

Semantic Externalism and the Causal Revolution

1. Psychologism

A variety of views we looked at either implicitly or explicitly endorse:

Psychologism: To know the meaning of a term is just to be in a certain psychological state.

- Perhaps the clearest example of psychologism is the *idea view of meaning*, which said that the meaning of a term is some idea in the head of the speaker.
- Arguably, Frege also endorsed a version of psychologism:
 - o Recall that senses are *ways in which the referent is presented to a speaker*. (In the case of proper name, this is something like a description of the referent of the name.) So even though Frege was adamant that senses are public, in that they can be shared, he thought that any who grasped a particular sense would be in a psychological state, which determined a referent.
- Same for Russell: the meaning of a name is some description, which must be grasped by the speaker.

2. From Psychologism to Externalism

Starting with Kripke, many philosophers of language came to repudiate psychologism in favor of a view on which meaning is determined – at least in part – by *extra-psychological* factors - factors that are outside the heads of speakers. This view is often known as “externalism”:

Semantic Externalism: The meanings of various terms are determined by factors that are external to the speaker.

In Putnam’s catchy slogan: “meanings just ain’t in the head!”

3. Motivations for Externalism

3.1 Kripke on Names

Judy and Trudy. Suppose Fred has met his next door neighbor, Judy. Unbeknownst to him, she has a completely identical twin: her twin looks just like her, has the same personality, etc. To make the two as similar as possible, let us suppose that this twin has also been named (quite cruelly!) “Judy”. However, Fred has never met this twin.

Question: Suppose Fred now makes a claim like, “I chatted with Judy the other day.” Which twin does his use of the term “Judy” refer to?

Kripke’s Answer: It refers to the Judy that he has met, not the identical twin he’s never met. Why? Not because of any description in his head – arguably, every description he associates with the name “Judy” applies equally well to both of the twins.¹ Rather, the reason why his use of the name, “Judy” refers to the person he’s met is that he stands in a certain causal relation with her: he has causally interacted with her, but he hasn’t causally interacted with her twin. This causal link is an *external factor* – it’s something outside his head.

¹ This is debatable. Recall from our discussion last week that there are some fancy descriptivist maneuvers. Maybe the relevant description should be “The person named ‘Judy’ that I’ve causally interacted with”. If this is the description, then this arguably tells the two twins apart.

3.2 Natural Kind Terms

Another motivation for semantic externalism comes from natural kind terms – that is, nouns that refer to natural substances.

Examples: “water”, “elm”, “gold”, “tigers”, “uranium”

Putnam on Twin Earth. Imagine a planet much like ours, Twin Earth. Just like on earth, the lakes, rivers, and streams of Twin Earth are filled with a clear liquid that people can drink (and in which fish swim, etc.). However, it turns out that this liquid has a different chemical structure than H₂O; let’s call it “XYZ”.

Putnam now asks us to imagine that in the 1700s – before anyone knew about the chemical composition of water – an Earthling named Oscar points to a glass and says:

(1) There is some water.

At the same time, Oscar’s Twin Earth counterpart – Twin Oscar – points to glass filled with XYZ and also utters (1). Putnam offers the following judgment about the case:

Judgment: In Oscar’s mouth, “water” refers to H₂O, whereas in Twin Oscar’s mouth, “water” refers to XYZ. This is true that even though neither Oscar nor Twin Oscar has the ability to distinguish H₂O from XYZ.

Following Kripke, Putnam holds that “water” is a rigid designator. Since Oscar’s use of “water” (on Earth) refers to H₂O, it refers to H₂O in all possible worlds. Similarly, Twin Oscar’s use of the word “water” (on Twin Earth) rigidly designates XYZ; it refers to XYZ in all possible worlds.²

Elms vs. Beeches. For a less far-fetched example, Putnam gives the example of different types of trees. He insists that he – like many ordinary people – doesn’t know how to distinguish elms from beeches. Putnam claims the following is the intuitive judgment:

Judgment: When the ordinary person says, “There’s an elm over there”, their use of the word “elm” refers to elms, and not beeches, even though they cannot distinguish between the two.

Burge on Arthritis. In the paper, “Individualism and the Mental” Burge introduces the example of a subject, Bert, who goes to a doctor. Bert has a number of true beliefs about arthritis (for example, that it is a potentially painful medical condition). However, he also has some false beliefs: in particular, he incorrectly believes that one can get arthritis in the thigh (rather than the joints). Complaining about a pain in his thigh, he says:

(2) I have arthritis.

² One important consequence of this is that “water is H₂O” (as uttered on Earth) expresses a necessary truth – there’s no world in which it is false. This is true even though we cannot discover it *a priori*. More generally, one important consequence of Kripke and Putnam’s work was to suggest that there can be necessary truths that are only knowable *a posteriori* - that is, through experience. (Previously, it was assumed that all necessary truths are *a priori*.)

Now imagine Bert has a twin, Ernie, who is physiologically the same as Bert, but belongs to a different linguistic community, one where everyone in the society uses the word “arthritis” to refer to the source of a pain that occurs in *both* the joints in muscles. Ernie also utters (2).

Judgment: When Bert utters (2), what he says is false. (We can imagine the doctor correcting him, “Actually, you don’t have arthritis.”) When Ernie utters (2), what he says is true.

Questions:

- Do you find these examples convincing?
- If so, how far should we extend semantic externalism? Should we leave it at names and natural kind terms? Or should we be semantic externalists about all referring terms in the language?

4. Externalism About Mental Content

So far we have focused on semantic externalism as a thesis about the meanings of certain *words*. But many writers also extend this thesis to our *concepts/thoughts*.

To motivate this, go back to the Twin Earth example. This time, imagine that instead of uttering the words, “There is some water”, Oscar just *thinks* these words. (And likewise, with Twin Oscar.) If you have the judgment that Oscar’s concept *water* is about H₂O, and that Twin Oscar’s concept *water* is about XYZ, then this motivates adopting an externalist picture of mental content as well. (Similarly points can be made with **Elms vs. Beech** and the **Arthritis** examples.)

This leads to a radical picture on which the contents of our very *thoughts* are determined – at least in part – by factors that are external to us!

5. Objections to Semantic Externalism

Some have objected to semantic externalism on the grounds that they do not share the relevant intuitions – they do not agree, say, that Oscar’s use of “water” only refers to H₂O.

Others have objected specifically to externalism about mental content. One of the main objections:

Privileged Access Objection: It’s often been held that we have some sort of privileged epistemological access to our own thoughts. On a traditional view – associated with Descartes – we can know our own thoughts directly and immediately, in a way that we can’t know facts about the external world.

“...if I mean only to talk of my sensation, or my consciously seeming to see or to walk, it becomes quite true because my sensation refers only to my mind.”
- Descartes

But if we embrace externalism about mental content, some have thought that we need to give up this view. After all, if externalism about mental content is true, then what I am thinking about is determined – in part – by factors that are external to me.

One way of making this worry vivid: suppose that while Oscar is sound asleep, he is abducted and transported to Twin Earth. For awhile he will go around using the word/concept “water”, all the while meaning H₂O. Eventually, as he spends enough time in Twin Earth, the meaning of his word/concept will shift, coming to mean XYZ.

But he won't be able to tell the difference: from the inside it will seem like he has been talking about the same thing this whole time.

Question: Do you find this objection compelling? Is there any plausible way of responding to it?

6. Advantages of Semantic Externalism?

We just looked at an epistemological objection to semantic externalism. However, Putnam actually thinks that semantic externalism carries an important epistemological advantage: it helps us refute skepticism about the external world!

A Classic Skeptical Argument:

A skeptical scenario: You aren't actually sitting in a classroom now. In fact, you don't even have a body. You're just a brain in a vat (BIV). Electrodes have been planted on your brain (that is, on you!); thanks to a computer, your brain is being bombarded with sensory stimuli that creates the illusion of being in a classroom. (This is basically the scenario from *The Matrix*, give or take. It's also very similar to Descartes' evil demon argument.)

According to the skeptic, you can't rule out the possibility that you're a BIV. Hence you can't know that you're not a BIV. But if you can't know that you're not a BIV, then you can't know that you're sitting in a classroom right now (or, more generally, that an external world exists).

Putnam's response relies on semantic externalism. Suppose a BIV has an image as of a tree, and consequently thinks, "There is a tree." Question: Is this thought true or false?

You might think the answer is obviously *false*. But wait! says Putnam. If semantic externalism is true, then the BIV's concept "tree" can only refer to some entity *x* if the BIV stands in an appropriate causal relation with *x*. But the BIV is not causally related to trees at all. So the BIV's concept, *tree*, isn't about trees. Instead, it's about something else – maybe certain states of the computer. But if it's about a certain state of the computer, then what the BIV thinks is true after all! (It's not true in English, but the BIV isn't speaking English! The BIV is speaking Vat-English, and the sentence, "There is a tree" is true in Vat-English.)

This is already a significant anti-skeptical result: we seem to have shown that the BIV is not massively deluded about the external world, contrary to what one might have initially thought. But Putnam thinks we can go even further, and give a proof that we're not BIVs after all.

- 1) Either I am a BIV or I am not.
- 2) If I am **not** a BIV, then when I say, "I am a BIV", what I say is false.
- 3) If I am a BIV, then when I say, "I am a BIV", I am speaking Vat-English, and my word "BIV" doesn't refer to brains in vats (it refers to states of the computer, or certain electric impulses). (from semantic externalism)
- 4) But if I am BIV, then it's false that I am state of the computer, or an electric impulse.
- 5) So if I am a BIV, then when I say, "I am BIV", what I say is false. (from 3, 4)
- 6) So when I say, "I am a BIV", what I say is false. (from 1, 2, 5)
- 7) I am not a BIV. (from 6)

Big Question: Have we successfully refuted the skeptic? Or is there something wrong with this argument?