

Semantics for Bayesians

Topics in Epistemology
Oxford, 14/6/12

Daniel Rothschild
daniel.rothschild@philosophy.ox.ac.uk

1 A classical view of semantics

Declarative sentences, in context, express propositions. E.g. ‘I’m tall’ said by Tom expresses the proposition that Tom is tall.

Assertion of a sentence s recommends belief in the proposition expressed by s . (At least if purposes of conversation are directed towards establishing truths, as opposed to say, speculation, fiction, bullshit, whatever.)

Semantics itself provides a mapping from sentences in the language to propositions (relative to contexts).

Conversation then allows the spread of propositional knowledge (or false beliefs) through assertion.

2 Communication graded belief

Bayesian perspective at the least adds a significant category of mental states beyond belief to the picture, graded beliefs. (Some argue for elimination of full belief as an interesting category.)

Note a disconnect between this category (graded belief) and the communication process: communication only allows (directly) for transference of propositional knowledge/belief, not graded beliefs. Clearly graded beliefs are transferred between individuals, however, with such sentences as:

(1) Tom is likely to come.

I can directly transfer full beliefs via assertion, but I cannot do so with graded beliefs. Yalcin (p.c.) calls this the “Bayesian bottleneck”.

3 Ways of relieving the bottleneck

Propositional surrogates We find a useful class of propositions, that play proxy for graded beliefs and facilitate their transfer through communication.

Novel communication model Some kind of new way of thinking of semantics/pragmatics of certain parts of language to allow direct transfer of graded beliefs.

4 Surrogates

We assume that (1) is a sentence which can often be used to transfer credences in conversation.

Physical propensities Of course, sometimes the propositions needed is right at hand.

Consider for instance, saying of a die that it has a $1/6$ th chance of landing heads. Clearly there is some fact about the world that makes this the case. Limited applicability.

Evidence relative objective chance If we have the view that a given body of evidence (however considered) always gives us an objective chance (whether mushy or fine). Then propositional content could be: relative to my evidence/shared evidence, probability is p .

Report of mental states In some sense the most obvious surrogate proposition is just one about one’s own mental state. In this case (1) simply expresses the proposition that one has a certain graded belief.

5 Problems for surrogates

5.1 Yalcin’s argument

Yalcin [2007] provides one argument against treating the second sentence (1) as expressing a sort of proposition. The argument goes as follows. Suppose (1) expresses a proposition, then it does not entail that Tom came since it seems strictly weaker. So that Tom didn’t come and that Tom is likely to come should both be able to be true. But then we should be able to say:

(2) Suppose it’s not raining and/but it’s probably raining. . .

or:

(3) If it’s probably raining and/but it’s not raining. . .

Both of these seem very odd.¹ Compare Moore-paradoxical sentences:

(4) It’s raining and it’s probably not raining.

(5) It’s raining and I don’t believe it.

Crucial test:

(6) Suppose it’s raining and I don’t believe it.

Note some claim that this bad, but my empirical studies with Emmanuel Chemla suggest it is not nearly as much so:

(7) Suppose it’s probably raining.

¹To see just how difficult a pragmatic explanation of this data proves to be, see the lengthy manuscript by Dorr and Hawthorne [2011] which tries to give a coherent pragmatic account of the contrast above.

5.2 Belief attributions

Suppose sentences like (1) expressed some proposition about one’s beliefs (or the beliefs of the conversational group). Then what would belief attributions plausibly mean?

(8) John believes it’s probably raining.

It seems implausible to think that this is really an attribution to John of a belief about his beliefs or information available in his conversational context.

Seems implausible in that it makes what seems like a first-order belief attribution seem higher-order.

What about with the objective probabilities account?

These problems lead Moss [forthcoming], Yalcin [2007], Swanson [2006] to give an EXPRESSIVIST account of (some) probability talk, according to which these statements do not directly express propositions, but rather put constraints on beliefs or the common ground.

Note on any surrogate view what John believes in (8) is not what is expressed by “it’s probably raining” since it must be relative to John’s evidence/beliefs, not the speakers. Yalcin [2007] argues that (8) should be true just in case John’s probabilities assign something high to the proposition that it’s raining.

5.3 Parallels with epistemic modals

Arguments above all originally deployed for epistemic modal ‘might’.

(9) Suppose it’s not raining but it might be.

(10) John believes it might be raining.

6 Expressivist picture

Following Yalcin, I use the term *expressivism* to label a class of views about the semantics and pragmatics of certain uses of probabilistic expressions. I use the term mostly negatively: expressivist views are ones in which sentences like (1) are not used to put forward a proposition.

In classical semantics the connection between semantic values and assertive content are tight: relative to a context anyway, they are identical.

Perhaps this explains why most writers about expressivism in the meta-ethical domain view it as a thesis about semantic content directly. Take these quotes (from Yalcin [forthcoming]):

The centerpiece of any quasi-realist [expressivist] “account” is what I shall call a psychologistic semantics for the region: a mapping from statements in the area to the mental states they “express” when uttered sincerely. The procedure is broadly recursive. Begin with an account of the basic states: the attitudes expressed by the simplest statements involving the regions characteristic vocabulary. Then assign operations on attitudes to the various constructions for generating complex statements in such a




Figure 1: Structure of semantic theory for declaratives

way as to determine an “expressive role” for each of the infinitely many statements in the area. Rosen (1998), pp. 387-8, discussing Blackburn (1993)

The basic idea of expressivism is that it is the job of a semantic theory to explain what a sentence, ‘P’, means, by saying what it is to think that P. Schroeder (2008a), pp. 704

These quotes both show a presupposition that an expressivist semantics in characterizing a language as not putting forward propositions but rather expressing certain kinds of non-propositional attitudes, must directly map in the semantics expressions into properties of mental states (or ‘expressive roles’ or whatever).

This effectively lumps conversational effect into semantics, whereas we might instead keep them separate. Figure 1 divides this into two stages, the first stage, semantics determines a semantic value for a sentence (an abstract object), whereas the second stage, the BRIDGING PRINCIPLES relates that object to its characteristic conversational effect.²

A rough semantics for probability statement based on Yalcin [2007, 2010], Swanson [2006], and Rothschild [forthcoming], goes as follows:

\mathcal{L} has a simple syntax of atomic, factual sentences $a, b, c \dots$, as well as complex sentences formed with a single application of the ‘probably’ operator, P (P -sentences). An interpretation is a pair $\langle W, \llbracket \cdot \rrbracket \rangle$ where W is the set of possible worlds, and $\llbracket \cdot \rrbracket$ is a function with domain \mathcal{L} such that for an atomic sentence s , $\llbracket s \rrbracket \subseteq W$ and for a P -sentence of the form Ps , $\llbracket Ps \rrbracket = \{m \in M : m(\llbracket s \rrbracket) > .5\}$, where M is the set of all probability

²Rothschild [2011, forthcoming] elaborates this picture.

measure functions over W .

We thus piggyback on a typical propositional semantics which assigns sets of worlds to sentences to get a semantics which assigns sets of probability functions to sentences. Sets of probability functions are themselves just that, certain kinds of abstract objects. Our semantics is only completed by assigning certain bridging principles between these and conversational effects.

This semantics is not itself non-factualist. We can have a bridging principle that gets us a proposition in the end:

Factual assertion: If s is a factual sentence, then an assertion of s is a suggestion to believe $\llbracket s \rrbracket$.

Factual P -assertion: If s is P -sentence, then an assertion of s by x is a suggestion to believe the proposition that x 's credal state is in $\llbracket s \rrbracket$.

These two principles give a factualist account of assertions of probability claims, even though the semantics assigns something non-propositional to sentences with the P operator.

Yalcin, Swanson, Moss and me (?) want to replace this with something less factual along the lines of this:³

Expressivist P -assertion: If s is P -sentence, then an assertion of s is a suggestion that the conversational participants adopt some credal state in $\llbracket s \rrbracket$.

This allows for direct transmission of probabilistic belief without the medium of a propositional assertion.

7 Frege-Geach/Earlier worries

The embedding problem for meta-ethical expressivism (often called/related to the FREGE-GEACH PROBLEM) is often expressed like this:

In many cases what the non-cognitivist says about the meanings of moral sentences used in simple predication cannot plausibly apply to the same sentences when embedded. For example, if a non-cognitivist says the meaning of 'Lying is wrong' is explained by the suggestion that it serves to express disfavor towards lying in the way that 'Boo lying!' might, that does not seem to be a good explanation of what the very same words are doing when they are used in many embedded contexts. [van Roojen, 2009]

In our framework above this problem merely points to the need to compose semantic values (but not assertive contents!). We can do so by adding rules like these to our semantics:

\mathcal{L}^O , introduces conjunctions and disjunctions of P -sentences or atomic sentences (but not combinations of P -sentences and atomic sentences).⁴
The semantics is straightforward: for any θ and ϕ in \mathcal{L}^O , $\llbracket \theta \wedge \phi \rrbracket = \llbracket \theta \rrbracket \cap \llbracket \phi \rrbracket$, and $\llbracket \theta \vee \phi \rrbracket = \llbracket \theta \rrbracket \cup \llbracket \phi \rrbracket$.

³All of these authors (except me) say considerably more subtle things than this, but crudeness has a certain pedagogical value.

⁴We could adopt the semantics of Yalcin [2007] to allow arbitrary combinations of P -sentences and non- P -sentences.

Note these are totally classical understandings of the connectives (we can do negation similarly).

This has good consequences for e.g. conjunctions of probabilistic statements:

(11) Bill is probably in Paris and David is probably in Berlin.

8 Why expressivism?

How do we handle belief attributions:

Belief transparency If p is a P -sentence then $\llbracket x$ thinks $p \rrbracket$ is the set of worlds where x 's credences are in $\llbracket p \rrbracket$.

9 Sharp credences and bridging principles

(12) John thinks it's probably raining or it's probably snowing.

If disjunction is narrow scope this gets a weird meaning.

Rothschild [forthcoming] suggests to respond to these problems (and similar ones for embeddings under quantifiers) by shifting to sets of probabilities as objects of beliefs, as in the mushy credence literature.

Mushy P -assertion: If s is P -sentence, then an assertion of s is a suggestion that the conversational participants adopt some credal set that is a subset of $\llbracket s \rrbracket$.

Mushy belief transparency If p is a P -sentence then $\llbracket x$ thinks $p \rrbracket$ is the set of worlds where x 's credal set is a subset of $\llbracket p \rrbracket$.

10 Motivation Expressivism?

Note some of the original arguments for expressivism to have gone away. We can handle both suppositional contexts and belief attributions in combination with factualism account of P assertion using this semantics.⁵ So in fact the arguments for expressivism (as opposed to a more sophisticated semantics are distinct for this).⁶

Expressivism of this sort seems coherent, but is it motivated? For our purposes the Bayesian bottleneck concern still holds. The factualist story still sits uneasily with Bayesian perspective as it requires communication of credences to go via the expression of propositions. Consider conversations like this:

(13) a. It's probably raining

⁵The suppositional contexts require treatment of conjunctions of sentences about probability with non-probability which requires a bit more machinery. For now we can just assume that a factual sentence s expressing p has as its semantic denotation the set of all probability functions that assign 1 to s .

⁶Yalcin [2007] is well aware of this, and provides different arguments, which are in fact quite a bit more subtle. However, many commentators on Yalcin seemed not to have noticed this point, which admittedly is rather subtle.

b. I doubt/believe it.

What does this ‘it’ pick out here? the proposition or the semantic value. It seems like the proposition expressed is not a viable option on the mental-state type view of what proposition is expressed. More plausible on the evidence-relative objective view. Yalcin [2011] argues that it is hard to get the right range of evidence to give a coherent account of disagreement. It cannot just be your evidence, or disagreement doesn’t make sense, but if it is more than your evidence than assertion is very hard to be justified. Literature on relativism about epistemic modals is relevant here. Yalcin [2011] also suggests there are conflicting intuitions that motivate non-factualism. He gives the following case:

Fat Tony secretly plants highly compelling evidence of his murder at the docks. The evidence is discovered by the authorities, and word gets out about his apparent death. The next evening, from his safehouse, Fat Tony watches a panel of experts on the news discussing the situation. Expert A has had a good look at the evidence found at the scene. “Fat Tony is dead,” he says. Expert B has also had a good look at the evidence, but his assessment is more cautious. “Fat Tony might be dead,” B says.

If we replace ‘might’ with ‘probably’ things do not seem very different. Yalcin says this about our intuitions here (based partly on a large empirical study with Knobe):

We can all agree that Expert A, however reasonable his speech act was in light of the information available to him, spoke falsely. Things are not as he says they are. Okay; what about Expert B? Is what he said true or false? Let me remind you that Fat Tony’s planted evidence was highly compelling. Let me remind you also that Fat Tony is definitely not dead. And, before you settle on an answer, let me ask you also to ponder whether Fat Tony himself should agree with your answer.

Now, what do you say about what B said—true or false?

It appears that, as a matter of empirical fact, intuitions are unclear about cases such as this—cases where an epistemic modal claim is assessed for truth from outside the discourse context. Some are inclined to say that B spoke truly; others are inclined to say that B spoke falsely; everyone else shrugs, or proposes to change the question to one with a clearer answer.

What needs explaining for eavesdropping cases such as this, I think, is not any univocal intuition we all have about the epistemic modal claim made in the case. In the absence of a systematic empirical study, there seems to be no single intuition there to explain. Rather, what needs explaining is the absence of agreement, by competent speakers of English, on what the right answer is. What needs explaining are the conflicting intuitions. My point for now is just that conflicting intuitions are not expected on descriptivist assumptions. If B’s utterance is in the business of representing the world as being a certain way, as A’s presumably is, then either the world is that way, or it isn’t. Other things being equal, we’d expect intuition concerning the truth of B’s utterance to be about as clear as it is with A’s. Descriptivists have work to do, then, explaining why things are not equal.

To summarize the concern: if there is a proposition being put forward by epistemic modals or probably claims, it is very unclear which one it is. Together with the Bayesian bottleneck this might motivate expressivism, but the arguments are much less definitive than those that motivated the basic semantic framework.

11 Knowledge attributions

Moss [forthcoming] notes that regardless of any strangeness in the status of statements like (1) they can be embedded in knowledge reports:

(14) John knows that it probably rained yesterday.

She argues that these attributions exhibit (intuitively) typical properties of knowledge attributions, in particular factivity and safety.

This presents a problem if we think that like (8) we cannot treat (14) as a relation between an individual and a proposition. After all, we commonly consider knowledge to be a PROPOSITIONAL ATTITUDE.

Factivity For any world w , if ‘ S knows that p ’ is true as evaluated at w , then ‘ p ’ is true as evaluated at w .

Safety For all cases α and β , if β is close to α and in α one knows that C obtains, then in β one does not falsely believe that C obtains.

What do we then say about:

(15) John falsely believes it’s probably raining.

Thus, she suggests we need to expand the category of knowledge to include properties of credal states as well as belief states.

References

- Cian Dorr and John Hawthorne. Embedding epistemic modals. manuscript, 2011.
- Sarah Moss. Epistemology formalized. *Philosophical Review*, forthcoming.
- Daniel Rothschild. A note on conditionals and restrictors. manuscript, 2011.
- Daniel Rothschild. Expressing credences. *Proceedings of the Aristotelian Society*, forthcoming.
- Eric Swanson. *Interactions with Context*. PhD thesis, MIT, 2006.
- Mark van Roojen. Moral cognitivism vs. non-cognitivism. *Stanford Encyclopedia of Philosophy*, 2009.
- Seth Yalcin. Epistemic modals. *Mind*, 116:983–1026, 2007.
- Seth Yalcin. Probability operators. *Philosophy Compass*, 5(11):916–937, 2010.

Seth Yalcin. Nonfactualism about epistemic modality. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*. OUP, 2011.

Seth Yalcin. Bayesian expressivism. *Proceedings of the Aristotelian Society*, forthcoming.