

# On Sense and Intension

David J. Chalmers

Department of Philosophy

University of Arizona

Tucson, AZ 85721.

[chalmers@arizona.edu](mailto:chalmers@arizona.edu)

## 1 Introduction

What is involved in the meaning of our expressions? Frege suggested that there is an aspect of an expression's meaning — the expression's *sense* — that is constitutively tied to the expression's role in reason and cognition. Many contemporary philosophers have argued that there is no such aspect of meaning. I think that Frege was closer to the truth: one can articulate an aspect of meaning with many, although not all, of the properties that Frege attributed to sense. This aspect of meaning is what I will call an expression's *epistemic intension*. I will argue that epistemic intensions can serve as quasi-Fregean semantic values, and that this claim is undefeated by the major contemporary arguments against Fregean sense.

The simplest aspect of an expression's meaning is its *extension*. We can stipulate that the extension of a sentence is its truth-value, and that the extension of a singular term is its referent. The extension of other expressions can be seen as associated entities that contribute to the truth-value of a sentence in a manner broadly analogous to the way in which the referent of a singular term contributes to the truth-value of a sentence. In many cases, the extension of an expression will be what we intuitively think of as its referent, although this need not hold in all cases, as the case of sentences illustrates. While Frege himself is often interpreted as holding that a sentence's referent is its truth-value, this claim is counterintuitive and widely disputed. We can avoid that issue in the present

framework by using the technical term ‘extension’; in this context, the claim that the extension of a sentence is its truth-value is a stipulation.

Different sorts of expressions have different sorts of extensions. By the stipulation above, the extension of a singular term is an individual: the extension of ‘France’ is a particular country (France), and the extension of ‘Bill Clinton’ is a particular person (Clinton). Analogously, the extension of a general term is plausibly seen as a class: the extension of ‘cat’ is a particular class of animals (the class of cats). The extension of a kind term can be seen as a kind: the extension of ‘water’ is a particular substance (water). The extension of a predicate can be seen as a property or perhaps as a class: the extension of ‘hot’ is a particular property (hotness) or a particular class (the class of hot things).

The extension of a complex expression usually depends on the extensions of the simpler expressions that compose it. This applies most obviously to the truth-values of sentences. For example, ‘Sydney is in Australia’ is true, and it is true because the extension of ‘Sydney’ (a particular city) is located in the extension of ‘Australia’ (a particular country). ‘Michael Jordan is short’ is false, and it is false because the individual who is the extension of ‘Michael Jordan’ does not have the property that is the extension of ‘short’. The same applies to typical complex expressions other than sentences: for example, it is not implausible that the complex singular term ‘the greatest cricket player’ has an extension (Don Bradman), and that this extension depends on the extensions of its parts.<sup>1</sup> Certain expressions (such as those involving belief attributions, to be discussed later) may provide an exception to this rule, but it is plausible that the dependence holds in a very wide range of cases.

There are various complexities here, and there are corresponding choices to be made. For example, some terms (e.g. ‘Santa Claus’) appear to have no referent: in such a case, one might say that they lack extension, or one might say that they have a null extension. In some cases (e.g. ‘greatest’ above), it appears that the extension of an expression can depend on context: for this reason, we may wish to assign extensions to

---

<sup>1</sup> In this paper I will assume for ease of discussion that descriptions are true singular terms. Nothing important turns on this: the main points of this paper carry over to a Russellian analysis on which definite descriptions are complex quantifiers. One simply needs to apply the framework to the appropriately regimented logical form.

expression tokens, or to expression types in contexts, rather than to expression types alone. In general, the truth-value of a sentence will be determined by the extensions of its parts within a regimented *logical form* (or perhaps some other underlying form), along with corresponding principles for determining truth-value of a regimented sentence from its logical form and the extensions of its parts; and the regimented sentence may look quite different from the original sentence, with different basic constituents and a more complex structure. Different semantic theories may assign extensions to expressions and logical forms to sentences in different ways.

Many of these complexities will not concern us here. The discussion that follows should be general over many specific proposals concerning logical form, extensions of simple terms, and compositional determination. I will simply take it for granted that sentences have a logical form and contain simple terms that have an extension; that these simple terms compose complex terms, which compose the sentence; and that the extension of a complex expression (including a sentence) is at least in many cases determined by its logical form and the extensions of its parts.

To clarify terminology: on my usage, an “expression” is any linguistic entity that has an extension or that is a candidate for extension. For ease of discussion, I will say that when an expression is of the sort that is a candidate for an extension, but appears to lack extension, it has a null extension. A “term” is any expression other than a sentence. Complex expressions are expressions (including sentences) that are composed of other expressions. Simple expressions are expressions that are not composed of other expressions.

## **2 Sense and extension**

A simple and attractive view of meaning ties the meaning of an expression to its extension. On such a view (e.g., Mill 1843; Salmon 1986), the meaning of a simple term is its extension, and the meaning of a complex expression is determined by the extensions of its parts. On the strongest version of this view, the meaning of a complex expression is its extension. On a slightly weaker version, the meaning of a complex expression is a

complex structure involving the extensions of the simple terms that are parts of the expression. Either way, all meaning is grounded in extension.

There are three traditional reasons for doubting this simple view of meaning. First: some simple terms (such as ‘Santa Claus’ and ‘phlogiston’) appear to lack extension. On the view above, these terms will lack meaning, or they will all have the same trivial meaning, corresponding to the null extension. But intuitively it seems clear that these terms have *some* meaning, and that their meanings differ from each others’. If so, then meaning is more than extension.

Second: in some sentences, the role of a word in determining a sentence’s truth-value appears to go beyond its extension. This applies especially to sentences about beliefs and related matters. For example, it is plausible that ‘John believes that Cary Grant is an actor’ could be true, while ‘John believes that Archie Leach is an actor’ is false, even though ‘Cary Grant’ and ‘Archie Leach’ have the same extension. If so, then either the truth-value of the sentence is not determined by the meanings of the terms, or there is more to meaning than extension.

Third: there is often more than one term referring to the same thing. In such cases, the terms often seem intuitively to have different meanings. Witness ‘Hesperus’, the ancients’ name for the evening star, and ‘Phosphorus’, their name for the morning star. Or take ‘water’ and ‘H<sub>2</sub>O’, both of which refer to the same substance. If ‘water’ and ‘H<sub>2</sub>O’ refer to the same thing, and if reference is all there is to meaning, then ‘water’ and ‘H<sub>2</sub>O’ refer to the same thing. But intuitively, ‘water’ and ‘H<sub>2</sub>O’ have *different* meanings. If that is so, then reference cannot be all there is to meaning.

In “Über Sinn und Bedeutung” (1892), Frege lays out the central issue roughly as follows. The sentence ‘Hesperus is Hesperus’ is *trivial*. It can be known a priori, or without any appeal to experience. The knowledge that Hesperus is Hesperus requires almost no cognitive work at all, and gives us no significant information about the world. By contrast, the sentence ‘Hesperus is Phosphorus’ is *nontrivial*. It can only be known a posteriori, by appeal to empirical evidence. The knowledge that Hesperus is Phosphorus requires much cognitive work, and gives us significant information about the world.

As Frege put it, ‘Hesperus is Phosphorus’ is *cognitively significant* whereas ‘Hesperus is Hesperus’ is not. Intuitively, this difference in cognitive significance reflects

a difference in the meanings of ‘Hesperus’ and ‘Phosphorus’. When a subject comes to know that ‘Hesperus is Phosphorus’, what she learns depends on what she means by ‘Hesperus’ and by ‘Phosphorus’. It seems plausible that the subject learns something different when she learns that Hesperus is Phosphorus than when she learns that Hesperus is Hesperus. If these two claims are correct, then ‘Hesperus’ and ‘Phosphorus’ have different meanings, and meaning involves more than extension.

If meaning involves more than extension, then what is the further element? Frege held that in addition to extension (or reference), an expression also has a *sense*. ‘Hesperus’ and ‘Phosphorus’ have the same referent, but different senses. ‘Water’ and ‘H<sub>2</sub>O’ have the same referent, but different senses. ‘Cary Grant’ and ‘Archie Leach’ have the same referent, but different senses. For all such cases, the intuitive difference in cognitive significance among pairs of terms such as these is reflected in a difference in the terms’ senses.

The notion of sense has a number of important features, which I discuss in what follows. The discussion is intended as a broad and informal outline of a Fregean view, rather than as a faithful representation of every feature of Frege’s own view. More precise versions of some of the following claims will be given later in this paper.

(1) *Every expression that has an extension has a sense.*

In “Über Sinn und Bedeutung”, Frege concentrated mostly on the senses of names, holding that all names have a sense. It is natural to hold that the same considerations apply to any expression that has an extension. Two general terms can have the same extension and different cognitive significance; two predicates can have the same extension and different cognitive significance; two sentences can have the same extension and different cognitive significance. So general terms, predicates, and sentences all have senses as well as extensions. The same goes for any expression that has an extension, or is a candidate for extension.

(2) *Sense reflects cognitive significance.*

The central feature of sense is that it is tied constitutively to cognitive significance. In the case of singular terms, Frege set out this connection as follows: two referring expressions 'a' and 'b' have different senses if and only if an identity statement 'a=b' is cognitively significant.<sup>2</sup> So 'Hesperus' and 'Phosphorus' have different senses, since 'Hesperus is Phosphorus' is cognitively significant. 'Hesperus is Hesperus', by contrast, is cognitively insignificant, and the two sides of the identity correspondingly have the same sense.

Frege's test for difference in sense is limited to singular terms, but one can naturally generalize it to other expressions, by suggesting that a pair of expressions of the same type have different senses when a statement of their coextensiveness is cognitively significant.<sup>3</sup> In the case of kind terms, one can apply the same test as before: 'a' and 'b' have different senses if and only if an identity statement 'a=b' is cognitively significant. So 'water' and 'H<sub>2</sub>O' have different senses, since 'water is H<sub>2</sub>O' is cognitively significant. In other cases, the test will be slightly different. For general terms, one can say that 'a' and 'b' have different senses when 'All a's are b's and all b's are a's' is cognitively significant: so 'renate' (creature with a kidney) and 'cordate' (creature with a heart) have different senses. For predicates, one can say that 'A' and 'B' have different senses when 'For all x, x is A iff x is B' is cognitively significant: so 'has a kidney' and 'has a heart' have different senses. For sentences, one might suggest that S and T have different senses when 'S iff T' is cognitively significant: so 'Hesperus is a planet' and 'Phosphorus is a planet' have different senses.

It is possible for two different expressions to have the same sense. When two words are intuitively synonymous — as with 'lawyer' and 'attorney', perhaps — an identity between them is cognitively insignificant. The truth of 'lawyers are attorneys' is arguably trivial: it is knowable a priori, requires no cognitive work, and gives no

---

<sup>2</sup> Throughout this paper, I use single quotes where some might use corner quotes, and I allow context to disambiguate whether symbols for linguistic expressions invoke use or mention of that expression.

<sup>3</sup> This extension beyond the case of singular terms goes beyond what is explicit in Frege. In various passages that touch on the equivalence of senses of sentences, Frege gives a number of criteria that are not obviously equivalent to each other. Some of these criteria closely resemble the criteria in the text, while others are related but are not obviously equivalent.

significant information about the world. If so, then ‘lawyer’ and ‘attorney’ have the same sense. In a similar way, it is plausible that the sentences ‘Vixens are rare’ and ‘Female foxes are rare’ are trivially equivalent. If so, the two sentences have the same sense.

We can think of the sense of an expression as mirroring the expression’s role in reason and cognition. When two expressions are trivially equivalent, they will play almost the same role in reason and cognition, and will have the same sense. When two expressions are not trivially equivalent, they will play different roles in reason and cognition, and will have different senses. In this way, we can think of an expression’s sense as capturing its cognitive significance, and as representing the “cognitive value” or “cognitive content” of the expression.

*(3) The sense of a complex expression depends on the senses of its parts.*

We saw before that the extension of a sentence (such as ‘John is hot and Mary is cold’) at least typically depends on the extension of the expressions it contains and on its logical form. In a similar way, the sense of a sentence at least typically depends on the senses of the expressions it contains and its logical form. The same goes for complex terms, such as ‘the greatest cricket player’: insofar as its extension depends on the extensions of its parts, its sense depends on the sense of its parts.

This dependence of an expression’s sense on the senses of its part may closely reflect the dependence of an expression’s extension on the extensions of its parts. To determine the sense of a complex expression, we first determine the logical form of a complex expression, then determine the senses of the basic terms involved, and then compose these senses in a way that depends on the logical form. Just how this composition works is not entirely clear, but I will say more about it in what follows.

*(4) Sense determines extension.*

Frege held that the extension of a word, a complex expression, or a sentence is determined in some way by its sense. The sense of an expression is not in general determined by its extension. ‘Hesperus’ and ‘Phosphorus’ have the same extension but

different senses, so it seems that there is no path from extension to sense. If the determination thesis is correct, however, then there is a path from sense to extension.

It is not entirely clear how the determination thesis is to be understood. We might say that sense *strongly* determines extension if sense determines extension on its own, without any further contribution from the world. In contemporary terms, we might say that sense strongly determines extension if any two possible expressions that have the same sense have the same extension. On this view, it seems that an expression's extension must somehow be present at least implicitly within its sense. While there are some indications of this sort of view in Frege, this idea arguably stands in tension with the idea that sense reflects cognitive significance. For example, the two terms 'the morning star' and 'the evening star' have the same extension, but this sameness of extension does not seem to be implicit in the cognitive roles of the terms. It is natural to suppose that someone in a different environment might use a term with the same cognitive role but a different extension. Similarly, it is not clear how the truth of a statement such as 'There are 90 chemical elements that occur in nature' could be determined by its cognitive role alone.

Alternatively, we can say that sense *weakly determines* extension if extension is determined by sense in conjunction with the world. It is natural to suppose that the sentence just mentioned is true not just because of its sense, but because of the way the world is. Likewise, it seems plausible that a term such as 'the morning star' refers to the planet Venus not just because of its sense, but because of the way the world is. Formulating the weak determination thesis so that it is both plausible and nontrivial is not easy: after all, is not everything determined by the way the world is? But there is at least an intuitive idea to keep in mind here, which we can return to later.

Frege also held some further theses concerning sense. These are not quite as crucial to a broadly Fregean view as the theses above, but they will be relevant to our discussion.

(5) *In indirect contexts, expressions refer to their customary senses.*



As we saw before, there are cases in which the truth-value of a sentence seems not to be determined by the extensions of its parts. This happens especially with sentences involving belief, and related ascriptions of attitudes. If ‘John believes that Cary Grant is an actor’ is true and ‘John believes that Archie Leach is an actor’ is false, and if ‘Cary Grant’ and ‘Archie Leach’ have the same extension, then the truth value of these sentence cannot be determined by the extensions of their parts. The same goes for many other constructions involving indirect contexts, a context where words appear inside a “that”-clause (such as ‘that Cary Grant is an actor’).

To deal with these cases, Frege suggested that in indirect contexts, an expression inside the “that”-clause does not have its usual extension. Instead, its extension is what is usually its *sense*. So inside such a clause, ‘Cary Grant’ does not refer to a person, but to the (customary) sense of ‘Cary Grant’. In this way, we can see that ‘Cary Grant’ and ‘Archie Leach’ have *different* extensions within these sentences, so the thesis that truth-value depends on extension is preserved.

(6) *The sense of a sentence has an absolute truth-value.*

For Frege, the sense of a sentence is a special sort of entity, a “thought”. A Fregean “thought” is not a mental entity. It is more like what many philosophers call a proposition, capturing the content that a sentence expresses, when stripped of the accidental clothing of a particular language. (I will use this terminology instead of Frege’s in what follows.) Just as a sentence can be true or false, a proposition can be true or false. Frege held that propositions are the primary bearers of truth, and that sentences are true or false derivatively: a sentence is true if and only if the proposition that it expresses is true. Further, Frege held that a proposition is true or false *absolutely*. On his view, it is not possible for the same proposition to be true or false, for example at different times. If two sentences, uttered by any subjects at any times, express the same proposition, they will have the same truth-value.

This has strong consequences for the notion of sense. Many sentences can be true when uttered on one occasion, and false when uttered on another. For example, ‘It is raining here now’ will be false if I utter it now, but it would have been true if I had

uttered it at this time yesterday. One might have been tempted to say that both of these utterances had the same sense. But if senses are propositions with absolute-truth value, this cannot be so. The two sentences must have different senses, and must express different propositions.

(7) *The sense of an expression can vary between occasions of use.*

It is tempting to hold that the sense of an expression is a *universal* feature of that expression: that is, that every token of an expression has the same sense. If this were right, then the sense of an expression could be seen as built into the language of which the expression is a part. On Frege's view, however, sense is not always universal in this fashion. One reason for this is given above: on Frege's view, the sense of a sentence such as 'It is raining here now' differs between different occasions of use. One can presumably trace this difference to differences in the sense of expressions such as 'here' and 'now' between occasions of use. Another quite different reason is tied to names. Frege (1892, second footnote) says that two different users of a name such as 'Aristotle' might associate a different sense with it. He says that this should not happen in a 'perfect language', but it does happen in natural languages.

For reasons like these, Frege's view entails that one cannot always attach sense to expression *types*. To handle cases like this, one has to attach sense to expression *tokens* (or to expression types as used in specific contexts, or to something else that is tied to an occasion of use). It follows that on Frege's understanding, the sense of an expression should not be identified with its *linguistic meaning*, where the latter is required to be common to all tokens of an expression type.

To sum up: on a Fregean view, expressions have senses that satisfy theses (1)-(7). In recent years, many philosophers have doubted this. It has been widely argued that expressions do not have Fregean senses, and there is no notion that can play the role that sense is intended to play. In particular it is widely believed that it is difficult to satisfy thesis (2), and that it is impossible to simultaneously satisfy theses (1)-(4).

To flag my conclusions: I think there is a viable notion of sense such that expressions have senses that satisfy slightly modified versions of the core requirements

(1)-(4). When sense is interpreted this way, theses (5) and (6) are rejected, but (7) is accepted. I think that such a view can vindicate the spirit, if not the letter, of Frege's view.

### 3 What are senses?

What are senses? I have outlined a number of features that Frege attributed to senses, but this is not yet to say what a sense is. Frege's own discussion leaves the matter somewhat unclear. He says that they are not mental entities, such as the idea or image associated with an expression, and he holds that they are abstract objects of some sort (inhabiting the 'third realm'). But this still leaves their nature open.

One natural suggestion is that senses are *descriptions*. Frege sometimes uses descriptions to specify senses. In talking about the sense of 'Aristotle', for example, Frege says 'It might, for instance, be taken to be the following: the pupil of Plato and teacher of Alexander the great'. One might similarly suggest that the sense of 'Hesperus' is something like 'the brightest object visible in the evening sky', and that the sense of 'Phosphorus' is something like 'the brightest object visible in the morning sky'.

Descriptions, on a natural understanding of the term, are linguistic entities. 'The pupil of Plato and teacher of Alexander the great' is a linguistic entity, a complex expression containing ten words. Senses obviously cannot be descriptions of this sort: descriptions have senses of their own, so we will be left either with descriptions that serve their own sense or endless chains or circles of descriptions. Furthermore, senses of this sort can never break out of the linguistic domain.

A more plausible suggestion is that the sense of an expression is the *sense* of an associated description. Even if this is true, however, it does nothing to tell us what senses *are*. It is also far from clear that associated descriptions of the right sort exist for all expressions. For example, 'knowledge' seems to be a paradigmatic term with a sense: some states qualify as knowledge but others do not, and one might think of the sense as encapsulating a criterion for knowledge. But there famously appears to be no description that captures such a criterion. Gettier showed that 'justified true belief' is inadequate, and all attempts at complex descriptions have failed. Nevertheless, even if there is no

linguistic description that captures what it takes to be knowledge, this does not show that ‘knowledge’ has no sense.

To find a better understanding of sense, it is useful to think about the work that descriptions are doing here. The role of a description is plausibly to give us a *condition on extension*: a condition that an entity in the world must satisfy in order to qualify as an expression’s extension, depending on how the world turns out. For example, if ‘Hesperus’ functions as above, then the associated description will give a condition on its extension. If we discover that the brightest object in the evening sky is Venus, then ‘Hesperus’ will refer to Venus. If we discover that the brightest object in the evening sky is Rigel Delta, then ‘Hesperus’ will refer to Rigel Delta. And so on.

Here the crucial property of a description is that it gives us a way of identifying an expression’s extension, given full knowledge of how the world turns out. It may be that for some expressions (such as ‘knowledge’), there is no description that can do this job. It is nevertheless not implausible that the expression’s extension depends in *some* fashion on how the world turns out, and in particular that full knowledge of how the world turns out puts a subject in a position to identify the expression’s extension. We can then generalize to think of an expression’s sense as the relevant condition on extension, whether or not this condition can be captured by a description.<sup>4</sup>

What do conditions on extension have to do with cognitive significance? An attractive idea is that when an expression plays a certain cognitive role for a speaker, then it will be associated with certain tacit criteria for identifying the extension of the expression, given sufficient information about the state of the world. It is natural to hope that these criteria will reflect the cognitive role of the expression in some deep respects. In order to tie a condition on extension to cognitive significance in this way, it is important that the relevant condition on extension be understood *epistemically*, in a manner that is closely connected to a speaker’s knowledge and cognition. To do this, it is helpful to think about possible states of the world in epistemic terms.

---

<sup>4</sup> Explicit discussion of the idea that a sense corresponds to a condition on extension is surprisingly rare in Frege, but it is present in his *Begriffsschrift*. This strand of Frege’s thought is emphasized by Dummett 1973.

For all we know a priori, there are many ways the world could be. We might live in a world with planets visible in the sky, or we might not. We might live in a world where people play cricket, or we might not. We might live in a world where some objects travel faster than light, or we might not. We might live in a world where the liquid in the oceans is a basic atomic substance, or we might not. We can put this intuitively by saying that there is a space of *scenarios* such that for all we know a priori, any one of these scenarios could be actual. To a first approximation, we can think of a scenario as something like a possible world, though some potential differences will arise later.

This a priori ignorance about the nature of our world reflects a corresponding a priori ignorance about the extensions of our expressions. For a typical expression, such as ‘Hesperus’, or ‘the greatest cricket player’, or ‘water’, we cannot know what the expression refers to without much observation of the world. For all we know a priori, it might be that Hesperus is Rigel Delta, or that the greatest cricket player is Dennis Lillee, or that water is a basic element. The extensions of our expressions depend on how our world turns out. That is, they depend on which scenario is actual.

Once we know enough about the nature of our world, we are usually in a position to know what our expressions refer to. Once we do enough astronomical work investigating the nature of the objects in the evening sky, we know that ‘Hesperus’ refers to Venus, not Rigel Delta. Once we know about the performances of cricket players throughout the history of the game, we know that ‘the greatest cricket player’ refers to Don Bradman, not Dennis Lillee. Once we know about the chemical makeup of the various substances in our environment, we know that ‘water’ refers to H<sub>2</sub>O, not to a basic element. And so on.

We can think of this as being part of what using a language involves. If a subject uses an expression, then given sufficient information about the world, the subject will be in a position to know the extension of the expression. Furthermore, something like this will be the case *however* the world turns out: for any world, given sufficient information about that world, the subject will be in a position to determine what the extension of the expression will be *if* that world is actual. Of course in some cases the extension may be indeterminate, as it sometimes is in the actual world; but in such a case, the subject will be in a position to determine that, too.

One could put forward a thesis holding that when a subject using an expression is given sufficient information about a world, the subject is in a position to know what the extension of the expression will be if that world is actual. A full and precise version of such a thesis would require careful attention that I will not give here, but some things can be said to clarify it.

First, what counts as “sufficient information”? If we allow too much, the thesis becomes trivial: given the information that Hesperus is Phosphorus, one can trivially know that Hesperus is Phosphorus. But it is clear that in this case, no such information is required; neutrally specified information suffices. It is plausible that at least in many cases, ‘qualitative’ information about the distribution of physical properties, appearances, and mental states (perhaps including some indexical information) in the actual world suffices to determine an expression’s extension. I will not give a precise account of the relevant information here (see Chalmers and Jackson 2001 for a more precise hypothesis). What matters is rather the general idea that there is *some* constrained sort of information such that information of this sort can suffice to determine an expression’s extension, and usually in a nontrivial way.

Second, when is a subject “in a position to know” an expression’s extension? We can that this is the case when *sufficient reasoning* from information available to the subject will allow the subject to know the expression’s extension. Here the reasoning is restricted to a priori reasoning (or to armchair reasoning, if one prefers), so further empirical observation is disallowed. And we idealize away from poor reasoning: it is not a subject’s actual reasoning that matters, but rather what the subject could know given unimpaired reasoning. For example, a subject possessing the relevant information might judge that 47 plus 59 is 116, due to a miscalculation, or that Alpha Centauri is the nearest star, because they overlook the sun. But these mistakes can be corrected by better reasoning, so they provide no counterexample to the thesis.

If something like this is right, then a subject using an expression is in a position to know the expression’s extension *given* relevant empirical information and sufficient reasoning. The second clause entails that there is a normative element here; but what matters is that the extension is within the reach of reason. This feature of language and thought is responsible for a deep link between meaning, possibility, and rationality.

## 4 Intensions

All this suggests that an expression's sense might be seen as an *intension*: a function from possibilities to extensions. This function takes a given possibility, and associates it with an extension relative to that possibility. The extension will be either an entity present in that possibility, or the null extension. We can say that the intension is *evaluated* at a possibility, and *returns* an extension in that possibility. In what follows, I will make a prima facie case that intensions can be associated with expressions in such a way that they can play much of the role of Fregean sense.

For the Fregean, the possibilities in the domain of an expression's intension will be thought of as *epistemic* possibilities, in a broad sense: ways the world could be, for all we know a priori. These epistemic possibilities are what I called scenarios above: for now, we can think of them as possible worlds. The intension of an expression can be thought of as an *epistemic* intension: it captures (very roughly) the way the extension of the expression depends on which epistemic possibility turns out to be actual. For a sentence S and a scenario W, for example, a useful heuristic is to ask: *if* W is actual, is it the case that S? Or to stress the epistemic nature of this conditional: if W *turns out* to be actual, will it turn out that S? If yes, then the intension of S is true at W.

Take an expression such as 'the greatest cricket player'. Let  $W_1$  be a scenario in which Don Bradman never plays cricket, and in which Phil Tufnell scores more runs and takes more wickets than any other cricket player. Let A be a scenario corresponding to the actual world, where Don Bradman's batting average is 99.94, and where the second highest batting average is around 61. For all we know a priori, scenario  $W_1$  could be actual. For all we know a priori, scenario A could be actual. If scenario  $W_1$  is actual, then the greatest cricket player is Phil Tufnell. So when evaluated at  $W_1$ , the intension of 'the greatest cricket player' returns Phil Tufnell. If scenario A is actual (as it is!), then the greatest cricket player is Don Bradman. So when evaluated at A, the intension of 'the greatest cricket player' returns Don Bradman.

What about a term such as 'Hesperus'? Here one can tell a Fregean story in a similar way. Let scenario  $W_2$  be one on which the brightest object visible in the evening

is Rigel Delta, 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 Rigel Delta. So when evaluated at  $W_2$ , the intension of 'Hesperus' returns Rigel Delta. If it turns out that A is actual, then it will turn out that Hesperus is Venus. So when evaluated at A, the intension of 'Hesperus' returns Venus.

The same applies to a term such as 'water'. 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  $H_2O$ .

One can do the same sort of thing for a whole sentence. Just as the extension of a term depends on the way the world turns out, so does the truth-value of a sentence. So a sentence will also be associated with an intension: this time, a function from scenarios to truth-values. This function takes a scenario, and returns a truth-value associated with that sentence in the scenario. This truth-value will be 'true', 'false', or perhaps neither.

Take a sentence such as 'Hesperus is Phosphorus'. One can tell a Fregean story here in a similar way. If it turns out that scenario  $W_2$  above (with Rigel Delta in the evening and Neptune in the morning) is actual, then it will turn out that Hesperus is not Phosphorus. So when evaluated at  $W_2$ , the intension of 'Hesperus is Phosphorus' returns 'false'. When evaluated at the actual scenario A, on the other hand, the intension of 'Hesperus is Phosphorus' returns true.

The intuitive characterization of epistemic intensions using the heuristics I have given here makes a strong prima facie case that expressions have epistemic intensions. Giving a truly precise definition of epistemic intensions involves complexities that I cannot go into here, but it may be useful to say a few words about the foundations of the idea, and about how a more precise definition (in the case of sentences) can be constructed.

The basis for epistemic intensions lies in our ability to describe and evaluate epistemic possibilities. Let us say that it is epistemically possible (in the broad sense) for a speaker that S when the speaker cannot know a priori that S is not the case. Then it is



epistemically possible that water is not H<sub>2</sub>O. It is also epistemically possible that our world is the XYZ-scenario: that is, that the clear liquid in the oceans and lakes (and so on) is XYZ. And when we reflect on the second, specific epistemic possibility, it reveals itself to us as an instance of the first epistemic possibility. That is, the epistemically possible hypothesis that the XYZ-scenario is actual is a specific version of the epistemically possible hypothesis that water is not H<sub>2</sub>O.

We can see, then, that we use language to describe and evaluate epistemic possibilities in a distinctive way. Quite generally, given a specific epistemically possible scenario W and some more general epistemic possible sentence S: a speaker can say, on reflection, that the epistemic possibility that W is actual is an instance of the epistemic possibility that S is the case, or an instance of the epistemic possibility that S is not the case, or is neither. If it is the first, then the epistemic intension of S is true in W. If it is the second, then the epistemic intension of S is false in W. If it is neither, then the epistemic intension of S is indeterminate in W.

Of course any specific scenario must be *described* in order for a speaker to be able to evaluate it as an epistemic possibility. When the speaker considers the epistemic possibility that W is actual, he or she really considers the epistemic possibility that D is the case, where D is a description (in some sense) of W. This raises the issues of what it is to be a description of an epistemically possible scenario, and of whether different descriptions of the same scenario might give different results. I discuss this matter in Chalmers (forthcoming a), isolating a class of canonical descriptions that give equivalent results.<sup>5</sup> For present purposes, the intuitive characterization of what it is for the epistemic intension of S to be true in W should suffice.

The epistemic intension of S at W corresponds to a speaker's judgment about whether the epistemic possibility that W is an instance of the epistemic possibility that S, or to the speaker's judgment about whether S is true if W is actual. But it is not the

---

<sup>5</sup> See Chalmers forthcoming a. A natural requirement is that a canonical description be epistemically complete, where D is epistemically complete if D is epistemically possible and there is no S such that both  $D \wedge S$  and  $D \wedge \neg S$  are epistemically possible. If scenarios are understood as possible worlds, one must also require that D use only semantically neutral expressions (roughly, those that are not vulnerable to Twin Earth thought experiments) supplemented by indexicals.

speaker's snap judgment that matters, nor any actual judgment of the speaker. Rather, it is the speaker's (potential) *rational* judgment that matters. Here we idealize away from poor reasoning, and consider judgments on ideal rational reflection.

An idealization like this can be made in a number of ways. The most obvious way is to invoke the idealization inherent in the notion of apriority. S is a priori not if a speaker knows that S, nor even if a speaker would be able to know that S on reflection, but rather if it would be possible for the speaker to know that S, using the concepts involved in S, on ideal rational reflection. So we can say that the epistemic intension of S is true in W if a priori reasoning by the speaker could reveal the epistemic possibility of W to be an instance of the epistemic possibility that S.

This approach suggests a natural characterization of the epistemic intension of S for a speaker. For a given scenario W, let D be any canonical description of W. Then the epistemic intension of S is true at W iff the material conditional ' $D \supset S$ ' is a priori; it is false at W if ' $D \supset \neg S$ ' is a priori; and it is indeterminate at W if neither is a priori.<sup>6</sup> Other forms of definition are also possible. If someone rejects the notion of apriority, for example, then one can appeal to a different construal of the relevant epistemic status, perhaps by appealing to whether an inference from the hypothesis that D is the case to the conclusion that S is the case would be rational. One also needs to say more about the nature of the relevant scenarios, descriptions, and epistemic status, in order to have a truly precise definition. (For a more precise treatment, and for expression of this treatment to expressions other than sentences, see Chalmers forthcoming a and b.)

It should be noted that to evaluate the epistemic intension of S at W, a token of S need not be present within W. None of the heuristics or definitions that I have given here rely on such a token; rather, they rely on first order claims about epistemic possibility and

---

<sup>6</sup> The notion of apriority invoked here will be relative to a speaker, for reasons discussed later in the paper. It is also important that apriority be understood so that sentences such as 'Hesperus is Phosphorus' are not a priori (contrary to some understandings). The easiest way to accommodate these requirements is roughly as follows: a sentence S is a priori relative to a speaker if the sentence as used by that speaker expresses a thought that can be justified independently of experience, on ideal rational reflection. Here a thought is a token propositional attitude such as a belief. This allows that the same sentence type can be a priori for one speaker but not another, and it entails that sentences such as 'Hesperus is Phosphorus' will not generally be a priori, as there will generally be no way to a priori justify the corresponding thought.

apriority using an expression present in the actual world. And one can often evaluate expressions at scenarios with no such tokens. For example, there are epistemically possible scenarios in which no-one uses language. When one considers such a scenario *W*, it reveals itself as an instance of the epistemic possibility that no words exist, that there are no novels, and so on. So the epistemic intension of ‘words exist’ is false at *W*, and so on. It is even arguably epistemically possible that no-one exists (as long as ‘I exist’ is not a priori). One can consider and evaluate that epistemic possibility in various ways; for example, the epistemic intension of ‘Someone exists’ will be false there. Of course some sentences may yield indeterminate results at these epistemic possibilities, but all that matters is that some sentences yield determinate results there.

This framework is grounded in the fact that when a speaker is given the right sort of information about the actual world, or about an epistemic possibility more generally, then conclusions about extension are within reach of reason, given appropriate reflection. We can use this to articulate a version of the Fregean notion of ‘grasping’ a sense. We can say that a subject grasps an intension when the subject is in a position to *evaluate* that intension: that is, when sufficient reasoning will allow the subject to determine the value of the intension at any world. Again, this does not require that the subject will actually determine the correct extension when a relevant epistemic possibility is specified, but it does require that the extension is within the grasp of reason. If what I have said here is right, then whenever a subject uses an expression, the expression will be associated with an epistemic intension that the subject grasps.

Occasionally, the epistemic intension of a term will be the same as that of an associated description. For example, the epistemic intension of ‘Neptune’ as used by Leverrier was arguably the same as that of ‘the object causing the perturbation of the orbit of Uranus’; and the epistemic intension of ‘bachelor’ is arguably the same as that of ‘unmarried man’. In these cases, there is a sense in which the epistemic intension can be “captured” by a description.

This does not hold in general. For many or most terms, there may be no description (and certainly no short description) with the same epistemic intension as the term. We saw this in the case of ‘knowledge’, and the same applies to most names. In these cases, the best one can hope for is a description whose epistemic intension

approximates that of the original term: as with ‘justified true belief’ for ‘knowledge’, or ‘the clear drinkable liquid in the oceans and lakes’ for water, and so on. These descriptions may give one a rough and ready sense of how a term’s epistemic intension functions, but they do no more than that. Usually there will be at least a small subset of epistemic possibilities (such as the Gettier cases, in the case of ‘knowledge’) in which the epistemic intension of the original term and of the description come apart.

So there is no reason to think that in general, an epistemic intension can be captured by a description. More generally, there is no reason to think that grasping an epistemic intension requires any sort of descriptive articulation of a concept by a subject. The epistemic intension is a function, not a description. It is revealed in a subject’s rational evaluation of specific epistemic possibilities, not in any sort of explicit definition. Even where such a definition exists, a subject need not be able to articulate it to grasp the epistemic intension. Indeed, we usually evaluate the plausibility of such definitions precisely by deploying our prior grasp of a term’s epistemic intension, to see how whether the definition gives the right results in specific cases. (Witness the literature on the definition of ‘knowledge’). So epistemic intensions are more basic than descriptions, and should not be assimilated with them.

Epistemic intensions can be associated with any expression that is a candidate for extension. Given the general type of the expression (singular term, general term, etc), it will be constrained to have a certain sort of extension (individual, class, etc). The intension of such an expression will be a function from scenarios to the appropriate sort of extension. So the intension of a singular term is a function from scenarios to individuals; the intension of a general term is a function from scenarios to classes; the intension of a sentence is a function from scenarios to truth-values; and so on.

## **5 Intensions as senses**

How can we connect intensions as discussed above to the Fregean theses outlined at the start of this paper? We could establish such a connection with the aid of the following four theses about intensions.

(1\*) Every expression has an intension, which returns at a scenario an extension of the type appropriate for the expression.

(2\*) A sentence is a priori iff its intension is true at all scenarios.

(3\*) If the extension of a complex expression E depends by some rule on the extensions of its parts, then at a scenario W, E's intension returns an extension that depends by the same rule on the extensions returned by the intensions of E's parts.

(4\*) At a scenario corresponding to the actual world, the intension of E returns the (actual) extension of E.

These theses all fit naturally with the understanding of epistemic intensions outlined above. (1\*) was discussed at the end of the last section. (3\*) follows from a natural extension of compositional semantics from the actual world to arbitrary epistemic possibilities. (4\*) is a consequence of the claim that given appropriate information about the actual world, a subject will be in a position to determine the expression's actual extension. (2\*) also fits naturally with the understanding above, but it raises some tricky issues that I will return to shortly.

Given the theses above and the equation of senses with epistemic intensions, it is not hard to see that versions of the Fregean theses (1), (3), and (4) follow. Requirement (1), that every expression has a sense, follows immediately from (1\*). Requirement (3), that the sense of a sentence depends on the sense of its parts, follows from (3\*) along with the thesis that the extension of a sentence depends on the extension of its parts. Note that even if there are some cases where the latter thesis fails (such as belief contexts), requirement (3) can still be satisfied as long as the extension of a sentence depends on the extensions *and the intensions* of its parts. For then it is plausible that the extension of a sentence at a scenario will depend on the extension of its parts at a scenario along with the intensions of the parts, so the intension of the sentence will depend on the intensions of the parts.

A version of requirement (4) follows naturally from thesis (4\*). Here the relevant version of (4) is the thesis involving weak determination, on which sense determines extension in combination with the world. To determine an expression's extension, one simply evaluates the expression's intension at the actual world. One can naturally think of the intension as supplying *criteria* for determination of extension: in combination with the actual world, these criteria will determine an extension.

What about the crucial requirement (2), that sense reflects cognitive significance? Here matters are complicated by the fact that we lack a precise definition of 'cognitive significance'. There is a natural understanding of cognitive significance that fits well with the intensional framework, however. We can say that a sentence S is cognitively insignificant (for a speaker) when S is knowable a priori (by that speaker, given ideal reflection): that is, when it is knowable with justification independent of experience. And S is cognitively significant when it is not knowable a priori. 'Hesperus is Hesperus' is knowable a priori, so it is cognitively insignificant. 'Hesperus is Phosphorus' is not knowable a priori, so it is cognitively significant.

This understanding of cognitive significance differs from Frege's. On Frege's account, a priori knowledge can be cognitively significant: the knowledge that  $59+46$  is 105 is cognitively significant, for example, because this knowledge requires some cognitive work. It is very hard to articulate this notion of cognitive significance precisely, however, and it is not clear that there is a useful precise notion nearby. The notion of apriority can serve at least as a useful substitute. It is clear that non-apriority entails cognitive significance in the Fregean sense (although not vice versa). Most of Frege's central cases of cognitive significance involve a posteriori knowledge, and are therefore accommodated well on this framework. The main difference is that certain pairs of expressions that Frege counts as cognitively inequivalent (e.g., "59+46" and "105") will count as cognitively equivalent on the current framework, because of the rational idealization inherent in the notion of apriority. But this understanding gives at least an approximation of a Fregean view (one that might later be refined), so it is this understanding that I will work with.

For senses to reflect cognitive significance in this sense, it was required (among other things) that an identity 'a=b' is cognitively significant if and only if 'a' and 'b' have

different senses, and that two sentences S and T have the same sense if and only if the material biconditionals 'S iff T' is cognitively insignificant. Recasting these claims in terms of intensions and apriority, we require the following: An identity 'a=b' is a priori if and only if 'a' and 'b' have the same intension; and two sentences S and T have the same intension if and only if the material biconditional 'S iff T' is a priori. For other expression types, such as predicates, general terms, and the like, analogous theses will be required.

It is not hard to see that the relevant requirements will all be entailed by theses (1\*)-(4\*), with the crucial work being done by thesis (2\*): that S is a priori iff its intension is true at all scenarios. This can be illustrated in the case of singular terms and in the case of sentences.

From (2\*), it follows that 'a=b' is a priori iff its intension is true at all scenarios. By (3\*) and the extensional semantics of identity, the intension of 'a=b' is true at a scenario iff the intensions of 'a' and 'b' return the same extension there. So 'a=b' is a priori iff the intensions of 'a' and 'b' return the same extensions at all scenarios. Equivalently, 'a=b' is a priori iff 'a' and 'b' have the same intensions.

From (2\*), it also follows that 'S iff T' is a priori iff its intension is true at all scenarios. By (3\*) and the semantics of material conditionals, the intensions of 'S iff T' is true at a scenario iff the intensions of S and T return the same truth-value there. So 'S iff T' is a priori iff the intensions of S and T return the same truth-values at all scenarios. Equivalently, 'S iff T' is a priori iff S and T have the same intensions.

We have seen that from (1\*)-(4\*), versions of the crucial Fregean requirements (1)-(4) can be satisfied. The question then becomes: is there reason to believe that sentences can be associated with intensions that satisfy (1\*)-(4\*)? The crucial claim is clearly (2\*). Is there reason to believe that expressions can be associated with intensions that satisfy (2\*)?

Here, the central reason stems from the Fregean understanding of these intensions and of the associated scenarios. As we saw above, the scenarios represent *epistemic* possibilities: ways the world might be, for all we know a priori. And the intensions are *epistemic* intensions: they capture a subject's idealized judgments about how an expression applies to an epistemic possibility, under the hypothesis that it is actual. On

this understanding, a strong connection between intensions and epistemic notions such as apriority is built into the framework.

So consider the left-to-right direction of (2\*): if a statement is a priori, its epistemic intension is true at all scenarios. This direction is straightforwardly plausible. If S is a priori, then for *any* epistemic possibility W: if W turns out to be actual, it will turn out that S. We might say that for all W, the epistemic possibility that W is actual is an instance of the epistemic possibility that S. So the epistemic intension of S is true in all scenarios.

This is brought out by cases such as ‘Hesperus is Hesperus (if it exists)’, ‘All bachelors are unmarried men’, and even ‘The meter stick in Paris is one meter long (if it exists)’, assuming that these are a priori. In each of these cases, there is no epistemic possibility of falsehood. No matter how the world turns out, it cannot turn out that Hesperus is not Hesperus, or that the meter stick in Paris is longer or shorter than one meter. So in every epistemic possibility W, the epistemic intensions of these sentences are true.

Something similar applies in the right-to-left direction, which we can consider in the contrapositive form: if S is not a priori, then there is some scenario in which its epistemic intension is not true. This fits the familiar cases. ‘Hesperus is Phosphorus’ is a posteriori, and its epistemic intension, as we saw above, is false in a scenario  $W_2$  where the objects visible in the morning and evening sky differ. ‘Water is  $H_2O$ ’ is a posteriori, and its epistemic intension is false in the ‘Twin Earth’ scenario  $W_3$ . And this is no accident: the aposteriority of these statements seems to be reflected in the existence of scenarios in which the epistemic intension is false.

If S is not a priori, it is clearly epistemically possible that S is not the case. So all we need for the right-to-left direction of (2\*) is the following principle: if it is epistemically possible that T, then there is an epistemically possible scenario W in which the epistemic intension of T is true. This principle is suggested very strongly by the above examples. We have seen that even in the case of a posteriori identities involving water, Hesperus, and the like, the principle appears to be satisfied.

The principle is not entirely trivial, however. Someone might hold that if scenarios are understood as possible worlds, there are counterexamples to the claim. For



example, it might be held that ‘a god exists’ is necessary but not a priori, for example. If this view is correct, then ‘no god exists’ is epistemically possible, but it will plausibly be verified by no possible world. Similarly, it might be held that the Continuum Hypothesis is necessary but not a priori. If this view is correct, then the negation of the Continuum Hypothesis is epistemically possible, but plausibly verified by no possible world. If these views are correct, and if scenarios are understood as possible worlds, then these cases (unlike cases involving water, Hesperus, and so on) provide a counterexample to the principle above, and so to principle (2\*).

These claimed counterexamples would be highly controversial, as they presuppose highly controversial views about gods and about mathematics respectively. A Fregean could simply deny these views as contrary to his or her principles, perhaps holding that they tacitly assume a false view of some of the crucial notions: apriority, possibility, intensions. It is useful to probe the basis for such a denial, however. The opposing arguments could be resisted in at least two ways, depending on the relevant conception of possibility.

There are two ways in which the Fregean might understand the space of epistemic possibilities over which epistemic intensions are defined. (See Chalmers forthcoming b.) First, this might be a space of worlds that is *independently characterized*, such as the space of metaphysically possible worlds. In this case, it becomes a substantive thesis that when S is epistemically possible, there is a world in which the epistemic intension of S is false. The thesis is plausible for familiar cases (including ‘water’ and ‘Hesperus’ cases), but an opponent might hold that in some cases, it is false. This might happen for reasons such as those in the case above, where an opponent could hold that there are not enough metaphysically possible worlds to go around.

On this understanding, for the Fregean to hold onto principles such as (2\*), she will have to deny the opponent’s analysis of the cases. She might hold that God’s existence cannot be necessary, for example, or perhaps that it is a priori; and she might hold that the Continuum Hypothesis is either a priori or indeterminate. This will require substantive argument. And the postulated connection between the independently characterized space of worlds and epistemic possibility will be a substantive thesis; it will have implications for just what is metaphysically possible, for example. Of course the

thesis may nevertheless be plausible. A Fregean might even argue for the thesis directly, perhaps by noting that there are no clear counterexamples, and by arguing that there are constitutive connections between notions of metaphysical and epistemic possibility.

On the second understanding, the possibilities involved are not independently characterized, but are understood as epistemic possibilities from the start. The Fregean might postulate or construct a space of scenarios understood as “maximal epistemic possibilities”, for example. (A construction of this sort of discussed in Chalmers forthcoming b.) These might intuitively correspond to maximally specific epistemically possible hypotheses about the way things are, hypotheses from which a priori reasoning can settle everything there is to be settled. This space might be defined quite independently of notions of ‘metaphysical possibility’.

On this understanding, a Fregean is free to accept much of the opponent’s analysis of the cases above. She is free to accept that ‘God exists’ is a posteriori but metaphysically necessary, for example. If she accepts this, she will simply insist that there is an epistemically possible scenario in which the epistemic intension of ‘God exists’ is false. This will follow naturally from the definition of those scenarios. Once epistemic possibility is separated from metaphysical possibility, the opponent will have no obvious reason to deny this. On this understanding, the truth of (2\*) does not involve commitment to the substantive theses about metaphysical possibility above, and it will be well-protected from counterexamples.

(A third possibility is that the Fregean could accept both the opponent’s analysis of the cases and an independently characterized notion of possibility, and accept that (2\*) is not true across the board. It might be held that there are certain special domains where it fails, but that it holds in the most domains, or in certain constrained domains. This view is not out of the question, but I will set it aside here.)

My view is that the Fregean can go in either of these two ways. I think that (2\*) is plausible even given an independently constrained notion of metaphysical possibility, and that it can be argued for directly (see Chalmers 2002a); but this is a nontrivial matter. If someone doubts this, however, the second option is available. This has the advantage of making fewer substantive commitments. Here, the main burden is making the case for the

relevant space of epistemically possible scenarios. This is also nontrivial, but I think that it can be done in a reasonably straightforward way (see Chalmers forthcoming b).

In any case, we can see that the Fregean is not being unreasonable in accepting (2\*). It appears to fit the clear cases, and there are natural ways to respond to one who doubts it. So it seems that the claim that there are intensions of some sort that satisfy (1\*)-(4\*) is well-motivated. If there are such intensions, then they can function as senses that satisfy (1)-(4).

To clarify what I have and have not done: I have not tried to precisely define epistemic intensions, in the sense of giving a precise recipe for evaluating a sentence S's intension at a scenario W. I have also not tried to give a knockdown argument that (1)-(4) can be satisfied. Rather, I have simply outlined a certain sort of Fregean view, and I have tried to make it plausible that intensions satisfying (1)-(4) exist.

If this is correct, a Fregean can hold that there is at least a *prima facie* case for a notion of sense that satisfies (1)-(4). But of course there have been numerous arguments *against* Fregean sense in recent years. So we now need to examine these arguments to see whether they have any force against the conception just outlined. I will focus on four main arguments: what we might call the argument from indexicality, the modal argument, the epistemological argument, and the argument from variability.

## **6 The argument from indexicality**

The first objection is not so much an explicit argument as a set of considerations put forward by a number of philosophers, especially John Perry (1977). These considerations revolve around Frege's treatment of indexicals, such as 'I', 'here', and 'now'.

Recall that Frege held that the sense of a sentence has an absolute truth-value. This entails that if two utterances of a sentence express the same sense, they must have the same truth-value. But it is clear that certain indexical sentences, such as 'It is now Saturday' can be uttered truly at one time and falsely at another time. So on Frege's picture, these two sentences must have different senses.

If the sense of the sentence depends on the sense of its parts, then some part of the sentence must have a sense that differs between the two occasions of utterance. The obvious source of the difference is the indexical expression 'now'. So Frege's view entails that such an indexical expression has different senses on different occasions of utterance. The same goes for many other indexicals: 'I', 'here', 'today', 'tomorrow', and so on.

It is very hard to see how this is supposed to work. One idea is that the sense of such an expression should build in its referent. If so, my utterance of 'now' today has a sense that builds in a particular Saturday, and my utterance of 'now' tomorrow has a sense that builds in a particular Sunday. Similarly, my utterances of 'I' have a sense that builds in a particular individual, David Chalmers, while your utterances have a sense that builds in a different individual.

It seems that something like this is required to preserve Frege's claims, but it is not clear that this is compatible with the overall spirit of Frege's framework. When I use an expression such as 'now' or 'today', the referent does not seem to be reflected in the cognitive significance of the expression for me. I might have no idea what day today is; and the day might change without it making any special difference to my cognition. So there seems to be some tension between this sort of claim and the thesis that sense reflects cognitive significance.

How will the intensional framework deal with these matters? Let us consider a sentence such as 'It is raining here now'. If such a sentence determines a function from worlds to referents, there is an immediate problem. I can utter the sentence truly today, and falsely yesterday. But both days, I inhabit the same world: at least, the same 'objective' world A. So if the intension of a sentence is a function from (objective) worlds to truth-values, then as before, the two utterances must have different intensions. And as before, the two tokens of 'now' must have different intensions.

It is once again quite unclear how this should work. One might suggest that the epistemic intension of 'now' should pick out a particular time (time  $t$ , the specific time of utterance) in all worlds. But then a sentence such as 'It is now time  $t$ ' will have a necessary epistemic intension. This will not do: the sentence is clearly a posteriori (and

cognitively significant), so its epistemic intension must be contingent. It seems that any other proposed intension will have a similar problem.

There is a natural way for the intensional framework to deal with this issue. A proponent of this framework should deny that intensions are functions from objective worlds to extensions. Rather, intensions can be seen as functions from *centered* worlds to extensions. Here, a centered world is a world marked with a ‘center’, where the center consists of an individual and a time present in that world.

This idea can be motivated in a natural way. On the Fregean intensional framework, scenarios are supposed to represent a sort of maximal epistemic possibility. But it is a familiar idea (this time from the work of Perry (1979), among others) that an objective description of the world leaves some matters epistemically open. When I lie awake in the middle of the night, then even if I had a full objective description of the world, I might still wonder ‘what time is it now?’, and I might not be able to settle this matter from the information available. Or I might have a full objective description of the world, but not know which individual in that world is *me*. So an objective description of the world is not an epistemically complete description of the world. To make it epistemically complete, the description also needs *locating information*: a ‘you are here’ marker, indicating which individual is *me* and what time is *now*. This sort of epistemic possibility is best represented by a centered world.

Once epistemic possible scenarios are represented as centered worlds, we can deal with the problem straightforwardly. We can say that the epistemic intension of ‘I’ picks out the individual marked at the center of any given scenario, and that the epistemic intension of ‘now’ picks out the time marked at the center. The epistemic intension of ‘today’ will pick out the day containing the time marked at the center of any given scenario, and the epistemic intension of ‘tomorrow’ will pick out the following day. These intensions will be common to all occasions of use of these expressions.

When a subject uses an expression, the actual extension of the expression will be given by evaluating the expression’s intension at the centered world inhabited by that subject: a scenario corresponding to the actual world centered on that subject and on the time of use. When two different subjects use an indexical expression such as ‘I’, they will inhabit two different centered worlds: one centered on the first subject, and one on the

second. So the epistemic intension of 'I' will pick out different actual extensions for each. For each subject, the intension will pick out himself or herself. Something similar applies to 'now': when this expression is used at different times, the intension will be evaluated at different centered worlds, and will always return the time of use.

Something like this can also help with terms that are not obviously indexical, such as 'water'. If I am given a full objective specification of an epistemic possibility, and am told that it contains a planet where the watery stuff is H<sub>2</sub>O and a planet where the watery stuff is XYZ, then I may not be in a position to know what the extension of 'water' is. To know that, I need to know which planet I am on. But if I am also given locating information, in which the center of my world is marked (e.g. on the planet with H<sub>2</sub>O), then there is no problem. I am now in a position to know which environment is *my* environment, to know which substances *I* am causally related to, and so on. So as long as I can derive the relevant objective information (about appearance, behavior, distribution of various substances), I will have no problem determining that if this centered world is actual, then water is H<sub>2</sub>O. So the epistemic intension of 'water' will return H<sub>2</sub>O at this world.

There are a couple of subtleties to the use of centered worlds. One is the following: it is arguably not a priori that I exist (I know this through experience), but the epistemic intension of 'I exist' is true in all centered worlds as defined. To deal with this, one should make the marking of a subject and time *optional*: some scenarios have no marked subject and time, or perhaps mark just one but not both. In a scenario without a marked subject, the epistemic intension of 'I exist' will be false. The other side of this coin is that there may be expressions that require the marking of further entities at the center of a scenario: some demonstratives ('that object') may require marked experiences, for example, and some special cases ('this thought') might require marked thoughts. These cases are not crucial to the current discussion, however, so we can mostly stay with scenarios in which at most a center and a time is marked.

The introduction of centered worlds to the Fregean framework has one major consequence. The sense of a sentence will no longer have an absolute truth-value. When I say 'It is raining here now' yesterday and today, my utterance has the same epistemic intension both times, but it is false yesterday and true today. So the intension is not true

absolutely, or false absolutely. It is true or false only relative to a subject and a time. So the Fregean requirement (6) fails.

It is not clear what Frege would have thought of this. The requirement that the sense of a sentence (a ‘thought’) be an absolute bearer of truth was very important to him. Still, it is widely held that Frege’s treatment of indexicals needs major repair. And it seems to me that giving up thesis (6) does not do any significant damage to a broadly Fregean framework, and it allows one to preserve the crucial connection with cognitive significance and the determination of reference. So it may be that giving up this thesis is the best way to preserve a framework that retains the broad spirit of Frege’s view.

This adjustment also entails that we must give up on Frege’s thesis (5), which holds that expressions in belief contexts refer to their senses. Consider a belief attribution such as ‘John believes that I am British’. Here the sense of ‘I’ is the epistemic intension that picks out the individual at the center of a given scenario, and the sense of ‘I am British’ is true only in scenarios where the individual at the center has a certain national origin. If John were to entertain a belief with that sense, then he would attribute that national origin to *himself*. But it is clear that this is not what John does when he believes that I am British. So on the current understanding of sense, thesis (5) entails that the wrong sort of belief is attributed to John. So we must give up on thesis (5).

I do not think that this is a high cost to pay. It is widely held that thesis (5) must be rejected, for a number of different reasons. The thesis yields an attractively elegant analysis of belief attributions, but on a close analysis, it seems that belief attributions are more subtle than the thesis suggests. The analysis of belief attributions was only a subsidiary element of Frege’s view, however, and it is clear that giving it up preserves the broad spirit of his view. Of course we still need a good analysis of belief attribution, but that is a subject that needs much discussion in its own right.<sup>7</sup>

## **7 The modal argument**

---

<sup>7</sup> I think that one can give an account of belief ascriptions that exploits epistemic intensions in a somewhat more indirect way. See Chalmers 2002b.

Perhaps the best-known argument against a Fregean view of language is the modal argument of Kripke (1980). Kripke's argument is concerned mostly with names, and also with natural-kind terms. He is arguing in the first instance against "descriptive" views of reference on which names are semantically akin to descriptions, but the arguments are generally taken to have force against any Fregean view.

Kripke argues as follows that names cannot be equivalent to descriptions. Take any name, such as 'Aristotle', and any description, such as 'the last great philosopher of antiquity'. Then it might have been that Aristotle was not the last great philosopher of antiquity; he might have died while an infant, for example. Something similar applies to any other description D that seems a likely candidate to capture the sense of the name: for any such description D, a judgment 'it might have been that Aristotle was not D' seems correct. So 'Aristotle' is not semantically equivalent to the description D.

Something similar applies in the case of 'Hesperus'. A Fregean might hold that 'Hesperus' is semantically equivalent to 'the brightest object visible in the evening sky', or some such. But Kripke argues that it might have been that Hesperus was not the brightest object visible in the evening; it might have been destroyed millennia ago, or it might have been struck by a comet and left the solar system. So again, it seems that a name is not semantically equivalent to a description.

Kripke uses these considerations to argue that it is not *necessary* that Hesperus is the brightest object visible in the evening sky, since it might have been otherwise. Or as it is sometimes put: it is not *metaphysically necessary* that Hesperus is visible in the evening, and it is *metaphysically possible* that it is not. In a similar way, he argues that there are *possible worlds* in which the evening star is not Hesperus: in a world where Venus was knocked off course by a comet and in which another object is visible in the evening, the evening star is not Hesperus but some other object.

He argues in a similar way that there are no possible worlds in which Hesperus is not Phosphorus. If Venus had been visible only in the morning with something else visible in the evening, this would have been a scenario in which Hesperus was not visible in the evening, and not a scenario in which Hesperus was not Phosphorus. He argues that Hesperus and Phosphorus are the same object (the planet Venus) in all worlds in which they exist. So 'Hesperus is Phosphorus (if they exist)' is necessary.



Kripke puts this by suggesting that names are *rigid designators*, picking out the same object in all possible worlds. Most descriptions, in contrast, are not rigid: they pick out different objects in different possible worlds. So names are quite unlike descriptions. Where names are concerned, Kripke's view is closer to the simple view, on which the meaning of a name is its referent, than to the Fregean view, on which the meaning of a name involves an associated sense.

Kripke gives related arguments concerning natural kind terms. These can be illustrated in the familiar case of 'water'. It can be argued that water might have behaved and appeared quite differently from the way in which it actually behaves and appears: it might never have appeared in liquid form, for example (witness the possibility of ice in the actual world). More generally, for any description D of water's macroscopic properties, it can be argued that if H<sub>2</sub>O had not satisfied D, water would not have satisfied D. So 'water is D' appears not to be necessary, and it seems that natural kind terms are not equivalent to descriptions.

In a similar way, it can be argued that something might have satisfied any such description D without being water. In Putnam's Twin Earth world (Putnam 1975), a different chemical substance XYZ has all the superficial properties of water. But Putnam argues that this substance is not water. That is, in a counterfactual scenario in which XYZ was watery, XYZ would not be water. Rather, water is necessarily H<sub>2</sub>O. If so, 'water' is akin to a rigid designator: it picks out H<sub>2</sub>O in all worlds.

Do Kripke's arguments have any force against the intensional framework I have outlined here? Do they show, for example, that names or natural kind terms do not have epistemic intensions that satisfy (1)-(4)? At first glance, one might think so. I have argued that Hesperus has an epistemic intension that picks out something like an object visible in the evening, in any given world. So in a world where Mars rather than Venus is prominent in the evening sky, the epistemic intension will pick out Mars. But Kripke argues that Hesperus picks out Venus in all worlds, and that it needs not pick out the evening star, or anything like it.

The conflict is only superficial, however. Kripke takes care to distinguish metaphysical possibility from epistemic possibility. And he allows that it is epistemically possible that Hesperus is not Venus; he simply denies that it is metaphysically possible.

Kripke allows that it might *turn out* that Hesperus is not Phosphorus; and he can allow that if it turns out that Venus was never visible in the evening but that Mars was, then it may turn out that Hesperus is not Phosphorus but Mars. So his argument is entirely compatible with ‘Hesperus’ having an *epistemic* intension that functions roughly as I have suggested. Much the same applies to ‘water’, and to other relevant terms.

This response is not ad hoc. The sort of possibility that is most relevant to a Fregean view is clearly epistemic possibility. When one thinks about sense in intensional terms, one thinks of it as giving criteria for the extension of an expression *depending on how the world turns out*. If the world turns out one way, it will turn out that water is H<sub>2</sub>O; if the world turns out another way, it will turn out that water is XYZ. Nothing in Kripke’s modal argument gives any reason to deny this. And the notion of sense was always tied to epistemic notions such as apriority, not to notions such as “metaphysical necessity”.<sup>8</sup> These connections are entirely preserved, even in the light of Kripke’s argument.

The other crucial property of sense was that it determines an expression’s extension in the actual world. Again, Kripke’s argument does nothing to suggest that this is not the case. The epistemic intension determines an expression’s actual extension when evaluated at the actual world. For all Kripke has said, it also determines an expression’s extension under all epistemically possible hypotheses about the actual world.

What an epistemic intension does not do, if Kripke’s arguments are correct, is determine an expression’s extension when evaluated in explicitly counterfactual scenarios. When we consider these scenarios, we are not considering them as epistemic possibilities: as ways things might be. Rather, we are acknowledging that the character of the actual world is fixed, and are considering these possibilities in the subjunctive mood: as ways things might have been. That is, rather than considering the possibilities *as actual* (as with epistemic possibilities), we are considering them *as counterfactual*. If Kripke is right, then evaluation in this sort of explicitly counterfactual context works quite differently from the evaluation of epistemic possibilities. This point still needs explaining.

It is striking that all of Kripke’s conclusions concerning modality are grounded in claims concerning what might have been the case, or what could have been the case, or

---

<sup>8</sup> Dummett 1973 makes a similar point in discussing Kripke’s argument.

what would have been the case had something else been the case. Kripke is explicit (1980, pp. 36-37) in tying his notion of necessity to these formulations, and almost all of his arguments for modal claims proceed via these claims. What all these formulations have in common is that they involve scenarios that are acknowledged not to be actual, and that are explicitly considered as counterfactual scenarios.

All these claims are *subjunctive* claims, not in the syntactic sense, but in the semantic sense: they involve hypothetical situations that are considered as counterfactual. The paradigm of such a claim is a subjunctive conditional: 'if P had been the case, Q would have been the case'. We can say that all these claims involve a *subjunctive context*, where a subjunctive context is one that invokes counterfactual consideration. Such contexts include those created by 'might have', 'would have', 'could have', or 'should have' (on the non-epistemic readings of these phrases), subjunctive conditionals involving 'if/were/would be' or 'if/had/would have', and other phrases. In Kripke's sense of 'possible' and 'necessary', where 'it is possible that P' is equivalent to 'it might have been the case that P', then modal contexts such as 'It is possible that' are themselves subjunctive contexts.

Kripke's central point against the description theory was that names and descriptions function differently in modal contexts: for a name 'N' and a description 'D', it is necessary that D is D (if it exists), but it is not necessary that N is D. We can put this somewhat more precisely by saying that names and descriptions function differently in subjunctive contexts. And more generally, names and descriptions seem to behave differently under subjunctive evaluation of hypothetical possibilities. How can a Fregean view handle this phenomenon?

I have not argued that names are equivalent to descriptions. I do allow, however that a name and a description can have the same sense, at least approximately and in some cases, as the case of 'Hesperus' suggests. If so, then if Kripke's point is accepted, we must explain the different truth-value of subjunctive sentences by appealing to something other than a difference in sense. So the question is: why is 'it might have been that Hesperus was not the evening star' true, while 'it might have been that the evening star was not the evening star' is false, given that 'Hesperus' and 'the evening star' have (roughly) the same sense?

There are a number of ways in which one might try to explain this. First, one might appeal to a difference in underlying logical form between subjunctive sentences containing names and descriptions. One could hold that in subjunctive contexts, names always take wide scope, so that the sentence involving ‘Hesperus’ above has the logical form of ‘Hesperus is such that it might not have been the evening star’. Or one could hold that names always involve an unarticulated ‘actual’, so that the sentence has the logical form of ‘it might have been that the actual Hesperus was not the evening star’. Once this logical form is in place, then substituting ‘the evening star’ for ‘Hesperus’ yields the same truth-value. On such a view, the difference in the modal status of the sentences is due to a difference in logical form, not a difference in the sentence’s semantic contents. These explanations merit considerable discussion, but I will set them aside here.<sup>9</sup>

The second sort of explanation appeals to a semantic difference between names and descriptions, in some aspect that goes beyond their sense, and that affects how they function in subjunctive contexts. On such a view, a name may have sense and extension, but sense and extension do not *exhaust* the meaning of a name. The simplest addition to the Fregean framework would be a semantic feature that is part of the meaning of all names but not part of the meaning of descriptions, indicating that in counterfactual contexts, the expression contributes its *actual* extension.<sup>10</sup> There will be a corresponding semantics of counterfactual contexts, such that the presence or absence of the feature is relevant to the truth-value of counterfactual statements. This will explain the difference between names and descriptions.

I favor a more general semantic explanation of the second sort. On this account, every expression is associated with *two* intensions: one governing its application to epistemic possibilities, and one governing its application to explicitly counterfactual possibilities. The first of these is the epistemic intension. The second is a *subjunctive intension*. Like the epistemic intension, this is a function from worlds to extensions. But here the worlds in question are seen as counterfactual metaphysical possibilities, and

---

<sup>9</sup> See Dummett 1973, Sosa 2001, and Stanley 1997 for views of this sort, and Soames 1998 for critical discussion.

<sup>10</sup> Recanati (1993) proposes a feature (labeled ‘REF’) that works something like this, though he conceives of it differently.

expressions are evaluated in these worlds in the way in which we evaluate counterfactual scenarios.

Just as the epistemic intension mirrors the way that we describe and evaluate epistemic possibilities, the subjunctive intension captures the way that we describe and evaluate subjunctive possibilities. To evaluate the subjunctive intension of a sentence *S* in a world *W*, one can ask questions such as: if *W* had obtained, would *S* be the case? For example, if *W* is a world in which Venus was knocked off course by a comet and in which Mars was prominent in the evening, we can say the following. If *W* had obtained, Hesperus would not be visible in the evening; if *W* had obtained, Hesperus would still be Phosphorus, and would still be Venus; if *W* had obtained, Hesperus would not be Mars. So the subjunctive intension of ‘Hesperus is Phosphorus’ is true in *W*.

The same sort of thing applies to a natural kind term, such as ‘water’. If *W* is the Twin Earth world, we can say: if *W* had obtained, water would still have been H<sub>2</sub>O, not XYZ. So the subjunctive intension of ‘water is XYZ’ is false in that world, and the subjunctive intension of ‘water’ picks out H<sub>2</sub>O in that world. One can tell a related story for a general term such as ‘cat’, whose subjunctive intension will pick out a class of members of a particular biological species in all worlds, and for a property term such as ‘hot’, which will arguably pick out a certain sort of molecular motion in all worlds. (Note that at least in the ‘cat’ case, the extension is a different class in each world, and is not always the actual extension. This can be handled straightforwardly by a subjunctive intension, but it is harder for the other accounts above to handle, as they rely on projecting the actual extension.)

Indexicals such as ‘I’ can also be accommodated naturally. The epistemic intension picks out the being at the center of a world, but the subjunctive intension does not. The subjunctive intension of my use of ‘I’ picks out me (David Chalmers) in all worlds. I can say ‘if David Chalmers were in a coma, then I would be in a coma’, and so on. In evaluating counterfactual scenarios, ‘I’ always picks out David Chalmers (though of course, he need not be *called* that). So we do not need a center to evaluate the subjunctive intension of ‘I’. More generally, the worlds in the domain of subjunctive intensions can be taken to be standard uncentered worlds, not centered worlds.

Just as every expression has an epistemic intension, every expression will have a subjunctive intension. This intension will be a function from worlds to extensions of the appropriate sort: individuals, truth-values, and so on. And as long as the extension of a sentence depends on the extensions of its parts, the subjunctive intension of a sentence will depend on the subjunctive intension of its parts.

A name and a description may have similar epistemic intensions and similar extension, but they may have very different subjunctive intensions. This can be seen by examining world W above. Here we can say that if W were the case, then Hesperus would still be Venus, but the evening star would be Mars. So the subjunctive intensions of 'Hesperus' and 'the evening star' are distinct. The former picks out Venus in all worlds, while the latter picks out something that is visible in the evening in all worlds. More generally, the subjunctive intension of a name picks out its actual extension in all worlds, while the subjunctive intension of a description picks out whichever object satisfies the description in a given world.

The subjunctive intension of a name depends directly on its actual extension. The extension of a name can usually not be known a priori, so the subjunctive intension cannot be known a priori, either. If we lack relevant empirical information about the actual world, we might be unable to evaluate an expression's subjunctive intension at a counterfactual world, even given a detailed specification of that world. For similar reasons, two names (such as 'Hesperus' and 'Phosphorus') may have the same subjunctive intension without the subject knowing that a priori. The same goes for natural kind terms and indexicals. So subjunctive intensions do not reflect the cognitive significance of the expressions involved. This contrasts with the subjunctive intensions of descriptions, which can often be evaluated a priori (at least if the descriptions contain no names, natural kind terms, or indexicals), and it contrasts with epistemic intensions, which can generally be evaluated a priori for any expression.

Subjunctive intensions are most directly relevant to the evaluation of subjunctive sentences. When an expression occurs in a subjunctive context ('it might have been that S'; 'if S had been the case, T would have been the case'; 'if A were B, C would have been D'; and so on), its subjunctive intension is used in evaluating the truth-value of the sentence. For example, 'it might have been that S' (at least in the relevant sense of 'might

have been') will be true iff the subjunctive intension of S is true at some world. 'If S had been the case, T would have been the case' is true roughly iff the nearest world that satisfies S's subjunctive intension satisfies T's subjunctive intension. And so on.

Subjunctive intensions are also relevant to the evaluation of modal contexts, at least when these are interpreted the Kripkean way. 'It is necessary that S' will be true iff S's subjunctive intension is true at all worlds. 'Possibly, T' will be true if T's subjunctive intension is true at some world. This is what we would expect, given that 'It is possible that T' will be true precisely when 'It might have been that T' is true.

By analogy, epistemic intensions are relevant to the analysis of corresponding epistemic contexts. The indicative conditionals 'If P is the case, then Q is the case' or 'If it turns out that P, it will turn out that Q' are true (or acceptable) for a subject approximately when the epistemically nearest world that satisfies P's epistemic intension also satisfies Q's epistemic intension, where "epistemically nearest" is defined in a way that depends on the subject's knowledge or perhaps beliefs. (See "The Tyranny of the Subjunctive".) And 'it is a priori that P' will be true if P has a necessary epistemic intension.

(The material in the remainder of this section is a somewhat more elaborate development of this framework, and can be skipped by those who prefer to move on.)

The subjunctive intension of an expression is always determined by the epistemic intension of the expression and by the character of the actual world. Because of this, it is possible in principle to associate a *two-dimensional intension* with an expression, which captures how its subjunctive intension will vary, depending on which epistemic possibility turns out to be actual. This two-dimensional intension can be thought of as a function from pairs (V, W) of epistemic possibilities and metaphysical possibilities to extensions.

To evaluate the two-dimensional intension of a statement S at a pair of worlds (V, W), one can ask: if V is actual, then if W were the case, would S be the case? This is reflected in some more natural English constructions: for example 'if water is XYZ, then water could not be H<sub>2</sub>O'. To determine the truth-value of statements like these, one needs the full two-dimensional intension, as neither epistemic intension nor the subjunctive intension of the terms involved carries information about how to evaluate subjunctive

statements under alternative epistemic possibilities. So for a fully general account of sentence's truth-values, one needs at least the full two-dimensional intension, although its full structure will be relevant only in rare cases.

Like an epistemic intension but unlike a subjunctive intension, a two-dimensional intension can be evaluated a priori. One needs no empirical information about the actual world, since all the relevant information is specified in the epistemic possibility. One might hold that the two-dimensional intension represents the true 'cognitive significance' of an expression, if one holds that the difference between the behavior of names and descriptions in counterfactual contexts represents a difference in cognitive significance; so one could hold that the two-dimensional intension is a Fregean 'sense' in some expanded understanding of the term. But we do not need to adjudicate that matter here.

An expression's actual subjunctive intension can be derived from the two-dimensional by using the speaker's actual scenario as the epistemic possibility parameter; this cannot be evaluated a priori precisely because we do not know which scenario is actual. It is tempting to reconstruct an expression's epistemic intension as the 'diagonal' of the two-dimensional intension, which results when the same possibility is used as the epistemic possibility parameter and (in an uncentered version) as the metaphysical possibility parameter. This depends on just how the possibilities are understood, however.

As before, there are two ways of understanding the class of epistemic possibilities. They can be understood in terms of the same worlds that function as metaphysical possibilities, with the addition of a center. Or they can be understood as a separately defined class of epistemically possible scenarios. The first option is more elegant, but requires a substantive philosophical thesis about possibility. The second option is more complex, but it requires fewer philosophical commitments. A philosopher who holds that the existence of a god is metaphysically necessary but not a priori can embrace the second option but not the first, for example.

If one takes the second option: then one has two distinct classes of possibilities: the epistemically possible scenarios and the metaphysically possible worlds. These possibilities may have certain relations to each other, but the two sets are nevertheless disjoint. One can evaluate a statement in an epistemically possible scenario, yielding the value of its epistemic intension there, and one can evaluate it in a metaphysically possible



world, yielding the value of its subjunctive intension there. But because the spaces are distinct, there is no natural way to map epistemic possible scenarios onto metaphysically possible worlds. As a result, one cannot reconstruct an epistemic intension as the ‘diagonal’ of the two-dimensional intension, and various other elegant properties are removed. This view also has the disadvantage of requiring a strong underlying modal dualism, with distinct modal primitives for each space of worlds.

If one takes the first option: then the worlds that function as epistemic possibilities are the same as the worlds that function as metaphysical possibilities, with the addition of a center. So epistemic intensions and metaphysical intensions are defined over (almost) the same space of worlds. To evaluate an expression’s epistemic intension at a world, we consider the world as an epistemic possibility: as a way our world might actually be. To evaluate an expression’s subjunctive intension at a world, we consider the world as an explicitly counterfactual possibility: as a way our world might have been, but (probably) is not. One can say that in the first case, we consider the world as actual; in the second case, we consider the world as counterfactual.

The first option is compatible with all the familiar cases. Take the case of ‘water’, and consider a Twin Earth world where the watery stuff is XYZ (near the center, if required), and where H<sub>2</sub>O is not watery. When we consider this world as actual, it is an instance of the epistemic possibility that water is XYZ: if W is actual, then water is XYZ. When we consider this world as counterfactual, it is an instance of the metaphysical possibility that water is H<sub>2</sub>O: if W had been actual, then water would still have been H<sub>2</sub>O, not XYZ. Something similar applies in all the familiar cases.

One might ask: is *W itself* a world where water is XYZ, or a world where water is H<sub>2</sub>O? On a standard philosophical view, it is the latter, not the former. This is because the phrase ‘a world in which S’ is almost always read as invoking a world W in which the subjunctive intension of S is true, or for which S is true when W is considered as counterfactual. My own view is that this reading is arbitrary, and that the phrase ‘a world in which S’ is ambiguous between readings that invoke epistemic and subjunctive intensions. If so, there is no determinate answer to the question above. To remove the ambiguity one can define new locutions: for example, when the epistemic intension of S is true in W, then W is a world *at* which S; when the subjunctive intension of S is true in

W, then W is a world *of* which S. In any case, this terminological issue does not matter too much for our purposes. As long as we are always clear about how we are evaluating statements, and have clear conventions for understanding the relevant phrases, no confusion should result.

If the first option is accepted, various matters become more straightforward. Because there is only one space of worlds (apart from the difference involving centers), one can reconstruct an epistemic intension from the two-dimensional intension as a ‘diagonal’ intension, where the value of the epistemic intension at W is the value of the two-dimensional intension evaluated at W (considered as actual) and an uncentered version of W (considered as counterfactual). This view is also much simpler metaphysically, since it is compatible with a deep underlying modal monism, with just one space of worlds. Of course there is still a dualism of epistemic and metaphysical possibility in language, but this simply arises from the dual nature of semantic evaluation over a single space of worlds.

Given that any sentence S has an epistemic intension and a subjunctive intension, someone might ask: which of these is the *content* of S? Which of these gives the *truth-conditions* of S? What is the *proposition* expressed by S? My view is that we need not settle these questions one way or another. We can say that the epistemic intension is S’s *epistemic content* and that the subjunctive intension is S’s *subjunctive content*. As for the content of S (unqualified), this is a complex content that subsumes both of these and possibly more. Similarly, S has *epistemic truth-conditions* (showing how S’s truth depends on how the world turns out), and *subjunctive truth-conditions* (showing how S’s truth varies in counterfactual possibilities).

As for propositions, I have avoided this terminology, as it is multiply ambiguous. If one’s conception of a proposition is a set of possible worlds (or something similar, such as a structure of intensions), then one could say that S expresses two propositions, an epistemic proposition and a subjunctive proposition. But if one’s conception of a proposition is more generally of what remains semantically of S once the arbitrary clothing of a given language is stripped away, then one could say that S expresses a complex proposition that involves a two-dimensional structure. One should not run these two conceptions together: for example, the fact that an utterance of S expresses two

propositions in the first sense in no way entails that the utterance is ambiguous, since ambiguity would involve expressing two propositions in the second sense. For my part, I prefer to use ‘proposition’ in the second, more general way; but I will largely avoid the expression here.

It is sometimes objected that the epistemic intension cannot be part of the content expressed by S at all, since it is really a matter of the content that S *would have* expressed had a token of S been present in a different context. This is a mistake. As we saw earlier, the epistemic intension of S can be evaluated in worlds that contain no token of S, and even if a world does contain such a token, its presence is usually irrelevant to evaluating an epistemic intension there. The epistemic intension of S is not defined in terms of counterfactual tokens of S at all. Rather, it is defined in terms of the first-order use of an *actual* expression S in evaluating various epistemic possibilities. This is precisely analogous to the way that the subjunctive intension of S is defined in terms of the first-order use of S in evaluating counterfactual possibilities. So epistemic intensions and subjunctive intensions are on a par here.

One could define the *contextual intension* of S as a function that maps centered worlds containing a token of S at the center to the extension of that token in that world. This might resemble an epistemic intension in some ways, but it would not be an epistemic intension. First, an epistemic intension will be defined over many more worlds. Second, it is unclear just what it takes to be a token of S, and depending on what we require, the contextual intension may give very different results from an epistemic intension. If only orthographic properties are required, then there will be worlds where a token of ‘water’ refers to horses or to the number two. If the same extension is required, then a token of a name such as ‘Hesperus’ will pick out the same individual (Venus) in all worlds. Perhaps there is an intermediate requirement that gives roughly the same results as an epistemic intension, but this is not obvious. The obvious suggestion is to presuppose the notion of an epistemic intension, and require that a token of S have the same epistemic intension of S. However one does things, there is not much point in defining such a notion, since we have epistemic intensions to do the job already.

This matter bears on the two-dimensional semantic frameworks developed by Kaplan (1989) and Stalnaker (1974). The framework I have developed here resembles

these in obvious respects, and it owes much to them. Epistemic intensions are analogous in certain respects to Stalnaker's "diagonal propositions" and to a version of Kaplan's "character", and subjunctive intensions are much the same as Stalnaker's "propositional content" and Kaplan's "content". But there are crucial differences. First: on Kaplan's and Stalnaker's frameworks, the analogs of epistemic intensions are defined in terms of the analogs of subjunctive intensions, whereas on the framework I have outlined, they are defined quite independently of subjunctive intensions. Second and most important: on Kaplan's and Stalnaker's frameworks, the worlds on the first dimension of evaluation are not considered as epistemic possibilities but as contexts of utterance, where evaluation requires the presence of a token of the expression within the context. As a result, this first dimension of evaluation yields a contextual intension, not an epistemic intension.

These features lead to strong limitations on using these frameworks for epistemic purposes, which both Kaplan and Stalnaker note. Kaplan's framework yields useful epistemic results only for indexicals and demonstratives, and not for names and natural kind terms. This is because the contextual intensions for indexicals and demonstratives behave much like their epistemic intensions, but the contextual intensions for names and natural kind terms behave very differently. (A name arguably picks out the same individual in every context, but it can apply to different individuals within epistemic possibilities.) Stalnaker (1999) argues that his "diagonal proposition" does not reflect matters of apriority directly, partly because of the problems involved in holding the meaning of an utterance constant across contexts. It seems that these frameworks are useful for epistemic purposes precisely to the degree that the notions involved resemble epistemic intensions.

A strong Fregean might criticize this framework from the other side, holding that only epistemic intensions, not subjunctive intensions, are part of the content of a sentence, and the subjunctive behavior can be accounted for by appealing to logical form or to a simpler semantic feature, as above. This matter is not cut and dried, but there are a number of advantages to including subjunctive intensions as an explicit semantic value.

First, subjunctive intensions allow a direct parallel between the treatment of apriority and of necessity, and between epistemic and subjunctive evaluation. Second, subjunctive intensions give a more general account of subjunctive behavior than accounts

that rely on actual extension, and consequently can more easily account for the distinctive subjunctive behavior of certain terms, such as ‘cat’. Third, we have at least one content (the subjunctive intension) associated with an expression that is “objective” in the way that Frege required. Fourth, people have intuitions about “what is said” by an expression, and sometimes these better reflect subjunctive intensions than epistemic intensions (e.g. when I say ‘I am Australian’ and you say to me ‘You are Australian’, there is a sense in which we say the same thing, although our epistemic intensions differ). Fifth, the presence of subjunctive intensions allows for a degree of continuity with current philosophical frameworks in which a central role is played by notion of “propositional content” closely tied to subjunctive intensions; this continuity allows a Fregean to retain some of the insights of this tradition rather than discarding them completely.

The addition of subjunctive intensions certainly goes beyond Frege’s view. But it is a supplement to the view rather than a radical overhaul. Senses, or epistemic intensions, are still present and playing the same role they always played. Sense and extension have merely been supplemented by a further semantic value in order to deal with various subjunctive and modal phenomena. Frege did not address these phenomena directly, so it is hard to know how he would have dealt with them. But I think this framework is compatible with the broad spirit of a Fregean view.

## **8 The epistemic argument**

Kripke’s second central argument against descriptive views of language is an epistemological argument.<sup>11</sup> Kripke recognizes that a description theorist might accept that names are not modally equivalent to descriptions, but might nevertheless hold that names are epistemically equivalent to descriptions. So he raises some quite different considerations to argue that this is not the case. Again, while these arguments are most explicitly aimed at a description theorist, they are generally taken to have force against any broadly Fregean view.

---

<sup>11</sup> Some discern a third “semantic argument” in Kripke, closely related to the epistemic argument. The considerations used here to reply to the epistemic argument can also be used to reply to the semantic argument.

The argument proceeds roughly as follows. The description theorist will hold that a name N (say, 'Gödel'), as used by a given speaker, is epistemically equivalent to some description or cluster of descriptions D (say, 'the person who proved the incompleteness of arithmetic'). On such a view, 'N (if it exists) is D' will be a priori for the speaker. But Kripke argues that for many names N, no such sentence is a priori. In effect, he does this by arguing that for any description D, 'N is not D' is epistemically possible for the speaker. If so, then N is not epistemically equivalent to D.

Take the name 'Gödel'. A speaker may associate a number of descriptions with her use of the name: 'the person who proved the incompleteness of arithmetic', and so on. But Kripke argues that there are (epistemically possible) scenarios in which the name 'Gödel' will turn out to refer to a person who does not satisfy any of the descriptions. So the speaker cannot know a priori that Gödel (if he exists) satisfies any of the descriptions.

One can put the argument strategy as follows. Take the name 'Gödel' and the description 'the man who discovered the incompleteness of arithmetic'. Then there is an epistemically possible scenario in which the incompleteness of arithmetic was discovered by a man named 'Schmidt', and in which the proof was stolen and published by a man named 'Gödel', to whom the proof was thereafter attributed. If this scenario is actual, then the speaker's term 'Gödel' refers to the second man, not to the first. So it is epistemically possible for the speaker that Gödel did not prove the incompleteness of arithmetic, and there is no a priori equivalence between the name and the description.

Kripke argues that the same can be done for any description associated with the name (such as 'the man to whom the discovery of the incompleteness of arithmetic is commonly attributed'). There will always be epistemically possible scenarios such that if they are actual, the speaker's term 'Gödel' refers to someone who does not satisfy the description. (For example, a scenario in which the discovery is now commonly attributed to Hilbert, without the speaker realizing.) If this is right, then the name cannot be equivalent a priori to any description.

Kripke also argues that there are actual cases in which most descriptions that a speaker associates with a name are false of the name's referent. This applies to 'Peano' and 'Jonah', for example: Peano did not discover the axioms associated with him, and Jonah was probably not swallowed by a whale. So here it is not even true, let alone a

priori, that the name and the associated description are coextensive. And he also argues that there are some names such that the speaker has *no* associated description that could fix reference: a speaker can use the term 'Feynman' to refer to Feynman while knowing nothing more than that he is a famous physicist, for example, where that description is satisfied by many individuals. So again, it seems that reference is not fixed descriptively.

Does this argument against the description theory 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 speaker 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 accommodate Kripke's 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 specified as 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.

In a way, what is going on here is analogous to what goes on in the analysis of a term such as ‘knowledge’, as discussed above. Someone might hold that ‘knowledge’ is equivalent a priori to some description, such as ‘justified true belief’. But then we come up with a scenario such that in that scenario, something falls within the extension of ‘justified true belief’ but not of ‘knowledge’. So we conclude that the two are not equivalent a priori. The process repeats itself for other descriptions, suggesting that ‘knowledge’ is not a priori equivalent to any such description. This suggests that the intension of ‘knowledge’ cannot be precisely captured in a linguistic description.

In a similar way, 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. And the epistemic intension still mirrors the cognitive significance of the name. The identity ‘N is D’ is a posteriori precisely because the two expressions have different epistemic intensions; that is, precisely because there is an epistemic possibility where they come apart. The subject has the ability to evaluate the name’s referent within this epistemic possibility, just as can be done with expressions in general.

The intension of an expression such as ‘knowledge’ can at least be approximated by certain linguistic descriptions, such as ‘justified true belief’, and by longer and longer versions that come gradually closer to the true intension. One might wonder whether something similar can be done with a name such as ‘Gödel’. Can the epistemic intension of the name in the case above at least be approximated by a linguistic description? This is not compulsory for the intensional framework, but it can at least be enlightening to look. A side benefit is that it provides some sort of at least approximate account of the features of the world in virtue of which the epistemic intension applies.

To answer this question, one needs to consider: when speakers use a name such as ‘Gödel’ or ‘Feynman’ in cases such as those above, 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 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’.<sup>12</sup> 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. So maybe the description above is at least a good first approximation.<sup>13</sup>

There are two sorts of objections that might be made to this sort of description. First, it might be held that it does not always give the right results. For example, it may be epistemically possible that the speaker misheard or misremembered the name, and that others were really using the name ‘Fireman’. In such an epistemic possibility, the description above will give the wrong results. But this is just the sort of thing that we should expect, given the imperfection of descriptions. As with ‘knowledge’, we could try to move to closer approximations. 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, just as with ‘knowledge’. But as in all these cases, the most this shows is that any such approximation

---

<sup>12</sup> Kripke considers some potential descriptions in this vicinity. He discusses ‘the man Jones calls ‘Gödel’ (1980, p. 92), where the speaker believes he acquired the name from Jones, and dismisses it on the ground that the belief may be false. And he discusses ‘the man called ‘Glumph’ by the people from whom I got it (whoever they are), provided that my present determination of the reference satisfies the conditions sketched in *Naming and Necessity* and whatever other conditions need be satisfied’ (1980, p. 162), and dismisses this on grounds connected to the second part of the description. Surprisingly, he never considers the obvious intermediate description (e.g. ‘the man called ‘Glumph’ by the people from whom I acquired the name’) that avoids both of these problems.

<sup>13</sup> Replies to Kripke’s epistemic arguments that appeal to metalinguistic descriptions of this general sort are given by Searle 1984, Fumerton 1989, Jackson 1998, and Lewis 1986, among others.

is imperfect. One refutes these approximations by evaluating the epistemic intension in certain epistemic possibilities and showing that the approximation gives the wrong results; so this sort of argument does nothing to show that the epistemic intension does not exist.

Second, it might be held that this sort of description is “circular”, perhaps because it appeals to the notion of reference. The circularity is not obvious, however. If one appealed to the notion of reference in a definition of ‘reference’, there would be a danger of circularity. But a definition of ‘reference’ is not being offered here. If the descriptivist were to offer a definition of ‘reference’ it might be something like ‘whatever satisfies a canonical description D associated with N’. This is a general, noncircular definition with no appeal to reference (except in the notion of satisfying a description, which presumably is to be accounted for separately). At worst there is a danger that if a description D involves ‘reference’, it will not yield a determinate result, since evaluating the description will require evaluating reference, which will require evaluating another description, which will require evaluating reference, and so on. But a descriptivist can reply that this is merely a recursive situation, not a circular one, and that the process will always eventually be grounded in a use of the name whose associated description does not involve ‘reference’ at all, such as an initial baptism. The descriptivist might even hold that this is a natural way to capture the insights of the causal theory of reference in a descriptivist framework: reference proceeds through recursive deference to others, ultimately grounded in an initial baptism.

In any case, this worry does not arise on the intensional framework. The epistemic intension of a name is simply a function from worlds to individuals that reflects a rational ability to determine a specific individual in a given epistemic possibility. It is not a description, and so makes no use of ‘reference’. Perhaps one might worry that if something like the picture above is right, then *evaluating* the epistemic intension of a term like ‘Gödel’ at a world will require having explicit information about others’ reference within that world, and that this would be circular. It is not clear why this would be circular, but in any case, explicit information about reference would not be required. Information about the epistemic intensions that others associate with various names would suffice (along with other information about causal relations, the properties of

various objects, and so on). This information might itself be derivable from information about other mental and/or physical states, or it might be some sort of mental primitive; that does not matter. With information about others' epistemic intensions, reference will be determined. And if there is a worry about *evaluating* others' epistemic intensions given that they are deferential also, then knowledge of epistemic intensions across a whole community (including its history) will suffice, since the deference will ultimately be grounded in a nondeferential use.

So given relevant information about the physical and mental states of individuals in the community, there will be no problem evaluating epistemic intensions. And the use of this information is in no way ad hoc; it corresponds to the information that we use in evaluating reference across various epistemic possibilities. In particular, when a name is used deferentially, information about the linguistic and cognitive practices of others will always be relevant.

Of course not every use of a name is a deferential use, so not every epistemic intension will function in this way. When a name is introduced in an initial baptism, there will be no deferential element involved; to evaluate reference in the actual world (and across epistemic possibilities), the speaker will not usually need information about the cognitive states of others. The same goes for some names used for very familiar referents. Say that a wife uses the name 'Fred' for her husband and has done so for years. In such a case, even were the speaker to discover (to her surprise) that no-one else in the community used that name for her husband, she would still reasonably hold that the name she uses refers to the spouse. If such a situation turned out to be actual, her utterance 'Fred is my husband' would plausibly be true, not false. This suggests that the epistemic intension of her use of the name has no deferential element.

There are also many intermediate cases, where a name is used with some mixture of deferential and nondeferential elements, so that for a speaker to determine reference of the name, relevant information will include both information about others' usage and independent information about properties of the referent (perhaps corresponding to some of the speakers' beliefs involving the name). This sort of intermediate case will be necessary to account for cases such as 'Madagascar' (Evans 1977), where the referent of our use of the name (an island) differs from the referent of the original use (part of

mainland Africa). If every use since the initial baptism was entirely deferential, this could not happen. So some uses in the causal chain must have been not entirely deferential, with epistemic intensions that were partly influenced by a speaker's beliefs. It is easy to imagine that even if the beliefs have only a small influence on a given speaker's epistemic intension, the effect of this influence would amplify as a causal chain proceeds. If every speaker in the chain has a small component of influence from the belief, then deference to a speaker whose epistemic intension is also influenced by the belief will increase the effect. In the case of 'Madagascar', the result might be that uses of the name would initially refer to the mainland location, would proceed through a period of divided reference, and would eventually emerge as referring to the island.

(One other subtlety: to evaluate a deferential use of the name in an epistemic possibility, the speaker may need the name itself (as used by her) to be present in an epistemic possibility, so that she can determine where that name was acquired. This suggests that deferential uses are exceptions to the principle that tokens of an expression need not be present within an epistemic possibility. I think that this sort of case is best handled by having one or more optional marked thoughts ('this thought') present at the center of a world, which be used to trace deferential reference in an epistemic possibility.)

Kripke also gives an epistemic argument against descriptive views of natural-kind terms (1980, pp. 116-23). He argues that it is not a priori that gold is yellow, or that tigers are striped, and so on, in effect because there are epistemic possibilities in which these statements could turn out to be false (if we were suffering from various illusions, for example). In a similar way, Putnam (1975) argues that it is not a priori that cats are animals, since there are epistemic possibilities in which it would turn out that they are robots from Mars. As before, these arguments do not apply to the intensional account: they proceed by evaluating an epistemic intension at various epistemic possibilities, and they show at most that the epistemic intension is not equivalent to the relevant description. They also suggest that a causal link between the term and the referent plays an important role; so any better descriptive approximation of the intension should give a significant role to this link.

I conclude that Kripke's epistemic arguments have no force against an intensional account of Fregean sense.

## 9 The argument from variability

Another argument against Fregean view is rarely articulated explicitly,<sup>14</sup> but it may motivate some opposition to the Fregean view. This argument notes that different speakers may associate *different* cognitive significance with the same name. For example, an identity such as "Bill Smith is William Smith" may be cognitively significant for one speaker (e.g. one who has heard of the same person under two names in two different contexts), and cognitively insignificant for another (e.g. the person's partner, who uses both names indiscriminately). So if Fregean sense is to reflect cognitive significance, then the sense of a name must vary between speakers. But if the sense of a name can vary between speakers, it is not part of the *meaning* of the name at all. The same may apply to other expressions, such as natural kind terms: the cognitive significance of such a term is variable, so its sense must be variable, so its sense cannot be part of its meaning.

We need not spend much time on the first part of this argument. On the account I have given, it is clear that the epistemic intension of a name can vary between speakers. We saw this in the discussion just concluded: the first user of a name may use the name with one epistemic intension, and later users may use it with a quite different epistemic intension. When Leverrier introduced the term 'Neptune' as a term for whatever planet was perturbing the orbit of Uranus, then the epistemic intension of his use of the term functioned roughly as described. But the next speaker — perhaps his wife, who knew only that Neptune was an astronomical object for which her husband was searching — might have used it with a different epistemic intension. And later users might well use it in a deferential way (as with 'Feynman' above), with an epistemic intension that reflects this. So different speakers can clearly have different epistemic intensions for the same name.

---

<sup>14</sup> Frege himself (in "The Thought") addresses objections based on variability, in the context of indexicals rather than names. Burge (1979) stresses that Frege's conception of sense differs from contemporary conceptions of linguistic meaning, in part because of variability.

Something similar may apply to natural kind terms, though this is not quite so clear. For example, two speakers might 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). It might be that for the first speaker, the epistemic intension of ‘water’ functions to pick out (roughly) a substance that takes on a certain liquid form, and that for the second speaker, the epistemic intension of ‘water’ functions to pick out a substance that takes on a certain solid form. In the actual world, these epistemic intensions both pick out the same substance, but in other scenarios, their extensions will differ. But arguably both are using the same word ‘water’. If so, there is no epistemic intension for ‘water’ that is common across all users of the term.

The same applies in an even stronger way to demonstratives such as ‘that’. Here different uses of the same term by the same speaker can have different epistemic intensions. It is plausible that the epistemic intension of ‘that’ depends at least in part on a speaker’s intentions, which may differ between uses of the term: on one occasion, the speaker may intend to refer to an object on her left, and on another occasion, to an object on her right. If so, the epistemic intensions on these different uses may differ. So there is no epistemic intension for ‘that’ that is common across all uses of the term, even for a single speaker.

This suggests that for a general account, epistemic intensions cannot be assigned to expression *types* but rather must be assigned to expression *tokens* (or perhaps to expression types in contexts of use). There will be some terms for which epistemic intensions are constant across all tokens of an expression type — some descriptive terms and indexicals, for example — in which case an epistemic intension can also be assigned to the type. And for names and natural kind terms, epistemic intensions might be constant at least across an individual speaker’s use of an expression (at least within a limited time frame), so there could be an assignment to more limited types. But a fully general account requires that epistemic intensions are assigned to tokens.<sup>15</sup>

---

<sup>15</sup> It should be noted that this difference in epistemic intensions between speakers goes along with differences in which sentences are a priori for a given speaker. For example, ‘Neptune (if it exists) affects the orbit of Uranus’ may be a priori for Leverrier, but not for his wife. Leverrier could use the sentence to

What of the second part of the argument: that if epistemic intensions can vary between tokens of a type, then they are not meanings? This issue is largely terminological. If it is stipulated that meanings are constant across all tokens of a type, then epistemic intensions are not meanings. If this is not stipulated, then epistemic intensions might be meanings. We could distinguish ‘type meanings’ and ‘token meanings’, and allow that epistemic intensions are not (in the general case) type meanings, but they are token meanings. Or we can use a different term, such as ‘content’, for the sort of meanings that can vary between tokens of an expression type. It is not clear that a substantive issue remains once the terminological issue is cleared up.

It might be insisted that if epistemic intensions can vary between tokens of a type, then they are not part of *language*, they are not an aspect of *linguistic* content, and that perhaps they do not fall within the domain of the philosophy of language at all. Again, this is a terminological issue. One might say that the epistemic intensions of names are not part of *a language* such as English, where this is considered as what is common between all English speakers. But this is no reason to deny that they are part of language in a broader sense, and they are in the domain of the philosophy of language. Of course these terminological issues are largely sociological in origin; so if someone resists on all these issues, it may help to point to some historical examples.

First, Frege. As we have seen, Frege himself held that the sense of an expression could vary between tokens of a type, in the case of indexicals and names. Frege’s theory of sense is generally taken to be one of the most important theories in the philosophy of language, and Frege himself is perhaps the most important figure in the field. To hold that variable semantic values are not part of the content of language would have the odd consequence of excluding Frege’s own theory of sense.

Second, Kripke. It is striking that in all his arguments against the Fregean view, Kripke (1980) never mentions an argument from variability. He appears to take it for

---

express a priori knowledge, but his wife could not. So insofar as the apriority of various sentences is invoked in the analysis of epistemic intensions, it will generally be a speaker-relative notion that is relevant. (This is reflected in Kripke’s own discussion (e.g. 1980, p. 73) which usually talks of a sentence being “a priori for a speaker”.) Of course the notion is still distinct from the notion of what the speaker knows a priori; it is a notion of what is knowable a priori for a speaker, given ideal rational reflection.

granted that on a Fregean theory, the descriptive content of a name may vary between speakers (his relevant formulations are all speaker-relative), and does not mention this as an objection. Given the number of other objections that are developed against a Fregean account, this suggests that Kripke assumes that variability alone does not rule out sense as an aspect of language. And it suggests that if the worst problem for a Fregean is variability, then a Fregean view of language would be broadly correct.<sup>16</sup>

Third, Kaplan. Kaplan's theory of character and content (for indexicals and demonstratives) is one of the most important theories in the philosophy of language in recent years. On Kaplan's account, it is clear that both the character and the content of a demonstrative such as 'that' will vary between uses of the term.<sup>17</sup> So if epistemic intensions are not an aspect of language, then neither are character or content.

Someone might object that the epistemic intension of an utterance is not always part of what that utterance *communicates*, since a speaker and a hearer may associate different epistemic intensions with the same expression, and the hearer may not know what the speaker's epistemic intension is. It is true for that this reason, epistemic intensions are not always communicated: a speaker's utterance with a given epistemic intension may cause the hearer to acquire a belief with a different epistemic intension. But it is not clear that this sort of communicative property is a *sine qua non* for any aspect of the content of language. Certainly it is not satisfied by Frege's sense or by

---

<sup>16</sup> This helps to defuse a natural objection: that opposition to Fregean theories was only ever intended as opposition to Fregean accounts of linguistic meaning, so that the existence of variable Fregean semantic values is irrelevant. It is clear that this is not how Kripke (the canonical source of recent opposition to the Fregean view) saw the matter. And if this objection were correct, then Kripke's modal and epistemic arguments would have been quite inessential to the anti-Fregean case, which could instead have proceeded directly via the argument from variability.

<sup>17</sup> Kaplan's formal language for demonstratives deals with this issue by holding that tokens of demonstratives that are associated with different intentions or demonstrations are in fact tokens of a different word. But it is extremely plausible that in a natural language such as English, different tokens of a demonstrative such as 'that' are tokens of the same word; so this formal stipulation still leaves character distinct from linguistic meaning in natural language. See Braun 1996 for critical discussions, and for alternative proposals regarding the linguistic meaning of demonstratives.



Kaplan's character or content, for example. At worst, epistemic intensions are in the same boat as these paradigmatic aspects of linguistic content.

It might also be objected that epistemic intensions are an aspect of "speaker meaning" rather than "semantic meaning". This objection appeals to a distinction that Kripke (1977) exploits to deal with Donnellan's (1966) distinction between the "referential" and "attributive" use of expressions such as 'the man in the corner drinking champagne', used by a speaker intended to refer to a man he is looking at in the corner. If the man in the corner is actually drinking a martini, there is a sense in which the expression refers to him (the referential sense), and a sense in which it does not (the attributive sense). Kripke argues that the attributive reading reflects the "semantic reference" of the expression, while the referential reading is merely an aspect of "speaker's reference". Perhaps something similar applies to epistemic intensions?

I think Kripke's analysis of Donnellan's cases is plausible, but it does not generalize to the case of epistemic intensions. As Kripke says:

The semantic referent of a designator is given by a *general* intention of the speaker to refer to a certain object whenever the designator is used. The speaker's referent is given by a *specific* intention, on a given occasion, to refer to a certain object.

It is clear that on this definition, the epistemic intension of a name or natural kind term is more akin to semantic reference than to speaker's reference, since it reflects a general intention on the speaker's part, not a specific intention. Indeed, the epistemic intension of 'the man in the corner drinking champagne' in the case above plausibly picks out the semantic referent (no-one), not the speaker's referent (the man in the corner). One might, if one wished, introduce a corresponding difference between "semantic epistemic intension" and "speaker's epistemic intension", where the latter picks out the man who is ostended irrespective of his other properties. But even then, epistemic intensions would be in the same boat as reference, and there would be a clear notion that falls on Kripke's "semantic" side of things.

Perhaps there is some *other* distinction that might be drawn between something one might call “semantic meaning” and something one might call “speaker meaning”: for example, if one stipulates that semantic meaning must be constant across all tokens of a type, while speaker meaning can vary between speakers. But that would be a very different distinction from Kripke’s, and it would do nothing to suggest that epistemic intensions are in the same boat as Donnellan’s “referential” uses of descriptions.

A final objection might be that epistemic intensions are not part of language, since they derive entirely from the contents of thought. On this view, it is the speaker’s *concept* of water that has an epistemic intension, and it is the conceptual content that varies between speakers, not any sort of linguistic content. One might respond by accepting that epistemic intensions are associated with concepts and thoughts<sup>18</sup> and that the epistemic intensions of linguistic expressions are derivative in some way on the epistemic intensions of thoughts, while denying that this entails that epistemic intensions are not also part of the content of language. It is not unreasonable to hold that all linguistic content is derivative in some way on mental content (this applies just as much to expressions whose content is constant between speakers as to expressions whose content varies). But if the objection in question were correct, then such a view would be false by definition, or at best, it would entail that there is no linguistic content. Perhaps the view is wrong, but if so, this is a substantive point, not a terminological one. It is also worth noting that Kripke’s own discussion above suggests that linguistic content can derive from the mental content inherent in a speaker’s intentions. So I think that there is no objection to epistemic intension as a sort of linguistic content here.

Ultimately, the best way to deal with any terminological issue is to reflect on the use to which a term is being put, and to determine which sense of the term is most relevant to a given purpose. It may be that there are some purposes for which the most relevant notion of the “meaning” or “content” of linguistic expressions is one on which meaning and content are required to be constant across all tokens of a type. Such purposes might include those of determining what is built into the semantic structure of a language such as English, giving an account of what is required to competently use an

---

<sup>18</sup> See Chalmers 2002b, in which epistemic intensions are used to give an account of the “narrow content” of thought.

expression of a given type, and perhaps addressing certain questions about what an expression will communicate between arbitrary speakers of a language. For these purposes, one can invoke a notion of meaning on which universality is required. On this conception of meaning, epistemic intensions may be part of the meanings of indexicals and some descriptive terms, but they will not be part of the meanings of names and natural kind terms.

For many other purposes, we do not need such a narrow notion of meaning or content, and we often need a broader notion. This arguably applies to most of the uses to which the philosophy of language and notions of meaning are put in other areas of philosophy: that is, the uses in virtue of which philosophy of language is sometimes said to be “first philosophy”. The same also applies to many uses of meaning and content within the philosophy of language itself.

One example: notions of meaning and content are often taken to be central in analyzing questions about necessity and possibility, which in turn play a crucial role in analyzing many metaphysical issues. But for this role, the question of whether meaning or content is constant across speakers is almost entirely irrelevant. It would make very little difference to the deepest issues if there were just one speaker of a language, or if different speakers used terms with different meanings or content. If I use an epistemic intension to reach a metaphysical conclusion about water, the worst possible consequence of variability will be that someone else will not be able to reach a similar conclusion using their term ‘water’. This seems unlikely to happen, due to the common referent. But even if it did, it would do nothing to invalidate my conclusion; it would just mean that someone else would have to express the conclusion differently. Once terminological issues are cleared up (as they often need to be), the substantive points will be as before. So if epistemic intensions are otherwise relevant to answering these questions, as I think they are, it is no objection that they can vary between speakers.

The same goes for many or most applications of the philosophy of language in metaphysics, epistemology, the philosophy of mind, the philosophy of science, and other areas. For most of these purposes, a requirement that meaning or content is constant over all tokens of an expression type is at best irrelevant, and is at worst harmful. If a semantic value can otherwise play a useful explanatory role in these domains, then the variability

of that semantic value will not be a serious objection. If this is right, then to stipulate that variable semantic values do not qualify as meanings or linguistic contents may be to severely constrain the explanatory role of meaning and linguistic content. To do so would be to compromise the status of the philosophy of language as first philosophy, or to send the burden upstairs to the philosophy of mind.

I think it is best instead to recognize that if variable semantic values can do much of the central work that notions of meaning and linguistic content are introduced to do, then they qualify straightforwardly as aspects of meaning or linguistic content. If so, then the argument from variability has little force against a Fregean view.

## **10 Conclusion**

I have argued that a broadly Fregean account of meaning is tenable. On this account, the notion of an epistemic intension plays the role of a Fregean notion of sense. Epistemic intensions are not the same as Fregean sense in all respects, but they are similar in many respects, and they allow versions of the core Fregean requirements on sense to be satisfied. It may be useful to summarize where the core Fregean theses stand in light of the preceding discussion.

Thesis (1), that every expression has a sense, has been preserved, by allowing that every expression has an epistemic intension. It has also been augmented by a thesis holding that every expression has a subjunctive intension in addition to its epistemic intension.

Thesis (2), that sense reflects cognitive significance has been preserved in a slightly modified form. In the modified form, cognitive significance is understood as non-apriority. So when an identity is not a priori, the expressions involved have different senses; when two sentences are not equivalent a priori, they have different senses; and so on.

Thesis (3), that the sense of a complex expression depends on the sense of its parts, has been preserved in a modified form. The thesis as it stands holds for all expressions except those involving modal and subjunctive contexts. To handle such expressions, one needs a slightly modified thesis: that the sense of an expression depends

on the semantic values of its parts, where the semantic value may include elements (such as subjunctive intension) that go beyond sense. It is also the case that all aspects of semantic value (extension, epistemic intension, subjunctive intension) of a complex expression depends on the semantic values of the parts.

Thesis (4), that sense determines extension, has been preserved in the “weak determination” version that holds that sense determines extension in combination with the world.

Thesis (5), that the sense of a sentence has an absolute truth-value, has been discarded. The sense of a sentence has a truth-value only relative to a subject and a time. Thesis (6), that expressions refer to their senses in indirect contexts, has also been discarded. A more complex account of these contexts is still required. Thesis (7), that the sense of an expression can vary between speakers and between occasions of use, has been preserved.

I have also argued that the most common objections to Fregean theories can be handled by such an account. At most, these objections show that (i) senses are indexical, (ii) senses should be supplemented by a further semantic value, a subjunctive intension, (iii) senses should be understood as intensions, not descriptions, and (iv) the sense of an expression type can vary between speakers and between occasions of use.

I have argued that extension, epistemic intension, and subjunctive intension are all part of the meaning and content of an expression, but I have not argued that these exhaust the meaning or content of an expression. In fact I think that they do not. First, there are plausibly aspects of meaning that have nothing to do with the determination of truth: the difference between ‘and’ and ‘but’ is an example. Second, there may be expressions that are a priori equivalent to each other, but that nevertheless have different meaning due to some more fine-grained cognitive difference: the difference between ‘equilateral triangle’ and ‘equiangular triangle’ is an example. To handle this last sort of difference, I think that one may need senses of a variety that are more fine-grained than epistemic intensions. The notion of an epistemic intension might be extended to do this (by moving to a more fine-grained space of epistemic possibilities), but this is a separate story (see Chalmers forthcoming b). It may also be that for some purposes, the meanings or

contents of complex expressions may need to be taken as structured complexes of extensions and/or intensions; this is quite compatible with the framework I have outlined.

I have also not given a conclusive demonstration that epistemic intensions of expressions exist and have all the properties I have attributed to them. A conclusive demonstration would require some more precise definitions, and a rebuttal of all counterarguments. I have argued that there is a strong prima facie case that epistemic intensions exist and have the properties I have attributed to them, however, and I have argued that the most obvious counterarguments can be rebutted. It may be that there are other arguments against the view; if so, I would be very interested to hear them. In the meantime, I think that a broadly Fregean approach to meaning holds considerable promise.

## **Acknowledgments**

I am grateful to audiences at ANU, Florida, Miami, Michigan, and at the “Language, Mind, and World” conference in Tlaxcala for discussion. Thanks especially to Chris Evans, Mike Harnish, Marga Reimer, and Jason Stanley for comments.

## **Bibliography**

- Braun, D. 1996. Demonstratives and their linguistic meanings. *Nous* 30:145-73.
- Burge, T. 1979. Sinning against Frege. *Philosophical Review* 58:398-432.
- Carnap, R. 1947. *Meaning and Necessity*. University of Chicago Press.
- Chalmers, D.J. 2002a. Does conceivability entail possibility? In (T. Gendler & J. Hawthorne, eds) *Conceivability and Possibility*. Oxford University Press.  
[consc.net/papers/conceivability.html]
- Chalmers, D.J. 2002b. The components of content. In (D. Chalmers, ed) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press.  
[consc.net/papers/content.html]
- Chalmers, D.J. forthcoming a. The foundations of two-dimensional semantics.  
[consc.net/papers/foundations.html]

- Chalmers, D.J. forthcoming b. The nature of epistemic space.  
[consc.net/papers/espace.html]
- Donnellan, K. 1966. Reference and definite descriptions. *Philosophical Review* 75:281-304.
- Dummett, M. 1973. *Frege: Philosophy of Language*. Duckworth.
- Evans, G. 1973. The causal theory of names. *Proceedings of the Aristotelian Society (Supplementary Volume)* 47:187-208.
- Frege, G. 1892. Über Sinn und Bedeutung. Translated in (P. Geach & M. Black, eds.) *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell, 1952.
- Fumerton, R. 1989. Russelling causal theories of reference. In (C. Savage and C. Anderson, eds.) *Rereading Russell*. University of Minnesota Press.
- Jackson, F. 1998. Reference and description revisited. *Philosophical Perspectives* 12:201-18.
- Kaplan, D. 1989. Demonstratives. In (J. Almog, J. Perry, and H. Wettstein, eds.) *Themes from Kaplan*. Oxford: Oxford University Press.
- Kripke, S.A. 1977. Speaker's reference and semantic reference. In (P. French, ed.) *Contemporary Perspectives in the Philosophy of Language*. University of Minnesota Press.
- Kripke, S.A. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lewis, D. 1979. Attitudes *de dicto* and *de se*. *Philosophical Review* 88: 513-43.
- Lewis, D. 1986. *Australasian Journal of Philosophy* 62:221-36.
- Mill, J.S. 1843. *A System of Logic*. London: Longmans.
- Perry, J. 1977. Frege on demonstratives. *Philosophical Review* 86:476-97.
- Perry, J. 1979. The problem of the essential indexical. *Nous* 13:3-21.
- Putnam, H. 1975. The meaning of 'meaning'. In (K. Gunderson, ed.) *Language, Mind, and Knowledge*. Minneapolis: University of Minnesota Press.
- Recanati, F. 1993. *Direct Reference*. Blackwell.
- Searle, J. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.

- Soames, S. 1998. The modal argument: Wide scope and rigidified descriptions. *Nous* 32:1-22.
- Sosa, D. 2001. Rigidity in the scope of Russell's theory. *Nous* 35:1-38.
- Stalnaker, R. 1978. Assertion. In (P. Cole, ed.) *Syntax and Semantics: Pragmatics, Vol. 9*. New York: Academic Press
- Stanley, J. 1997. Names and rigid designation. In (C. Wright and R. Hale, eds) *Companion to the Philosophy of Language*. Blackwell.

## **Notes**