Indexicals, Context, and Implicit Content

Antony Eagle

antonyeagle.org

Philosophy of Language » Lecture 8

Contents

Indexicals and Context

A problem for compositionality?

- In the last lecture I discussed **compositionality**: the claim that the meaning of a natural language sentence is a function of the meanings of the constituent words and its syntactic structure.
- I examined some challenges to compositionality; here is another. Suppose I utter the words *I am Antony*, and Sylvester utters those words too. We **say the same thing**, in some obvious sense.
- But since what I say is true, and what Sylvester says is false, we must **not** be saying the same thing, in some equally obvious sense. Our utterances must mean different things, if one is true and the other not true.
- So one and the same sentence, with the same syntactic structure, has two different meanings. Farewell compositionality!

The many-but-only-one meanings of *I*

- Well, maybe not. For isn't it obvious that the word *I*, when it occurs in Sylvester's utterance and mine, functions to denote **different** people in the different utterances?
- So the many meanings of *I am Antony* – meanings we witness when that sentence is spoken by different people – are generated by different **extensions** assigned to *I*.
- But this is **not** a case of **lexical ambiguity**, for the word *I* is not ambiguous.
  - Speakers don't need two entries in their mental lexicon to be competent with this expression.
  - More importantly, there does seem to be a common **rule** that governs how to assign an extension to *I*, which suggests that – again, in some sense – there is a common meaning.

Context-sensitivity

- When used, *I* almost always denotes the same thing as the description *the speaker of this sentence*, were it uttered at the same time by the same speaker.
  - *I am the speaker of this sentence* is **trivially true** whenever it is uttered.
- That description comes close to capturing the **rule** speakers and hearers need to be competent with – that, in some sense of 'meaning', is a meaning of *I*.
- Note this **isn't the intension** of *I*, since the description varies in its referent from place to place within a given possibility, while an intension maps each expression to its referent in a possibility – intensions capture variability of referent across possibilities, but not within them.
- This behaviour of *I* is shared with a number of other expressions, including *you*, *this*, *here*, *now*, *tomorrow*,....
  - All of these vary in extension between utterances within a possibility.
- Such expressions are said to be **context-sensitive**.
- Obviously a **context** must be part of a possibility – many contexts exist in the same possible world. But what exactly are they, and how can an expression vary in meaning depending on context?

Context

- A context is something like an **utterance situation**.
  - This includes the concrete **surroundings** of an utterance – the speaker, the audience, when and where, what other things are around and being pointed at, etc.
  - But also the **conversational context**: the content of the foregoing conversation, what has been talked about.
- A context-sensitive expression is one such that its extension on some particular occasion of use depends on some feature of this **context of utterance**.
  - Uses of *I* are sensitive to the speaker of the context; an utterance of *now* is typically sensitive to the time of the context, etc.
- For an interesting class of context-sensitive expressions, we may represent contexts by a collection of parameters, or **indices** (Kaplan 1979: 81), standing for features of the utterance situation.
  - Typical indices include the **agent or speaker** of the context, the **time** of the context, the **position** of the context, and the **world** of the context (Kaplan 1979: 82); some advocate further indices (Kaplan 1979: 97).
- The class of expressions including *I*, *here*, *now* are known as **indexicals**, because of their index-sensitivity.

The Semantics of Indexicals: Context and Circumstance

First Attempt to Connect Context and Truth

- A first attempt to use context to capture the behaviour of the indexical *I* might go something like this:

  I-Indexical *I*, when evaluated relative to an index, has as extension the speaker of that index.

  - So when I utter *I am Antony*, this is to be evaluated related to the actual index (the actual sequence of contextual parameters). Since I am the speaker, the speaker parameter is me, hence the extension of *I* is me; when Sylvester utters it in a different context c' represented by a different sequences of indices, *I* denotes them.
- When is a sentence, possibly containing context-sensitive expressions, **true**?
- The original thought was that contextual variation was rather like modal and temporal variation (Kaplan 1979: 81–82) – that truth relative to an index was like truth relative to a possible world, but there were a lot more indices than possible worlds:

  Contextual Truth† A sentence *S* is **true** relative to index *c* iff *S* says something true when the extensions of indexicals in *S* are fixed by features of *c*.
- So when I actually assert *I am Antony*, I am the speaker of this context, and the sentence then expresses a truth about the actual world, namely, that Antony is Antony. When Sylvester says it, is expresses the falsehood that Sylvester is Antony.

Making and Reporting Context-Sensitive Utterances

- Consider my current true utterance of this sentence:
  (1) I am here.
- Since a is the speaker in c, and p is the place in c, then when we apply Contextual Truth† we get this:
  (2) *I am here* is true relative to c iff the speaker in c is at the place of c at the present time of c.
    - Note the time parameter is implicit in the original sentence in the present tense verb, but is explicitly indexical in the truth conditions (2).
- And since the speaker is actually at the place of utterance at the time of utterance, this predicts that (1) is actually true – which of course it is.
- As Kaplan notes, this makes (1) come out roughly to say the same thing as
  (3) Antony is in his office/Napier 711.
- This looks desirable; these seem to make much the same contribution to a conversation and to say much the same thing.
    - Note how someone might report my utterance of (1): *Antony said he is in his office/Napier 711*.

Triviality and Embedding

- But (3) seems to differ from (1) in some respects – in particular, (1) is completely **trivial**, and (3) is not.
- No one could ever or anywhere say something false by an utterance of (1). But someone could easily say something false by saying (3). How should we capture these facts?
- The most natural idea is to note that (1) is true relative to every utterance situation, which corresponds to what Kaplan calls a 'proper' index (1979: 83) – i.e., an index where the parameter values derive from some real possibility as to whether a given speaker might be located at a given place and time.
  - E.g., this is a proper index $\langle w_@$, Antony, Princeton, September 11, 2001$\rangle$, since I was actually there then. But $\langle w_@$, Antony, Melbourne, September 11 2001$\rangle$ is improper, as is $\langle w_@$, Antony, Princeton, January 1 1066$\rangle$, etc.
- *I am here* is true at every proper index, unlike *Antony is in Napier 711*, which is false at the proper index $\langle w_@$, Antony, Princeton, September 11 2001$\rangle$.

Necessity and Triviality

- But now we've rendered ourselves unable to accommodate the contingency of (1). Consider
  - (4) Necessarily, I am here.
- Applying Contextual Truth†, assuming *necessarily* quantifies over indices:
  - (5) *Necessarily, I am here* is true relative to c iff, for every proper index c': *I am here* is true relative to c'.
- Since (1) **is** true at every proper index, (5) predicts that (4) will be true. But it's not. While I am in fact here (in this place), I certainly didn't need to be (Kaplan 1979: 83).
  - Compare: *Antony is speaking* which I cannot **falsely utter** and yet it can be **false** – it's false whenever I'm not speaking, and though I can make it true just by saying it, it's not necessary.
- While every proper index is such that its speaker is at its place, and so make *I am here* true, that shouldn't be enough to make it **necessarily** true.

What's gone wrong?

- The standard diagnosis is this: we have used indices in **two different roles**, when we should have distinguished them.
  1. We used the context of utterance to generate proper indices which **fix the contents of indexical expressions**; and
  2. We used proper indices to proxy for possible worlds, that is, to serve as points at which to **evaluate the truth of the sentence** after its contents have been fixed.
- The problem is that if you use one and the same entity to do both things, then anything which **cannot be falsely uttered** becomes a **necessary truth** – and yet these are not equivalent (though in many cases they overlap).
  - Since there is no context in which the speaker of the context isn't at the place of the context, there is no single proper index that can both fix the referents of indexicals *I* and *here* and falsify the claim *I am here* under the same interpretation.

> there are difficulties in the attempt to assimilate the role of a context in a logic of demonstratives to that of a possible world in the familiar modal logics…. (Kaplan 1979: 83)

Double Indexing (Speaks 2024: §2.1.4)

- Kaplan's solution is that **contexts of utterance** should determine the indices with respect to which indexical expressions have contents, and possible worlds should serve as **circumstances of evaluation**, with respect to which contents have referents/truth values.
  - There is accordingly a two-step process to assign a truth value to a sentence; fix a content relative to one index (the context), and evaluate that content for truth with respect to a potentially different index (the world).
- So this gives us a revised truth rule:

  Contextual Truth A sentence S is **true** in w relative to c iff, when the contents of context-sensitive expressions in S are fixed by features of c, the proposition then expressed by S is true at w.
- And here the world of the utterance context can **diverge** from the world which is fixed by the circumstance of evaluation.

Modal Truth in the Double Indexing Framework

- Much of the time, context and circumstance coincide.
  - If there are no indexical elements to a sentence, then the content is invariant across all indices; and even if there are indexical elements, we typically want to assess the truth of what is said in a context with respect to the circumstances surrounding that context.
- But what our example of the contingency of *I am here* shows is that **modal operators** like *necessarily* can **shift the circumstance of evaluation while keeping the context of utterance fixed**.
  - With *Necessarily, I am here*, we use the context of utterance to fix *I* as denoting Antony, and *here* as denoting Napier 711. This gives us what is said (the proposition expressed) by **this use** of *I* and *here*.
  - Then we treat *necessarily* as a quantifier over circumstances of evaluation, such that *necessarily* ϕ is true (relative to c) iff what ϕ says as used in c is true at every possible circumstance.
  - And since *Antony is in Napier 711* is not true at every possible world, *Necessarily, I am here* is false.
    - It is false at just those possibilities where I cannot express that Antony is in Napier 711 by an utterance of *I am here* – though whatever I **do** express by an utterance of that claim will still be true.

Two Kinds of Meaning: Content and Character

Content and Character

> The content of an expression is always taken *with respect to* a given context of use. Thus when I say 'I was insulted yesterday' a specific content – *what I said* – is expressed. Your utterance of the same sentence, or mine on another day, would not express the same content. … This content … has been often referred to as a 'proposition'. So my theory is that different contexts … produce not just different truth values, but different propositions.
>
> … I call that component of the sense of an expression which determines how the content is determined by the context, the 'character'…. Just as content … can be represented by functions from possible worlds to extensions, so characters can be represented by functions from contexts to contents. The character of *I* would then be represented by *the function … which assigns to each context the content which is represented by the constant function from possible worlds to the agent of the context*. (Kaplan 1979: 83–84)

Illustrating the Standard Picture

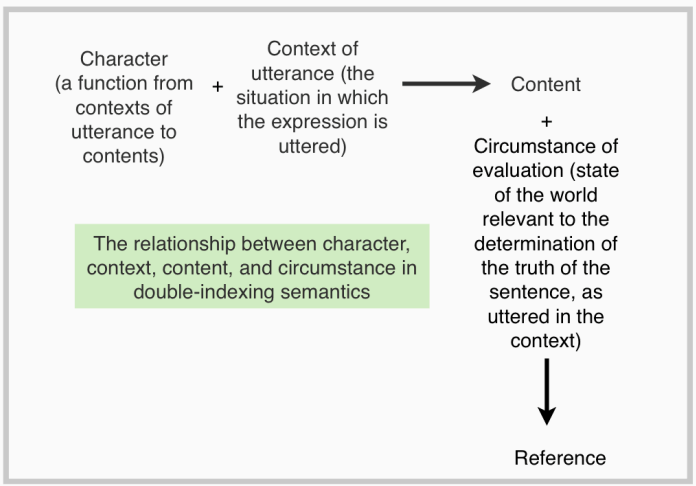

Figure 1: The Relationship of Character, Context, Content, and Circumstance (Speaks 2024: §2.1.4).

The standard picture

- The picture Kaplan gives us is this (see fig. 1). Consider *I am speaking*, relative to c. We look at the **character** of the expressions:
  1. The character of *I* is the function which assigns to each context c the intension $[\![p]\!]$, where p rigidly denotes the speaker in c. So it assigns to the present context the intension of *Antony*, though it assigns the intension of *Sylvester* to *I* in other contexts.
  2. The character of the non-context sensitive expression *is speaking* is the function which assigns to each context the intension of *speaking*.
- Given a **context**, character determines **content**: the intensions of *Antony* and *is speaking*. Combining these, the content of the sentence is this intension: $\{w : \text{Antony speaks in } w\}$.
- We evaluate this at the **circumstance of evaluation**, by checking whether the circumstance world is in the intension — obviously, this one will come out true, but only contingently so.

Context-Sensitivity and Rigidity

- The character Kaplan offers for *I* treats it as rigid: if it has a content in a context, it has that content in every possible world.
  - So *I might not have been speaking* is true, because *I am not speaking*, relative to my present context, is true in any possible world in which Antony does not speak.
- That means it has a different intension than ⟦*the speaker at such-and-such a place on such-and-such date*⟧, which is a context-insensitive but non-rigid designator.
  - That describes different things in different possibilities, but it has no indexical elements – all the indexical parameters are explicitly spelled out.
- Other possibilities are available too; see tbl. 1.

Table 1: Combinations of rigidity and context-sensitivity.

	Rigid	Non-Rigid
Context-Sensitive	Indexicals: *I*, *here*	*The man next to me now*
Invariant	Names: *Antony*, *Gödel*	*The man next to Antony on April 7, 1993*

Analyticity

> What is it that a competent speaker of English knows about the word 'I'? … It is the character of 'I' …. Thus that component of sense which I call 'character' is best identified with what might naturally be called 'meaning'.…
>
> (c) In all contexts, an utterance of ['I am here'] expresses a true proposition [i.e., when evaluated at the context itself].
>
> On the basis of (c), we might claim that ['I am here'] is analytic (i.e., it is true solely in virtue of its meaning). Although ['I am here'] rarely or never expresses a necessary proposition. This separation of analyticity from necessity is made possible – even, I hope, plausible – by distinguishing the kinds of entities of which 'is analytic' and 'is necessary' are properly predicated: characters (meanings) are analytic, contents (propositions) are necessary. (Kaplan 1979: 84–85)

A priority and Necessity: Stalnaker's proposal

- Kaplan's proposal makes some analytic sentences contingent, which seems wrong.
  - Recall Kripke's suggestion that 'analytic' should be reserved for sentences which are both *a priori* and necessary (Kripke 1980: 39).
- Explicitly inspired by Kripkean ideas, Stalnaker suggests that sentences like *I am here*, which are true whenever they are uttered, are examples of the **contingent *a priori***:

> An a priori truth is a statement that, while perhaps not expressing a necessary proposition, expresses a truth in every context. (Stalnaker 1978: 83)

  - Such sentences are counterpart to the cases Kripke discusses of the **necessary *a posteriori***, cases like *water is* H_2O.
- A famous example from the history of philosophy is from Descartes:

> I must finally conclude that this proposition, *I am*, *I exist*, is necessarily true whenever it is put forward by me or conceived in my mind. (Descartes 1641, AT 25)

  - Here he seems to conclude that *I exist* cannot be falsely uttered; this makes it *a priori*, though clearly contingent.

Worlds, Propositions and Circumstances

- Our working hypothesis identifies the propositional content of a sentence with a **function from possible worlds to truth values**.
- Given this, our circumstances of evaluation only need to include a possible world, since that is the only thing that a proposition's truth value **varies** with respect to.
- But some have argued that the circumstance needs other indices too:

> The most popular candidate for a second index is a time. The view that propositions can have different truth-values with respect to different times – and hence that we need a time index – is often called 'temporalism'. The negation of temporalism is eternalism. (Speaks 2024: §2.3.2)

- Kaplan himself includes time amongst the indices of the context to fix the referents of temporal indexicals like *now*, but also thinks that contents can vary in truth value over time, and so is a temporalist (Kaplan 1979: 91–92).
- By contrast, Richard (1981) argues for eternalism; Zimmerman (2005) offers a comprehensive overview.

Where is Context-Sensitivity to be Found?

Tests for Context-Sensitivity

- So much for the framework; we now have the capacity to model words which are sensitive to features of the context.
- But how can we **tell** whether a word is context-sensitive, so that we might add an index for it in our models of contexts?
- One standard test is the **indirect report** test:

  IDR If reporters can easily and truly indirectly disquotationally report an utterance of a sentence *S* by an agent *A*, i.e., with '*A* said that *S*', despite indifference about, or ignorance of, its original context of utterance, then it is unlikely *S* is context-sensitive. (Donaldson and Lepore 2012: 121)

- This test obviously counts *I*, *now*, and the easy cases as context-sensitive – **you** cannot report my utterance of *I am Antony* by saying *Antony said that I am Antony*.

Unreliability of IDR test: *On the left* (Hawthorne 2006)

> Terms for relative directions, like 'left', seem to be almost as obviously context-sensitive as 'I'; the direction picked out by simple uses of 'left' depends on the orientation of the speaker of the context. But we can typically use 'left' in disquotational 'says' reports of the relevant sort. Suppose, for example, that Mary says
>
> > The coffee machine is to the left.
>
> Sam can later truly report Mary's speech by saying
>
> > Mary said that the coffee machine was to the left.
>
> despite the fact that Sam's orientation in the context of the ascription differs from Mary's orientation in the context of the reported utterance. Hence our test seems to lead to the absurd result that 'left' is not context-sensitive. (Speaks 2024: §2.3.1)

Another test: Agreement

> If A and B both utter S and can be reported as agreeing, say, with 'A and B agree that S', then that is evidence S is semantically invariant across its distinct utterances. If, on the contrary, distinct utterances cannot be so reported, this is evidence S is not semantically invariant across its distinct context of utterance. (Cappelen and Hawthorne 2009: 54–55)

- A variant agreement test:

  A2 If A utters α *is* F and B utters β *is* F, and can be reported as agreeing that α **and** β **have something in common** by someone indifferent to the original contexts of utterance, that is evidence that F is semantically invariant across distinct utterance contexts.

Some more possible examples

Table 2: Verdicts of tests for context-sensitivity; '+' means 'is judged context-sensitive by the test'

Expression	IDR	Agreement	A2
actually	+	+	+
flat	?/+	?	+
tasty	?	-	?/+

- The judgements for *actually* are most easily elicited when considering **fictional discourse**; e.g., if
- For *flat*, think about this case for the IDR test: we are discussing whether the ground is flat enough so our drinks don't spill; it may not be appropriate to say *Mary truly said that South Australia was flat, and this park is in South Australia, so it must be flat too.*
- *Tasty* is trickier – lots of people think it has something to do with the tastes of the speaker, but there also seems to be the possibility of genuine (not merely verbal) (dis)agreement.
  - Lots of controversy about this example (Lasersohn 2005-12)!

Is *Knows* an Indexical?

- One interesting thesis in contemporary epistemology is the thesis that *knows* is context-sensitive (DeRose 1995), and that it is sensitive to the **standards** in play in the context.
  - These standards vary with the alternative possibilities to actuality we need to consider.
  - So consider the actual fact that I have hands. If we are considering the BIV alternative, we need lots of evidence to rule it out, and if we are considering just the ordinary alternative (I don't have hands but I just didn't notice) that doesn't need lots of evidence to rule it out.
- The central question then is: **is *knows* context-sensitive** (Stanley 2004)?
- Both of our tests predict **no**.
  - It seems we can take an utterance by *A* in a high standards context of *Jane knows that she has hands*, and in a low standards context by *B* of the same sentence, and report them as agreeing.
  - And I can report *A*'s claim accurately using *A*'s words.

Pronouns

Pronouns and Context-Sensitivity

- Another interesting class of context-sensitive expressions are **pronouns**: *you*, *she*, *they*, *it*,....
  - These certainly get their referent from the context; and they certainly are predicted to be context-sensitive by our tests. (Jane says *she was happy*, unambiguously about someone other than Jane; I cannot report this unambiguously by *Jane said that she was happy*.)
- But the way that pronouns get their reference fixed is not always by picking up on features of the real-world context in which they are uttered (Evans 1980).
- Sometimes they pick up on **what was already said in the conversation**:

  (6) Bob walked over to the couch. He sat down.

  (7) Joe walked over to the couch. He sat down.

  - Here, the referent of *he* is parasitic upon the previously introduced name; we say that the pronoun depends **anaphorically** on the foregoing content, specifically the name.
- Sometimes the pronoun picks up on more general features of the context:

  (8) Hey you!

  - Here, the referent of *you* is given by whoever is physically in the surroundings of the utterance; this is sometimes known as **deixis**.

Pronouns and Binding

- **Bound variable pronouns** are governed by quantifiers:
  - (9) Every boy thinks he deserves ice cream. (Elbourne 2011: 115)
  - (10) Only Sally loves her mother.
- (9) can mean: Every boy x is such that x believes x deserves ice cream.
- (10) can mean: Only Sally is an x such that x loves x's mother, **or** only Sally is an x such that x loves Sally's mother.
  - Note these can have non-bound readings too: *Maxwell has been so well behaved today that every boy thinks he deserves ice cream*.
- A general treatment: the meaning of a pronoun in context is a temporary **assignment** of an extension (Portner 2005: 103–8).
  - In (9) and (10), the assignment is determined by the quantifier phrases *every boy* and *Only Sally* – in particular, *Every boy (… he …)* is true iff *(… he …)* is true under **every assignment** of a boy to be the extension of *he*.
  - When a pronoun doesn't occur in the scope of some assignment-determining expression like a quantifier? Then **context** determines the assignment, as in (6)/(7)/(8).

Donkey Anaphora

(11) Each man who owns a donkey beats it.

> Here, again, in the most natural reading, *it* cannot be referential. A speaker of this sentence would not be referring to some particular donkey, Flossy, and saying that every man who owns a donkey beats Flossy. There is no consensus, however, on what the pronoun here does mean. (Elbourne 2011: 116)

- The problem: in (11), it looks like the logical form should be something like this:
  (12) For every man, if there exists a donkey that he owns, then he beats it:
  $\forall x(Mx \rightarrow (\exists y(Dy \wedge Oxy) \rightarrow Bxy))$.
- The problem is evident in the logical form: the last occurrence of y is unbound. What we want is for the quantifier *a donkey* to somehow bind that final *it*, but it would seem to be unable to 'reach out' of its clause to bind the pronoun.
- The same phenomena is seen in explicitly conditional donkey anaphora (13); the inability of a quantifier to bind pronouns out of scope is illustrated in the otherwise parallel (14):
  (13) If Sarah owns a donkey, she beats it.
  (14) *If Sarah owns every donkey, she beats it. (King and Lewis 2021: §2)

Compositionality and Donkey Anaphora

- There is no issue with understanding donkey anaphora, or seeing what the logical form should be. The issue is understanding how to put the apparent meanings of the parts **compositionally** together into generate that meaning/logical form. The following illustrates this problem.
- It is evident that (13) requires Sarah to beat every donkey she owns; if she only beat one of her herd, it would be false.
- So maybe the pronoun *it* is bound by *a donkey*, but the quantifier actually has **universal** force.
  - So (11) means *Every man x is such that every donkey y is such that if x owns y, x beats y*, and (13) means *Every donkey is such that: if Sarah owns it, she beats it*.
- But in similar occurrences, *a donkey* clearly is existential.
  - E.g., in *Each man owns a donkey*, *a donkey* is an existential.
  - Consider even *If a man owns a donkey, he beats it*.
- So can we assign a consistent and coherent meaning to *a donkey* that yields the right meaning for (11) compositionally?

E-type pronouns

(15) Amy owns a donkey. She cares for it.

- In (15), the pronoun *it* is anaphoric on the material in the preceding sentence; it means something **descriptive** like: *the donkey Amy owns*. This is a so-called **e-type pronoun** (Evans 1980).
- Parasitic on this usage, perhaps *it* in (11)/(13) likewise has the force of a **description** (Heim and Kratzer 1998: 295; King and Lewis 2021: §3.3).
  - Some say that e-type pronouns in donkey anaphora are not definite descriptions, but rather something 'numberless' and with universal force, like *any donkey he/she owns* (Neale 1993).
  - Others say that *it* in (11) means something like *the donkey*, but that the referent of the definite description varies along with the individuals quantified over, so (11) would say something like *Each man$_x$: if s is a minimal situation in which x owns a donkey, in s, x beats the donkey* (Elbourne 2005).
- A problem for the theories sketched above comes from cases like this:

  (16) If a bishop meets a bishop, he blesses him. (King and Lewis 2021: §3.3)

  This is a problem because, whether *him* has definite or numberless force, there are two bishops here, and incorrect truth conditions will apparently be predicted.

Implicit Content and Unarticulated Constituents

Context and Content Again

(17) It's raining.

- An utterance of (17) in a context *c* expresses (more or less) the proposition *At the time and place of c, the weather is rainy*.
- This is context-sensitive; the **agreement test**, in particular, demonstrates it, since when Alice in Palo Alto says *It's raining*, and Bob in Oxford says at the same time *It's not raining*, they are not disagreeing.
- Can we explain this, using Kaplan's content/character distinction? The only way would be something like this:

> The word 'rain', when used in a context *c*, refers to the property of raining at the location *contextually salient at c*. (Donaldson and Lepore 2012: 124)

Supplanting the Contextual Location

- A problem arises with this sentence:

  (18) It's raining in Oxford.

- Suppose (18) were uttered in Adelaide. Go by the quasi-Kaplanian context-sensitive truth conditions we just offered, plus compositionality, and we get this:

  (19) It's raining at the location contextually salient in Adelaide in Oxford.

- (19) is borderline **ungrammatical**; and even if it is grammatical, what it seems to express is **it is raining in Adelaide**, not that it is raining in Oxford.

Unarticulated constituents

> In order to assign a truth-value to [*It's raining*], as I just did, I needed a place. But no component of [this] statement stood for a place. The verb 'raining' supplied the relation *rains*(t, p) – a dyadic relation between times and places, as we have just noted. The tensed auxiliary 'is' supplies a time, the time at which the statement was made. 'It' doesn't supply anything, but is just syntactic filler. ([Perry's footnote:] Note that if we took 'It' to be something like an indexical that stood for the location of the speaker, we would expect 'It is raining here' to be redundant and 'It is raining in Cincinnati but not here' to be inconsistent.) So Palo Alto is a constituent of the content of my son's remark, which no component of his statement designated; it is an *unarticulated* constituent. Where did it come from? (Perry 1986: 138)

Related examples: quantifier domain restriction

(20) Everyone is smoking.

- (20) can be truly uttered in a context where **everyone around the speaker is smoking**, even while it is false that **absolutely everyone** everywhere is smoking.
- So we need context to supply a **restriction** on the domain of quantification, so that an utterance of (20) is only about a certain restricted domain:

  Everyone The word *everyone* when used in a context c quantifies over the domain of **people in** c.
- One problem is that the speaker of (20) **need not be smoking** for them to utter (20) truly. Compare:

  (21) Everyone over there is smoking.

Three Approaches to Unarticulated Constituents

Approach 1: Propositional Radicals

- The puzzle is this: (17) (*It's raining*) seems to be location-sensitive, but no expression in (17) is location-sensitive, given that (18) which embeds it is not.
  - Similar observations apply to (20) and (21).
- The natural thought is that (17) might be context-sensitive in some new way. Namely, what (17) expresses depends on its surrounding **linguistic context**.
- Maybe: what (17) and (20) express **aren't propositions**, but **propositional radicals** (Bach 1994); like propositions, but with a 'gap':

  (22) It's raining at *t* in _____.
  (23) Everyone in _____ is smoking.
- What's crucial is that the gap can be filled in either by context, as with (17) and (20), **or** by further explicit content as in (18) and (21).

How to Fill a Gappy Proposition

- Note that this gap in the proposition **is** filled, but how it is to be filled is **not** triggered by any overt word in the sentence; no word has a **character** that specifies how context generates a content here.
- Kaplan's picture is overturned; this proposal claims that a sentence in a context determines a propositional radical; and then a **propositional radical in a context** determines what is said. Context plays two roles:
  1. It generates a content, as Kaplan claims;
  2. It combines with a content to express **another content**, more like a traditional proposition. (Which latter content determines a truth value, given a circumstance of evaluation.)
- But how does this second step work? Since the sentence and context have already contributed all they have to offer when the propositional radical is determined, what principles govern how a propositional radical expresses a specific proposition in a context?

A Global Solution?

- The most obvious solution is this: **when someone expresses a propositional radical in c, the 'gaps' are all filled by** c. (Elbourne 2011: 122–24)
- But there is a problem. It's easiest to see in the case of quantifier domain restriction.
  - (24) Everyone is asleep and being monitored by a research assistant. (Soames 1986: 357)
  - (25) Every sailor waved to every sailor. (Stanley and Williamson 1995; Stanley and Szabó 2000: 249)
  - (26) Exactly two people irritated everyone.
- Suppose Alice and Bob are gatecrashers, who arrived uninvited and behaved obnoxiously to all the invited guests. Intuitively, there is a reading of (26) which is true even if Alice and Bob did not irritate themselves (or each other). A framework in which context supplies a single domain cannot explain the truth of (26) in this case:

> contextual supplementation works at the level of constituents of sentences or utterances, rather than the level of the sentences or utterances themselves. (Soames 1986: 357)

Gaps Filled New Ways

- A sentence like (26) seems to have two gaps, on the propositional radical view, made explicit as follows:

  (27) Exactly two people in _____ irritated everyone in _____.

- We can't have both gaps filled by the global context, since that gives the wrong results for (24)–(26).
- So **how** do those gaps get filled? We need some mechanism, otherwise the theory will be consistent with all sorts of bizarre gap-fillers – e.g., when the first gap is filled by people with red hair, and the second by people with blue eyes – which aren't intended or plausible readings of the sentence at all.
  - **In the absence of character, how is the specific rule mapping context to gap-filling content implemented?**
- It seems like we need some theory of how this works without postulating novel and unattested context-sensitive mechanisms.

Approach 2: Covert Variables (Stanley 2005)

- To deal with (24)–(26), we need context to supply domains to **constituents** of the sentence:

  (28) Exactly two people [in domain d_1] irritated everyone [in domain d_2].

- As long as context supplies to d_1 the set of all party-goers, and to d_2 the set of just the invited party-goers, this semantics will get the intended reading of the utterance of (26).
- This is the **covert variable** approach: it says that in the syntax of *people*/*one* is an unpronounced domain variable that behaves in context somewhat like *in this/that place*.
  - It behaves therefore like a **locative pronoun** would, but English doesn't have them – which is maybe *why* it is covert since there is no overt equivalent.

Stanley and Szabó on Covert Indexicals

> A sentence [like *every bottle is on the shelf*] can communicate a proposition concerning a restricted domain of bottles, because, relative to certain contexts, it expresses such a proposition. It expresses such a proposition relative to certain contexts because common nouns such as 'bottle' always occur with a domain index. It follows that, in the logical form of quantified sentences, there are variables whose values, relative to a context, are (often restricted) quantifier domains. (Stanley and Szabó 2000: 258)

Applying covert variables to *It's raining* I

- We can apply this same sort of approach to *It's raining*.
- The approach would say that the real form of *It's raining* has a covert free variable, as depicted in the tree to the right.
- There is a covert variable expression *in l* associated with *rains*, which behaves like *here*, because the location variable is contextually assigned.

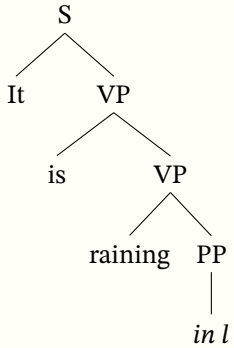

Applying covert variables to *It's raining* II

- In (7), by contrast, the covert variable is saturated by the **overt** locative expression *in Oxford*, as depicted in the tree to the right.
- The value assigned to the variable/pronoun *l* is anaphoric on the explicitly mentioned place, just like in the assignment to the pronoun *him* in *Regarding Bob: Mary likes him* is anaphoric on the explicitly mentioned man.

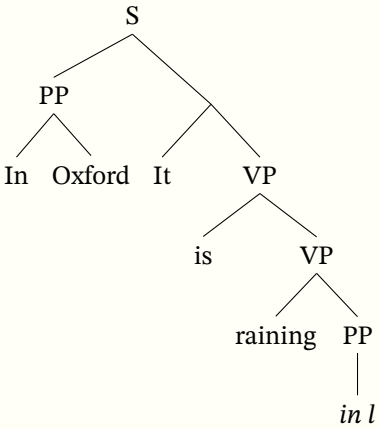

The Binding Argument

- Stanley points out that there is a reason to posit covert pronouns/variables in the syntax: that those variables can be **bound** by quantifier phrases (Stanley 2005: 240).
  - (29) Every time John lights a cigarette, it rains.
  - (30) Always, every bottle is on the shelf.
- If (29) is uttered in *c*, it expresses the proposition that on every occasion when John lights a cigarette, it rains at **the occasion of the lighting**.
  - Namely: *Every e of John lighting a cigarette is such that it rains at the place of e*.
  - Here we assume that the time parameter is explicitly set by the tense, so the place variable is what is bound.
- So, Stanley says, *It rains* has the syntax *it rains at p*; the covert pronoun *p* is assigned a content by context when it is not overly bound as in (29).

Approach 3: Pragmatic Enrichment

- The third way denies the presence of covert variable elements in the syntax. The meaning of *it's raining* is just the generic proposition that IT'S RAINING [SOMEWHERE].
- But what an utterance of (17) communicates is not its meaning, but a more specific 'proposition that the speaker wished to express' (Elbourne 2011: 127).
- How do we get to that proposition? By **pragmatic enrichment**: by applying general principles – psycho-social background beliefs about intention, communication, and relevance – to the utterance to figure out what it means.
- This is not rule-based – it violates compositionality, since the **real meaning** of (17) is the output after pragmatic enrichment, not the impoverished thing spat out by the compositional semantics.
- Hard to evaluate this until we know **what** might be pragmatically assigned as the meaning.

References

References

Bach, Kent (1994) 'Conversational Impliciture', *Mind & Language* **9**: 124–62. doi:10.1111/j.1468-0017.1994.tb00220.x.

Cappelen, Herman and John Hawthorne (2009) *Relativism and Monadic Truth*. Oxford University Press.

DeRose, Keith (1995) 'Solving the Skeptical Problem', *Philosophical Review* **104**: 1–52. doi:10.2307/2186011.

Descartes, René (1641/1996) *Meditations on First Philosophy: With Selections from the Objections and Replies*, John Cottingham, ed., trans. Cambridge University Press.

Donaldson, Tom and Ernie Lepore (2012) 'Context-Sensitivity', in Gillian Russell and Delia Graff Fara, eds., *The Routledge Companion to the Philosophy of Language*: 116–31. Routledge.

Elbourne, Paul (2005) *Situations and Individuals*. MIT Press.

Elbourne, Paul (2011) *Meaning: A Slim Guide to Semantics*. Oxford University Press.

References (cont.)

Evans, Gareth (1980) 'Pronouns', *Linguistic Inquiry* **11**: 337–62.

Hawthorne, John (2006) 'Causal Structuralism', in *Metaphysical Essays*: 211–27. Oxford University Press.

Heim, Irene and Angelika Kratzer (1998) *Semantics in Generative Grammar*. Blackwell.

Kaplan, David (1979) 'On the Logic of Demonstratives', *Journal of Philosophical Logic* **8**. doi:10.1007/bf00258420.

King, Jeffrey C and Karen S Lewis (2021) 'Anaphora', in Edward N Zalta, ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2021/entries/anaphora/.

Kripke, Saul (1980) *Naming and Necessity*. Harvard University Press.

Lasersohn, Peter (2005-12) 'Context Dependence, Disagreement, and Predicates of Personal Taste*', *Linguistics and Philosophy* **28**: 643–86. doi:10.1007/s10988-005-0596-x.

References (cont.)

Neale, Stephen (1993) *Descriptions*, A Bradford Book. MIT Press.

Perry, John (1986) 'Thought Without Representation', *Proceedings of the Aristotelian Society, Supplementary Volumes* **60**: 137–51. doi:10.1093/aristoteliansupp/60.1.137.

Portner, Paul H (2005) *What Is Meaning? Fundamentals of Formal Semantics*. Wiley-Blackwell.

Richard, Mark (1981) 'Temporalism and Eternalism', *Philosophical Studies* **39**: 1–13. doi:10.1007/bf00354808.

Soames, Scott (1986) 'Incomplete Definite Descriptions', *Notre Dame Journal of Formal Logic* **27**: 349–75. doi:10.1305/ndjfl/1093636680.

Speaks, Jeff (2024) 'Theories of Meaning', in Edward N Zalta and Uri Nodelman, eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2024/entries/meaning/.

Stalnaker, Robert C (1978/1998) 'Assertion', in *Context and Content*: 78–95. Oxford University Press.

References (cont.)

Stanley, Jason (2004) 'On the Linguistic Basis for Contextualism', *Philosophical Studies* **119**: 119–46. doi:10.1023/b:phil.0000029353.14987.34.

Stanley, Jason (2005) 'Semantics in Context', in Gerhard Preyer and Georg Peter, eds., *Contextualism in Philosophy*: 221–53. Oxford University Press.

Stanley, Jason and Zoltán Gendler Szabó (2000) 'On Quantifier Domain Restriction', *Mind and Language* **15**: 219–61. doi:10.1111/1468-0017.00130.

Stanley, Jason and Timothy Williamson (1995) 'Quantifiers and Context-Dependence', *Analysis* **55**: 291–95. doi:10.1093/analys/55.4.291.

Zimmerman, Dean W (2005) 'The A-theory of Time, the B-theory of Time, and "Taking Tense Seriously"', *Dialectica* **59**: 401–57. doi:10.1111/j.1746-8361.2005.01041.x.