

Understanding as immersion

by R. M. Sainsbury (University of Texas at Austin and King's College London)

Abstract: Understanding has often been regarded as a kind of knowledge. This paper argues that this view is very implausible for understanding words. Instead, a proper account will be of the “analytic-genetic” variety: it will describe immersion in the practice of using a word in such a way that even those not previously equipped with the concepts the word expresses can become immersed. Meeting this condition requires attention to findings in developmental psychology. If you understand a declarative utterance, you thereby know what the speaker said in uttering the words she did. The converse is close to true: if you come to know what a speaker of a declarative utterance has said, then normally you understand the utterance. An example of an abnormal situation would be one in which the speaker’s utterance is in a language you do not understand, but you are authoritatively informed, and so know, what the speaker said. The connection between understanding a whole utterance and knowing something can be preserved by alluding to the basis of the knowledge: it should arise from understanding the words in the utterance, and how they are put together. This leads to a fairly uncontroversial conditional equivalence:

- (1) If an utterer of an utterance, U, thereby says that p then:
(X knows this, on the basis of understanding the words in U and how they are put together) iff (X understands U).

The relevant knowledge has no mysteries: it is propositional, typically fully explicit, and is invoked in memory and in reports of speech.

The question to be addressed in this paper is whether we can apply this kind of connection between understanding and knowledge to explain what is taken for granted in the equivalence

above, namely, the understanding of words. I will suggest that we cannot. Under some idealizations, one can formulate necessarily true equivalences between understanding a word and knowing something, but these equivalences are unilluminating. The knowledge is not of the straightforward kind invoked in connection with whole utterances, and it is mysterious how one might come to possess it. An account of understanding words requires an “analytic-genetic” story, an analysis which offers some philosophical illumination while being consistent with what is known about the genesis of participation in linguistic practices.

Before beginning to offer this positive account (§3), I describe some ways in which one might attempt to explain the understanding of words in terms of knowledge (§1), and how Dummett has objected to some such accounts (§2).

I

Knowing what a word means.

The following understanding/knowing equivalence for words cannot be doubted:

X understands a word, W, iff X knows what W means.

What can be doubted is whether this offers any illumination of what it is to be in a state of understanding with respect to a word. The right hand side contains an indirect question: it holds of X iff X can answer the question “what does W mean?”. X can be counted as providing an answer by acting in an appropriate way, for example by stating some fact X knows about W.

Answers of the latter kind will be discussed in the remaining two subsections of this section.

What other kind of action might constitute showing one knows the answer to the question? Using the word as it should be used, that is, with the meaning it has, is indubitably one kind, but this as

it stands offers very little insight into the nature of understanding. The positive account of this paper could be categorized as an attempt to flesh out this answer in a more illuminating way.

Non-homophonic knowledge, e.g. of definitions, or that red is the color of ripe tomatoes.

Homophonic knowledge is knowledge expressed by re-using the very concept expressed by the target word, for example, the knowledge that “red” is true just of red things. That case, with its appearance of some kind of circularity, is discussed two paragraphs hence, and also in §2. The present topic is non-homophonic knowledge. In one kind of case, it is overtly metalinguistic. Such knowledge is arguably available for defined terms: to understand “bachelor” is to know that, by definition, “bachelor” is true of all and only unmarried adult males. Such cases are exceptional. There could be a language with no redundancy (no words definable in terms of others) and the question of what it is to understand a word in such a language would remain. Knowing a definition is obviously unavailable for words with no definition.

We often use words to introduce someone to a new word. We say that *Aristotle* was a famous Greek philosopher, a pupil of Plato, or that *red* is the color of ripe tomatoes, or that Ford now makes the Focus in *vermilion*. In these example, the target word (in italics) appears to be used rather than mentioned: at any rate, overt marks of mention, like inverted commas, are typically absent. In fortunate circumstances, encountering such utterances can result in someone understanding a word they had not understood before. The surrounding words cannot be thought of as definitions, since the target words could be understood even by one who did not know the facts stated, and the facts are not necessary. Moreover, this way of instilling understanding is available only to students who already possess a language, so it could not be the right story in

every case. In short, even if some such knowledge is sufficient for understanding, it is not necessary. A knowledge/understanding equivalence cannot be obtained in these ways.

Homophonic knowledge, e.g. that “Hesperus” stands for Hesperus

In the Davidsonian tradition, the aim is to throw light on the nature of meaning by saying what demands are imposed on a theory of meaning if knowledge of it is to be sufficient for understanding utterances. Knowledge of a theory (which is a set of sentences closed under a relation of entailment) concerns the facts it states: at a first approximation, it is knowledge of those facts. For the knowledge to be genuinely theoretical, it needs to be derived from the theory's axiomatic facts, the facts its axioms state. One who knows the axiomatic facts, and how to derive other theory-endorsed facts from them, can be credited with knowledge of the theory, even if he has not explicitly effected every (perhaps any) such derivation.

In this tradition, a test for the correctness of a theory of meaning is that it should provide the wherewithal to attain correct indirect reports of speech. If someone utters “Nothing travels faster than light”, one who encounters the utterance, and knows a correct theory of meaning for English, should be able to reach the verdict that the speaker has said that nothing travels faster than light. This inevitably gives primacy to the conception of homophony sketched above. One who utters the words “Hesperus is visible” should be reported as having said that Hesperus is visible, not that Phosphorus is visible or that the evening star is visible. Suitable axioms need to take the form “‘Hesperus’ stands for Hesperus”, rather than the form “‘Hesperus’ stands for Phosphorus” or “‘Hesperus’ stands for the evening star”. Axioms of the preferred form are

standardly called homophonic; but we need to refine that notion, since we are concerned with knowledge of facts, and not with knowledge of the truth of sentences.

A *homophonic sentence* is one in which the mentioned expression is reused in effecting its semantic description. In the preferred axiom of the previous paragraph, “Hesperus” was reused in giving the semantic description of this very expression. When we consider knowledge, this notion of homophony is inapplicable. Everyone agrees that both an initially wholly monolingual German speaker, and an initially wholly monolingual English speaker, could come to learn that “rouge”, in French, is true just of red things. The English and German speakers would tend to express what they know in different sentences, neither of which would be homophonic.¹ But what they know is the same as what French speakers know, and French speakers would tend to express the knowledge homophonically. If we wish to apply the notion of homophony to what is known, we cannot make do with a definition that is essentially related to sentences.

If we use a finely-individuated notion of a fact as a potential object of knowledge, it brings with it a notion of content (the same objects of knowledge have the same content). I suggest that a *homophonic fact* is one whose content re-uses the concept expressed by the target expression. English, German and French speakers can have knowledge of the homophonic fact that “rouge” stands for red things. That “Fido” stands for this dog is non-homophonic; that “Fido” stands for Fido is homophonic.

Our question is: can we give an illuminating description of word-understanding in terms of knowledge of homophonic facts? With appropriate restrictions on the kind of word W is², it is hard to deny the following equivalence:

(2) If W expresses the concept C then:

(X knows that W is true just of things that are C) iff (X understands W).

If “Hesperus” expresses the concept of *being Hesperus*, then (X knows that “Hesperus” is true just of things that are Hesperus) iff (X understands “Hesperus”). Can we not regard word-understanding as a species of knowledge, just as we regard utterance-understanding? Though (1) tends to strike people as correct and mildly illuminating, (2) tends to invoke negative reactions: it is “circular”, we never state such knowledge, we cannot use it to teach people the use of expressions, it is at best implicit. I consider one such reaction in the next section.

II

Dummett (1975 “What is a theory of meaning?—I”) has claimed, in effect, that appeal to homophonic knowledge “does not fully display in what an understanding of the object-language consists”, since it presupposes mastery of some language. The idea seems to be that we can give no account of the source of knowledge of homophonic axiomatic facts. This interpretation is supported by a related article, which makes a possibly stronger claim:

the fundamental aim of the theoretical representation [of the practical ability to use a language] is to explain what it is that someone who does not yet know any language has to acquire in order to come to know the given language (Dummett 1976: 70 “What is a theory of meaning?—II”).³

This constraint could be interpreted in a way which makes it too strong: knowledge of a language arises gradually, and what has been learnt in earlier stages is available to help acquire further linguistic knowledge in later stages. But a weaker interpretation of the constraint is plausible, and it targets axiomatic and theoretical knowledge differentially: on the supposition that at least some indefinable concepts can be possessed only as words expressing them come to be understood, at least some axiomatic knowledge needs to be non-homophonic, if acquisition of this knowledge is to be part of acquiring language. This allows for the knowledge involved in understanding whole utterances to be homophonic, provided we have an adequate account of the axiomatic knowledge from which it flows.

If I know that that dog is Fido, or is called Fido, I have non-homophonic knowledge, since “that dog” expresses a concept other than the concept of being Fido. If I know that Fido is called Fido, my knowledge is homophonic. One could well doubt whether the concept of being Fido, which the homophonic content contains, could be acquired by one ignorant of the meaning of “Fido” or some synonym. If it could not, there appears to be no room for an account of how the homophonic axiomatic knowledge could be acquired: how someone without the concept of being Fido could come to know that Fido is called Fido. No parallel doubt arises for the envisaged non-homophonic knowledge: a child could well identify a dog, exploiting some concept other than that of being Fido, and come to learn its name; in the process, coming to possess the concept of *being Fido*. There is no general worry that the learner would first have to understand the target expression in order to have the concept which needs to be exercised in the relevant piece of knowledge.

The general argument could be formulated thus: if understanding is to be illuminatingly described in terms of knowledge, there must be an account of how the relevant knowledge is acquired. If the knowledge requires exercise of the concept expressed by the target expression, that concept needs to be available to the learner before acquiring the knowledge. But in some cases, to acquire a new expression is to acquire a new concept. Hence if understanding is to be illuminatingly described in terms of knowledge, some of this knowledge needs to be non-homophonic.

Dummett's positive account takes him in the direction of a "rich" or "full-blooded" theory of meaning, one which would contain only or mostly theorems stating non-homophonic facts. Such a theory could not satisfy the prime Davidsonian demand on a theory of meaning, that it should deliver optimal reports of what speakers have said by their utterances. We report sayings better the closer we approach homophony, for if we attain it the content reported will not involve concepts other than those exercised by the speaker in her utterance. If we side with Davidson's requirement that a theory of meaning should deliver optimal speech reports, we cannot adopt Dummett's own proposal to fix the problem to which he is drawing attention. I propose that, instead, we should stop focusing on knowledge in our attempt to give an illuminating account of what is involved in understanding words. Instead, we should use the notion of a linguistic practice: what it is for there to be such a thing, and what it is to join it. This will enable us to tell an "analytic-genetic" story, one which explicitly does justice to what is involved in learning a language.

To understand a word is to be immersed in the practice of using it, where the crucial propositional attitude which ensures immersion is intention (the intention to conform to the way others use it). If this basic idea is correct, the prospects are good for the following dependent theses:

- a. Immersion is the basis for a not necessarily inferential account of the source of the knowledge in which understanding complete utterances consists.
- b. The actual language of a population is one whose words have the meanings they have in the practices in which the members of the population actually engage.
- c. Practices are the “use” (the subvenient psychological facts) which determine whether or not a semantic description is correct for a population.

(a) registers an idea that deserves more exploration: perhaps understanding expressions and modes of construction relates to knowledge of what speakers say rather as perceptual experiences relate to knowledge of the perceivable state of one’s environment. In each case, the former both causes and justifies the latter; it does not follow that the former states are states of knowledge. Whereas there is some prospect of putting inferential activities on hold, there is no possibility of not coming to know things about one’s environment if one is perceiving it, or not coming to know what others have said if their speech is audible and in a language one knows.

Each word (semantically individuated) first came into use in some specific act, which I call a baptism (no ceremonial connotations are intended⁴). Other uses in the practice descend from the baptism in a tree-like way: new users are recruited to the practice by exposure to the uses of existing users. I will assume that for each practice, its baptism fixes its semantics, and that this is constant throughout the life of the practice. (A full account would drop this idealization, in order

to speak to the possibility of a word changing its meaning within what, by some criteria, counts as a single practice.) It can be left open at this stage to what extent each practice also encodes information relevant to composition.

As I see it, a practice involves four basic elements:

- baptism; the first event in a practice, and determinant of the word's semantics
- initiation; the introduction of a new user into the practice
- continuation; an act by an existing participant in the practice which belongs to the practice
- departure; one who was at one point a participant is no longer

I present an account of initiation, which I think is the hardest case.

My suggestion is that the following conditions are necessary and sufficient for initiation into a practice, P, of using the word, W:

1. Exposure to one or more correct uses of W in P. (In a correct use of W, the user understands W.)
2. Recognition on the part of the subject that W is a *word*: a reidentifiable element which is not a tool proprietary to the source-user, but is available for use by the subject.
3. A non-accidentally successfully realized conformist intention, which may be initially characterized thus: *to use W as that person does* (or as those people do). The intention is to be caused by exposure to W, via a route of the kind which permits the flow of knowledge.

Most of the rest of the paper concentrates on the content of the conformist intention, showing that there are progressively richer ways of having it, depending upon the level of conceptual and linguistic sophistication of the agent. The idea is to ensure that the analysis makes room for an appropriate genetic story. Fully mature subjects normally have a conformist intention by intending to use a word with the meaning it has in the mouths of those from whom they learnt it. Although this intention refers to the meaning of the word, it can be possessed even when the word is not understood. Immature subjects can have the conformist intention without exercising the concept of meaning.

I encounter a new word in a foreign language that I am studying, and intend to use it as my source did. I summon up the will: *now* I will use it just as the source did, I hereby *resolve* to do so. Do I meet the third condition? There is proper room for doubt whether my activity so much as counts as intending. Merely to wish or hope is not to intend. If I intend to A, arguably A must either be a basic action that is available to me, or else an action for whose successful accomplishment I have at least an outline of a plan which stretches back to basic actions.⁵

However that issue is resolved, the mere willing certainly cannot add up to the non-accidental successful realization of a conformist intention. That would involve, for example, a correct translational hypothesis, put to some kind of test (though this is no guide to how to have a successful conformist intention in learning a first language).

Bringing a conformist intention to a successful conclusion may take time. Psychologists debate whether young children genuinely have concepts like *baby*, given their apparent unwillingness to extend the notion beyond humans (Carey 1988). The debate could as well be put in terms of

whether the children have fully grasped the meaning of “baby”. One view would be that they have not, and that it is acquired gradually. If that is right, the successful realization of the conformist intentions must be regarded as taking some time. Another example, for which there are detailed chronological claims, concerns the definite article. On a plausible view, full mastery of the use of “the F” requires the user to appreciate that two distinct conditions must be fulfilled, which Tomasello (2003: 210) refers to as givenness and specificity. Some F needs to be “given” in the conversational context, in the sense that the speaker knows that it is already familiar to the hearer, and it should be “specific” in that it should be the only potentially relevant F. You should not ask for *the* biscuit if there is more than one biscuit (specificity), or if your audience is not already familiar with the unique biscuit in question (givenness). It seems that children appreciate the specificity of “the” relatively early (3 years), but still make errors relating to givenness at age 4 and 5 (2003: 212). This suggests that months, or even possibly years, may be required from the moment a conformist intention is formed to the time at which it is successfully achieved.

An objection to the necessity of the proposed analysis is that, among both children and adults, there are cases of passive understanding, understanding unaccompanied by any intention to use the expression, and so unaccompanied by a conformist intention. It is said that for some children their first utterance is a complex sentence, revealing that they already understood the component expressions even though they had not used them. Most adults understand some taboo words they would never utter. And one can imagine someone who understands but is determined never to contribute to the relevant practices (perhaps a prisoner of war). For these cases, a weaker intention seems appropriate: the subject intends to be in a position in which she could conform

her use to the others, should she wish to contribute to the practice. The weaker intention seems enough in the more usual cases as well.

An objection to the sufficiency of the proposed analysis is that a subject might conform her intention to a model, in the sense of using the word for just the things the model in fact uses it for, so far as her experience goes, without thereby understanding it. A typical case would be excessive restriction. As far as I know, Daddy uses “toy” just for our toys; conforming my use to his is doing likewise. This would be to under-extend “toy”. One cannot object to the form of the counterexample: it might be that Daddy has limited experience, and never encounters toys other than ours, so his use is itself in a sense under-extended. This sort of case reveals a useful indeterminacy in the notion of *using as he does*. The counterexample adopts a narrow construal, one which may play a role developmentally (that is, at some stages, this is all the conformity aims at). What the analysis also makes room for is a wider construal, one which takes into account possible as well as actual situations, and the community of users, rather than a single representative.

There lurks here a major issue, a contrast between two pictures of language. At one extreme is the picture which reaches its fullest development in classical model theory, shared by people as diverse as Frege, Russell and Quine. Setting aside vagueness, common nouns have determinate extensions: for every object, either the noun determinately applies to it or it determinately does not. That is guaranteed by the determinacy of its meaning (and failures marked as due to vagueness may also be construed as failures of determinacy of meaning). Meaning is seen as “rails invisibly laid to infinity” which fix application for every possible case (Wittgenstein 1953:

§218). At the other extreme is a picture owed to the later Wittgenstein, according to which a serious implementation of the idea that meaning is use must reject the classical picture, and see use itself as something plastic and potentially indeterminate, especially with respect to possibilities not yet envisaged. The present account is more in the mould of this second picture.

To illustrate, consider different approaches to the question whether a swampman (a microphysical duplicate of, say, Donald Davidson, except formed in a lightning flash in a swamp a few seconds ago) is a human being. The classical picture insists that the existing meaning of “human being”, assuming this not to be relevantly vague, must settle an answer to that question. The debate is then whether cladistic approaches have become incorporated into that meaning, or whether the meaning of “human being” is still, perhaps unscientifically, determined by non-historical features. On the Wittgensteinian picture, whether “human being” applies to swampmen or not waits upon how people would come to use that expression, if confronted with swampmen. We all know what is likely to happen. That this conflicts with cladistic criteria suggests that scientists are likely to introduce a special word for their purposes (so that “human being” will be to it as “weight” is to “mass”). In all this, there is no call to say that there has been a change in the meaning of “human being”. Meanings are not fixed by criteria.

The present approach to initiation is more sympathetic to the Wittgensteinian outlook. A conformist intention defers to the community of users in relation to both actual and possible situations. This aim is in one way broader and in another way more restrictive than matching the applications of whatever users one encounters. It is more restrictive, for you are entitled to

disagree even with the person from whom you learnt the word. It is broader, since it relates to other users and to possible as well as actual uses.

IV

In explaining language acquisition, imitation is basic; it gives rise to the most primitive form of the conformist intention. In affirming this, I appear to be in conflict with Chomskyans, like Steven Pinker:

[ungrammatical] sentences show clearly that language acquisition cannot be explained as a kind of imitation. (Pinker 1994: 45)

The idea is that since children will typically not have encountered ungrammatical sentences, we cannot explain their production of such sentences as their simply reduplicating something they have heard. This point cannot be contested, but it is not enough to show that there is no significant room for imitation in an account of language acquisition, notably in the acquisition of the lexicon. Imitation may occur with or without reduplication, and reduplication with or without imitation.

Imitation is a complex notion, to be distinguished from various masqueraders.⁶ An obvious one is mimicry. X mimics Y if X does something similar to Y, and in response to Y's action. This does not require X to have recognized Y's action as of some specific type. A standard example of mimicry is birdsong: there is no need to suppose that a mimicking bird recognizes the intentions of the bird it mimics (the model), or even that the model has intentions. The tendency of human neonates to reproduce certain facial expressions, upon suitable exposure, is also generally classified as mimicry, rather than imitation. Y speaks on the phone, and X, an 11

month old, mimics this behaviour, making babbling noises into a toy phone. X does not speak on the phone. Mimicking speaking on the phone may not be speaking on the phone; this contrasts with imitation.

An imitator, as opposed to a mimic, must associate something the model does with some intentional action type, must intend to produce an action of that type, and this intention must be realized for the imitation to be successful.⁷ The imitation relation holds between an agent, a model, and an action-type: in doing A, X imitates Y; this entails X does A. X can imitate Y with respect to A even if Y does not token A but merely visibly tries unsuccessfully to token it. There is no analog for mimicry. Imitation varies in its degree of sophistication and specificity, depending upon how sophisticated and specific is the action-type in question. This spectrum of specificity is what allows a convincing genetic story: prelinguistic children imitate simple actions, and this enables us to see simple-minded conformist intentions as a route into a linguistic practice. A child can imitate a model's use of an expression without already understanding it, and without possessing semantic concepts.

The genetic story I offer has four chapters. (1) role-switching imitation; (2) goal-identifying imitation; (3) identification of an action type which is not tokened; (3) identification of intended referent. By this last stage, a child is manifesting conformist intentions.

(1) You touch your ear and I (the baby) touch my ear (as opposed to yours). Imitation of this kind appears very early (well before language), but, assuming that it is genuine imitation, and not merely the triggering of an innate motor routine, it is quite sophisticated. The baby needs to

distinguish the action type *touching your ear* from the type *touching own ear*. Only the second type is one we both token, and which I token in imitation of you tokening it. The reflexive action type structures a switch of roles: your ear is the object acted on by you, but does not have this role in my action. Underlying such a switch is appreciation of the fact that I can control my limbs in a way that I cannot control yours. In a standard linguistic situation, roles of speaker and hearer switch. Suppose the adult (the model) draws the baby's attention to something while uttering a word. An imitation of this involves drawing the adult's attention to that thing using (another token of) that word. The relevant action type is *drawing hearer's attention to O using W*, rather than *drawing my attention to O using W*, though of course the adult's action is also of the latter type. In imitation, the baby needs to switch the role played by herself and the adult. Self-ear-touching suggests that this skill is present prelinguistically, and so can be recruited for language acquisition.

(2) In an experimental situation, 14-month old babies are allowed to observe an adult pressing a horizontal surface with her head, upon which a light goes on.⁸ One week later, children are re-introduced to the experimental situation. "67% of the infants duplicated the act" (Melzoff 2005: 59). What act was that? One possibility is that the head-touching was not "duplication", but simply a "natural" response to the apparatus; but 0% of the control group, who had not been allowed to observe the original head-touching, touched the surface with their heads (even when their attention had been drawn to the apparatus). One action type that might have been duplicated is simply that of *touching the surface with own head*. Follow-up studies showed that the relevant action type involved the light going on. In experiments reported by Carpenter et al (1998), there was a delay before the light came on. Most of the 12-month-old and older infants looked towards

the light in anticipation before it came on (Tomasello and Carpenter 2005: 138). This suggests that the relevant action-type was *touching the surface with own head to get light on*. In related experiments, reported by Gergely et al (2002), the 14-month old subjects were divided into two groups (see Tomasello and Carpenter 2005: 140–1). Infants in one group observed the original model situation. Infants in the other group observed an adult with both hands occupied (by a blanket around her shoulders) perform the head-touching routine. Children from the second group were much more likely to touch the horizontal surface with their hands only, and not their heads. It is as if they reasoned: “the model cannot touch the surface with her hands because they are full; she is using her head because that is the easiest thing under the circumstances; but I am more fortunate and can use my hands to do the same thing”. These children can be seen as tokening the action type *touching the surface to get light on*, but not as tokening either of the previous types (in italics). In many of these imitations, there is sensitivity to the goal; this needs to be mentioned in a descriptively appropriate action type.

(3) Meltzoff (2005: 62ff., reporting Meltzoff 1995, 1999 and Bellagamba and Tomasello 1999) found that “young toddlers [18 months] can understand our goals, even if we do not fulfill them”. In one of his experiments, children were placed in three groups: those in one were shown an adult prising apart a dumbbell, those in another were shown an adult attempting unsuccessfully to prise apart a dumbbell, and the third group, the controls, were shown neither actions. Children in the first two groups both prised apart the dumbbell significantly more often than those in the control group. The act type needed in a proper statement of imitation is something like *prising apart a dumbbell*, a type which some of the children did not see tokened in the experimental situation, rather than *trying unsuccessfully to prise apart a dumbbell* or *sliding one’s hands in*

outwards motions over each end of the dumbbell. The conclusion is reinforced by other experiments, in which Meltzoff showed children an adult apparently trying to prise apart an outsize dumbbell, too big for their tiny hands. They used alternative means, for example, placing one end between their knees (Meltzoff 2005: 64). In the relevant action type, goal rather than means must predominate: *getting a dumbbell prised apart*, rather than *using one's hands to prise apart a dumbbell*.

The results were not duplicated when, instead of an adult acting on a dumbbell, an inanimate apparatus was caused to produce similar movements. In these cases, children would engage in prising apart in response to the apparatus producing examples of this, but would not respond in this way to the apparatus making movements similar to unsuccessful attempts to prise apart. Meltzoff explains this by saying that the children can “understand successes, but not failures” (Meltzoff 2005: 64). An alternative explanation is that this is a case of what Tomasello calls emulation: the device reveals to the child a causal potential of the object which the child then exploits. We do not need to describe the case as one in which children try to produce an action type to which they have been exposed (whether by successful or unsuccessful tokenings). This suggests that imitation of actions whose tokenings have not been observed is critical as evidence of imitation. When observed action types are reduplicated, careful controls are needed to ensure that this is not emulation as opposed to imitation.

In the three phases considered so far, children who are at best at the very earliest stages of language learning display language-independent skills of recognizing intentions, and imitating action types essentially containing intention-related content. This should pave the way for them

to bring their sensitivity to intentions to bear in learning language. They need to appreciate that when an adult is directing some stretch of language at them, there is typically some object or feature of their shared environment which the adult is trying to get them to attend to. The object or feature is the object of the adult's intention, and the action-type in the relevant imitation will have a corresponding content.

(4) A standard view is that to learn a word is to associate it with an object. Rico is a border collie with a reported vocabulary of over 200 words: say "duck", and he will search around his owner's apartment and return with the toy duck; say "ball", and he will return with the ball; and so on. There is no doubt that he has associated words with objects; moreover, he apparently learns new word-object pairs very quickly, in a way suggestive of a child's "fast mapping" (Kaminsky et al., 2004). Is Rico like a young child with a vocabulary in the region of 200 words?

A Rico language might suggest the picture that Wittgenstein called Augustinian (Wittgenstein 1953: §1–10). Augustine's own words suggest a further element:

When they (my elders) named some object, and accordingly moved towards something, I saw this and grasped that the thing was called by the sound they uttered when they meant to point it out (cum eam vellent ostendere). (Augustine, *Confessions*, I.8; quoted and translated in Wittgenstein 1953: 2).

Augustine stresses that the learner identifies referential intentions on the part of the model, whereas there is no need to suppose that Rico does this. The role of referential intentions has been borne out experimentally. The idea was to test the hypothesis that a child would associate a word with the most perceptually salient object in the environment when it was uttered. In the

experiment, the child is presented with a new and exciting toy, while the experimenter looks into a bucket, which the child cannot see into, and introduces a new word: “Wow! There’s a modi! What a pretty modi!”. When the child is asked to find the modi, he or she looks into the bucket; he or she does not think that the new toy is a modi (Bloom 2000: 63–4, reporting work by Dare Baldwin (1991, 1993)). I am not aware of evidence that this is how Rico would behave. In looking in the bucket, the child is showing sensitivity to an action type on the lines of *responding to “modi” on the basis of its having been used for whatever the adult had in mind*. Perhaps there is also sensitivity to an action type *responding to “modi” on the basis of its having been used for whatever is in the bucket*, but we would need to see this action type as in play only in virtue of an appreciation of the fact that whatever the adult had in mind is whatever is in the bucket, else we would have no sensible explanation of the fact that the child will behave similarly in a wide range of cases in which she cannot see the intended referent, regardless of local peculiarities like the bucket.

As Tomasello puts it, learning language involves mind-reading. It is reassuring that, as we have seen in (1)–(3), this is a capacity that children have independently of language. It is something they can bring to bear upon the language-learning task.

V

We now have the elements of both the analytic story (from §3) and the genetic one (from §4).

We must put them together, and show that appropriate rules of non-circularity are not infringed.

The analysis requires that entering a word-using practice involves intending to use the word as the others do, and achieving what is intended. Having this intention requires appreciating that the others engage in intentional acts with words: they *use* them. The genetic story assures us that an appreciation of this general kind is present in prelinguistic children, and indeed that more complex language-relevant appreciations, concerning referential intentions, are at work very early in language-learning. Hence the analysis is one which can make room for, and exploit, developmental facts; the same could not be said for accounts of understanding in terms, for example, of homophonic knowledge.

I conclude by considering two versions of the accusation that the analysis is circular. The first clause of the conditions for initiation requires exposure to *correct* uses of the target word. So why not simply short-circuit the complexities, and say that to understand is to come to use correctly? However, “correct” in this context can be explained away using a recursive structure. A baptismal use counts as correct by definition, since, as idealized, it is simply a stipulation concerning what will count as correct use. So a use can count as correct thanks to being derived (by various applications of the conditions for initiation or continuation) from a baptismal one.

Another accusation of circularity takes the form of a dilemma. Either the content of the conformist intention contains explicitly semantic material (e.g. the initiate must succeed in conforming his use to the *correct* use of the word) or else it will be inadequate. The first option induces circularity. The inadequacy mentioned in the second option would be exemplified by someone who took using “dog” as the others do to amount to do more than using “dog” for the

family pet. The criticism is that there is no space between the options: recoiling from either option lands one with the other, and both are unsatisfactory.⁹

By contrast, I suggest that it is not inadequate to have conformist intentions which have no explicit semantic content, and not damagingly circular to allow for intentions with explicitly semantic content.

As I use the notion of explicit content, an intention whose content includes *the color of fire-engines* does not thereby have *red* as part of its explicit content. The distinction has, in effect, been exploited in the contrast between homophonic and non-homophonic knowledge. One not already possessing the concept *red* could come to possess it if he knew the non-homophonic fact that it applied just to things the color of fire-engines. In a similar way, a learner can intend to have intentions with the content of the intentions of a model, while being hazy about what the content is of those intentions. Such cases have already been suggested in the genetic part of the account. A child may identify an adult's intention as to speak of the object in the bucket, while not supposing that the dominant intention guiding that action has as content *the object in the bucket*. A conformer's aim is to act as the model does; that is, to token actions of the most significant of the action-types under which the model's action falls. Action-types are determined by intentional content. So a conformer, in wishing to token whatever is the most relevant action-type with which he is confronted, aims to have the appropriate intentions, whatever they may be. It is thus possible for a conformer to be guided by semantic content, without the intention explicitly possessing semantic content. If the conformer is successfully guided by such content, one has no grounds for saying that he has not done enough.

Many conformist intentions, especially for learners of second languages, will have explicitly semantic content. Circularity in the analysis would exploit that fact, which the present analysis does not. But even if it did, it is not clear that the analysis would be worthless. Consider, for example, the claim that something is red iff it looks red, or that various people form a social group iff they think of themselves as a group. In both cases, we have a concept on the left side of a biconditional, and on the right that same concept figures within some kind of non-extensional operator (“looks”, “think”). These claims may be mistaken, but they are highly substantive, and their truth would have a marked impact on our conception of colors and of the social world, inspiring some kind of “response-dependent” view. Perhaps they should not be thought of as “analyses”, but that would not prevent them being philosophically significant. We could say the same about semantic concepts. To keep to the same style of Euthyphronic oversimplification, the claim that an expression means such-and-such iff people intend it to mean such-and-such may well be (in fact clearly is) incorrect as it stands. But the link it suggests between intentions and meaning may be on the right lines, and could be the beginning of a detailed intention-based account of meaning. So I conclude that even if the present analysis were driven to recognize explicit semantic content in the conformist intentions, that would not in itself show the account to be worthless.¹⁰

¹ Those who hold that all thought, and so all knowledge, occurs in an inner language (“Mentalese”) will not regard that language as English; for these theorists, a sentence which represents knowledge acquired through an encounter with a homophonic English axiom will not be homophonic.

² We should also allow for words like “of”, “if” and “very”, which do not fit the pattern displayed in the text. Everyone knows how to extend something similar to “if”, the extension to “very” may not be problematic, but the wide ranges of ways in which “of” can occur makes extending the account to this case highly problematic.

³ Dummett’s two papers contain a great deal of rich and complex material. The points I select should not be thought of as a proper representation of his views. In particular, his endorsement of a molecular approach (1976: 72) might well lead to a relaxation of constraints on axiomatic knowledge.

⁴ Hughes (2004: 39) shows that the first use of a name might, through a misunderstanding, be governed by conformist intentions, so would not count as a baptism in even a relaxed version of the way this word is ordinarily used. My departure from the ordinary meaning of “baptism” goes further: it is just the first use.

⁵ By a basic action, I mean one which can be performed just by willing. For a normal person in normal circumstances (though not for a paralyzed person) raising one’s arm is an available basic action, but switching on the light is not.

⁶ My account is meant to be sympathetic to, and indeed in large part shaped by, approaches in the psychological literature like Tomasello’s. Tomasello and Carpenter (2005: 134–5) argue for a fourfold classification: in *mimicry*, “an organism reproduces the body movements of another, without an understanding of any goal that might be structuring those movements”; in *emulation learning*, an observer learns about objects through observing another organism interacting with them; and in *goal emulation* “the learner parses the observed action into end and means, but then chooses to ignore the means used by the demonstrator”; in *imitative learning*, the learner classifies the demonstrator’s action under a means-end action type taken to be intended (and not

necessarily realized), and produces a token of that type. As Tomasello stresses, the crucial point is that even prelinguistic children are imitative learners, and my account follows this lead.

⁷ “Imitate” and “successfully imitate” are probably equivalent in ordinary English, but I will make the element of success explicit when this is important.

⁸ The original study was by Melzoff (1988). The present account draws on Melzoff 2005, and Tomasello and Carpenter (2005). 14-months is about the age that children start to acquire language.

⁹ Thanks to Carrie Jenkins for this criticism.

¹⁰ Earlier versions of this paper were presented at the Linguistics Department at the University of Texas at Austin, and at Philosophy Departments at the University of Florida (Gainesville), King’s College London, and the Universities of Bristol, York and St Andrews. My thanks to participants for their comments. I recall in particular suggestions by Michael Jubien, Kirk Ludwig, Greg Ray, Gene Witmer, Keith Hossack, Richard Samuels, Gabriel Segal, Mark Textor, Hilary Kornblith, Finn Spicer, Timothy Williamson, Andrew Woodfield, Barry Lee, Tom Stoneham, Kit Fine, Frank Jackson, Carrie Jenkins, and Crispin Wright.

References

- Baldwin, D. A. (1991). "Infants' contribution to the achievement of joint reference." *Child Development* 62: 1289–1306.
- Baldwin, D. A. (1993). "Early referential understanding: Infants' ability to recognize referential acts for what they are." *Developmental Psychology* 29: 832–843.
- Bellagamba, F. and M. Tomasello (1999). "Re-enacting intended acts: Comparing 12- and 18-month-olds." *Infant Behaviour and Development* 22: 277–282.
- Bloom, P. (2000). *How Children Learn the Meanings of Words*. Cambridge, Mass., MIT Press.
- Carey, S. (1988). "Conceptual differences between children and adults." *Mind and Language* 3: 167–181.
- Carpenter, M., K. Nagell, et al. (1998). "Social cognition, joint attention, and communicative competence from 9 to 15 months of age." *Monograph of the Society for Research in Child Development* 63: serial # 255.
- Gergely, G., H. Bekkering, et al. (2002). "Rational imitation in preverbal infants." *Nature* 415: 755.
- Hughes, C. (2004). *Kripke. Names, Necessity and Identity*. Oxford, Oxford University Press.
- Kaminski, J., J. Call, et al. (2004). "Word learning in a domestic dog: evidence for 'fast mapping'." *Science* 304: 1682–1683.
- Meltzoff, A. N. (1988). "Infant imitation after a one-week delay: Long term memory for novel acts and multiple stimuli." *Developmental Psychology* 24: 470–476.
- Meltzoff, A. N. (1995). "Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children." *Developmental Psychology* 31: 838–850.

Meltzoff, A. N. (1999). “Origins of theory of mind, cognition, and communication.” *Journal of Communication Disorders* 32: 251–269.

Meltzoff, A. N. (2005). “Imitation and other minds: the 'like me' hypothesis”. In S. Hurley and N. Chater (ed). *Perspectives on Imitation* Cambridge Mass, MIT Press, 55–77.

Pinker, S. (1994). *The Language Instinct*. London, Penguin.

Tomasello, M. and M. Carpenter (2005). “Intention reading and imitative learning”. In S. Hurley and N. Chater (ed). *Perspectives on Imitation* Cambridge Mass, MIT Press, 133–148.

Wittgenstein, L. (1953). *Philosophical Investigations*. Oxford, Basil Blackwell.