

Faculty of Philosophy Part II paper 2 Philosophy of Mind

Intentionality

These lectures will cover four topics:

1. The concept of intentionality
2. The propositional attitudes
3. The non-existent
4. Externalism and internalism about intentionality

Lecture 1 The concept of intentionality

Terminology & history

Intentionality is a philosophical term of the 'aboutness' or 'directedness' of mental states: the fact that some (or maybe all) mental states represent things. So it is also called mental representation.

The term 'intentionality' derives from the Latin 'intentio', which literally means *stretching* or *tension*; but in the scholastic Aristotelian philosophy of the middle ages was used as a technical term for *concept* or *notion*. (See the entry 'Intentionality' in the REP online.)

Franz Brentano (1838-1917) revived the terminology of intentionality to characterize the essence of the mental (he used the term 'intentional inexistence' (German: *Intentionale Inexistenz*) rather than 'intentionality', but he meant (NB 'inexistence' means 'existence-in', it does not mean 'non-existence').

Brentano's famous statement was:

Every mental phenomenon is characterised by what the scholastics of the Middle Ages referred to as the intentional inexistence of the mental, and what we would call relation to a content, direction upon an object or immanent objectivity... Every mental phenomenon includes something as object within itself. In presentation something is presented, in judgement something is affirmed or denied, in love loved, in hate hated, in desire desired and so on.

Franz Brentano, *Psychology from an Empirical Standpoint* (1874) p.88

Brentano here says that every mental phenomenon exhibits intentional inexistence (intentionality), and gives the examples of how in judgement something is affirmed or denied, in love something is loved etc. (NB 'Presentation' means what we might mean by 'conscious experience'.)

The general idea of an intentional state S is that when a subject is in S, there is always something which is S'd.

The point of Brentano's definition was to distinguish psychology from the other sciences. Brentano's view was that psychology is distinguished from the other sciences by its subject matter: mental phenomena. And mental phenomena are distinguished from physical phenomena by their intentionality.

It is clear from this that intentionality in our sense is not the same as *intending* to do something. Intending in this sense is an intentional state — one cannot intend without intending something — but intentionality as such has no special connection to intending.

Another terminological caution: we must distinguish intentionality, considered as a property of mental states, and *intensionality* which is a property of sentences or linguistic contexts. A linguistic context is intensional when it is not extensional, and a context is not extensional when either or both of the following principles fail to apply:

(1) Substitution of co-referring terms

From '....a...' and 'a=b' infer '...b...'

(For example: from 'Vladimir is taller than George Orwell' and 'George Orwell=Eric Blair' infer 'Vladimir is taller than Eric Blair'.)

(2) *Existential generalization*

From '...a...', infer '∃x...x...'

(For example: from 'George Orwell is shorter than Vladimir' infer 'There is/exists someone who is shorter than Vladimir'.)

These principles should be familiar to you from elementary logic: the logic you learned in 1A is all extensional logic. Many linguistic contexts which describe intentionality are intensional. But not all of them are (e.g. 'x sees y'). See Crane, *Elements of Mind* §§4-6 for the idea of intensionality and its relationship to intentionality.

The significance of intentionality

Why is intentionality philosophically interesting? One reason is that it provides us with a way of answering Richard Rorty's question: what unifies the concept of the mental? What is it that pains and beliefs have in common that justifies us in calling them 'mental'? (See Rorty, *Philosophy and the Mirror of Nature*, chapter 1). Brentano gives an answer: all mental phenomena are intentional (this is called 'Brentano's thesis': see appendix below, for this topic which was not discussed in the lectures).

But even if we are not interested in answering Rorty's question, we should be troubled by the problem of intentionality: that is, the problem of the representation of the non-existent.

Two things are commonly called the problem of intentionality: (i) how to explain the place of intentionality in the natural or physical world; (ii) how to explain thought about the non-existent.

Here we shall be interested in the second of these problems. It can be expressed in terms of the conflict between the following three propositions:

- (A) Intentional states are relations between thinkers and the things thought about
- (B) Relations entail the existence of their relata
- (C) Intentional states can be about things that do not exist

It seems that these cannot all be true together, so which one should we give up?

Some views think that (A) is a non-negotiable commitment: intentionality is essentially relational. An example of this is a common view inspired by Russell's distinction between acquaintance and knowledge by description: we are related to **particulars** (in Russell's case, sense-data) by acquaintance, and to **properties** by description. But strictly speaking, none of our thoughts are about things that do not exist.

But there seem to be many cases of thinking about things which cannot fit into the relational model: in particular, things that are neither cases of relations to properties nor relations to existing objects.

There are two kinds of case to consider: (i) fiction, where we know that what we are talking about does not exist; and (ii) error, where we mistakenly think that something exists when it does not.

Fiction: Consider Hamlet, who did not exist, and contrast him with Richard III, who did. One of the things you are thinking about exists, and one does not.

Error: consider Vulcan, the planet postulated by Le Verrier to explain the perturbation in Mercury's orbit, and contrast it with Neptune, the planet postulated in a similar way. Neptune does exist, Vulcan does not.

Both Hamlet and Vulcan seem to be particular things that do not exist.

This formulation of the problem of intentionality presupposes a metaphysical thesis: that relations can hold only between existing things. Someone might doubt this. But the intuitive idea behind it is: just as a property (like having a mass of 5kg, for instance) can only be had by something real, so a relation can only hold between something real. Otherwise we have no reason to think the relation is instantiated at all.

APPENDIX: BRENTANO'S THESIS

What is normally called 'Brentano's thesis' is the claim that all and only mental states are intentional.

This can be divided into two claims:

- ALL:** if something is mental, then it is intentional
- ONLY:** if something is intentional, then it is mental

It is often said that the following mental states present obvious counter-examples to **ALL**: bodily sensations, emotions of a certain kind (Searle: 'undirected anxiety and depression'). If this were right, then Brentano's thesis is false.

But there are also problems for **ONLY**: it is sometimes said that there is something like intentionality in the natural world: that natural non-mental phenomena can 'carry information' about other things, and so in a sense are 'about' or 'represent' those things. For example, the rings of a tree carry information about the age of the tree; clouds carry information about rain; smoke carries information about fire. This kind of intentionality (what Grice calls 'natural meaning' and Dretske calls 'information') is also claimed to be the natural or physical basis of our mental states. See Grice, 'Meaning' and Dretske, *Knowledge and the Flow of Information*.

Physicalism and Brentano's thesis

Sometimes the name 'Brentano's thesis' is given to certain other views too: for example, to the view that nothing physical is intentional. See Field, 'Mental Representation'. This view is, however, distinct from the view that all mental phenomena are intentional. For that all mental phenomena are intentional does not imply (against Dennett 1969) that nothing physical is intentional; since if physicalism were true, then the mental itself would be physical. What I am concerned with here, however, is the idea that all mental states are intentional, regardless of (a) whether anything else is, or (b) whether anything physical is. So I will put natural meaning/information and physicalism to one side for the time being.

Intentionalism

Intentionalists hold that all mental states are intentional: mentality is sufficient for intentionality. Intentionalists about *consciousness* hold that all conscious states are intentional. (Intentionalists about consciousness are also called 'representationalists'. Examples: Gilbert Harman, Michael Tye, Alex Byrne, David Chalmers.)

Intentionalism about consciousness has been characterised in terms of a number of different theses:

The phenomenal character of an intentional state is identical with its intentional content ('strong intentionalism': Tye)

The phenomenal character of an intentional state supervenes on its intentional content ('weak intentionalism': Tye)

The phenomenal character of an intentional state is determined by its intentional content plus its other intentional features ('impure intentionalism': Chalmers)

Question: how could you be an intentionalist about consciousness without being an intentionalist about every mental state? One answer: the higher-order thought (HOT) theory of David Rosenthal. On this view, a state is conscious when it is the object of a higher-order state. Not all the states that are the objects of higher-order states need be intentional: on Rosenthal's view, sensations are not intentional. But consciousness is explained in terms of the representation of other mental states.

It might be objected that intentionalism has no *explanation* of consciousness, since it assumes that some intentional states are conscious and some are not. But in this respect it is in the same position as non-intentionalism, which says that some mental states are conscious and some are not.

Proposed counter-examples to intentionalism

The typical attempted counter-examples to intentionalism are (a) bodily sensations; (b) emotions and moods.

(a) What should an intentionalist say about bodily sensation? (For example: pain.) Philosophers like Searle say that they are examples of states which are mental, but not intentional. Tye (1995) responds by arguing that pains have propositional content: they represent that the body is damaged in some way. See also DM Armstrong, *Bodily Sensation*.

We should compare this thesis with a simpler intentionalist thesis about sensation: that sensations have intentional objects which are their apparent locations in the body. See MGF Martin, 'Bodily Awareness: a sense of ownership'.

(b) The other apparent counter-example to intentionalism (the '**ALL**' part of Brentano's thesis) is the case of so-called 'objectless emotions' like 'undirected anxiety'. Obviously an intentionalist cannot accept this description of these states! So what should they say? See Crane, 'Intentionalism' §5.

Faculty of Philosophy
Part II paper 2, Philosophy of Mind
Intentionality

Lecture 2: The propositional attitudes

Intentional object and intentional content

The object of an intentional state – the ‘intentional object’ – is what the intentional state is about. This is a definition. Since we can think or talk about things of all sorts of ontological categories – properties, facts, particular events etc. – we should not restrict intentional objects to *particular material objects*.

Intentional states are not simply about objects as such – they often involve representing that an object is a certain way. For example, the belief that Hesperus shines in the sky represents the object (Venus, the evening star, call it what you like) as shining in the sky. It also represents it as *Hesperus* rather than as *Venus* etc.

So to specify an intentional state, it is not enough to specify its object, one also has to specify *the way it represents its object*. This is what I call the *content* of the mental state. It is a consequence of this definition that a state can have an existing content without having an existing object

So the object of an intentional state is **what it is (either in the world or not) that the state is directed on or about**. The content of an intentional state is how its object is represented. These are my definitions of two of the central ideas in the theory of intentionality. (NB not everyone defines these terms in this way, but I claim that my definitions capture what is simple and uncontroversial about their connotations.)

Notice that the object of a belief is what the belief is about; *what* you believe is the content of the belief. The object of seeing, by contrast, is *what is seen*. See Crane, *The Objects of Thought*, chapter 4, for more on this.

What more can be said in general about the notion of content? One popular idea in the 20th century is the idea that content is *propositional* – that is, something that can be true or false. Many philosophers treat intentional states as ‘propositional attitudes’ (the term derives from Bertrand Russell’s 1921 book, *The Analysis of Mind*).

What are propositional attitudes?

Propositional attitudes are those states whose contents are attributed in whole sentences: the belief that Labour will win the next election, the hope that Labour will win the next election, the expectation that Labour will win the next election, and so on.

Propositional attitudes are normally considered to be relations, but relations to **propositions**, not to **objects** (i.e. what the attitude is about). A belief that Pegasus is a winged horse is about Pegasus, but it is not a relation to Pegasus because Pegasus does not exist. But consistently with this, it could be a relation to *the proposition that Pegasus is a winged horse*.

Jerry Fodor (‘Propositional attitudes’) gives two arguments why we should hold that propositional attitudes are relations to propositions:

(i) the first relies on the inferential properties of mental states. If someone believes that p & q, then they will normally believe that p. How could this be explained on a completely non-relational view of mental states, which treats believing-that-p and believing-that-p-&-q as distinct intrinsic properties?

(ii) the second relies on the fact that there are simple inferences we make from ascriptions of attitudes. For example: Jane believes that p, John believes that p, therefore there is something that they both believe. This inference would be validated by a relational structure.

What are propositions?

No particular view of propositions is presupposed by the idea of a propositional attitude, except that: propositions are the bearers of truth value and the meanings of indicative sentences.

Two dominant views of propositions are the Russellian (or neo-Russellian) view, according to which they are composed of objects and properties and relations; and the Fregean view according to which they are composed of what Frege called 'senses': modes of presentation of objects and properties.

Propositionalism

Propositionalism is the view that all intentional states are propositional attitudes.

Propositionalism is a popular and widespread view; but is it true? On the face of it, we often talk as if our thoughts were directed on objects, and this directedness is not described in a propositional way. For example: we talk about wanting an apple, or desiring a long holiday, being frightened of dogs. An apple, a long holiday, dogs – these are either particulars (objects or events) or kinds of things, not propositions, not things that can be true or false.

A propositionalist will typically respond by saying that 'an apple' underspecifies what it is you want: you don't just want *an apple*, but you want (e.g.) *to eat an apple*, and this is an abbreviation for a propositional description of a state of affairs: *that you eat an apple*. This propositionalist 'reduction' of desire is the model for the way we should understand apparently non-propositional intentional states.

But some states of mind resist such 'reductions'. Consider the 'object-directed' emotions such as love and hate. On the face of it, these are attitudes to objects – for example, people. Can there be a propositionalist 'reduction' of love and hate – i.e. is loving someone just a matter of having certain propositional attitudes? No-one has been able to do this in any kind of plausible way.

Love and hate seem to be relations – if so, then you cannot love what does not exist (an atheist should not say, then, that a believer loves God; only that the believer thinks they love God). But there are also non-relational intentional states which seem to resist propositionalist translations. These are those states described by what are known as 'intensional transitive' verbs.

For example (see Graeme Forbes's SEP article 'Intensional Transitive Verbs'):

1. Verbs of depiction or representation: imagine, portray, visualize, write (about), belief (in);
2. Verbs of anticipation: anticipate, expect, fear, foresee, plan;
3. Verbs of desire: prefer, hope (for);
4. Verbs of evaluation: fear, worship, scorn, respect
5. Verbs of requirement: need, require, deserve

These are transitive verbs because they take a direct object (e.g. like 'hit' and not like 'walk'), but they are intensional because they exhibit the marks of intensionality or non-extensionality (see lecture 1).

The point about these verbs is that (unlike with wanting or desiring) there is no obvious propositional reduction of many of them: what are the propositional attitudes which constitute a fear of dogs, for example?

If this is right, there are three kinds of intentional state:

- (a) The propositional attitudes: belief, hope, maybe desire too
- (b) Object-directed relations: love, hate, dislike
- (c) States described by intensional transitives: fear, imagine, expect (etc.)

Faculty of Philosophy
Part II paper 2, Philosophy of Mind
Intentionality

Lecture 3: The non-existent

Propositionalism and the problem of intentionality

One of the supposed advantages of propositionalism is that it gives a straightforward unified account of the central intentional phenomenon of thinking about something that does not exist. Our thoughts are directed on not just real things, but also on things that do not exist – things mistakenly supposed to exist (e.g. by false theories of the world) or things which are pretended to exist or where we know they do not exist (e.g. fictional characters).

This brings us back to the problem of intentionality:

- (A) Intentional states are relations between thinkers and the things thought about
- (B) Relations entail the existence of their relata
- (C) Intentional states can be about things that do not exist

It seems that these cannot all be true together, so which one should we give up?

If you give up (A) you say that intentional states are not relations to their objects. But how can thinking about x fail to be a relation between you and x ?

Some views think that (A) is a non-negotiable commitment: intentionality is essentially relational. An example of this is a common view inspired by Russell's distinction between acquaintance and knowledge by description: we are related to **particulars** (in Russell's case, sense-data) by acquaintance, and to **properties** by description. But strictly speaking, none of our thoughts are about things that do not exist. See Crane, *The Objects of Thought* chapter 1.

But if (C) and (B) are true, then (A) cannot be. And there seem to be many cases of thinking about things which cannot fit into the relational model: in particular, things that are neither cases of relations to properties nor relations to existing objects.

There are two kinds of case to consider: (i) fiction, where we know that what we are talking about does not exist; and (ii) error, where we mistakenly think that something exists when it does not.

Fiction: Consider Hamlet, who did not exist, and contrast him with Richard III, who did. One of the things you are thinking about exists, and one does not.

Error: consider Vulcan, the planet postulated by Le Verrier to explain the perturbation in Mercury's orbit, and contrast it with Neptune, the planet postulated in a similar way. Neptune does exist, Vulcan does not.

Both Hamlet and Vulcan seem to be particular things that do not exist. It seems that claim (C) must be true.

What about (B)? This formulation of the problem of intentionality presupposes the metaphysical thesis that relations can hold only between existing things. Someone might doubt this. But the intuitive idea behind it is: just as a property (like having a mass of 5kg, for instance) can only be had by something real, so a relation can only hold between something real. Otherwise we have no reason to think the relation is instantiated at all.

Could a relation hold between an existing thing and a non-existing thing? To understand this idea we need to say more about existence and non-existence.

Non-existence and existence

What do we mean when we say that something does not exist? Well, what do we mean when we say that something exists? Here I take existence as unanalysable: something exists when it is in the world, part of reality, or real. But you won't understand what I mean by 'in the world' or 'real' unless you also understood what I meant by 'exist'. So I am not *defining* 'exists' in these terms, since all these ideas are too closely related.

In particular, we should not assume that 'exists' or 'in the world' means *in space and time* or that ideas or mental things are not in the world. To say that only spatio-temporal things exist is a thesis about *what exists* not a thesis about *what the nature of existence is*.

We should distinguish, though, between the claim that Pegasus exists in space and time, and the claim that the *idea* of Pegasus exists. When it is denied that Pegasus exists, it is the first of these that is meant, not the second. Of course the idea of Pegasus exists, if any ideas do. But when someone denies that Pegasus exists, they do not deny that the idea exists; they deny that Pegasus does. (See Quine, 'On what there is'.)

Negative existentials

One of the traditional problems about non-existence is the problem of negative existential statements – e.g. 'The round square does not exist' Why is this a problem?

A traditional statement of the problem is given by Moore: 'In saying that there is no such thing as a round square, I seem to imply that there is such a thing. It seems as if there must be such a thing, merely in order that it may have the property of not-being. It seems, therefore, that to say of anything whatever that we can mention that it absolutely is not, were to contradict ourselves: as if everything we can mention must be, must have some kind of being.' (Moore 1953: 289)

Russell says a similar thing: non-existent objects 'all have being, for if they were not entities of a kind, we could make no propositions about them' (1903: §427).

But why believe this? Why should we believe that in order to 'make a proposition' (i.e. say something intelligible about something) that thing has to have some kind of being? A simple view of truth and predication, of the sort we employ in elementary logic, might suggest that this is so: 'Fa' is true just in case the thing referred to by 'a' has the property referred to by 'F' (or is a member of the set of Fs).

But the applicability of this kind of view to our language and thought is exactly the question we are considering! So we cannot assume this view of truth in posing the problem of negative existentials.

Recent developments in theories of names might suggest that there is a specific problem with negative existentials which employ names. If the meaning of a name is just the object referred to (the 'direct reference' or 'Millian' view of names: see Kripke 'Naming and Necessity'), then how can one say something does not exist when using a name for it? For example, how can we meaningfully say that Pegasus does not exist?

Other truths about the non-existent

But negative existential propositions are not the only problematic propositions about the non-existent.

Fiction seems to provide us with one kind of case. Is it true that Sherlock Holmes is a detective, for example? On the one hand, this is what the story says – it does not say that he is an aerobics instructor. But on the other hand, if you listed all the detectives who have ever lived (people who have solved crimes, or who have attempted to solve crimes, or who are employed to solve crimes etc.) then Sherlock Holmes will not be among them.

A natural response to this is to say that the claim that SH is a detective is 'really' a claim that contains a hidden reference to the SH stories. In other words, the sentence 'SH is a detective' should be *analysed* as 'In the SH stories, SH is a detective'. This is clearly true. Can this 'operator' approach to statements about fictional characters give a general account of their truth and falsehood?

The view is initially plausible, but there are problems. One is the problem of how to recognize such 'hidden' operators – how do we know when a sentence requires such an 'analysis'? After all, names in fiction do not bear any obvious signs that make them easily distinguishable from other names.

But there are more specific worries about the operator view. Many claims we make about fictional characters are so-called 'external' claims: i.e. they are not claims made in the context

of telling a story, but in talking *about* the story. 'Napoleon is a character in *War and Peace*'; 'Harry Potter is a character invented by JK Rowling'; 'Sherlock Holmes (For more on fictional characters, see Stacie Friend's excellent *Philosophy Compass* article)

In any case, the representation of the non-existent is not just a problem about fiction – i.e. about cases where we *know* that something does not exist. There are also cases where we mistakenly think something exists when it does not. For example, in the case of 'things' postulated by mistaken scientific theories: the ether, the humours, phlogiston, the planet Vulcan... etc. In these cases, we cannot appeal to facts about 'fictions' in order to explain why claims about these things are true.

How does this leave the problem of intentionality?

Faculty of Philosophy
Part II paper 2, Philosophy of Mind
Intentionality

Lecture 4: Internalism and externalism about intentionality

What does it mean to say that intentionality is not a relation?

It looks as if the response to the problem of intentionality is to deny (A): intentional states are relations between thinkers and intentional objects (call this the thesis that intentionality is a relation).

We need to distinguish two ways of denying that intentionality is a relation:

Weak thesis: not all intentionality is relational; some is, some isn't

Strong thesis: no intentionality is relational

If no intentionality is relational (the strong thesis) then it seems that any thought you have could be had by a brain in a vat (i.e. your brain with the input and output nerves hooked up to a *Matrix*-like set-up which simulates the appearance of a real world outside).

But if the weaker thesis is true, then some intentionality is relational: that is, there are thoughts which you could not have unless the objects of your thought existed. This is a version of the doctrine of *externalism*, the idea that mental states essentially involve objects outside the brain and body of the thinker.

If externalism were true, then it would not be possible for all your thoughts to be had by a brain in a vat. This has significant consequences: for example, certain kinds of global scepticism are undermined; also, since one would not be able to tell my introspecting whether one is having a certain thought, the nature of one's thoughts is not transparent to one's own mind.

Externalism about intentionality

Externalism is *not* the view that the objects of our thought (the things we think about) exist outside our brains and bodies. Almost every philosopher of mind thinks this!

Rather it is the view that the content of our thoughts, what distinguishes one thought from another, *what we are thinking*, is something that is essentially determined by what is outside our brains and bodies. Meaning or mental content is 'not in the head' (but remember Davidson's sunburn analogy, described in the lectures). It is this that the opponents of externalism deny.

Again, the idea is not just that, as a matter of fact, the things outside us cause us to think what we do; internalists accept that there are things outside us, and that they cause us to think about them. The question is, whether we could have had all those thoughts even if none of those things existed? This is why externalism is expressed as the view that thoughts *essentially depend* on what is outside our brains and bodies.

Given all this, why should we believe in externalism?

The argument for Externalism

The central argument for externalism (though note that there are others) is the so-called Twin Earth argument invented by Hilary Putnam in 1975 (see 'The Meaning of "Meaning"').

The argument is based on a thought experiment: suppose that somewhere in this universe there is a planet which is exactly the same as this one, 'Twin Earth', except that what they call 'water' has a complex constitution we can abbreviate to XYZ, rather than H₂O. Each of us has a 'twin' on Twin earth (a molecule for molecule replica), and suppose my twin and I are thinking about the thing we call 'water'. Then the argument goes:

- (a) My twin and I are atom-for-atom identical
- (b) The object of my twin's thoughts is XYZ, the object of my thoughts is H₂O
- (c) XYZ ≠ H₂O
- (d) Content determines object

(e) Therefore: my twin and I have different thoughts

The scope of externalist arguments

How general is Putnam's argument? Is it just an argument about natural kind terms? Some have thought so (Mellor, 'Natural Kinds' in *BJPS* 1978). But in fact, it is more general than this. If you can create a Twin Earth scenario where things seem exactly the same to two physical duplicates but what they think about is different, then you can have an argument for externalism.

Here Tyler Burge's 'social externalism' should be mentioned, since the argument is slightly different. Burge (see 'Individualism and the Mental') considers someone who goes to the doctor because he thinks he has arthritis in his thigh. In fact, this is impossible. But consider a duplicate in another possible world where the word *arthritis* applies to diseases of muscles as well as joints. Then, Burge argues, our poor chap will be thinking something true rather than something impossible; and this is because his thought has a different content.

Burge's argument complicates matters because it appeals to the way we would attribute thoughts and attitudes; but it is not clear that all attributions would go Burge's way. See the section on Burge's argument in Tim Crane, 'All the Difference in the World' *Philosophical Quarterly* 1991

Assessing Putnam's argument

Premises (a), (b) and (c) are just parts of the story (though some have doubted (c), however for the reasons given in the lectures, this is a mistake).

The crucial premise is (d). This is a claim about the relationship between what we think (content) and what we think about (object). It says that what we think determines what we think about. It is a consequence of this that if two people are thinking the same thing then they are thinking about the same thing.

Frege held a similar principle: that the sense ('mode of presentation') of word determines its reference (the thing in the world it is about).

Principle (d) can be disputed. For example, it does not seem to be true of *indexical* thoughts, those expressed using indexical words like 'I' 'here' and 'now'. If I say in London 'it's hot here' and you say in Cambridge 'it's hot here' it seems like there is a clear sense in which we have said the same thing or expressed the same thoughts. But we are thinking about different places. So it seems that what we are thinking does not determine what we are thinking about after all.

There are other objections to externalism: for example, that it makes authoritative self-knowledge impossible. But externalists will see this as question-begging (see Farkas 'What is externalism?').

But is indexical thought eliminable?

One possible externalist response is to say that indexical thought is eliminable in favour of non-indexical thought (or 'reducible' to non-indexical thought). So, for example, 'here' might be eliminable in favour of something like 'the place of this utterance' and '*this* utterance' eliminable in favour of 'the utterance of sentence S at place P and time T' --- or something along those lines.

But such a reduction or elimination is hopeless, for at least two reasons. First, a purely non-indexical (sometimes called 'descriptive') account of reality in thought and language would not be able to distinguish the part of the universe which we are attempting to refer to from a qualitatively indistinguishable section of the universe. If all you had were qualitative descriptions, nothing would make it the case that your descriptions apply to the duplicated things on the other side of the universe. This problem (called 'massive reduplication') was invented by PF Strawson.

The second reason that indexical thought cannot be eliminated is that indexicals seem to be essential for action. This point has been made in a famous series of papers by John Perry. Perry considers the case of a shopper following a trail of sugar around a supermarket, hoping to alert the shopper who is making a mess with the leaking bag. Suddenly he stops, and

realizes that it is *he himself* who is making a mess! The indexical judgement *I am making a mess* (as opposed to the merely descriptive judgement that *the shopper with the leaky bag is making a mess*) is essential to explain why he acts as he does. Indexical thought or belief is essential to action.

Indexicality and externalism

What has been said so far suggests that the phenomenon of indexicality is a fundamentally internalist-friendly phenomenon. Unfortunately, the relationship between indexicality and externalism is actually more complex than this.

Two confusing features: (i) Given the objection above, it is rather ironic that Putnam described his own view as the view that words like 'water' have an 'unnoticed indexical component'. (ii) Note that some philosophers (e.g. Evans, McDowell) claim that indexical thoughts are those thoughts that involve the environment more than any others! What is going on? A tricky issue.

Broadly speaking, there are two kinds of view about indexicality. One view is inspired by Frege, and holds that the sense of an indexical determines its reference. This means, as Frege said in 'The Thought', that the sense associated with the indexical 'I' is different for everyone who utters it. (This view is defended by Gareth Evans, 'Understanding Demonstratives'.)

The second view is that there are two aspects to the meaning of an indexical: (1) the constant linguistic meaning, common to all utterances of a specific kind of indexical (e.g. 'I' = the speaker of this utterance), and (2) the actual contribution to the truth-conditions (or the reference) given by the particular utterance on a particular occasion (e.g. a particular person).

Kaplan calls (1) the 'character' of the indexical, and (2) the 'content'. Characters can be thought of as functions from contexts to contents: uttered in a particular context, the character of 'I' delivers the speaker as its content.

So are indexical thoughts internalist or externalist? If Evans is right, they have externalist propositional Fregean content (recall the principle that sense determines reference).

But if Kaplan and Perry are right, there is an element to the meaning of an indexical utterance which is internal (the 'character'), and one which is external (the 'content' in Kaplan's sense). This suggests a 'two-dimensional' approach to the intentionality of indexical thought.

And in fact, this approach can be generalized to apply to the Twin Earth case: the internalist ('narrow') content is the same across Earth and Twin Earth, but the context is different, yielding different externalist ('broad') contents. See Fodor, *Psychosemantics* chapter 2.

Tim Crane
March 2014