EXTERNALISM AND INTERNALISM

1. Two theses about meaning

    (i)   *Understanding/meaning*: If a speaker understands a sentence in a language, she knows that the sentence has a certain meaning in that language.

    (ii)   *Meaning/truth-conditions*: A meaning of a sentence determines a set of conditions under which it is true.

These give rise to *Bloomfield's problem*: a theory of meaning that satisfies both seems unreasonably ambitious. Burge and Putnam brought the problem to focus.

1.1. *The arthritis argument.* Felix has pain is his thigh and sincerely utters 'I have arthritis.' As uttered by Felix, the sentence is false. Felix has a duplicate who lives in a linguistic community that uses 'arthritis' to talk about all instances of joint or bone inflammation. As uttered by Twin Felix, the sentence is true. By (ii) the sentence has a different meaning in English and in Twin English. Felix and his twin, being duplicates, know the same truths. So, assuming they both understand 'I have arthritis' in their own language, they both understand it in the language of the other. However, this is implausible.

What if we insist that it is analytic that arthritis is an inflammation of the joints, and thus, that Felix does not understand 'I have arthritis'? Since similar examples can be constructed left and right, this sort of response acknowledges that a semantic theory must come close to being the theory of everything.

1.2. *The water argument.* Oscar is pointing at a glass of water and sincerely utters 'This is water.' As uttered by Oscar, the sentence is true. Oscar has a duplicate who lives in a world where the watery stuff is largely made up of XYZ molecules. As uttered by Twin Oscar (pointing at a glass of water), the sentence is false. By (ii) the sentence has a different meaning in English and Twin English. Oscar and his twin, being duplicates, know the same truths. So, assuming they both understand 'I have arthritis' in their own language, they both understand it in the language of the other. However, this is implausible.

What if we insist that the watery stuff on Twin Earth is really water, and thus, that 'This is water' as uttered by Twin Oscar is true? Since similar examples can be constructed left and right, this sort of response rules out the possibility of learning that things are fundamentally different from what they appear to be.

2. Two unsuccessful answers

The arguments rest on two crucial assumptions. The first is that duplicates know the same truths (i.e. if *p* is true and one of them knows it then the other does too). The second is that 'meaning' is used in the same sense in (i) and (ii).

2.1. *The language response*: the first assumption if false because duplicates differ in what they know about English. For example, Felix knows what is expressed by 'Arthritis is an English word' but his twin (not being a speaker of English) does not.

The answer is unsatisfactory for two reasons. First, unlike Twin Felix, Twin Oscar is arguably a speaker of English, and so, the language response does not help to disarm the water argument. Second, the arthritis argument can be run with a much weaker assumption, to wit, that both Felix and his twin know that 'I have arthritis' has a meaning and neither believes it has another one.

2.2. *The indexical response*: the second assumption is false because it ignores the distinction between character and content. In particular, both Oscar and his Twin know the character of 'This is water' but in their different environments this character determines different contents.

The answer is unsatisfactory for two reasons. First, unlike 'water', 'arthritis' is not a natural kind term, and so, it presumably has a constant character. Second, the interpretation of 'water' does not depend on the context of utterance, but the context of common use. If Oscar is transported to Twin Earth his use of 'water' will still refer to water, not twater – at least for a while.

3. Externalism

Externalists opt for a more radical version on the language response. They maintain that duplicates can differ in their mental states that do not concern language. In particular, Twin Felix has no beliefs about arthritis; twin Oscar has no beliefs about water. This allows them to block both the arthritis and water arguments and hold on to (i) and (ii).

Internalism: necessarily, intrinsic duplicates are mentally the same.
Externalism: internalism is false.

3.1. A standard skeptical argument goes like this: "You know you have hands only if you know that the thought you would express by the sentence 'I have hands' is true. However, for all you know, you could be a brain in the vat, in which case this thought would be false. So, you don't know whether you have hands."

If externalists are right, this argument fails: if I am a brain in a vat the thought I would express by the sentence 'I have hands' would be different but it would still be true.

3.2. Here is a version of the skeptical argument that is compatible with externalism: "You know you have hands only if you know that the thought you would express by the sentence 'I have hands' is about your hands (and not about your own experiences, electrical stimuli, or aspects of a computer program, etc.). However, for all you know, you could be a brain in the vat, in which case this thought would not be about your hands. So, you don't know whether you have hands."

Transparency: we know *a priori* the contents of our thoughts.
If externalists are right and transparency is true this argument also fails: I believe that I have hands, I know the content of my belief, so I know that my belief is about my hand.

(Of course, if I am a brain in the vat, I have an equally compelling argument proving something I would express by saying 'I am not a brain in the vat.' The conclusion would be impeccable, but it would not say what this sentence says in English. If I were a brain in the vat, I would not have beliefs about brains and vats.)

(Of course, this argument will not convince a resolute skeptic who no longer believes that he has hands.)

3.3. Some philosophers think the response to the skeptic fails because externalism and transparency are incompatible. Imagine that one night while he is sound asleep, secret agents move Oscar to Twin Earth, and they dispose of his twin. When he wakes up Oscar will notice no difference and keep saying things like 'There is water in that glass over there' but his claims will be false: what's over there is XYZ and XYZ is not water. However, over time, as he interacts with XYZ, the word 'water' is his mouth will eventually come to refer to XYZ (perhaps after a period of referential indeterminacy). Back on Earth before the *slow switch*, at time *t* Oscar believed that water is wet at *t*, and given transparency, he knew that he believed this. Given externalism, after the switch, at time *t'* Oscar does not have any beliefs about water, so he does not believe that water is wet at *t*, and hence, he does not know that he believes this. So, it looks as if Oscar knew something before the switch that after the switch he no longer knows. And yet, it seems strange to say that there is anything that Oscar has forgotten.

4. Narrow content

Proponents of narrow content opt for a more radical version of the indexicalist response. They do not merely argue that we need to distinguish between the character and the content of 'water' – they claim that we need a distinction between two sorts of content. Narrow content is what speakers who understand an expression associate with it in a given context of utterance, but it is wide content that determines truth-conditions. They hold that (i) and (ii) are both true but employ different notions of meaning.

4.1. One argument for narrow content appeals to *causation*: when Oscar and his twin are both reaching out for a glass of water we want to give the same explanation for their behavior. But if they don't have the same beliefs and desires (as externalists insist) the explanation must be different. Externalists object that this confuses causation and causal explanation. Causation is plausibly local: features of the environment can affect an agent's behavior only by affecting her intrinsic properties. But in providing causal explanations we must describe causes and effects, and it is here that the environment makes a difference.

4.2. Another argument for narrow content appeals to *rationality*: we might say that Pierre associates different narrow contents with 'Londres' and 'London' even though the wide contents of these words are the same. Perhaps rationality is a matter of narrow content, which is why Pierre can be rational even though he affirms both 'Londres est jolie' and 'London is not beautiful.' But this is problematic. Suppose someone were to tell Pierre that his beliefs concern the same city. Then Pierre would be rationally required to abandon at least one of these beliefs. But if rationality is a matter of narrow content, this would be hard to explain: narrow content

does not determine reference, so learning that 'Londres' and 'London' refer to different cities should *prima facie* make no difference to the rationality of holding both beliefs.

4.3. Chomsky thinks that the *linguistically relevant* notion of content is narrow. If we assume that anaphoric dependence requires identical semantic value, we are driven to an absurd metaphysics.

(1)   London is near the ocean but if sea levels rise it may have to move.
(2)   The book I just read hardly fits onto my shelf, and yet it is sold at every airport.
(3)   The door I just passed through was painted green.

> Suppose cup_1 is filled from the tap. It is a cup of water, but if a tea bag is dipped into it, that is no longer the case. It is now a cup of tea, something different. Suppose cup_2 is filled from a tap connected to a reservoir in which tea has been dumped (say, as a new kind of purifier.) What is in cup_2 is water, not tea, even if a chemist could not distinguish it from the present contents of cup_1. The cups contain the same thing from one point of view, different things from another; but in either case cup_2 contains only water and cup_1 only tea.

Externalists would challenge this diagnosis by stressing that the two points of view do not have the same status: the perspective of the chemist trumps that of ordinary folk when it comes to the issue of what is and what isn't water.

At the heart of Chomsky's resistance to externalism is his rejection of semantic deference. We should defer to chemists when it comes to the nature of water but they hold no special authority when it comes to the meaning of 'water.' For Chomsky, understanding a natural language is a strictly individual capacity, one that is independent not only of the speaker's social embedding but even of her communicative aims.

I-language:

(i)   intensional   (individuated by generative procedures – function in intension)
(ii)   idiolect   (independent of social relations)
(iii)   internalist   (independent of environment)