which temporal reference is powered by a grammatical system of morphemes. For example, Italian necessity (*dovere*) and possibility (*potere*) modal auxiliaries, can inflect for past tense (*dovetti*/*potetti*), present (*devo*/*posso*), and future (*dovrò*/*potrò*).

Another relatively clear observation is that aspectual marking occupies a different position, lower than TP. By "aspectual marking", we mean the grammatical marking of perfective, imperfective, and pro-

[10] After all, in earlier theories, such as that of Chomsky (1981), what we had was an undifferentiated inflectional layer IP. Pollock (1989) was central to showing that we have multiple inflectional layers, with TP near the top.

[11] See Ogihara (1996) and Abusch (1997). Apparently the label *woll* was suggested by Mats Rooth.

gressive.[12] In English, the progressive appears below tense, as we see with:

(16) a. (i) Sarah wrote a dissertation.
(ii) [PAST [Sarah write dissertation]]
b. (i) Bill was in love with Sarah.
(ii) [PAST [Bill in love with Sarah]]
c. (i) Sarah was writing a dissertation.
(ii) [PAST [PROG [Sarah write dissertation]]]

So far, we have TP appearing very high, and modals and aspectuals appearing below it, but above the vP. Beyond this, the situation gets even more complicated, and the evidence typically involves cross-linguistic comparisons. Much current work implies, or assumes, the existence of a stable hierarchy. The seminal paper for this is Cinque (1999). Setting aside some complications involving non-root modals, this hierarchy looks like this:

(17) Tense > Aspect > $\text{Modal}_{\text{root}}$

(Note that we have not justified the position of root modals with respect to aspect, nor the restriction to root and not epistemic modals (see Cinque 1999; Hacquard 2010).) Assuming this ordering indicates syntactic positions Asp and Mod and associated phrases AspP and ModP, the syntactic picture that emerges, closely enough, includes structure like:

(18)

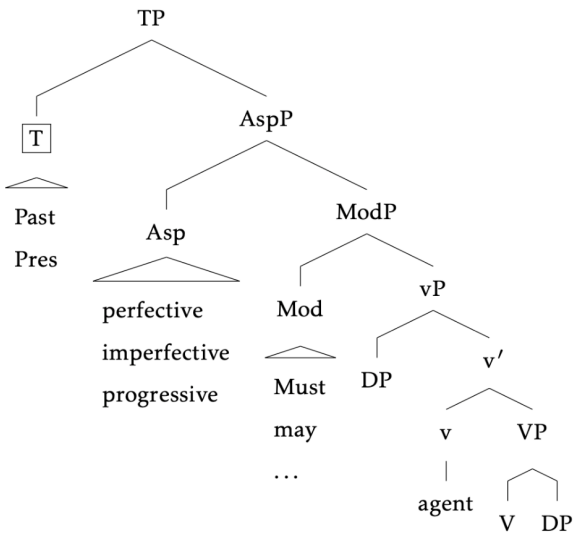

[12] This kind of aspect is sometimes called 'viewpoint aspect' (Smith 1997). Semantically it indicates whether we see an event as completed or ongoing. For surveys of the grammar of aspect, see de Swart (2012), Zagona (2013), and the comprehensive Smith (1997).

We hasten to add that this is oversimplified and some of the claims built into it are controversial. It is common to see epistemic modals as occupying a position above T (e.g. Cinque 1999; Hacquard 2010), and yet there is a vivid debate concerning whether they may themselves embed under tense. Question forms are almost always assumed to project a layer above TP, usually called CP. Many theories in the 'cartographic' tradition posit much more above TP (e.g. Rizzi 1997).[13] We have not tried to say where negation fits in. It is traditionally placed below T (Pollock 1989), but the issue remains controversial (Zanuttini 2001). Furthermore, there is more syntactic complexity to T. It is often seen as central to issues of case, agreement, finiteness, and so on.

So, here is a plausible idea about what tense is: it is a syntactic position. And if we want to distinguish genuine tenses in English from *will*, we have syntactic resources to do it. Genuine tenses live in T, while *will* lives in Mod, along with other root modals. We could go even further, and claim that, at least in English, what lives in T has the kind of semantics we reviewed in the previous section, while what lives in Mod has a different, distinctively modal semantics. One of the main theses of TMF is that *will* has a very particular modal semantics. But even if all of that argument failed, we would still find clean distinctions both syntactically and semantically between genuine tenses and *will*. So perhaps our syntax and semantics give us independent ways of narrowing down on the same core phenomenon.

This outlook might appear very satisfying. It seems well-justified for English, as well as other languages that are relevantly like it. And it builds on the Cinque hierarchy which, together with its relatives, seems well-supported by cross-linguistic evidence. But there remains great room for caution. Our goal in asking about the nature of tense was more ambitious than to simply ask how we can spot tenses in English, German, and some other languages. We wanted something more fundamental. Whether we have that is much less clear.

One way to press this concern is to ask about the extent to which we have latched onto a phenomenon that is linguistically universal. Here, the situation is not so clear. One point that will become more central in the rest of our discussion is that many languages lack overt tense morphology altogether. That can make the question of whether a language has a T head very complicated. It is all the more complicated by the many different jobs we have asked T to do. What we are considering is in effect a proposal discussed by von Fintel and Matthewson (2008: 170), who put it like this: "All languages possess a syntactic head T whose function is to locate the reference time with respect to the utterance time." They quickly conclude that this is "probably false," though they then note that what is really needed is more work, and that there

[13] This is the enterprise of mapping the functional structure of languages, often relying on extensive cross-linguistic investigation. The already-mentioned Rizzi (1997) and Cinque (1999) are good examples.

may yet be generalizations like this to be uncovered, even if this one, as stated, is likely false.

Here is one illustration of the concerns that drive von Fintel and Matthewson. It is well known that Mandarin Chinese shows no overt tense morphology (and little morphology of any kind). One might claim that in spite of this, it has a phonologically null T position (Sybesma 2007). This would support the proposed universal. But there are other options. According to influential analysis by Lin (2006), there is no T node, and viewpoint aspect, syntactically AspP, does the work in Mandarin that tense does in English. If this analysis is on the right track, it might undercut the claim that a highest tense layer is a universal. If Lin is right, the work done by T in many languages might be shifted to AsP in others. This is not to say that the analysis is certainly correct,[14] but it illustrates the reasons for von Fintel and Matthewson's caution.

So, to put it over-simply, one possibility is that a different syntactic position, perhaps AspP, can do the work that TP does in English. Here is another interesting possibility, put forth by Matthewson herself (Matthewson 2006) for St'át'imcets (Lillooet Salish). Like Mandarin Chinese, St'át'imcets lacks tense morphology, and so might appear tenseless. In this case, and in contrast to Lin, Matthewson argues that St'át'imcets does have a T head. Unlike English, semantically what occupies that head more or less expresses being non-future. That means we can have a T head, but not have it filled by what we normally think of as tense in languages like English. Again, we have lost the simple identification of tense through overlapping semantic and syntactic properties, if we are seeking full linguistic universality.

We recommend Matthewson's discussion (in her section 7) of the various complexities of talk of languages being tenseless. For our purposes, we can leave the matter with the observation that, given the number of roles TP is asked to play in many theories, it is not a huge surprise that we can find detailed analyses of specific languages that divide up those roles differently, both syntactically and semantically. Thus, if in fact the particular combination of roles we find in English turns out not to be universal, that would not be very surprising. This illustrates the point that there are, we believe, important semantic and syntactic properties that go with tense, but perhaps not a straightforward standard for what a tense is that is cross-linguistically universal.[15]

We do note with satisfaction Matthewson's speculations at the end of her paper, where she suggests that it may be universal that the future is different from present or past. This is in keeping with the mo-

[14] There might be some other complications here: lexical aspect ('aktionsart') and scope and temporal adverbs are also important. Additionally, we would need to know what hosts subjects if AspP is the highest layer. (Lin (2010) addresses this issue in more detail.)

[15] Many other works substantiate this. Among them: Ritter and Wiltschko (2009), which discusses Blackfoot (Algonquin) and Halkomelem (Salish); and Bittner (2014), which discusses both Mandarin Chinese and Kalaaisut (Greenlandic).

tivating view of TMF. It is suggestive that major cross-linguistic work points in the same direction, albeit admittedly in a discussion which is explicitly labeled as speculative.

Also, we should note that the Cinque-style placement of T in a rigid hierarchy raises some very abstract questions about what good explanations in syntax should be. We will not go into detail, as the foundations of linguistic theory is not our topic here, but we can simply note that the Cinque hierarchy is data-driven but has seemed to many to be stipulated. (Much the same is claimed for other exercises in syntactic cartography.) One might wonder if other explanations can be found. This is of particular concern in the current literature, as it relates to some of the key goals of the 'minimalist program' in syntax (e.g. Chomsky 1995). We recommend the discussion of Ramchand and Svenonius (2014), which though opinionated in its conclusions, is judicious in its overview and gives a good sense of the issues. We also mention Ramchand and Svenonius as it raises the possibility that one might opt for a more coarse-grained functional hierarchy than Cinque offers, that will be less discriminating between tense and other inflectional elements. If that turns out right, the robust distinction between tenses and modals in syntax might vanish. We have already considered reasons that might be so cross-linguistically, but we should be aware that the theoretical situation is complex even for one language—English.[16]

Where does this leave our question about what a tense is, and whether a modal is a tense? As with semantics, syntax offers us important insights but nothing like a definitive criterion. Different analyses put the work of temporal modification in different positions, and the idea that there is a hard-wired and robust demarcation between tense positions and other high positions remains contentious, and may be more like a helpful theoretical idealization than a robust fact about language.

Yet there still seems to be sufficient evidence to distinguish English *will* from tenses. Our conclusion in the semantics discussion was that though semantic behavior alone does not cleanly demarcate tense, there are substantial differences between tenses and *will*. The analogous conclusion here is that even if we cannot say there is one universal syntactic position for tense, there are substantial syntactic differences between the positions of tense and those of operators like *will*. We cannot say for certain that these differences are universal, but they do appear to have *some* cross-linguistic robustness.

What, then, is a tense? Well, we can say with some specificity what a tense is in English, and others have said what it is in Japanese, Mandarin Chinese, St'át'imcets, and so on. These show some common elements, and some variation. The analysis from Lin (2006) gives a good

[16] It might be that there are syntactic generalities here yet to be found, as Matthewson hints at but does not claim. It might be that a more abstract level of description might yield better results. See, for instance, Wiltschko (2014).

illustration. According to this analysis, much of the standard semantics of tense is written into the semantics of aspectual markers. One is left wondering if the right way to describe things is there is no tense, or rather if aspect kindly absorbed the job tense might have done. We can say the same about Matthewson's analysis of what might occupy a TP position. So we have a bundle of features of tense, and jobs that it does, both semantically and syntactically. These seem to pattern strongly together. But the various parts of the bundle can be divided in somewhat different ways, as Lin's analysis illustrates. We do not think it is a great surprise that languages might divide up such a bundle in somewhat different ways.

4. *The morphological account of tense*

We have so far looked at tense in semantics and syntax. We have been cautious to avoid making global claims, but we have continued to offer two linked ideas. The distinctive properties of tense may come together differently cross-linguistically, but they show strong distinction between tenses and modals.

Another equally influential conception of tense focuses on the morphology. Tense is a grammatical system whose job it is to anchor situations to certain times, defined by their relation to the utterance time, in the process of fixing their truth conditions. According to Comrie's (1985: 9) extremely influential definition, tense is "grammaticalised location in time." What counts as a grammatical system is itself a vexed question, but plausibly a system of bound morphemes counts as such (a bound morpheme is one that only occurs as a proper part of a word). Comrie again says:

> The English past/non-past opposition is a clear instance of a grammaticalised opposition. It is quite impossible to construct an English sentence containing a finite verb that is neutral as between the two poles of this opposition, i.e. *John runs* is clearly non-past, and *John ran* is clearly past, and there is no third term that is neither. Moreover the expression of the distinction is by means of bound morphemes (taken to include morphophonemic alternation, i.e. anything that does not involve a separate word). (1985: 10)

Under this view, what it is for the sentence *I played soccer* to be tensed is that it features the English bound morpheme *-ed* with the verb; the semantic role of this morpheme is to locate the situation emerging from the whole verb phrase in the past.

Like the previous analyses, the idea that tense is part of a grammatical system of bound morphemes directly implies that the devices that achieve future reference in English—auxiliaries like *will* and phrases like *going to*—are not tenses. As we have noticed, *will* is inflected for tense, and furthermore it appears in complementary distribution with other modals:

(19) a. Enkidou will free animals.
b. Enkidou might free animals
c. Enkidou may free animals.

At the same time, the morphological approach turns the question whether a language has tenses—as well as the question which tenses a language has—into a rather brittle, language-variant matter. English has a bound morpheme for past tense (*–ed*); more controversially, English can be viewed as having the bound morpheme *–s* for third person singular present (otherwise it does not appear to mark the present tense).[17]

Romance languages typically offer of inflectional paradigms for past, present and future. Thus, French and Italian have a simple future:

(20) a. Nager / nuotare ("to swim")
b. Je nage / Io nuoto ("I swim")
c. Je nagerai / Io nuoteró ("I will swim")

Some languages, such as Mandarin Chinese, lack these bound morphemes entirely.[18]

We have already seen a few options for how to approach languages without bound tense morphemes. In some cases, as Matthewson (2006) argued for St'át'imcets, it might be there are unpronounced morphemes, that occupy syntactic T heads. Or it might be, as Lin (2006) argued for Mandarin Chinese, that there are other aspectual markers that do the work of tense. Now, one might take Lin's proposal, and the proposal of TMF as counterexamples to the morphological conception of tense. Tenses do not have to be realized by systems of bound morphemes, and if a word like *will* wants to behave like a tense, we should not deny it tense status just because it is not a bound morpheme. This will either send us back to the semantic conception of tense (thus to characterizing tenses as items with temporal meanings) or to the syntactic conception; or as we suggested above, to a view that looks for multiple features and how they are divided up in a given language.

There is, however, another way here. One may insist that the morphological conception of tense is roughly correct. The somewhat radical conclusion would be that when it comes to theorizing about items with temporal meanings, "tense" is an unhelpful category, because it only latches on a incomplete subset of the whole panoply of temporally significant expression. Such a category might serve an important purpose for typological investigations. It can be a useful one. But our discussion of semantics and syntax suggests it may miss some important underlying commonalities in languages that differ substantially in morphology.

Another unexpected conclusion one would draw here is that tense and modality are not incompatible categories (Cariani forthcoming). The very same item, say the Italian or French future tense may be

[17] For a descriptively oriented discussion of English, see Huddleston and Pullum (2002).

[18] In addition to the references above, see also Bochnak (2019).

both a tense because it satisfies certain morphological criteria and a modal because it bears semantic properties that naturally group it with modals. As we have seen, there are multiple ways one can examine tense, and it is not all that surprising that they can cross-cut eachother in some cases.

5. *Discussion and conclusions*

It may be that our main conclusion does not *have to* be stated out loud. But it's probably a good idea to do so anyway. There is not *one* clear answer with regards to the question *what is a tense?* Consequently, the question of whether tense and modality can overlap does not have a unified, fully determinate answer. What we can do, however, is explore the different things that are called "tense" in the context of linguistic research, articulate multiple precise conceptions and answer our motivating question against each of them.

Yet we have also suggested that even with multiple, partly overlapping notions of tense, patterns may still emerge. We, with Bochnak and Matthewson, suspect that a robust tense versus modal distinction can be found within the many overlapping ideas about tense. We recognize that this remains speculative, especially when it comes to the rich and confusing range of cross-linguistic data and theories available. But we think it an appealing speculation.

References

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

Adger, D. 2003. *Core Syntax.* Oxford: Oxford University Press.

van Benthem, J. 1983. *The Logic of Time.* Dordrecht: Kluwer.

Bhatt, R. and Pancheva, R. 2017. "Implicit arguments." In M. Everaert and H. van Riemsdijk (eds.). *The Wiley Blackwell Companion to Syntax. Second Edition.* Hoboken: John Wiley and Sons, 1–35.

Bittner, M. 2014. *Temporality: Universals and* Variation. Oxford: Wiley-Blackwell.

Bochnak, M. R. 2019. "Future reference with and without future marking." *Language and Linguistics Compass* 13: 1–22.

Boylan, D. 2023. "Review of *The Modal Future." The Philosophical Review* 132: 320–325.

Cariani, F. 2021. *The Modal Future.* Cambridge: Cambridge University Press.

Cariani, F. Forthcoming. "Future displacement and modality." In E. Lepore and U. Stojnić (eds.). *Oxford Handbook of Philosophy of Language: Volume 2.*

Cariani, F. and Santorio, P. 2018. "*Will* done better: Selection semantics, future credence, and indeterminacy." *Mind* 127 (505): 129–165.

Carnie, A. 2021. *Syntax: A Generative Introduction.* Fourth edition. Malden: Wiley-Blackwell.

Chomsky, N. 1957. *Syntactic Structures*. The Hague: Mouton.
Chomsky, N. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.
Chomsky, N. 1986. *Knowledge of Language*. New York: Praeger.
Chomsky, N. 1995. *The Minimalist Program*. Cambridge: MIT Press.
Cinque, G. 1999. *Adverbs and Functional Heads*. Oxford: Oxford University Press.
Comrie, B. 1985. *Tense*. Cambridge: Cambridge University Press.
Condoravdi, C. 2002. "Temporal interpretation of modals: Modals for the present and for the past." In D. Beaver, L. Casillas, B. Clark and S. Kaufmann (eds.). *The Construction of Meaning*. Palo Alto: CSLI Publications.
Copley, B. 2009. *The Semantics of the Future*. Routledge.
Cresswell, M. J. 1990. *Entities and Indices*. Dordrecht: Kluwer.
Edmonds, J. 1980. "Word order in generative grammar." *Journal of Linguistic Research* 1: 33–54.
Enç, M. 1986. "Towards a referential analysis of temporal expressions." *Linguistics and Philosophy* 9: 405–426.
Enç, M. 1996. "Tense and modality." In S. Lappin (ed.). *Handbook of Contemporary Semantic Theory*. Oxford: Blackwell, 345–358.
von Fintel, K. and Heim, I. 2011. "Notes on intensional semantics." MIT.
von Fintel, K. and Matthewson, L. 2008. "Universals in semantics." *Linguistics Review* 25: 139–201.
Glanzberg, M. 2011. "More on operators and tense." *Analysis* 71: 112–123.
Glanzberg, M. and King, J. C. 2020. "Binding, compositionality, and semantic values." *Philosophers' Imprint* 20: 1–29.
Hacquard, V. 2010. "On the event relativity of modal auxiliaries." *Natural Language Semantics* 18: 79–114.
Heim, I. 1992. "Presupposition projection and the semantics of attitude verbs." *Journal of Semantics* 9: 183–221.
Heim, I. 1994. "Comments on Abusch's theory of tense." In H. Kamp (ed.). *Ellipsis, Tense, and Questions*. DYANA Deliverable R2.2.B. Amsterdam: University of Amsterdam, 143–170.
Heim, I. 2008. "Features on bound pronouns." In D. Adger and S. Bejar (eds.). *Phi Theory: Phi-Features Across Modules and Interfaces*. Oxford: Oxford University Press, 35–56.
Heim, I. and Kratzer, A. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.
Huddleston, R. and Pullum, G. K. 2002. *The Cambridge Grammar of the English Language*. Cambridge: Cambridge University Press.
Kaufmann, S. 2005. "Conditional truth and future reference." *Journal of Semantics* 223: 231–280.
King, J. C. 2003. "Tense, modality, and semantic value." *Philosophical Perspectives* 17: 195–245.
Klecha, P. 2014. "Diagnosing modality in predictive expressions." *Journal of Semantics* 313: 443–455.
Klein, W. 1994. *Time in Language*. London: Routledge.
Kratzer, A. 1998. "More structural analogies between pronouns and tense." *SALT* 8: 92–110.

Kratzer, A. 2012. *Modals and Conditionals*. Oxford: Oxford University Press.

Kusumoto, K. 1999. *Tense in Embedded Contexts*. Ph.D. dissertation. University of Massachusetts at Amherst.

Kusumoto, K. 2005. "On the quantification over times in natural language." *Natural Language Semantics* 13: 317–357.

Larson, R. K. 1988. "Implicit arguments in situation semantics." *Linguistics and Philosophy* 11: 169–201.

Lin, J.-W. 2006. "Time in a language without tense: The case of Chinese." *Journal of semantics* 23: 1–53.

Lin, J.-W. 2010. "A tenseless analysis of Mandarin Chinese revisited: A response to Sybesma 2007." *Linguistic Inquiry* 41: 305–329.

Matthewson, L. 2006. "Temporal semantics in a supposedly tenseless language." *Linguistics and Philosophy* 29: 673–713.

McCloskey, J. 1997. "Subjecthood and subject positions." In L. Haegeman (ed.). *Elements of Grammar*. Dordrecht: Kluwer, 197–235.

Montague, R. 1970. "English as a formal language." In B. Visentini (ed.). *Linguaggi nella Società e nella Tecnica*. Reprinted in Montague 1974. Milan: Edizioni di Comunità, 189–224.

Montague, R. 1974. *Formal Philosophy: Selected Papers of Richard Montague*. Edited by R. H. Thomason. New Haven: Yale University Press.

Ogihara, T. 1995. "The semantics of tense in embedded clauses." *Linguistic Inquiry* 26: 663–679.

Ogihara, T. 1996. *Tense, Attitudes, and* Scope. Berlin: Springer.

Ogihara, T. 2007. "Tense and aspect in truth-conditional semantics." *Lingua* 1172: 392–418.

Ogihara, T. and Sharvit, Y. 2012. "Embedded tense." In R. I. Binnick (ed.). *Oxford Handbook of Tense and Aspect*. Oxford: Oxford University Press, 638–668.

Partee, B. H. 1973. "Some structural analogies between tenses and pronouns in English." *Journal of Philosophy* 70: 601-609.

Pollock, J.-Y. 1989. "Verb movement, universal grammar and the structure of IP." *Linguistic Inquiry* 20: 365–424.

Portner, P. 2011. "Perfect and progressive." In K. von Heusinger, C. Maienborn and P. Portner (eds.). *Semantics: An International Handbook of Natural Language Meaning*. Volume 2. Berlin: de Gruyter, 1217–1261.

Prior, A. N. 1957. *Time and Modality*. Oxford: Oxford University Press.

Prior, A. N. 1967. *Past, Present, and Future*. Oxford: Oxford University Press.

Ramchand, G. and Svenonius, P. 2014. "Deriving the functional hierarchy." *Langauge Sciences* 46: 152–174.

Reichenbach, H. 1947. *Elements of Symbolic Logic*. New York: Dover.

Richard, M. 1981. "Temporalism and eternalism." *Philosophical Studies* 39: 1–13.

Ritter, E. and Wiltschko, M. 2009. "Varieties of INFL: TENSE, LOCATION, and PERSON." In J. van Craenenbroeck (ed.). *Alternatives to Cartography*. Berlin: de Gruyter, 153–202.

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

Smith, C. S. 1997. *The Parameter of Aspect*. Second edition. Dordrecht: Kluwer.

Stowell, T. A. 1981. *The Origins of Phrase Structure*. Ph.D. dissertation. MIT.

Sudo, Y. 2012. *On the Semantics of Phi Features on Pronouns*. Ph.D. dissertation. MIT.

de Swart, H. 2012. "Verbal aspect." In R. I. Binnick (ed.). *Oxford Handbook of Tense and Aspect*. Oxford: Oxford University Press, 752–780.

Sybesma, R. 2007. "Whether we tense-agree overtly or not." *Linguistic Inquiry* 38: 580–587.

Wiltschko, M. 2014. *The Universal Structure of Categories*. Cambridge: Cambridge University Press.

Zagona, K. 2013. "Tense, aspect, and modality." In M. den Dikken (ed.). *Cambridge Handbook of Generative Syntax*. Cambridge: Cambridge University Press, 746–792.

Zanuttini, R. 2001. "Sentential negation." In M. Baltin and C. Collins (eds.). *Handbook of Contemporary Syntactic Theory*. Malden: Blackwell, 511–535.

Acknowledgement to Referees

The Editorial Board would like to thank our colleagues listed below who served as referees of manuscripts submitted to the *Croatian Journal of Philosophy* during the past two years.

Peter Baumann, *Swarthmore College, USA*
Simon Beck, *University of the Western Cap, South Africa*
Nicola Bertoldi, *Catholic University of Louvain, Belgium*
Krystyna Bielecka, *University of Białystok, Poland*
Sandy Boucher, *University of New England, Australia*
T.Y. Branch, *University of Cologne, Germany*
Zdenka Brzović, *University of Rijeka, Croatia*
Fausto Corvino, *University of Gothenburg, Sweden*
Tamara Crnko, *University of Rijeka, Croatia*
Michael Devitt, *Graduate Center of the City University of New York, USA*
Jason Dockstader, *University College Cork, Ireland*
Dušan Dožudić, *Institute of Philosophy, Zagreb, Croatia*
Callum Duguid, *University of Leeds, UK*
Aleksandar Fatić, *University of Belgrade, Serbia*
Harriet Fagerberg, *Hunter College and The Graduate Center, City University of New York, USA*
José Andrés Forero-Mora, *Corporación Universitaria Minuto de Dios, Colombia*
Anne-Marie Gagné-Julien, *McGill University, Canada*
Adam Gjesdal, *Fordham University, USA*
Ramiro Glauer, *University of Applied Sciences Potsdam, Germany*
Nathalie Gontier, *University of Lisbon, Portugal*
Keith Green, *East Tennessee State University, USA*
Ljudevit Hanžek, *University of Split, Croatia*
Clare Hay, *University of Reading, UK*
Katja Hettich, *The Bauhaus-Universität Weimar, Germany*
Sanja Ivić, *Institute for European Studies, Belgrade, Serbia*
Marek Jakubiec, *Jagiellonian University, Poland*
Tvrtko Jolić, *Institute of Philosophy, Zagreb, Croatia*
Marko Jurjako, *University of Rijeka, Croatia*
Dunja Jutronić, *University of Split, Croatia*
Faik Kurtulmus, *Sabancı University, Turkey*
Gregory Landini, *University of Iowa, USA*

Giorgio Lando, *University of L'Aquila, Italy*
Alexander Linsbichler, *Johannes Kepler University Linz and University of Vienna, Austria*
Goran Lojkić, *Institute of Philosophy, Zagreb, Croatia*
Jakub Mácha, *Masaryk University, Czechia*
lIan MacLean-Evans, *York University, Canada*
P. D. Magnus, *University at Albany and State University of New York, USA*
Filippo Mancini, *University of Bonn, Germany and University of Padua, Italy*
Poppy Mankowitz, *University of Bristol, UK*
Igor Martinjak, *University of Zagreb, Croatia*
Gordan Matas, *University of Split, Croatia*
Nenad Miščević, *University of Maribor, Slovenia*
Kathleen Murphy-Hollies, *University of Birmingham, UK*
Ryan Nefdt, *University of Cape Town, South Africa*
Davor Pećnjak, *Institute of Philosophy, Zagreb, Croatia*
Paloma Pérez-Ilzarbe, *University of Navarra, Spain*
Michal Polák, *University of West Bohemia, Czechia*
Nenad Popovic, *Miami Dade College, USA*
Aleta Quinn, *University of Idaho, USA*
Jiří Raclavský, *Masaryk University, Czechia*
Ivan Restović, *Institute of Philosophy, Zagreb, Croatia*
Alberto L. Siani, *University of Pisa, Italy*
Nenad Smokrović, *University of Rijeka, Croatia*
Stephen Snyder, *Bogazici University, Turkey*
Jussi Suikkanen, *University of Birmingham, UK*
Adam Świeżyński, *Cardinal Stefan Wyszyński University in Warsaw, Poland*
Iris Vidmar Jovanović, *University of Rijeka, Croatia*
Adrian Walsh, *University of New England, Australia*
Brandon Yip, *Australian National University, Australia*
Elia Zardini, *Complutense University of Madrid, Spain*
Paweł Zięba, *Jagiellonian University*, *Poland*

Table of Contents of Vol. XXIII