

Bad Intentions*

Alex Byrne and James Pryor

1. Three Roles for Associated Properties

Let us say that a speaker *associates property P with word T* iff the speaker believes that the referent of T (if it exists) has P.¹ Here are three roles that associated properties might fill.

First, a speaker might be able to know that the referent of word T has certain properties (if it exists), armed only with her understanding of T and a bit of a priori reflection. If so, then let us say that those properties fill *the a priori role* (for word T). For instance, perhaps anyone who understands the word *water* is able to know, without appeal to any further a posteriori information, that *water* refers to the clear, drinkable natural kind whose instances are predominant in our oceans and lakes (if *water* refers at all—we will suppress this qualification from here on). Or, less controversially, perhaps anyone who understands *water* is able to know that *water* refers to a natural kind, or at least that it doesn't refer to an abstract object like a number. Or, almost *uncontroversially*, perhaps anyone who understands *water* is able to know

* Thanks to David Chalmers, Mike Nelson, Scott Soames, an audience in Barcelona, and two anonymous referees for helpful comments.

¹ Two points of clarification. First, the beliefs may be implicit, in the sense that the speaker would only judge that the referent of T (if it exists) has P upon ideal *a priori* reflection. More on this later. Second, for simplicity we will concentrate on singular terms, although the semantic theory (“two-dimensionalism”) that is the topic of this paper is not so restricted. We will treat *water* as a singular term referring to a chemical kind. (We ignore predicative uses, as in *O’Leary has some water in his basement.*)

that it refers to *water*. This last example shows that, plausibly, there will always be *some* property filling the a priori role for word T that its referent uniquely possesses—*being water*, in the case of *water*. What is entirely unobvious is whether speakers have more interesting kinds of identifying knowledge about the referents of words: say, that *water* refers to the clear, drinkable natural kind predominant in our oceans and lakes. At first glance, such cases seem to be the exception, not the rule.

Frege's puzzle provides the second role for associated properties. As Frege pointed out in "On Sense and Reference," sentences like *Bob Dylan is Robert Zimmerman*, unlike the sentence *Bob Dylan is Bob Dylan*, "often contain very valuable extensions of our knowledge." The "cognitive significance" (or "informativeness") of these sentences differ, and this is evidently because the cognitive significance of the name *Bob Dylan* differs from that of the coreferential name *Robert Zimmerman*. To explain these differences in cognitive significance, many philosophers appeal to differences in the properties that speakers associate with the names *Bob Dylan* and *Robert Zimmerman*. When the explanation of why T differs in cognitive significance from other coreferential words appeals to properties that the speaker associates with T, we will say that those properties fill *the Frege role* (for T).

Notice that properties that fill the a priori role need not fill the Frege role. The property *being Bob Dylan* (which is the same as the property *being Robert Zimmerman*), and arguably also the property *being sentient*, fill the a priori role for both *Bob Dylan* and *Robert Zimmerman*. Since these properties are associated with both names, they cannot help explain the difference in cognitive significance between *Bob Dylan is Bob Dylan* and *Bob Dylan is Robert Zimmerman*; accordingly they do not fill the Frege role.

Notice also that properties that fill the Frege role need not fill the a priori role. *Being the author of Mr. Tambourine Man* for example, might fill the Frege role for **Bob Dylan** simply because it is a very well-known a posteriori fact that Dylan wrote *Mr. Tambourine Man*. Alternatively, *being the author of Blow Ye Winds of Morning*, might—at least in principle!—fill the Frege role for **Bob Dylan**, for some speakers. But a speaker cannot *know* that the referent of **Bob Dylan** has this property, because Dylan *didn't* write *Blow Ye Winds of Morning*.

The question of *reference-fixing* provides the third and final role for associated properties. What makes it the case that the name **Bob Dylan**, as we use it, refers to a certain person, namely Robert Zimmerman? (We may assume that this question has a non-trivial answer: it is not a brute fact that **Bob Dylan** refers to Robert Zimmerman.) The much-maligned *description theory of reference* gives one answer to this question. According to this theory, a word T (as used by a particular speaker) refers to an object *o* because the speaker gives a certain kind of reference-fixing authority to some properties P_1, \dots, P_n . This makes T refer to whatever uniquely possesses P_1, \dots, P_n —and that happens to be object *o*. When a speaker gives some of the properties she associates with T this kind of reference-fixing authority, we will say that those properties fill *the reference-fixing role* (for T).

Notice that it does not suffice, for some associated properties P_1, \dots, P_n to fill the reference-fixing role for T, that the referent of T is the unique possessor of P_1, \dots, P_n . For properties to fill the reference-fixing role, the speaker has to (somehow) give them the special reference-fixing authority. (Of course, it is no easy matter to say exactly how a speaker might do this; for present purposes we can leave this tricky question aside.) Nor does it suffice, for P_1, \dots, P_n to fill the reference-fixing role for T, that the referent of T is the unique possessor of P_1, \dots, P_n *and* that P_1, \dots, P_n fill the a priori role for T. Properties can fill the a priori role for T

without the speaker giving them reference-fixing authority. For example, the property *being water* fills the a priori role for **water**, and water uniquely possesses it, but the speaker need not have fixed the reference of **water** to be whatever uniquely possesses this property. For present purposes, though, we can allow the converse. We can assume that speakers have some sort of privileged access to the facts about what properties they have given reference-fixing authority to; and, hence, that any property that fills the reference-fixing role for T also fills the a priori role for T.

Notice that properties that fill the Frege role need not fill the reference-fixing role. We have already seen that a property that fills the Frege role need not be possessed by the referent (for example, *being the author of Blow Ye Winds of Morning*, in the case of **Bob Dylan**). In addition, a property that fills the Frege role need not be uniquely identifying. (For example, perhaps *being a raspy-voiced singer* fills the Frege role for **Bob Dylan**.)

Also notice that properties that fill the reference-fixing role need not fill the Frege role. Presumably someone could introduce **Raspy** as a nickname for Bob Dylan by giving the appropriate reference-fixing authority to the property *being Bob Dylan*. But, as we have seen, this property is associated with any name for Bob Dylan, and so does not fill the Frege role. We will mention another way of making the same point at the end of the paper.

So, with the one exception noted a few paragraphs back, there are no entailments (or, at any rate, no uncontroversial entailments) from filling one role to filling another. Moreover, for a given word T, although we may grant that *some* properties fill the a priori role for T, and that *some* (possibly distinct) properties fill the Frege role for T, it will often be controversial whether *any* properties fill the reference-fixing role for T.

Take *water*, for example. Well-known arguments due to Kripke and Putnam appear to eliminate all the interesting candidates for filling the reference-fixing role for *water*, for example *being the clear, drinkable natural kind predominant in our oceans and lakes*. All that remains are rather unexciting candidates like *being water*. And it is not at all obvious that even this property fills the reference-fixing role for *water*. Of course, there will be *some* story to be told about why *water* has the referent it does; but the reference-fixing story we've been discussing is just one way this might be accomplished.

Given what we've said so far, it should seem rather implausible that a single set of associated properties could fill all three roles for a word. However, according to a sophisticated revival of the classical description theory—the semantic theory known as *two-dimensionalism*—this implausible claim is actually *true*. For any word T, there are associated properties that simultaneously fill the a priori role, the Frege role, and the reference-fixing role. These properties are represented by a word's "primary" or "epistemic" intension: a certain function from possibilities to referents. Many proponents of two-dimensionalism take the theory to be something of a philosophical panacea, resolving a host of puzzles about language and thought—and posing a formidable challenge to physicalism into the bargain.

We think this enthusiasm is misplaced. Two-dimensionalism is incorrect basically for the reasons Kripke and Putnam gave thirty years ago, or so we will argue.

We will proceed as follows. Section 2 sets out the two-dimensionalists' central explanatory apparatus. We focus on David Chalmers' version of two-dimensionalism, in

particular his notion of “epistemic intensions.”² Section 3 examines some considerations Chalmers gives for believing that words have epistemic intensions. We do not think that these considerations are persuasive. Section 4 briefly recapitulates part of the old, familiar case against the classical description theory, which can readily be adapted to apply to two-dimensionalism: Kripke’s arguments from ignorance and error. Section 5 criticizes Chalmers’ response to Kripke; and section 6 examines a second response to Kripke, which we think also fails.

2. Epistemic Intensions

We now give a nuts-and-bolts summary of Chalmers’ version of two-dimensionalism, making a number of simplifications for the sake of brevity.³ In particular, we will ignore complications due to indexicals like *I* and *now*.

An *epistemic possibility* is a hypothesis about how the actual world is, in respects that are left open by all one can know a priori. So, since the population of Barcelona is not an a priori matter, there is an epistemic possibility in which Barcelona has 1.1m inhabitants, another in which it has 1.2m, and so on. On the face of it, epistemic possibilities are distinct from the more common sort of metaphysical possibilities. Since it is not a priori that water is H₂O, there is an epistemic possibility in which water is, say, XYZ, and not H₂O, even though there is no such metaphysical possibility. In fact, Chalmers argues that the metaphysical possibilities and the epistemic possibilities are the *same* (minor qualifications aside); we will not be discussing this part of his view.

² Two-dimensionalism has also been defended recently by Frank Jackson (see especially his 1998a). See Byrne 1999 for some discussion of Jackson’s account. It has much in common with Chalmers’ account, although there are some differences. For reasons of space, we cannot examine the differences here.

An *epistemically possible world* or *scenario* is a “maximal” epistemic possibility: an epistemic possibility E^* that a priori implies all the other epistemic possibilities that are compossible with it.⁴ (Henceforth, when we speak of “epistemic possibilities” we mean these “maximal” epistemic possibilities.)

The *epistemic intension* of a word T is a function from epistemic possibilities to objects that exist “in” or according to those epistemic possibilities. According to Chalmers, the value of T ’s epistemic intension at some epistemic possibility E may be determined by considering instances of the following schema (where t is replaced by the word T , and n is replaced by a singular term that appears in the specification of E):

(Turns-Out) If E “turns out to be actual”—that is, if it correctly represents how the world really is—then t will turn out to be n .⁵

³ For more careful expositions, see Chalmers **XXXX**, Stalnaker 2001, and Pryor 2003.

⁴ In other words: E^* does not leave any facts a priori open. For any epistemic possibility E , it is either (i) a priori that if E^* is correct, then E is correct; or (ii) a priori that if E^* is correct, then not- E is correct; or (an arguable qualification) (iii) a priori that if E^* is correct, there is no determinate fact of the matter whether E is correct.

As will become clear shortly, the epistemic possibilities Chalmers officially defines his intensions over are specified in a very limited vocabulary (roughly: that of physics and phenomenology). Accordingly, it is entirely unobvious that these official epistemic possibilities are maximal in the sense just explained (not that Chalmers thinks otherwise).

⁵ We assume that the conditional in this schema is the material conditional. We also assume that whenever E a priori implies that n exists, n appears in the specification of E . (Compare the “identifying descriptions” in Chalmers and Jackson 2001, p. 318.)

If (and only if) anyone who understands this conditional can know it to be true, perhaps after a bit of a priori reflection, then T's epistemic intension will be a function that maps E to the object n.⁶ We will say that a speaker can *identify the referent of T in E* if and only if the speaker can know some instance of this schematic conditional to be true, in the way just described. In general, Chalmers supposes that for any word T, and any epistemic possibility E, anyone who understands T can identify its referent in E. As Chalmers and Jackson put it: an understanding of T by "a suitably rational subject bestows an ability to evaluate certain conditionals of the form $E \square C$, where E contains sufficient information about an epistemic possibility and where C is a statement using [T] and characterizing its extension, for arbitrary epistemic possibilities" (2001, p. 324, footnote omitted).⁷

Here are two examples Chalmers gives of identifying the referent of a word in an epistemic possibility:

What about a term such as 'Hesperus'?...Let scenario W_2 be one on which the brightest object visible in the evening is Jupiter, and where the brightest object visible in the morning is Neptune. For all we know a priori, W_2 is actual. If it turns out that W_2 is actual, then it will turn out that Hesperus is Jupiter. So when

⁶ On this formalization, n would always have to exist, because it is the value of a function that exists. Epistemic possibilities can however *say* that certain objects exist, which do not and indeed could not exist. This raises interesting questions about the ontology of epistemically possible objects. We cannot pursue those questions here, so we will assume for the sake of argument that they can be answered in a way that makes the notion of an epistemic intension coherent.

evaluated at W_2 , the intension of ‘Hesperus’ returns Jupiter. If it turns out that A [the epistemically possible world that happens to describe the actual world correctly] is actual, then it will turn out that Hesperus is Venus. So when evaluated at A, the intension of ‘Hesperus’ returns Venus. (2002b, pp. 145-6)

And similarly:

Let W_3 be a ‘Twin Earth’ scenario, where the clear, drinkable liquid in the oceans and lakes is XYZ. For all we know a priori, W_3 is actual. If it turns out that W_3 is actual, then it will turn out that water is XYZ. So when evaluated at W_3 , the intension of ‘water’ returns XYZ. If it turns out that A is actual, then it will turn out that water is H_2O . So when evaluated at A, the intension of ‘water’ returns H_2O . (2002b, p. 146)

(These reflections about what will turn out to be the case are supposed to be a priori.)

So, according to Chalmers, the epistemic intension of *Hesperus* differs from that of *Phosphorus*, and the epistemic intension of *water* differs from that of *H₂O*. He thinks that, in general, two words T_1 and T_2 have the same epistemic intension if and only if a speaker competent with these words can know that they are coreferential, armed only with her understanding of the words and a bit of a priori reflection. Since Chalmers takes synonyms to be words with the same epistemic intension, he also holds that if a speaker understands a pair of

⁷ The quotation actually concerns “concepts,” rather than words, but clearly Chalmers and Jackson would allow the substitution. (See their footnote 7, p. 323.)

synonyms T_1 and T_2 , she can know that they are coreferential. This claim is controversial, but we will not discuss it further here.

The apparatus of epistemic intensions is not supposed to be the whole semantic story, of course. *Two* “semantic dimensions” are required, because a word T also has a more familiar sort of intension: the function that takes a metaphysically possible world w to the referent of T at w . (That is, the function that delivers T 's referent in possibilities taken to be *ways the world could, counterfactually, have been*, not *ways the world may be, for all one knows a priori*.) Since, necessarily, *Hesperus is Phosphorus*, and *water is H_2O* , the “metaphysical” or “counterfactual” intension of *Hesperus* is the *same* as that of *Phosphorus*, and similarly for *water* and *H_2O* .

We said that the epistemic intension of a word is determined by which instances of the schematic conditional like (Turns-Out) a speaker will be able to know a priori. What enables a speaker to know which of these conditionals are true, and which are false? We can think of matters like this. For any word T a speaker understands, there are some properties P_1, \dots, P_n that the speaker associates with T . More precisely, the speaker believes that the referent of T possesses P_1, \dots, P_n in the following sense: upon ideal a priori reflection, the speaker would judge that the referent of T possesses P_1, \dots, P_n . These properties are such that the value of T 's epistemic intension at epistemic possibility E is the object described by E as being the unique possessor of P_1, \dots, P_n (if there is such an object). According to Chalmers, any such properties will fill all three of the roles we mentioned earlier: the a priori role, the Frege role, and the reference-fixing role.

To illustrate these points, take *water*. Going by the previous quotation, the associated properties are something like: *being clear, being drinkable, being in the oceans and lakes*. Since these properties fill the a priori role for *water*, someone who understands *water* doesn't need any

further a posteriori knowledge to know that the referent of *water* is clear, drinkable, and found in the oceans and lakes.

Since these properties fill the Frege role, the cognitive significance of a sentence like *Water is H₂O* derives from the fact that *being clear, being drinkable, being in the oceans and lakes* are associated with *water*, and some other properties are associated with *H₂O*. We can also put this point in terms of the epistemic intensions of *sentences* (functions from epistemic possibilities to truth values): *Water is water* is cognitively *insignificant* because its epistemic intension is the constant function that take every epistemic possibility to the True; *Water is H₂O* is cognitively significant because its epistemic intension takes certain epistemic possibilities to the True and others to the False.

Lastly, since these properties fill the reference-fixing role for *water*, *water* refers to the unique clear, drinkable stuff found in the oceans and lakes. If some epistemic possibility says that XYZ is the unique stuff with these properties, then the epistemic intension of *water* will map that epistemic possibility to XYZ.

As is apparent from the above quotations, a competent speaker is supposed to be able to identify the referent of a word like *water* in an epistemic possibility E that is specified *without using the word water* (or cognate expressions): for example, a possibility in which “the clear, drinkable liquid in the oceans and lakes is XYZ.” So competent speakers are not supposed simply to know that *water* refers to *water*. Likewise, a competent speaker is supposed to be able to identify the referent of *Bob Dylan* in epistemic possibilities that are specified without using the name *Bob Dylan*. Let us put this point by saying that speakers are supposed to have *substantial identifying knowledge* of the referents of *water* and *Bob Dylan*. This amounts to having an ability

to evaluate, upon ideal a priori reflection, all instances of the schematic conditional (Turns-Out), where E is specified without using the word T (or any of its cognates).⁸

In fact, Chalmers thinks that speakers will be able to identify the referents of their words in epistemic possibilities specified in *strongly* reductive terms. The only expressive resources required, he thinks, are the language of a complete fundamental physics and a language suitable for describing “the phenomenal states and properties instantiated by every subject bearing such states and properties, at every time” (Chalmers and Jackson 2001, p. 319), plus a few other bells and whistles.⁹ Given an epistemic possibility E specified using only these vocabularies, speakers who understand *Bob Dylan* and *water* are supposed to be able to identify the referents of *Bob Dylan* and *water* in E.

For our purposes, though, two-dimensionalism need not be viewed as having such strong reductive aspirations. We will just take the two-dimensionalist to be employing *some kind of* “reductive” specification of epistemic possibilities, leaving the details open.

⁸ Three points of clarification. First, substantial identifying knowledge is intended to be *nothing more* than the ability to evaluate these conditionals. Chalmers and Jackson stress that this ability need not always be underwritten by the subject’s explicit judgments about what properties T’s referent possesses; often, they think, the ability will precede and explain any such judgments (see Chalmers and Jackson 2001 §3, and Jackson 1998b, pp. 211-12).

Second, at the beginning of this paper we said that a speaker “associates P with T” iff the speaker believes that the referent of T (if it exists) has P. We should emphasize (again) that these beliefs may be ones that the subject has only “implicitly,” in virtue of having the ability to evaluate these conditionals.

Finally, for our purposes, nothing turns on exactly how the notion of “ideal a priori reflection” is to be understood.

⁹ The additions are a “‘that’s all’ statement” (Chalmers and Jackson 2001, p. 317) and a “‘you are here’ marker” (p. 318).

So far we have given the impression that a word has a *unique* epistemic intension. However, the two-dimensionalist can and typically will allow that a word's epistemic intension often varies from speaker to speaker. For example, Chalmers considers the case of two speakers, who "have been exposed to different forms of water: one has only been exposed to water in liquid form (knowing nothing of a solid form), and the other has been exposed only to water in solid form (knowing nothing of a liquid form)" (2002b, p. 174). It might be, he says, that the epistemic intension of *water* as used by the first speaker differs from the epistemic intension of *water* as used by the second, although of course both intensions return the same referent at the actual world, viz. H₂O. Again, to accommodate Putnam's *elm/beechn* example, Chalmers says that the epistemic intension of *elm* as used by the botanical ignoramus is (roughly) given by the description *The tree the experts call 'elm'* while the epistemic intension of *elm* as used by the experts is something quite different (2002a, pp. 617-8). Since none of our arguments turns on the assumption that words have unique epistemic intensions, for convenience we will mostly ignore this kind of alleged variation.

It is a strong and unobvious claim that speakers have substantial identifying knowledge of the referents of words like *water* and *Bob Dylan*. Why think that they do? If speakers *don't* have this identifying knowledge, then they won't be in a position to know what these words refer to in the two-dimensionalist's reductively specified epistemic possibilities, and hence the corresponding epistemic intensions won't be well-defined. So another way of asking our question is: why think that words like *water* and *Bob Dylan* have epistemic intensions of the sort we have described?

3. Chalmers' Argument from Examples

There are various arguments for two-dimensionalism in the literature. Some of these are of an indirect sort: two-dimensionalism should be accepted because it neatly solves some theoretical puzzles—for example, puzzles about the necessary a posteriori.

Other arguments are more direct. For example, Chalmers says that two-dimensionalism is suggested naturally by armchair reflection on what speakers would say if the world turned out one way rather than another. In this way speakers can manifest their alleged abilities to identify the referents of words in different epistemic possibilities.

In section 2, we quoted a few passages from Chalmers that are intended to exhibit a fragment of the epistemic intensions of *Hesperus* and *water*. In the second of those passages, Chalmers suggests that the following conditional is a priori (that is, it can be known to be true by anyone who understands it, after a priori reflection):

(CDL) If it turns out that XYZ is the clear, drinkable liquid in the oceans and lakes, then it will turn out that water is XYZ.

The *Argument from Examples*, as we will call it, starts with a discussion of *water* and other examples, and concludes that “[t]he intuitive characterization of epistemic intensions using the heuristics I have given here makes a strong prima facie case that expressions have epistemic intensions” (2002b, p. 146).

Now if (CDL) really is a priori, then this would help support a crucial part of the two-dimensional package, namely that speakers have what we called *substantial identifying*

knowledge of the referents of their words. And perhaps with further argument, it can be used to support all the main two-dimensional claims. So, is (CDL) a priori?

Offhand, it can appear that way. Admittedly, given the present state of chemical knowledge, it would be somewhat deviant to utter (CDL) assertively. But we can imagine some chemical ignoramus justifiably doing so, and it seems that the sentence she utters is *true*. (After all, it *hasn't* turned out that XYZ is the clear, drinkable liquid, etc. When the ignoramus discovers that water is H₂O, she does not have to *retract* her earlier assertion of (CDL).) Presumably the ignoramus could even *know* that (CDL) is true. And since she is ignorant, it might seem that in order to know that (CDL) is true, she only needs to understand it.

But consider an obvious fact about water, for instance that it is the liquid that comes out of taps in Barcelona, and consider the conditional:

(TAPS) If it turns out that XYZ is the liquid that comes out of taps in Barcelona, then it will turn out that water is XYZ.

Just as before, it would be somewhat deviant to utter (TAPS) assertively, but a chemical ignoramus might well do so. Again as before, it seems that (TAPS) is true, and that the ignoramus could know this to be so.

However—we may safely presume—(TAPS) is not a priori. When we imagine the ignoramus assertively uttering (TAPS), we are tacitly assuming that she knows some obvious a posteriori facts about water, in particular that it comes out of taps in Barcelona. If we imagine instead that the ignoramus has never heard of Barcelona, or that she believes that wine comes out of Barcelona taps, then she will have no justification for uttering (TAPS).

This should raise considerable suspicion concerning the status of (CDL). The fact that it is easy to imagine a scientific ignoramus knowing (CDL) to be true does not support the claim that the conditional is a priori. For in imagining the ignoramus to know (CDL), we may be tacitly assuming that she knows some obvious a posteriori facts about water, in particular that it is the clear, drinkable liquid in the oceans and lakes.

Now it might be insisted that even if we explicitly *stipulate* that the ignoramus has no a posteriori knowledge (beyond that conferred by her knowledge of English), it will *still* be plausible that she would be justified in accepting (CDL). Well, perhaps. Our only point at present is that one tempting but superficial reason for thinking that (CDL) is a priori collapses on further examination.

Having made this defensive point, it is time to go on the offensive. We think that familiar arguments from Kripke and Putnam show quite conclusively that no conditional like (CDL) is a priori. More-or-less equivalently, they show that speakers do not ordinarily have substantial identifying knowledge of the referents of words. Let us turn then to these arguments; in particular, to Kripke's arguments from ignorance and error.

4. Kripke's Arguments from Ignorance and Error

Kripke's examples of the names *Cicero* and *Feynman* support the view that a speaker can be a competent user of a name despite lacking substantial identifying knowledge *because of ignorance*. "[M]ost people," Kripke says, "when they think of Cicero, just think of *a famous Roman orator*, without any pretension to think that either there was only one famous Roman orator or that one must know something else about Cicero to have a referent for the name" (1980, p. 81). Similarly, the man in the street may use the name *Feynman* to refer to Feynman, even

though “[w]hen asked he will say: well, he’s a physicist or something. He may not think this picks out anyone uniquely” (1980, p. 81).

Kripke’s story about Gödel and Schmidt supports the view that a speaker can be a competent user of a name despite lacking substantial identifying knowledge *because of error*. In Kripke’s story, speakers use the name *Gödel* to refer to Gödel, even though the achievements they ascribe to Gödel—discovering the incompleteness of arithmetic—were really performed by the unfortunate Schmidt. The properties that speakers associate with the name *Gödel* are rich enough to uniquely identify someone, but the person they uniquely identify is not the name’s referent.

Notice that the Gödel/Schmidt story doesn’t *just* teach us something about speakers who have false beliefs. It teaches us something stronger, namely that for *any* speaker (not just speakers in error), the properties the speaker associates with the name *Gödel* do not fill the reference-fixing role (with possible exceptions for those who named Gödel in the first place, or for the property *being Gödel*). For consider some competent user of the name *Gödel* who *knows* that it refers to the individual having such-and-such properties—say, the property of discovering the incompleteness of arithmetic. Since the speaker knows that the referent of *Gödel* has this property, she believes it does, and hence she *associates* this property with the name. However, if this property filled the reference-fixing role, then in a nearby possible world in which the Schmidt story is true, and the speaker uses the name *Gödel* with the same semantic intentions, we should find that *Gödel* refers in her mouth to Schmidt. But for typical speakers, this is just what we don’t find. Typical speakers may know that Gödel discovered the incompleteness of arithmetic, but they don’t give that property reference-fixing authority.

Similarly with *water*. Most competent users of this word *do* know that it refers to the kind that has certain properties, for instance the kind many instances of which are clear, drinkable, liquid, and found in the oceans and lakes. But considerations just like those in the Gödel /Schmidt case show that these associated properties don't fill the reference-fixing role for *water* (as the word is used by these speakers).

The arguments from ignorance and error are concerned with a typical user of a name who has picked it up from someone else. It might be argued that associated properties will at least be needed to fill the reference-fixing role in the special case where a speaker explicitly introduces a name. This is an issue too large to be properly discussed here, but it is worth noting that the matter is not at all straightforward. Take the case of ostensive definition. Suppose a speaker sees a dog, and dubs him *Checkers*. There will be many properties that pick out the dog (e.g., *being the dog the speaker is looking at*). But it is unclear whether the speaker needs to associate any such properties with the word, and a fortiori unclear whether any such properties fill the reference-fixing role. And even if an associated property *does* fill the reference-fixing role, it might be the unexciting property of *being this particular dog, Checkers*. The speaker may be able to name the dog *Checkers* simply because she stipulates, *of the dog she is seeing, that it is the referent of Checkers*. These sorts of associated properties seem ill-suited for Chalmers' purposes; they will not provide the kind of substantial identifying knowledge that he is looking for.

In any case, concentrating on typical speakers, the arguments from ignorance and error seem to show that associated properties do not fill the reference-fixing role for words like *water* and *Bob Dylan*. Therefore these words, as used by typical speakers, do not have epistemic

intensions. Chalmers, though, is quite unimpressed by these arguments, for reasons that we will now examine.

5. Chalmers' Response to Kripke

The core of Chalmers' response to the arguments from ignorance and error is expressed in the following passage:

Does this argument against the description theory [i.e., Kripke's arguments from ignorance and error] yield an argument against the intensional framework I have been outlining? It seems clear that it does not. This argument works with a conception of descriptions on which they correspond to linguistic expressions. When Kripke argues that the descriptions that the speaker "associates with" the name cannot fix reference, he always invokes linguistic descriptions that the speaker associates with the name, or at least explicit descriptive beliefs of the speaker. But the intensional framework is not committed to the idea that descriptions always correspond to linguistic expressions; in fact, at least part of the motivation of the framework comes from an independent rejection of this idea. And the intensional framework is not even committed to the idea that the intensions associated with a name correspond to explicit beliefs of the speaker. So there is no clear argument against the intensional framework here.

In fact, Kripke's central method of argument seems to be obviously compatible with the intensional framework. A proponent of this framework could cast the argument strategy as follows. We want to show that for a given name N and description D, 'N is D' is not a priori. To do this, we consider a specific

epistemically possible scenario *W*. We then reflect on a question such as the following: ‘if *W* turns out to be actual, will it turn out that *N* is *D*?’ And we find that the answer is no. If so, the epistemic intension of ‘*N* is *D*’ is false in *W*. So ‘*N* is *D*’ is not a priori.

On this interpretation, when we think about the Gödel/Schmidt case, for example, we are tacitly evaluating the epistemic intension of ‘Gödel’ at a world as specified in the example. When we consider that world as an epistemic possibility, it reveals itself as an instance of the epistemic possibility that Gödel did not discover incompleteness. That is, we find that the epistemic intension of ‘Gödel’ does not pick out the prover in this world, it picks out the publisher. If so, the epistemic intensions of ‘Gödel’ and of ‘the man who discovered the incompleteness of arithmetic’ are distinct. (2002b, p. 169)

There are three main points in this passage. First, as Chalmers puts it a little later, “Kripke’s arguments suggest that the epistemic intension of a name such as ‘Gödel’ cannot be precisely captured in a linguistic description. But they do nothing to suggest that the epistemic intension does not exist” (2002b, p. 170). Second, even if the description is linguistically expressible, the speaker might associate it with the name only tacitly or implicitly—if asked for an explicit statement of what properties she was using to identify the referent of *Gödel* in various epistemic possibilities, she might be at a loss. Third, Kripke’s own methodology is best viewed as a way of *revealing* or *articulating* a name’s epistemic intension, rather than as demonstrating that the name *has no* epistemic intension. (See Chalmers 2002a, n. 11; Chalmers and Jackson 2001, pp. 326-7; and Jackson 1998b, pp. 212-14.)

Take the first point first. Suppose that a speaker has seen a proof of the first incompleteness theorem, and retains a capacity to recognize the proof visually. Let us further suppose that the speaker associates some properties with the name *Gödel* that she cannot fully articulate in English. The best she can do is something like *the man who discovered this proof*, uttered while demonstrating the appropriate pages in *From Frege to Gödel: A Source Book in Mathematical Logic*. But *what* she has in mind is an essentially visual way of thinking of the proof; her demonstrative utterance (let us suppose) doesn't fully articulate it.

Kripke's story about Schmidt straightforwardly shows that these associated properties do not fill the reference-fixing role for *Gödel*, as it is used by this speaker. For, in the story, the person who possesses these properties is Schmidt; yet the speaker's word *Gödel* refers to Gödel. Further, any other linguistically inexpressible properties that a speaker might associate with *Gödel* would also appear to be subject to a Schmidt-type objection. So although Chalmers is right to claim that the properties that a speaker associates with *Gödel* need not be linguistically expressible, this does not seem to help at all in fending off Kripke's argument from error.

Neither does the first point help in fending off Kripke's argument from ignorance. Perhaps many ordinary speakers have some complex idea of ancient Rome, derived from *Ben Hur* and *Gladiator*, that resists complete articulation in English. The properties they associate with *Cicero* might be gestured at with phrases like *a famous orator from that place*, while demonstrating various sword-and-sandal scenes. So the properties they associate with *Cicero*, let us suppose, are also not linguistically expressible. But obviously these properties do not pick out the referent of *Cicero* uniquely. And since there is no reason to suppose that ordinary speakers associate other linguistically inexpressible properties with *Cicero* that *do* pick out the referent uniquely, the argument from ignorance stands.

Turn now to Chalmers' second point, the one about explicitness. This point does indicate a need for caution: we should not conclude that an ordinary speaker does not have substantial identifying knowledge of the referent of *Cicero* just because the speaker herself cannot explicitly state it. Substantial identifying knowledge might make its presence known through the speaker's disposition to apply the name, rather than through her verbal reports (see note 8). But it seems clear that even when we take this into account, ordinary speakers are often impressively ignorant about the referents of names like *Cicero*. Their poor performance on history exams is due to their *lack* of knowledge of Cicero's life and times, not to its *implicitness*. And in any case, even if we found that speakers did associate properties with *Cicero* that were both suitably reductive and uniquely identifying, the Gödel/Schmidt example shows that they usually won't fill the reference-fixing role.

Finally, let us turn to Chalmers' third point, that Kripke's examples help to *reveal* or *articulate* a name's epistemic intension, rather than demonstrate that it doesn't have one. Recall that the epistemic intension of *Gödel* is supposed to represent a speaker's ability to identify the referent of *Gödel* in some *reductively specified* epistemic possibility—her *substantial identifying knowledge*. Possibilities specified as ones containing *Gödel* don't count. So Chalmers seems to be saying that evaluating Kripke's example involves identifying the referent of *Gödel* in a reductively specified epistemic possibility. Kripke gives us an epistemic possibility in which certain people (bearing the names *Schmidt* and *Gödel*) do certain things, and given that epistemic possibility, "the epistemic intension of 'Gödel' does not pick out the prover [of the theorem] in this world, it picks out the publisher." If that is the right account of Kripke's Gödel/Schmidt example, then it would *not* show that *Gödel* lacks an epistemic intension. Rather, the

example would presuppose that *Gödel* has an epistemic intension, and it would help us to articulate what that intension is.

However, we think this is a misrepresentation of Kripke's example. *Kripke* does not offer any reductive specification of the Gödel/Schmidt possibility. As we read Kripke, he is asking us, in effect, to imagine a situation in which a speaker who falsely believes that *Gödel* refers to the man who discovered the incompleteness of arithmetic, nonetheless *uses Gödel to refer to Gödel*. The situation is specified in terms of properties that *Gödel* does and does not have. ("A man named 'Schmidt'...actually did the work in question. His friend Gödel somehow got hold of the manuscript..." (Kripke 1980, p. 84).) Kripke's point is that that situation is perfectly coherent, which makes it plausible that the referent of the name *Gödel* is not fixed by properties that the speaker associates with it. There is nothing at all in Kripke's description of the example to support the view that a competent user of the name *Gödel* can identify its referent in some *reductively specified* epistemic possibility. So there are no grounds here for thinking that anyone who understands *Gödel* has substantial identifying knowledge about its referent.¹⁰

Kripke *could have* presented his story about Gödel and Schmidt without using the name *Gödel*, but by using instead an expression his readers knew to apply to Gödel, such as *the member of the Institute for Advanced Study who starved himself to death*. If he had done so, though, he would have been exploiting shared a posteriori identifying knowledge about Gödel, rather than identifying knowledge that we all have just in virtue of understanding *Gödel*.

It may be that anyone who understands *Gödel* will know some substantial conditions that are *necessary* for being the referent of *Gödel*. (Conditions, that is, that can be specified without using *Gödel* or its cognates.) For example, perhaps anyone who understands *Gödel* knows that it

refers to a sentient being, if it refers at all. If so, the conditional *If it turns out that there are no sentient beings, then it will turn out that Gödel does not exist* will be a priori. Competent speakers may also know some interesting *sufficient* conditions for being the referent of *Gödel*. For example, if competent speakers know the necessary condition just mentioned, then they will also know that if there is exactly one sentient being and if *Gödel* refers, then it refers to this sentient being. If so, the conditional *If it turns out that Gödel exists and there is exactly one sentient being, then it will turn out that Gödel is this sentient being* will be a priori.

It is however a considerably stronger claim that competent speakers know substantial conditions that are *both necessary and sufficient* for being the referents of their terms; that is, substantial identifying knowledge. We do not think that examples like Kripke's provide any support for this strong claim—even if the knowledge is allowed to be linguistically inexpressible and implicit. Our ability to identify referents in such examples typically owes to the fact that the examples are specified in *non-substantial* terms, or are specified using descriptions that the referents are *known a posteriori* to satisfy, or both. So these examples do not give us reason to attribute substantial identifying knowledge. Rather, as Kripke says, they show that competent speakers do *not* typically need to have such knowledge.

It seems to us, then, that Chalmers' three points do not deflect the force of Kripke's *Cicero* and *Gödel* examples.

6. The Metalinguistic Response to Kripke

After responding to Kripke's arguments, Chalmers turns to the question of whether the epistemic intension of a name like *Gödel* can “at least be approximated by a linguistic description.” “This

¹⁰ Soames forthcoming offers similar objections.

is not compulsory for the intensional framework,” he says, “but it can at least be enlightening to look” (2002b, p. 170).

To answer this question, one needs to consider: when speakers use a name such as ‘Gödel’ or ‘Feynman’ in cases such as those above [i.e. when they are mistaken or ignorant], how do they determine the referent of the name, given sufficient information about the world? For example, if someone knows only that Feynman is a famous physicist and that Gell-Mann is a famous physicist, how will external information allow her to identify the distinct referents of ‘Feynman’ and ‘Gell-Mann’? The answer seems clear: she will look to *others’* use of the name. Further information will allow her to determine that members of their community use ‘Feynman’ to refer to a certain individual, and that they use ‘Gell-Mann’ to refer to a different individual. Once she has this information, she will have no problem determining that her own use of ‘Feynman’ refers to the first, and that her own use of ‘Gell-Mann’ refers to the second.

This suggests that if we want to approximate the epistemic intension of the speaker’s use of “Feynman” in a description, one might start with something like ‘the person called ‘Feynman’ by those from whom I acquired the name.’ It certainly seems that if relevant information about others’ uses is specified in an epistemic possibility, then this sort of description will usually give the right results. The same goes for the ‘Gödel’ epistemic possibility. In all these cases, it seems that a name is being used *deferentially*: in using a name, the speaker defers to others who use the name. (2002b, pp. 170-1, note omitted).

We think that this should be Chalmers' official response to the epistemological arguments, not the three points discussed above. The moral of the arguments from ignorance and error is that if two-dimensional account of names is to be workable, then the epistemic intension of a name like *Gödel* cannot be given by any sort of "famous deeds" description, like *the man who discovered the incompleteness of arithmetic*. Instead, the epistemic intension has to be given by something like the description *the person called 'Gödel' by those from whom I acquired that name*.¹¹

As Chalmers notes, Kripke discusses various proposals along these lines, for example "By 'Gödel' I mean the man *Jones* calls 'Gödel'" (1980, p. 92). These proposals are said either to fall to a Gödel/Schmidt type objection, or else to violate Kripke's noncircularity requirement.

And, again as Chalmers notes, his own proposal seems to be vulnerable to Gödel/Schmidt type objections. To accommodate cases where the speaker mishears or misremembers the name, Chalmers tries a "closer approximation":

Perhaps 'The referent of the relevant name used by the person from whom I acquired the antecedent of my current term 'Gödel'' would do a better job. But no doubt there would be further counterexamples...But as in all these cases, the most this shows is that any such approximation is imperfect. One refutes these approximations by evaluating the epistemic intension in certain epistemic

¹¹ Note that Chalmers is allowing *semantic* specifications of epistemic possibilities here: for example, descriptions of the referential history of *Gödel* as used by a certain speaker. On his official reductive account, these are dispensable. See also Jackson 1998b, pp. 209ff. There is a large literature discussing metalinguistic proposals of this sort. Nelson 2002 gives a useful overview.

possibilities and showing that the approximation give the wrong results; so this sort of argument does nothing to show that the epistemic intension does not exist. (2002b, p. 171)

Indeed there are further counterexamples. Suppose a speaker baptizes Gödel with the name *Gödel*, and so doesn't acquire the name from someone else. Further suppose she forgets that this is so. Her use of *Gödel* still refers to Gödel. But if the property *being the referent of the relevant name used by the person from whom she acquired the antecedent of her current term Gödel* filled the reference-fixing role, then—since she never acquired *Gödel* from anyone—her use of *Gödel* would not refer.

Even if we assume that the metalinguistic proposal can be fixed up to avoid obvious counterexamples, at least three objections remain.

The first objection is that the metalinguistic proposal imposes unreasonable demands on understanding a word. Admittedly, the proposal does not require speakers to have explicit metalinguistic beliefs (see note 8 and the preceding section). But it does require competent speakers to have *an ability to evaluate conditionals* whose antecedents contain sophisticated semantic vocabulary, like *the antecedent of my current term n, the referent of a term as used by speaker S*, and so on. One would have thought, on the contrary, that the ability to speak and understand a language comes *first*: understanding words is a precondition of such conceptually sophisticated abilities, not the other way around.¹²

The second objection is that metalinguistic properties, even if they do fill the reference-fixing role, will not generally fill the Frege role. Consider an example. Imagine that Rosa Zola

was taken to the high school prom by Robert Zimmerman; despite having a wonderful evening, they lost touch after graduation. One day many years later Rosa hears an assertive utterance of *Bob Dylan is Robert Zimmerman*. She is utterly astonished and delighted. The information she gains is highly non-trivial, and it leads her to contrive a reunion with her old prom date. Two-dimensionalism promises an account of this: the cognitively significant information Rosa gains is the contingent proposition that the D is the Z, where *being D* determines the epistemic intension of *Bob Dylan*, and *being Z* determines the epistemic intension of *Robert Zimmerman*. However, on the metalinguistic proposal this contingent proposition is something of the following sort:

The referent of *Bob Dylan* as used by those from whom Rosa acquired that name is the referent of *Robert Zimmerman* as used by those from whom Rosa acquired that name.

And this information is patently *not* the news that excited Rosa and moved her to action. What excited her, we may suppose, is the information that the singer of *Mr. Tambourine Man* is the person she dated in high school. Rosa gained this information by hearing *Bob Dylan is Robert Zimmerman* because she associates *Bob Dylan* with the property *being the singer of Mr. Tambourine Man*. The associated properties that play the Frege role will be “famous deeds” properties like this one, not metalinguistic properties. And as we’ve already argued, these “famous deeds” properties will typically be ill-suited to play the reference-fixing role. For typical

¹² For further discussion, see Braun 1995 and Soames forthcoming.

speakers, those kinds of properties will always be vulnerable to Gödel/Schmidt-type counterexamples.¹³

So, adopting the metalinguistic proposal prevents epistemic intensions from solving Frege's problem, and thus removes one of the advertised advantages of two-dimensionalism. (See Chalmers 2002a, pp. 622-4; cf. Jackson 1998a, p. 76.)

The third objection is both the simplest and the most fundamental: the metalinguistic proposal is unmotivated. Before trying to make the metalinguistic proposal work, better reason is needed for thinking that the referent of a word *must* always be determined by the speaker's giving reference-fixing authority to some associated properties. In our opinion, no adequate case for this assumption has yet been supplied.

References

Braun, D. 1995. Katz on names without bearers. *Philosophical Review* 104: 553-76.

Byrne, A. 1999. Cosmic hermeneutics. *Philosophical Perspectives* 13: 347-83.

<<http://mit.edu/abyrne/www/CosmicHermeneutics.html>>

Chalmers, D. 2002a. The components of content (revised version). In *Philosophy of Mind: Classical and Contemporary Readings*, ed. D. Chalmers. Oxford University Press.

<<http://consc.net/papers/content.html>>

Chalmers, D. 2002b. On sense and intension. *Philosophical Perspectives* 16: 135-82.

<<http://consc.net/papers/intension.html>>

¹³ This is the "other way," alluded to in section 1 above, of making the point that properties that fill the Frege role need not fill the reference-fixing role. Here we are indebted to Thau (2002, Ch. 3).

Chalmers, D. 2004. The foundations of two-dimensional semantics. *This volume, chapter XX*.

Chalmers, D., and F. Jackson. 2001. Conceptual analysis and reductive explanation.

Philosophical Review 110: 315-60. <<http://consc.net/papers/analysis.html>>

Jackson, F. 1998a. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press.

Jackson, F. 1998b. Reference and description revisited. *Philosophical Perspectives* 12: 201-18.

Kripke, S. 1980. *Naming and Necessity*. Blackwell.

Nelson, M. 2002. Descriptivism defended. *Noûs* 36: 408-35.

Pryor, J. 2003. Varieties of two-dimensionalism.

<<http://www.princeton.edu/~jimpryor/papers/Varieties2D.pdf>>

Soames, S. forthcoming. Reference and descriptions. In *The Oxford Handbook of Contemporary Analytic Philosophy*, ed. F. Jackson and M. Smith. Oxford University Press.

Stalnaker, R. 2001. On considering a possible world as actual. *Aristotelian Society Supplemental Volume* 75: 141-56.

Thau, M. 2002. *Consciousness and Cognition*. Oxford University Press.