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.