

FRANK JACKSON

WHY WE NEED A-INTENSIONS

ABSTRACT. I think recent discussions of content and reference have not paid enough attention to the role of language as a convention-governed system of communication.¹ With this as a background theme, I explain the role of A-intensions in elucidating one important notion of content and correlative notions of reference.

LANGUAGE, INFORMATION, REPRESENTATION, TRUTH: BACKGROUND

There are many ways of using our bodies to convey putative information about what our world is like. Shrugging and drawing pictures are two examples. An especially powerful and effective way is by using our bodies to produce words and sentences, perhaps as patterns of sound, or as patterns on paper, or increasingly as patterns on computer monitors. Evidence of the power and effectiveness of this method is the fact that a great deal of our knowledge of what our world is like comes from interacting with sentences, most especially from testimony in one form or another. Moreover, our grasp of the meanings of the relevant words and sentences is essential to the way they transmit information. Although we learn some things from the production of words we do not understand – for example, that a certain person is awake or is using words we do not understand, we learn a great deal more from the production of words we do understand. That is why we buy books written in languages we understand and listen to news broadcasts in languages we understand.

But to learn is to narrow the possibilities for how things might be. When I give you information about when the train leaves by pointing to an entry in the timetable, I select that time ahead of the others as the one when, according to me, the train leaves. My gesture represents the time of departure by marking it off from alternatives. Something similar happens when I draw a circle on a map to indicate



Philosophical Studies **118**: 257–277, 2004.

© 2004 Kluwer Academic Publishers. Printed in the Netherlands.

where the treasure is. I give you information about where I think the treasure is by reducing the possibilities for where it might be from anywhere-on-the-map to somewhere-inside-the-circle. More generally, language conveys putative information through being a system of representation that divides how things are being represented to be from the other ways they might be, and our observation about the role of understanding and knowing meanings in greatly facilitating the conveying of information tells us that how words and sentences represent things as being is a semantic property of them – a semantic property put to a highly pragmatic task but a semantic property all the same.

I will use the term “representational content” for how things are being represented to be by a sentence in the communicative role it possesses in consequence of what it means. For brevity, I will use “content” without the prefix when no harm is done.

The connection with truth is made by noting that if a sentence represents things as being a certain way, then things being as they are represented to be is what it takes for the sentence to be true. From this perspective,

“Snow is white” is true iff snow is white

is *a priori* true because the LHS is true iff things are as the sentence “Snow is white” says they are, and the sentence itself, the RHS, is a good way of saying what that is. We can, that is, from the representationalist perspective see that the sentence is true simply by knowing that the sentence named in the LHS is the sentence on the RHS, something that follows from quotation mark conventions.

We have now arrived at the familiar picture of the content of a sentence as the set of possible worlds where the sentence is true. Before we look at some issues that show the need for *A*-intensions, it may be helpful to set a very old issue aside.

Many insist that the content of “There are equilateral triangles” differs from that of “There are equiangular triangles” and yet the two sentences are true at the very same possible worlds. One response is that this objection confuses representational content with meaning; it is the meanings and not the contents that differ. Although content depends *inter alia* on meaning, it is not the same as meaning. An expression in polar coordinates and another in cartesian coordinates can agree in representing that a certain figure is a circle of a given

size and location, can agree in content, but, by virtue of doing the representational job in different ways, count as differing in meaning. But we do not need to adjudicate this issue. For our purposes, what matters is that the right notion of content for understanding how language conveys putative information about how things are is the division among possible worlds one (with a caveat to come shortly).

WORLDS AND REPRESENTATION: THREE COMPLICATIONS

I now come to three complications; three cases where we know that we have to complicate the simple picture of content as a set of possible worlds. Our main interest is in the last case, but the first two are needed to set the stage and for later reference.

1. Much of what we say and believe is egocentric in the sense of having a kind of self-reference built into it. Suppose that I believe and say that there is a blowfly above my head. My belief and my saying are about what my world is like *inasmuch* as it is about how things are with *me*. The contrast is with my belief and saying that there are, at some time or other, blowflies somewhere or other. This example is the tip of an iceberg – we say and believe that it is raining near where we *ourselves* are, that a train will arrive soon in the sense that it will soon be in front of *us*, that *you* (in the sense of the person in front of *us*) have a beard, that *our* great, great grandfathers were not born in Australia, and so on. These claims all concern *inter alia* how things are with speakers and believers, and concern the world inasmuch as speakers and believers need to be in worlds of a certain kind for things to be that way with them.

Sets of worlds in themselves are not enough to capture egocentric content. When I believe that there is a blowfly above my head, I believe something about what my world is like: any world in accord with what I believe will have to contain at least one head and one blowfly. But I also believe something about how things are vis-à-vis myself and this cannot be captured without remainder in terms of what the world I live in is like.²

Although egocentric content cannot be captured in terms of divisions among possible worlds, it is not a challenge to approaching content and representation in terms of divisions among possibilities.

All the plausible treatments of egocentric content on the market work in terms of divisions among possibilities in one way or another. However, the possibilities are not possible worlds. For example, David Lewis's (1979) treatment is in terms of self-ascription of properties where properties are sets of possible individuals. Others prefer to use centred worlds.³ The debate over the best way to handle egocentric content does not matter for our purposes but it will be important later that we be alert to the correlative phenomenon of egocentric properties and this is why I mention the complication.

2. The second complication arises from sentences containing "actually" or equivalents, where the role of "actually" is to shift truth evaluation and reference at a world w back, in one way or another, to the actual world. "Actually P " is true at w iff " P " is true at the actual world. "The actual F is G " is true at w iff the thing which is the F in the actual world is G in w . " a is actually G " is true at w iff a is in w and is G in the actual world.⁴

What do I say about how things are when I say "Actually, there are electrons"? Surely, the very same as when I say "There are electrons". The presence or absence of the word "actually" makes no difference. However, the set of worlds at which the first sentence is true is very different from the set at which the second is true. The first set is the set of all possible worlds; the second is the proper sub-set of that set consisting of those with electrons in them. The same is true for the pair: "The actual President of the USA in 2002 is the son of a former President", and "The President of the USA in 2002 is the son of a former President". They are alike in what they say about how things are but differ markedly in the worlds at which they are true. The same is true for the pair: "The President of the USA in 2002 is the actual son of a former President", and "The President of the USA in 2002 is the son of a former President". They are alike in what they say about how things are but differ markedly in the worlds at which they are true. (The first is true at every world where the President in 2002 in that world is, in the actual world, the son of a former President; the second is true at every world where the President in 2002 in that world is the son of a former President in that same world.)

This means that although there clearly are senses of “content” on which “Actually, there are electrons” and “There are electrons” differ in content, on the sense we are giving the word, the sense we called “representational content”, the sense concerned with elucidating what is being claimed and communicated about how things are, the two sentences have the same content. The same goes for the other pairs.

The solution to the second complication is to construe (representational) content as the set of worlds w which are such that the sentence is true at w under the supposition that w is actual. In the case of “There are electrons”, this set is one and the same as the set at which the sentence is true. In the case of “Actually, there are electrons”, this set is likewise one and the same as the set at which “There are electrons” is true – exactly what we want. Similar points apply to the other examples.

I call the set of worlds w such that S is true at w under the supposition that w is actual, the A -intension of S – “ A ” for “actual”. It makes good intuitive sense that A -intensions should deliver content in our sense. When I communicate how things are using S , I am communicating what kind of world might be the one we are in, might be, that is, the actual one according to S , and the worlds that might be actual according to S are the worlds w such that w 's being actual is consistent with S 's being true. I use C -intension for the set of worlds w such that S is true at w – “ C ” because all but one of these worlds will be counterfactual. A -intensions are like the primary (or notional) intensions and C -intensions are like the secondary (or relational) intensions in David Chalmers.⁵

How is truth at a world under the supposition that that world is the actual world related to truth at a world *simpliciter*? It would be good to have an assurance that there are no problems special to the former, as Ned Block convinced me (which is not to say that he will be happy with what follows). For some sentences, their A -intension is one and the same as their C -intension. Examples are: “There are electrons” and “Some things are square”. For them, truth at a world and truth at a world under the supposition it is the actual world are one and the same. There is a difference between a sentence's A - and C -intensions if and only if the evaluation of the sentence at a world requires reference back to the way the actual world is as a result of

some explicit or implicit appearance of “actually”, or an equivalent rigidification device, in the sentence.⁶ But when this happens, the role of worlds in settling truth values is the *standard* one, the one that applies when it is *C*-intensions that are in question. The only difference is that the value at every world but one depends in part or in whole on how things are at another world. There is no difference in the role of how things are at worlds in settling truth values; the difference is in which worlds are in play. To put the point in terms of a simple example: (a) “The actual *F* is *G*” is true at *w* under the supposition that *w* is the actual world iff “The *F* is *G*” is true at *w*; and (b) what follows “iff” in (a) contains “is true at *w*” and not “is true at *w* under the supposition that *w* is the actual world”.

3. The final complication is one that amounts to a major issue. It can be raised using an old warhorse: water and H₂O. In my view, similar points apply to “Hesperus = Phosphorus” but I will not be discussing proper names here, and one might accept what will be said about water and H₂O and similar cases involving kind terms while insisting that proper names are different.

What we say about what the world is like using the sentence “There is water” is different from what we say about what the world is like using the sentence “There is H₂O”. Otherwise it would not have been a *discovery* that water is H₂O. But the worlds at which the sentences are true are one and the same. (I know that some bite the bullet here. They insist that the claims about how things are, in the sense of what we communicate about what our world is like, when we use the word “water” is one and the same as when we use “H₂O”. On their view, years before Lavoisier people were conveying information to each other about, for example, where there was H₂O when they used the word “water” despite not knowing that water is H₂O. This is not an attractive position. Moreover, the only way to give the position any attraction is to insist on a distinction between what is conveyed and what is knowingly conveyed: before Lavoisier, we conveyed lots of information about H₂O by using the word “water” without knowing that we were. But now we have the problem of giving an account of the content of what is knowingly conveyed: we would merely have shifted the bump in the carpet.)

I think we can resolve the issue by drawing on what we said about sentences containing “actually”. To obtain how things are being represented as being by the use of “water” sentences and “H₂O” sentences, we look to the set of worlds w such that under the supposition that w is the actual world, they are true at w : we look, that is, to their *A*-intensions. And the *A*-intension of “There is water” is very different from the *A*-intension of “There is H₂O”.

The rest of this paper elaborates this suggestion.

IF THINGS WERE EASY

It would be easy if everyone agreed that “water” was a descriptive name or a rigidified definite description, or a term that reference-fixes via some description. (Some say these are three ways of saying the same thing, some make distinctions. For our purposes, the differences if any will not matter.) The relevant description might well be something like David Chalmers’s “watery stuff”, spelt out in the usual way: the clear liquid of our acquaintance that falls from the sky, fills the oceans and so on (Chalmers, 1996, p. 57). This would make “water” equivalent to “the actual watery stuff”, and because the watery stuff is H₂O, and “H₂O” is rigid, this would make “There is water” and “There is H₂O” true at exactly the same worlds. It would give them, that is, the same *C*-intension. We would then explain the difference in representational contents by noting that “There is water” is equivalent to “There is actual watery stuff” and that the representational content of that sentence is its *A*-intension, that is, the set of worlds where there is watery stuff, whereas the *A*- (and *C*-) intension of “There is H₂O” is the set of worlds where there is H₂O.

However, it is far from true that everyone agrees that “water” is equivalent to “the actual watery stuff” or some such. There are three levels of disagreement. One is over one or another spelling out in words of the relevant description. For example, some argue that it is wrong to include being a liquid (at room temperature) in the spelling out on the ground that it might turn out that water is not a liquid. It might have been that although most things that flow have loosely linked molecules, water is one of the exceptions in being made of very small rigid molecular aggregations that slide over each other

easily. In that conceptually possible case, water is no more a liquid than is very fine sand (see White, 1982).

I think we can set this kind of problem aside. Provided that there is some suitable description to be had, we can apply the solution that appeals to *A*-intensions, using that description, whatever it is. Equally, we do not need to worry about suggestions that it may be vague what the right description is, or that it may vary over time or from person to person, or may change with change in conversational context. Points like these would tell us that it may be vague what the right *A*-intension is, that it may vary with time or person, or may change depending on conversational context. They do not tell us that *A*-intensions fail to do the needed job. Indeed, because what we convey about how things are when we say “There is water” is to some extent vague, varies with person and time, changes with conversational context, and so on, it had better be the case that *A*-intensions do likewise.

The second level of disagreement is over the availability of words to do the job, rather than over one or another choice of words. Perhaps “water” reference-fixes on that which has such and such properties – we apply the word rigidly to that which has the properties – but we lack words for the properties. In that sense, there is no question of finding an “actually” definite description or a descriptive name which can be regarded as semantically equivalent to the word “water”. This lack might be regarded as a sort of accident. It so happens that we lack the words to do the job. Or it might be regarded as something that goes deeper; there is some kind of barrier to acquiring the needed words. In either case, we could not give the *A*-intension of “There is water” in words.⁷ Nevertheless, on this second view, it would still be the case that *A*-intensions do the needed job of capturing content. The *A*-intension of “There is water” would be the set of worlds where there is stuff with the properties – the ones we don’t have words for, on the view in question. This set would be different from the corresponding set for “There is H₂O” and would give the content of “There is water”. The existence of an appropriate *A*-intension does not require that we have words for the properties that do the reference fixing.

The third level of disagreement is the one that would bite. On this view, there is no such thing as the property (I’ll use the singular

from now on but “the” property on any plausible view would be a vague disjunction of conjunctive properties) that fixes the reference of “water”. And the point is intended to apply to other similar cases; there is nothing special about the word “water”. This, runs the disagreement at the third level, is something we learn from Saul Kripke (1980) and Hilary Putnam (1975) and, in particular, from the famous arguments from ignorance and error, and from Twin Earth, against the description theory of reference. The message of these attacks on the description theory of reference is that the references of words like “water” and “gold”, as well as that of proper names, do not go by known associated properties; they go by something we users of the words need know nothing about – for instance, a certain kind of causal link. If this is right, then the suggestion that we use *A*-intensions to capture content does not get off the ground. There are no suitable properties with which to construct the required *A*-intensions.

I now turn to arguing that there are such properties. I start with some general remarks designed to create a presumption in favour of the existence of the needed properties.

WORDS AND ASSOCIATED PROPERTIES

In olden days, ships flew flags to let other ships and harbourmasters know about diseases on board. If I remember aright, flying a yellow flag meant that there were cases of yellow fever. This simple system of communication requires a known association between yellow flags flown in the relevant circumstances and yellow fever. The knowledge part of the story is important. Associations we don't know about are no use.⁸ In the case of the number of tree rings telling us about age of a tree, we know of the association between number of rings and age through botanical research. In the case of flags and diseases, we know about the association by knowing about the agreements entered into by seafaring communities to convey information about diseases by using flags.

We noted at the beginning, the great value of language in communicating information about how things are. This requires that we can use language to tell about the distribution of properties much as flags were once used to tell about the distribution of diseases. We

have just noted that the latter requires known associations between flags and diseases arising from agreements entered into by users of the flags. It makes eminent sense that the former equally requires known associations between, in its case, words and properties. How could I successfully use the word “square” to convey to you what something’s shape is if you and I did not know about the association between the word and the relevant shape property? It is also plausible that our knowledge of the associations between words and properties, like that between flags and diseases, arises from agreements to use the words for the properties. The word “square” might have been used for the property we in fact use the word “round” for, but we agreed otherwise and we know that we agreed (implicitly) otherwise.

There are many hard questions in this area – questions that go back to Locke and beyond, and are discussed by, for instance, H.P. Grice (1957), Lewis (1969) and Jonathan Bennett (1976). But three themes stand out from the mass of contentious detail.

First, our ability to use language to convey putative information about our world shows that very many words have known associations with properties. To suppose otherwise would make a nonsense of our ability to communicate views about the distribution of properties using words. Of course, our efforts at communication are not always successful. The confusions that can arise between speakers of English English and US English with words like “trunk” and “take away” are examples. But hypothetico-deductive considerations allow us to have a good deal of confidence in the by-and-large correctness of many of our views about which words are associated with which properties. For example, by using words we succeed in co-ordinating our actions to a remarkable degree. The antecedent probability of the people at this seminar being together in this room at this time is minuscule. The best explanation of our being together in this room at this time requires that we, by and large, agree in the properties we associate with the many words that passed between us in conversation, in notices, in e mails, and so on.

The second theme is that some of these associations are highly resilient. We are able to talk and write sensibly about situations we know for certain do not exist. You and I know that every object to which the word “square” applies has the property of being square.

You and I know that every object to which the word “square” applies has the property of not being a diamond bigger than the Ritz. All the same, we can sensibly discuss, using words, the possible square diamond that is bigger than the Ritz. This is possible because the word “square” retains its association with the property of being square while losing its association with not being a diamond bigger than the Ritz. Only that way could we know what it is whose existence is being denied in the sense of how things would have to be for there to be such a thing. The association with squareness is, that is, resilient.

We take resilience for granted all the time. We discuss all sorts of possibilities, many in the full knowledge that they do not obtain. We wonder how things would be if we won the lottery; how things would be if various laws had been inverse cube ones instead of inverse square ones; how things would be with swampman; and so on. The fact that this involves breaking connections between properties we know for certain obtain does not faze us. We are able to discuss sensibly these known-not-to-obtain possibilities using words, and that means that many, many words keep certain known associations with properties despite our placing them in very unlikely surroundings. The ones they keep are the resilient ones. These are the associations that allow these words to do the important job of keeping track of various properties as we discuss in words this, that, and the next possibility including *very* unlikely ones. This property is very important for philosophers. More than any group in the community, we discuss possibilities we know for sure do not obtain. If it were not for resilience, philosophers would not know what they were talking about much of the time. And even those who insist that our intuitions about swampman, say, are of little value take it for granted that they know what the case is.

Finally, the resilient, known associations are *semantic* properties of the words in the sense that the knowledge in question is part of what’s involved in understanding a language that contains the words. People who understand a language like English are able to use it to report on, learn about and discuss a vast range of possible concatenations of properties. What we learn *inter alia* when we learn English is that the word “square” in English is a word to use when we want to discuss the existence of something with a certain property, debate

whether or not that property is ever found together with being a diamond bigger than the Ritz, or tell someone what their new desk is like, and so on.

Of course, the fact that some words have resilient, “known through the process of coming to understand them”, associations with properties does not imply that all do, and the words “or” and “some” are among the obvious exceptions. But it is plausible that, as rule, adjectives and common nouns do have resilient, known associations with properties, and, in particular, it is *very* plausible that a word like “square” has a resilient, known association with a property – just about anyone reading or hearing this paper knows which property it is, and also knows the answer for words like “round”, “flat” and so on. The claim about adjectives and common nouns more generally is plausible because of the role played by words in these syntactic categories in making claims and conveying information about what things are like.

Of course, the resilience is very far from absolute. Famously, context can shift the property the word “flat” is associated with; the property “nice” is associated with has changed over the years; and the phenomenon of ambiguity means that many words have a range of associated properties and we select, fallibly but with some success, using context, principles of charity and humanity, and so on, the one that is right on some given occasion. The important topic of precisely how to handle the many complexities is not crucial for our concerns. What is important is that, somehow or other, we have spotted the patterns in the midst of all the complexity. If we had not, words would not be as useful as they manifestly are for passing on information about what our world is like, might have been like, would have been like if such and such, and so on.

For any word W that has, for users of language L , a resilient known association with property P , there is the relation between W and any x which has P . This relation might well be called “reference” but does not have to be. It might, for example, be called “denotation”, with “reference” reserved for the relation between W and the set, or perhaps the aggregation, of things that have P , or, if it comes to that, reserved for certain causal relations that users know nothing about.

It is worth noting though that when Kripke objected to the description theory of reference, he did not object to the word “reference” for the relation we have just identified. Consider, for example, the following passage from *Naming and Necessity* (Kripke, 1980, p. 91, my emphasis):

The picture that leads to the cluster-of-descriptions theory is something like this: . . . one determines the reference for himself by saying – “By Gödel I shall mean the man, whoever he is, who proved the incompleteness of arithmetic”. *Now you can do this if you want to. There’s nothing really preventing it.* You can just stick to that determination. If that’s what you do, then if Schmidt discovered the incompleteness of arithmetic you *do refer* to him when you say “Gödel did such and such”.

Kripke is here granting, to put the matter in our terms, that “Gödel” might have had, as a result of sticking to the determination he refers to, a resilient known association with the property of being the person who proved the incompleteness of arithmetic. And if it had, “Gödel” would be a good word to use if one wanted to make a claim about the person that has the property of proving the incompleteness of arithmetic. Kripke is not objecting to the use of the word “reference” in this counterfactual case for the relation between the word and the bearer of the property.

Why do I talk of properties and not descriptions? To identify the issue that has, it seems to me, always been on the table. No-one doubts that “water” refers to water. No-one doubts that “water” is a description in one proper sense. No-one doubts that “water” applies to water and only to water. So, in one proper sense of “description”, *of course* the reference of “water” goes by the satisfaction of a resilient known associated description. Would it then be right to urge that the description theory of reference was beyond question true for “water”, “gold”, “elm”, and so on, and all that’s at issue is its truth for proper names on the ground that they are not, syntactically speaking, descriptions? Clearly not. The issue is about whether the reference of some class of words or other goes by resilient, known associated properties, whether or not the words are called “descriptions” in a grammar book and whether or not we have words for the properties themselves.

I said that we know the resilient known associated properties for some words and gave as examples the words “square”, “flat” and

“round”, but what about the word “water”? Many hold that Twin Earth tells us that there are no such properties for “water”. And what about the words “elm” and “beech”, and more generally words used by experts as well as the folk? What, if any, are the resilient known associated properties for them?

Because of the role of the properties we are discussing in allowing us to make claims about how things are, I am going to talk of them from now on as the *representational* properties. Being square is the representational property for the word “square” in English. Equally, it is the representational property for “actually square”. Our observations about the very valuable role of adjectives in helping us pass on information about what the world is like strongly suggest that, as a rule, adjectives have representational properties. The same goes for common nouns. Words like “water”, “gold”, “elm”, “quark”, and so on, are very useful for saying how things are, co-ordinating behaviour, transmitting information, and so on. The rhetorical question we asked a moment ago now becomes the question of what to say about the famous arguments that suggest that these words lack representational properties? Too much for this paper. I will restrict myself to saying something about how recognising *inter alia* implicit knowledge and egocentric properties shows us how to respond to two of the arguments.

THE LANGUAGE OF THE EXPERTS AND IMPLICIT KNOWLEDGE

Hilary Putnam (1975, p. 226) claims that he does not know what separates beeches from elms but that he succeeds in referring to beeches when he says, say, that he does not know how beeches differ from elms. In our terms, his claim is that he lacks a representational property for the word “beech”. My reply is the far from original one that he does know how they differ from elms: only they are called “beeches” by the experts in his language community. However, the example, and the many cases where we folk co-refer with the experts by drawing on their word usages, raise a number of matters we need to address.

First, we need to distinguish the representational properties of a word for different members of a given language community. My representational property for the word “quark” is very different

from that of a leading physicist. I communicate putative information about the distribution of the property of having the representational property physicists use “quark” for, whatever it is, whereas physicists are communicating putative information about the property, call it, Q , itself.

Secondly, it may be unclear who the experts are (see further, Jackson, 1998). Many folk refer to quarks when they use the word “quark” despite the fact that they do not know whether it is a term from physics or from biology. However, these folk do know that there are experts somewhere or other, and that these experts lie at one end of an information-preserving causal chain that has whoever they themselves borrowed the term from at the other. But now we can specify the representational property for these folk for the word “quark”. It is the property of having the representational property, whatever it is, of the word “quark” for the users that they are borrowing from. And what is the representational property for this group of users for the term “quark”? Either they are the experts, in which case it is property Q (whatever it is), or they are not the experts, in which case it is the property of having the representational property, whatever it is, of the word “quark” for the users that they in turn are borrowing from. And so on.

We have, that is, a recursive story. “But suppose there is no chain of the kind in question?” In that case it is no longer one where it is plausible to say that these folk refer to quarks. When someone tells a story about a chain that ends up in the “wrong” place, say, with the German cheese, they tell a story about how one might use the word “quark” and yet fail to refer to the quarks of physics.

Finally, there is a question raised by the requirement that we know representational properties. It will be urged by some that the kinds of properties we have just been invoking to secure the right reference for “quark” and “beech” in the mouths of the relatively ignorant are not properties the relatively ignorant among folk speakers knowingly associate with “beech” and “quark”. They are too “fancy”; they cannot be representational properties in our sense.

I reply, first, that the folk often say things that make it clear they are aware of these properties and which strongly suggest that they are relying on them to secure reference. People who do not know much physics, and know that they do not know much physics, often

ask questions like, Is it established for sure that quarks exist? And when asked precisely what question they are asking, they answer that they are asking about the things physicists use the word “quark” for. Secondly, often the knowledge the folk have of representational properties is implicit. This is true quite generally, not just of the representational properties in cases of reference borrowing, and explains why the folk (and philosophers) often have trouble finding words for representational properties.

What do I mean by implicit knowledge?⁹ Consider the situation logic students are in before they are given the recursive definition of a wff. Although they cannot specify what it is to be a wff, they typically can reliably classify formulae into wffs and non-wffs. Moreover, they can say for any ill-formed formula what triggers their judgement that it is ill-formed. When presented with “(p v q”, they do not say that they can see that it is ill-formed but cannot say where the problem is. They know exactly where the problem is and how to fix it – add a RH bracket after the “q”. Similarly, they know what changes to a particular wff would make it ill-formed. They are in the following position: for each particular example (of reasonable length), they can say whether or not it is a wff and why, but they cannot give in words a story that covers all cases. The same is true for nearly all of us in our judgements of grammaticality. We can say, for particular examples, whether and why they are or are not grammatical – this is why “behaviourism” about our grasp of grammar is a mistake – but we cannot give the general story in words. Or consider the situation of many bridge players. They cannot state in detail in a way that goes anywhere near covering all the cases, the rules of bridge. At the same time, for any given stage of the game, they can correctly identify the legal moves and what changes to a given legal move would make it illegal and *vice versa*, and in principle (and in practice for the more able ones) critical reflection on their classifications would allow them to write down the rules. In this sense, they know the rules implicitly.

I think we should say the same about the sense in which speakers know the representational properties for words like “water” and “life”. Consider, for example, the discussions engendered by Twin Earth scenarios. There is considerable agreement about what to call “water” and what not to call “water” in these various scenarios. The

impact and importance of the writings of critics of the description theory of reference derive from this fact, and the same goes for the considerable agreement about what to call “gold” – and what to call “Gödel”, “Aristotle” or “life”, if it comes to that. But if speakers can say what to call “water” when various possibilities are described to them, we can identify the representational property for the word “water”: it is the property that, often implicitly, guides them when they say which stuff, if any, in each possibility to call “water” when presented with the various scenarios. When the guidance is implicit, the pattern that underlies the various verdicts will be one they cannot state in words.

This is not to say that, after reflection on their verdicts in the various possible cases, they won’t sometimes be able to make good stab at stating the pattern. I think we can in fact do this for the word “water” when we allow for the vaguenesses. Whether we can do the job for the word “knowledge” is more controversial.

TWIN EARTH AND THE EGOCENTRIC

I said that there is a representational property for the word “water”, and that intuitions about what counts as water in various possible cases do not make the case for the opposite view; instead, they are part of the exercise of making explicit our in part implicit knowledge of the property. Why then have so many thought that the Twin Earth argument shows that there is no such property for the word “water”? I speculate that the key role of the egocentric has been overlooked. Twin Earth arguments seek to show that water, the stuff the word “water” refers to, might lack any and every plausible candidate to be the representational property for the word “water”.¹⁰

The usual list includes: being potable, liquid (or flowing) at room temperature, being clear and odourless in its pure form, and filling the lakes and oceans. But, runs the argument, when a suitable version of the Twin Earth scenario is described to us, we readily agree that the H₂O on Twin Earth might lack every one of these properties and still be water. If this is right, how could appeal to implicit knowledge help one iota?

Now, as it happens, I do not think the agreement is that ready. I think we might equally well have decided that the right thing to say

is that H₂O comes in a water form and a non-water form. We might, that is, have treated the relation between H₂O and water as like that between carbon and diamond. I think, following Lewis (2002, p. 95, n2), that Putnam's advocacy of the answer that H₂O on Twin Earth is water, more or less no matter what form it takes, had the effect of shifting the usage of philosophers away from the folk's pre-analytic indecision between saying that H₂O is water no matter what, and saying that H₂O must have some of the properties we listed in order to count as water, towards a no-matter-what resolution.¹¹ But the question remains, what is the representational property for the word "water" on the no-matter-what disambiguation, the one that holds sway among most contemporary analytic philosophers?

On that disambiguation, the intuition that H₂O on Twin Earth is water comes from the fact that H₂O is the kind that has the listed properties on Earth. If we came to doubt that H₂O is the watery stuff on Earth, if we returned to our pre-Lavoisier epistemic state, we would correspondingly doubt that H₂O on Twin Earth was water. The two opinions go hand in hand. Now this kind is a kind that *we* Earthians are acquainted with and interact with; among its properties is the egocentric one of being the watery-on-Earth kind of *our* acquaintance.¹² It is this egocentric property that secures the reference of the word "water" in our mouths to the H₂O on Twin Earth. Similarly, it is XYZ's lack of this egocentric property that debars it from being water. If we overlooked the fact, noted near the beginning, that much of our thought and talk is egocentric, we would be at a loss to find the representational property for the word "water" on the "no matter what" resolution.

CONCLUSION

You now have before you the case for holding that "water" (on the resolution of vagueness that holds sway among philosophers) is a term that reference-fixes via a property, the property we have been calling the representational property. This completes the case for the solution bruited earlier to why the representational content of, for example, "There is water" differs from that of "There is H₂O" despite the fact that water is necessarily H₂O. The *A*-intensions of the sentences differ, and it is the *A*-intensions that give represen-

tational content. The A-intension of “There is water” is the set of worlds where that which has the representational property exists, and this is plausibly what we convey about how things are when we say that there is water, and it is, in any case, a different set from the set where there is H₂O. Of course, although we have focussed on one example, it is obvious how to generalise.

Coda: The Words That Do Not Have Representational Properties

I have argued that certain adjectives and common nouns have representational properties. To avoid misunderstanding, I should emphasise that it is crucial that these words are in a public language. The adjectives and common nouns of languages of thought (if such exist) cannot possibly have representational properties in our sense. This is because we do not know what the words of the language of thought are. Known associations between words and properties require that we know both the words and the properties. Also, the known associations we have been talking about arise from agreements to use certain words for certain communicative purposes and we did not make agreements, not even implicit ones, about how to use the words of the language of thought. And if we had, we would have had to hire brain surgeons. No doubt the words of the language of thought (if such there be) represent and thus can be paired with properties, but the pairings are not the result of agreements to use them in certain ways.

I mention the matter to highlight that our defence of a kind of description theory of reference for certain adjectives and common nouns is no defence of a description theory for the adjectives and common nouns of the language of thought.

NOTES

¹ See, e.g., Grice, 1957; Lewis, 1969; Bennett, 1976; Locke, 1690, Book III, Ch. II, §2.

² Many have argued this in one form or another; see, e.g., Perry, 1982; Castañeda, 1966; Lewis, 1979.

³ But sometimes centered worlds figure as ways of capturing how content may be a function of location of assertion, whereas we are concerned with the related but distinct question of how to capture the content *per se*.

⁴ The discussion in this section is indebted to Gareth Evans (1979/1985, p. 210) on the epistemic equivalence of “Actually *Q*” with “*Q*”, and to Jason Stanley (1997). Stanley should not be held responsible for the conclusions I draw.

⁵ Chalmers, 1996, §2.4. I originally thought of *A*-intensions as essentially the same as diagonal propositions in Robert Stalnaker but I am now not so sure in view of Stalnaker (2001); what is, in any case, clear is that he would not agree with my use of *A*-intensions. Chalmers’s most recent account of his views (2002) differs from mine in that I am working with a single space of possibilities whereas Chalmers now prefers for certain purposes, *not* all (see 2002, fn. 10) to explicate his distinction using different kinds of possibility.

⁶ In cases that call for centered worlds, the reference back will be to an actual context in the actual world.

⁷ It may be urged that we have at least one word to do the job, namely, “water” itself, for we know that anything “water” applies to has the property of being water. But our problem is that the word “water” tracks H₂O and we are looking for a description that tracks what it is we use the word “water” to say about how things are, and which conjoined with the term “actually” delivers a rigid designator of H₂O.

⁸ In discussions it has sometimes been suggested that knowledge is not essential. True justified belief would do. Nothing here hangs on the difference.

⁹ I use the term with reluctance as it has another meaning in linguistics. My excuse is that I think the sense I give it is pretty much the folk sense.

¹⁰ I mean Twin Earth arguments in their original “remote place in our world (galaxy)” manifestation. Twin Earth arguments in their “other possible world” manifestation tell us about rigidity.

¹¹ *Mutatis mutandis* for XYZ.

¹² And, of course, Putnam (1975) talks of ostension and indexicality.

REFERENCES

- Bennett, J. (1976): *Linguistic Behaviour*, Cambridge: Cambridge University Press.
- Castañeda, H-N. (1966): ‘He*’: A Study in the Logic of Self-Consciousness’, *Ratio* 8, 130–157.
- Chalmers, D.J. (1996): *The Conscious Mind*, New York: Oxford University Press.
- Chalmers, D.J. (2002): ‘The Components of Content’, in DJ. Chalmers (ed.), *The Philosophy of Mind* (pp. 608–633), New York: Oxford University Press.
- Evans, G. (1979): ‘Reference and Contingency’, *The Monist*, 62, 161–189. Reprinted in his *Collected Papers*, Oxford: Clarendon Press, 1985, pp. 178–213.
- Grice, H.P. (1957): ‘Meaning’, *Philosophical Review* 66, 377–388.
- Jackson, F. (1998): ‘Reference and Description Revisited’, in J.E. Tomberlin (ed.), *Philosophical Perspectives*, Vol. 12 (pp. 201–218), *Language, Mind, and Ontology*, Cambridge, MA: Blackwell.
- Kripke, S. (1980): *Naming and Necessity*, Oxford: Basil Blackwell.

- Lewis, D. (1969): *Convention*, Cambridge, M: Harvard University Press.
- Lewis, D. (1979): 'Attitudes *De Dicto* and *De Se*', *Philosophical Review*, 88, 513–543.
- Lewis, D. (2002): 'Tharp's Third Theorem', *Analysis* 62.2, 95–97.
- Locke, J. (1690): *An Essay Concerning Human Understanding*.
- Perry, J. (1982): 'The Problem of the Essential Indexical', *Noûs* 13, 3–21.
- Putnam, H. (1975): 'The Meaning of "Meaning"', in *Language, Mind and Reality* (pp. 215–271), Cambridge: Cambridge University Press.
- Stalnaker, R.C. (2001): 'On Considering a Possible World as Actual', *Proceedings of the Aristotelian Society* 75(suppl.), 141–156.
- Stanley, J. (1997): 'Rigidity and Content', in R.G. Heck Jr (ed.), *Language, Thought, and Logic: Essays in Honor of Michael Dummett* (pp. 131–156), Oxford: Clarendon Press.
- White, S.L. (1982): 'Partial Character and the Language of Thought', *Pacific Philosophical Quarterly* 63, 347–365.

Philosophy Program
Research School of Social Sciences
Australian National University
Canberra, ACT 0200
Australia
E-mail: frank.jackson@anu.edu.au

