

EXPRESSIVISM

Philosophy of Language · Bob Beddor

Representation & Possible Worlds Semantics

Suppose a speaker says:

- (1) It's raining outside.

It's natural to think that they are trying to *represent* the world as being a particular way—namely, one where it is raining outside. Similarly, if a speaker says:

- (2) China is more populous than Japan.
- (3) There's traffic on the AYE.

It's natural to interpret them as trying to represent the world as being a certain way—as being a world where China is more populous than Japan, or as being a world where there's traffic on the AYE.

One way of implementing this idea is in terms of a possible worlds semantics. According to possible worlds semantics, the meaning of a sentence is the set of possible worlds in which it is true. For example, the meaning of (1) will be the set:

$\{w \mid \text{It's raining outside at } w\}$.

Similarly, the meaning of (2) will be the set:

$\{w \mid \text{China is more populous than Japan at } w\}$.

Is All Language Representational?

It seems plausible, then, that much of ordinary language aims to represent what the world is like. But does all language aim to do so?

To see why one might doubt this, consider 'expressives' such as:

- 'ouch', 'oops', 'yum!', 'wheeee!', 'yay!', 'boo!'

Arguably, these expressions don't purport to represent the world as being a particular way. Rather, they serve to directly express some emotions/attitudes of the speaker.

More generally, for any class of linguistic expressions e we can distinguish between two positions:

Representationalism Utterances involving e aim to represent the world.

Expressivism Utterances involving e do not aim to represent the world; rather, they aim to express the speaker's attitudes/emotions.

Equipped with this distinction, we can go on to ask two big questions:

- Which classes of expressions fall in the expressivist camp? That is, which—if any—bits of language aim to express attitudes rather than represent the world?
- How should we analyze the meaning of expressive discourse?

There is considerable controversy regarding the answers to both of these questions.

Taste Predicates

One controversial class of expressions is *taste predicates*—that is, terms like ‘tasty’, ‘disgusting’, ‘delicious’, ‘fun’, etc. Do such expressions purport to represent the world, or do they serve to express the speaker’s attitudes/emotions?

This is controversial. One position is that taste predicates *do* purport to represent the world. In particular, they purport to represent facts *about* the speaker’s emotions/attitudes. On this view a sentence such as:

(4) Chili is tasty!

means roughly: *I (the speaker) enjoy the taste of chili*. On this view, if (4) is uttered by Bob, then the meaning of this utterance corresponds to the set of worlds:

$\{w \mid \text{Bob enjoys the taste of chili at } w\}$.

However, this proposal—while natural—runs into a problem explaining *disagreement*. Consider the following exchange:

- (5) a. *A*: The chili is tasty.
b. *B*: No, it’s not tasty!

In this case, it’s natural to view *A* and *B* as disagreeing. Such a disagreement seems to be perfectly natural. Contrast this with the following exchange:

- (6) a. *A*: I enjoy the taste of the chili.
b. *B*: ?? No, I don’t enjoy the taste of the chili!

This exchange seems to be very weird. Why? The answer seems to be that *B*’s claim doesn’t actually disagree with *A*’s claim, which is why it’s odd for *B* to use a negation marker (*No*). Note that this difference between (5) and (6) would be hard to explain if *The chili is tasty* simply meant *I (the speaker) enjoy the taste of the chili*.

Some have taken this motivate an expressivist treatment of taste predicates. On this view, when I utter (4), I am not representing myself (or anyone else, for that matter) as enjoying the taste of chili. Rather, I am directly *expressing* my enjoyment of chili.

An analogy:

What does it mean for a sentence to directly express a state of enjoyment? Here an analogy may be helpful. Contrast the following two sentences:

- (7) It’s raining outside.
(8) I believe it’s raining outside.

It would be implausible to suggest that these two sentences mean the same thing. After all, (7) makes a claim about the weather, but (8) makes a claim about the speaker’s beliefs. What, then, is the relation between the

two? One answer: someone who utters (7) does not represent themselves as believing that it's raining outside. Nonetheless, they manage to *express* this belief.

Similarly, expressivists about taste predicates maintain that someone who utters (4) does not represent themselves as enjoying the taste of chili. Nonetheless, they manage to express this state of enjoyment.

Moral Discourse

Thus far most of the work developing expressivism has been focused not on taste predicates but instead on *moral discourse*, such as:

(9) Killing innocent people is wrong.

(10) Helping the less fortunate is good.

According to expressivists, someone who utters a moral sentence is expressing their desire-like attitudes (desires, states of disapproval, preferences, and the like). For example, someone who says (9) is expressing their disapproval of killing. Someone who says (10) is expressing their approval of helping the less fortunate.

The debate between representationalism and expressivism about moral discourse is one of the major debates in metaethics. But why would anyone want to be a moral expressivist in the first place? Two main arguments have been given in the literature:

First Argument: Evades Traditional Problems in Metaethics
Morality gives rise to some tricky puzzles:

- *Metaphysical puzzles:* Are moral properties located in space and time? If so, where are they? Can they stand in causal relations with us? (If not, can we even refer to them?)
- *Epistemological puzzles:* Most of us think that moral knowledge is possible. We know, say, that torturing innocent people is wrong. But how is it that we come to know these things?
 - Contrast with perceptual knowledge, which seems comparatively unmysterious. Do we perceive moral properties/facts?
- *Motivational puzzle:* Some metaethicists have argued that there's a particularly close connection between moral judgment and motivation. According to these *motivational internalists*, if someone judges that they morally ought to e.g., giving to charity, then they will necessarily be at least somewhat motivated to give to charity. On reflection, this motivational pull can seem somewhat mysterious. How can a judgment directly motivate action in this way?

Some have thought that expressivism helps avoid these puzzles.

- Q: is this right?

The Frege-Geach Problem

Despite its apparent advantages, expressivism faces serious obstacles. (Here I'll focus on expressivism about morality—it is worth thinking about whether these obstacles also apply to expressivist theories of other areas of discourse, such as taste predicates.). Probably the most significant problem is what's known as the 'Frege-Geach Problem'.

The Frege-Geach Problem arises from the fact that moral discourse seems to behave like ordinary representational discourse in many respects.

In particular, we can embed ordinary representational discourse in a variety of complex constructions:

- (11) It's not raining outside. (negation)
- (12) Either it's raining outside or it's sunny outside. (disjunction)
- (13) It's raining outside and it's cold outside. (conjunction)
- (14) If it's raining outside, you should grab an umbrella. (conditionals)

We can embed moral discourse in the exact same ways. For example:

- (15) Lying is not wrong. (negation)
- (16) Either lying is wrong or lying isn't wrong. (disjunction)
- (17) Lying is wrong and I enjoy it nonetheless. (conjunction)
- (18) If lying is wrong, then I'm in trouble. (conditionals)

In this regard, moral discourse seems to closely resemble representational discourse, and seems to differ from 'yay' and 'boo' talk:

- (19) ?? Not boo lying. (negation)
- (20) ?? Either boo lying or not boo lying. (disjunction)
- (21) ?? Boo lying and I enjoy it nonetheless. (conjunction)
- (22) ?? If boo lying, then I'm in trouble. (conditionals)

Compositionality & the Joys of Truth Conditions

Why is this a *special* problem for expressivism?

To see the answer, we need to first see how to explain the embedding behavior of representational sentences. Note that the underlying structure of the sentences (11)—(31) can be represented as follows:

- (23) not p . (negation)
- (24) p or q . (disjunction)
- (25) p and q . (conjunction)
- (26) If p , then q . (conditionals)

It's natural to try to give a general recipe for analyzing these sentences. In particular, it's natural to try to give a general formula for analyzing the meaning of *Not* p in terms of the meaning of p , and the meaning of p or q in terms of the meaning of p and the meaning of q , etc.

This is precisely the approach that's adopted by a traditional descriptivist semantics, and which you may have encountered in a logic class. The simplest version of this idea is that we can define the meaning of (23)—(26) using a truth-table. For example, the truth table for negation will be:

p	$\neg p$
0	1
1	0

Here the basic insight is that negation reverses the truth-value of a sentence: if p is true, not- p will be false, and vice versa. Put in terms of possible worlds talk: the meaning of not- p will be the set of possible worlds in which p is false.

The truth table for conjunction is:

p	q	$p \wedge q$
0	0	0
0	1	0
1	0	0
1	1	1

The idea is that a sentence of the form *p and q* is true iff *both p* is true and *q* is true. Put in possible worlds talk: the meaning of *p and q* will be the set of possible worlds in which both *p* is true and *q* is true (i.e., $p \cap q$).

Finally, here is the truth table for the conditional:

<i>p</i>	<i>q</i>	$p \rightarrow q$
0	0	1
0	1	1
1	0	0
1	1	1

The idea here is that a conditional such as *If it's raining outside, you should grab an umbrella* is falsified by a situation in which it's raining but you should **not** grab an umbrella. However, it isn't falsified by the mere fact that it isn't raining. In possible world terms: the mean of *If p, then q* is the set of worlds in which either *p* is false or *q* is true (equivalently: it's the set of worlds in which it's not the case that both *p* is true and *q* is false).

More generally, a representational semantics allows us to define the meaning of a complex sentence in terms of the meaning of its constituent sentences. This is a very important advantage. (Recall the notion of *compositionality*.)

Why Expressivists Have a Problem

So if you have a truth conditional semantics, you have a nice general recipe for getting from the meaning of constituent sentences to the meanings of complex sentences. But it's not at all clear that expressivists have a correspondingly nice general recipe for doing so. According to a very simple way of developing an expressivist semantics, the meaning of:

(27) Lying is wrong.

is the mental state it expresses—specifically, *disapproval of lying*. Presumably, then, the meanings of the following complex sentences should also be the mental states they express:

(28) Lying is not wrong. (negation)

(29) Either lying is wrong or lying isn't wrong. (disjunction)

(30) Lying is wrong and I enjoy it nonetheless. (conjunction)

(31) If lying is wrong, then I'm in trouble. (conditionals)

But what are these mental states? And in what sense can they be derived as a function of the meaning of the constituent sentence, (27)?

There are really two distinct questions here. The first is:

- *What mental states do these complex sentences express?*

The second is:

- *How can we derive the mental states that these complex sentences express as a function from the mental state expressed by (27)?*

Blackburn's Higher Order Account

Here is one strategy that Blackburn uses to try to make progress on the Frege-Geach problem. Take a conditional sentence such as:

(32) If lying is wrong, then tricking someone else into lying is wrong.

Blackburn's suggestion is that the conditional expresses disapproval of being in the following complex state:

- Disapproving of lying & not disapproving of tricking someone else into lying

Here's one way of representing this. Let $[S]$ be the mental state expressed by an utterance of S (which, on the expressivist's picture, is the meaning of S). And let DIS represent the state of disapproval. Then we can analyze an arbitrary conditional sentence as follows:

$$[\text{If } S_1, \text{ then } S_2] = \text{DIS}([S_1] \& \neg[S_2])$$

Some questions:

- What does the higher-order account say is the meaning of the following sentences?

(33) If giving to charity is good, then stealing from charities is wrong.

(34) If lying is wrong, then I'm in trouble.

- Is the higher-order account's analysis of these sentences plausible? Why or why not?

Modus Ponens and Blackburn's Account

One important constraint on a theory of conditionals is that it should explain why it's irrational to accept the premises of the following argument but not the conclusion:

1. p
2. If p , then q
3. q

For example:

1. It's raining.
2. If it's raining, you should take an umbrella.
3. You should take an umbrella.

Lying Argument

1. Lying is wrong.
2. If lying is wrong, then convincing others to lie is also wrong.
3. Convincing others to lie is wrong.

Blackburn's explanation of this is that *lying is wrong* expresses disapproval of lying, and *If lying is wrong, then convincing others to lie is wrong* expresses disapproval of being in the state: disapproving of lying and not disapproving of convincing others to lie. And so someone who accepted both of the premises without accepting the conclusion would be committed to disapproving of the very state that they're in.

van Roojen's objection: This doesn't explain the difference between Lying Argument and the following:

Variant Lying Argument

1. Lying is wrong.
2. It is wrong to think that lying is wrong without also thinking that convincing others to lie is wrong.
3. Convincing others to lie is wrong.

Lying Argument is valid, whereas Variant Lying Argument is not valid. Blackburn's account seems unable to explain the difference.