Language and context

Book or Report Section

Accepted Version

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.

www.reading.ac.uk/centaur

Emma Borg: Language and Context

David Edmonds: *Imagine that I'd said to Nigel after an interview, 'that was the best Philosophy Bites interview we've ever conducted!' Perhaps I was being ironic or sarcastic, but perhaps I meant it literally. If you had more information, you might be able to work out my meaning—say, if you'd heard my tone of voice, or if you'd been sitting in on the interview too. So what is the relationship between meaning and context? We spoke to Emma Borg—to put her into context, she's an expert on the philosophy of language.*

Nigel Warburton: *The topic we're going to focus on is language and context. Now, what's context got to do with language?*

Emma Borg: Let's first think about a question that comes slightly before that—let's think about what it would be to understand a language. How is it that when you learn a language, like French or German, you understand the sentences or the expressions that you are being presented with? One possibility would be that what you've learnt when you're learning a language is something that pairs up the sentences of that language with the meaning in that language. So when someone presents you with a bit of this new language, what you do is you go to your big look-up list and you find the sentence in question and you pair it with a meaning—and that's it, you've got the meaning! That might be one way in which language learning could go. But it seems that after a moment's reflection, that can't be quite right. And it can't be quite right because language looks like it's got this quite surprising feature, namely that we can generate lots and lots—maybe even an indefinite or infinite number—of sentences, within just one natural language. So when you were learning French, you didn't just

learn a few sentences—rather, you learned a potentially infinite number of sentences. You learned how to say 'snow is white', but you also learnt how to say 'snow is white and grass is green', or 'snow is white and grass is green and two plus two is not five'—and we could just go on doing that, we could go on generating brand-new sentences. And that seems to show us that you haven't just got a big look-up list, because one thing we know about the mind (or about the brain) is that it's not infinite—you can't have an infinite list of sentences in that way.

NW: *Well, in my case of learning French, I was actually mapping the French onto English, so it was parasitic on my knowledge of English.*

EB: You might think in that case it was parasitic, but let's think about the first language that you learnt. Say you were learning English as your first language—it can't be the case that what you learnt was an infinite set of sentences, because, smart as we all are, we still haven't got infinite minds in that way. And so then we face the question as to what you know and how you understand your language. A suggestion that's been made (and it seems quite appealing) is that what you know are the meanings of some finite number of things (like the meanings of certain words) and you know some strategies or some rules for putting those words together in meaningful ways. So maybe that's what understanding a language consists in—just knowing word meaning and structures.

NW: *By structures, do you mean something like grammar?*

EB: Yeah, something like grammar, something like the structure of a sentence for how bits relate to one another.

NW: *Okay. So you've got this grammar, you've got this vocabulary—where does context come in?*

EB: Well, now we can begin to think about the worry of context. Because that picture I just described doesn't seem to have a place for context. But again, as soon as we start thinking about our language, we can see that there has to be a role for context. For example, when I say, 'I'm happy', I say something different to what you say when you say, 'I'm happy'. In the one case, I say, 'Emma's happy'; in the other case, you say, 'Nigel is happy'. So what's going on in that kind of contextual shift? Or again, if I say, 'Nigel is sitting', it looks like what I'm saying is 'Nigel is sitting at a particular time'. So it looks like in order to get at meaning we've got to have some kind of appeal to the context of utterance—we've got to know things like who's speaking, where they're speaking, when they're speaking, etc. As such, something in our knowledge of language has to capture that context sensitivity.

NW: *Presumably, you have to know even more than that—you have to know if the speaker is playing a game, or if they're joking with a friend, or if they're deadly serious, or being ironic, and so on—there must be quite a lot of contextual things that feed into our interpretation of what a word or a sentence means.*

EB: That's a good point. Someone might think that context sensitivity is quite a limited feature of our language. They might suppose that we need to think about context *only* when we're dealing with words like 'I' and 'this' and 'that' and 'today' and 'tomorrow', and tense markers. And some very smart philosophers, people like David Kaplan and Robert Stalnaker and John Perry, gave us a way of dealing with those kinds of expressions. They said that all you need to know when you're thinking about 'I' or 'this' or 'that' is just a very simple contextual rule which will take you from that expression to the thing it's picking out in the context. So, for 'I', what you need to know is that an utterance of 'I' always refers to the speaker. And then you have to go to context to find out who the speaker is, but it's quite a rule-based kind of

activity. But your question just now was picking up on lots of other ways in which context gets in, and that's precisely the way that the debate has gone in recent years. Philosophers, in particular Charles Travis, have raised a number of examples which seem to show that context sensitivity is pervasive in language.

NW: *Before we get onto Travis, I was just thinking about a phrase as simple as 'cat's eyes'—if you're talking about that in relation to your pet, that means one thing, and if you're talking about it as you're driving down the road, its meaning is completely different.*

EB: However, we might genuinely think that 'cat's eyes' has two different dictionary entries—the cat's eyes you find on the road may well be actually recorded in a different way in a good dictionary. So we might think that that expression is genuinely ambiguous. But there are other cases where it's clear that we do get a shift in meaning to do with the expressions that we're putting together. For example, there's a case that gets discussed a lot in the literature about the difference between 'baking a potato' and 'baking a cake'. You do something different when you're baking a potato (in that you're actually cooking the thing), whereas baking a cake means creating it (in that there was no cake before you baked it). But you might think that's not such a weird kind of context sensitivity because it's emerging from the expressions and the way they're put together—it's because in one case we've got a potato and in the other case we've got a cake, and they're different kinds of things, and as such could give rise to a difference of meaning in 'baking' in each case.

NW: *You mentioned Charles Travis's examples that are supposed to demonstrate that context is incredibly important for understanding language. What are they?*

EB: He has this one very famous case, which is probably a good place to start. Imagine this girl called Pia, and Pia has a Japanese maple tree, and as we all know, Japanese maple trees have red leaves. But one morning, Pia looks out of her window and sees her tree and she comes to think that red is the wrong colour for leaves. She thinks leaves should really be green. So she goes out and she meticulously paints every single leaf on her Japanese maple tree in green. And then Travis asks us to consider two different scenarios that Pia might be involved in. In the first scenario, a friend of hers who's an artist rings up and says, 'I'm about to paint a picture, but I need something green to balance out the composition. Have you got anything green I could borrow?'. And Pia says, looking at her painted leaves, 'yes, you can have these, these leaves are green'. And Travis says, in that context, our intuitions are that Pia said something true—speaking to the artist about the painted leaves, it's true to say that they're green. In the second scenario, a botanist friend of Pia's calls up, and she says, 'I'm about to run an experiment on photosynthesis. Have you got any green leaves that I can have for my experiment?'. And Pia, looking at her painted leaves, says, 'yes, you can have these, these leaves are green'. Travis says our intuitions in this second case are that Pia has no longer said something true. Speaking of the painted leaves to the botanist, it's no longer correct to say of these leaves that they're green. Yet we've got a single sentence—'these leaves are green'—where the circumstances don't change and we're just talking about the same painted leaves, yet in one context we feel something true has been said, and in another context we feel something false has been said. And that shows us that we've got context sensitivity where we didn't think we had it.

NW: *So that's supposed to show that context is everything for language. But does it really show that?*

EB: Well, these kinds of cases have given rise to many different positions in the literature. Some of those positions absolutely endorse Travis's position that such cases show us that context is all-pervasive and we should think about language in a much more Wittgensteinian way. We shouldn't think of getting at meaning as combining stable meanings in certain kinds of repeatable ways; really we should think about language as a tool that we use, and to understand meaning we need to know about the purpose for which that tool is being deployed, and so context is everything. So that's one position—that's Travis's kind of position in the literature. I've adopted a position at the other extreme, where I try and argue that these cases don't show us that context is all-pervasive. And to get at that kind of position, I think we need to do a couple of things. Namely, we need to think very carefully about each of the cases Travis gives us, because it's not clear that the same thing will be going on in every single case. For instance, there's been quite a lot of work in linguistics about whether colour terms are actually ambiguous, and maybe that's what's coming out in that Pia example. And there might be other things going on in some of the other cases that Travis gives us. But the general motto of my type of position, 'minimalism', is that these sorts of thought experiments aren't the right place to be looking when we're trying to think about literal meaning, because they're just too complicated and too infected by all the different features we bring to bear when we're trying to interpret one another to abstract out literal meaning.

NW: *So minimalism is the position that you can have a literal meaning of a sentence and that meaning is given by the words and their arrangement within a grammatical order, and that the context can be relevant but isn't the determining factor of meaning—is that the right way to describe it?*

EB: Yeah, absolutely. So minimalists are driven by those kinds of concerns we talked about earlier—about learnability and systematicity—and they think that if you haven't got something fixed and stable and unchanging, then you're not going to be able to explain those features of language. And combined with that, they think back to a very influential philosopher in this area—Paul Grice of Oxford—who drew a really clear distinction between sentence meaning and what speakers can say when they produce those sentences. So a famous case from Grice was when he asked us to consider a letter of reference that he is writing for someone. He's writing a letter of reference for one of his students, Jones, and he just writes one sentence: 'Jones has nice handwriting'. And the thought is that if Jones has been applying for a philosophy job and that's the letter of reference Grice writes for him, then the message conveyed to somebody is going to be something like 'Jones is not a good philosopher' or 'Jones shouldn't be given this job'. But Grice says, even if that's right, we're not tempted to think that the sentence 'Jones has nice handwriting' could ever literally mean 'Jones has no philosophical talent'! There's a difference between what a sentence means and what a speaker conveys by it. And what minimalists want to say is that these Travis cases are very nice, but at least some of them are just trading on that old distinction between sentence meaning and speaker meaning.

NW: *Why do you call this minimalism?*

EB: The main reason I call it minimalism is that it adopts a very minimal job description for semantics. So there used to be an assumption, which maybe wasn't always spelt out, that what we wanted from a semantic theory, a theory of literal meaning, was that it captured our intuitions about speech act content—that it would tell us what someone has said when they produce one of these sentences. And a minimalist for me is someone who says that's too rich a job description. A semantic

theory couldn't do that at the same time that it was trying to capture facts about learnability and systematicity, because speech act content is affected by a whole range of other things—you need encyclopaedic knowledge, you need to know about socio-economic status, you need to know all kinds of things about the person you're talking to, in order to really get at what message they're trying to get across. Instead, the minimalist says we should think about semantics as just a very limited, constrained kind of activity—it is just trying to get at literal meaning. And then that literal meaning plays a part in working out what someone has said, but it certainly doesn't exhaust it.

NW: *But then what is the role for literal meaning?*

EB: That's a very good question because one of the objections that is made to my position is that I can have the minimal literal meanings, but that they're completely redundant, and there is no reason to want them. However, I think there are a range of ways in which minimal meanings or literal meanings are going to be useful. And I think it's right that minimal meanings don't do all the really nice things that we might have thought semantic content was going to do for us, but they do still have some role to play. This can be demonstrated when thinking about examples of when things go wrong. So sometimes you're having a conversation with somebody and you suddenly realize that you don't really know what they're saying to you, or there's some aspect of the context that you don't know about and it's stopping you from getting the whole meaning. In that case, what can you do? Well, the only thing you *can* do is fall back on the literal meaning. So I can think to myself that I don't really know what Nigel said when he uttered that sentence, and I don't really know what he meant or what he was trying to get across to me, yet I can still know that he said a given sentence with a given literal meaning. Therefore, I can fall back to that kind of content. And if you

spend much time with children or philosophers or law-makers, you will find they're very good at accessing that level of content. For example, when I say to my seven-year-old child, 'give some of that cake to your brother', he is very good at passing just the tiniest fragment over. And what that shows is that he is exploiting the literal meaning—he knows full well that's not what I meant, that I meant for him to give some reasonable portion of it to his brother—but he's got a nice grip on what the minimal content was and he's going to exploit it. So minimal contents do have a role to play in that sort of exchange. And if we think about legal exchanges, there are some cases where it's very important that we just get at what someone is committed to by the words that she says rather than the way she says them or the background context. And then the final point where I think minimalism is important is that these different theories about the role of context might tell us something about the nature of the mind as well—they might tell us about the relationship between language understanding and other kinds of cognitive activity that we engage in.

NW: *But what does your understanding of language have to say about the mind?*

EB: My view of language has some quite clear repercussions for what we want to say about the mind. Minimalism claims that when you understand a language, what you understand are the meanings of these primitives—the words—and a way of putting those words together to generate meanings for complexes. But if that's all that's involved in understanding a language, then we could have some quite dedicated part of the mind that is responsible for such understanding. We could have what Jerry Fodor has called a module—a module is supposed to be an encapsulated computational system which doesn't look to everything that you as an individual know, it just looks to a very small subset: a set of rules, and a library of information. And minimalism says that maybe language understanding is like that and so could be

underpinned by a dedicated module. And note that this might fit nicely with some cases where things go wrong. We know that in some kinds of cases of impairment, what happens is that your language understanding either gets impaired or preserved, while your general understanding—your pragmatic understanding—is either preserved or lost in a different way. If we think about some cases on the autistic spectrum, namely those with Asperger's Syndrome, patients on the high-functioning autistic spectrum, they have a very well-preserved language ability and they tend to be over-literal—they access the literal meanings of sentences but they're not so good at picking up on the conveyed communicated content. And what the minimalist takes those kinds of cases to show is that maybe this is reflecting a structure of the mind. We could preserve the module that underpins language understanding when there's damage occurring in other areas of the mind.

NW: *Academics are often asked to defend what they do to a wider public. How could you do that when you're just reflecting on the nature of language?*

EB: We do need to justify why we're doing what we're doing, and I think this issue is one that we can definitely find reason for studying. One reason is its inherent interest. These things are activities that we engage in every day—we talk to one another, we get the message, but how does that work? These things are very 'up close and personal' and it would be good to understand what's going on. Another reason that these things matter is that if we think about the impairment cases we were talking about a moment ago, if we're going to be able to design the best kind of interventions for those sorts of difficulties, we need first of all to know what's going on when things go right. If we want to understand what's going on when things go wrong, we need to know how they compare with typical language development. So a good theoretical understanding of language matters for treating those kinds of concerns.

And if we think about all the developments that are going on at the moment with brain scans and advances in science, it seems to me that none of those results are going to be fully robust unless we've got a good theoretical model to plug them into. If we want to know why it's useful to understand that this bit of your brain lights up when you're listening to a language and this bit of your brain lights up when you're doing something else, we need to know about the theoretical structures that they plug into. And that's what philosophy delivers for you.

Further Resources

Emma Borg (2004) *Minimal Semantics*, Oxford University Press.

Emma Borg (2007) 'Minimalism versus Contextualism in Semantics', in G. Preyer and G. Peter (eds.) *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*, Oxford University Press, 546–571.

Kepa Korta and John Perry (2020) 'Pragmatics', *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta:

http://plato.stanford.edu/archives/spr2020/entries/pragmatics.

François Recanati (2004) *Literal Meaning*, Cambridge University Press.

Charles Travis (1997) 'Pragmatics', in B. Hale and C. Wright (eds.) *Companion to the Philosophy of Language*, Blackwell, 87–107.