

There's Glory for You!

Jeffrey Ketland
Pembroke College, Oxford

Abstract:

This dialogue concerns metasemantics and language cognition. It defends a Lewisian conception of languages as abstract entities (Lewis 1975), arguing that semantic facts are necessities (Soames 1984), and therefore not naturalistically reducible. It identifies spoken languages as idiolects, in line roughly with Chomskyan I-languages. It relocates traditional metasemantic indeterminacy arguments as indeterminacies of what language an agent speaks or cognizes. Finally, it aims to provide a theoretical analysis of the cognizing relation in terms of the agent's *assigning* certain meanings to strings.

Alice and Humpty Dumpty are talking ...

Alice: I don't know what you mean by 'glory'.

Humpty Dumpty: Of course you don't—till I tell you. I meant 'there's a nice knock-down argument for you!'

A: But 'glory' doesn't mean 'a nice knock-down argument'.

HD: When I use a word, it means just what I choose it to mean—neither more nor less.

A: But what *you* mean by 'glory' is different from what *I* mean by 'glory'. Which one is what 'glory' *really* means?

HD: There is no such thing as what 'glory' means *simpliciter*.

A: And what does 'simpliciter' mean?

HD: Hmm. 'simpliciter' means 'absolutely', 'without qualification'.

A: So, 'simpliciter' *does* in fact have a meaning! And surely so does 'glory'. And so on. Words *must* have a meaning. For, otherwise, the sounds we make are all meaningless!

HD: I didn't mean to say that 'glory' has *no* meaning simpliciter either! Rather what I think is that its meaning is both *mind-dependent* and *language-dependent*. Perhaps I should explain ...

A: I hope it won't take long.

HD: First of all, a *person* can mean something; can *use* a word to mean something. For example, we might ask: what did *you* mean, or what were *you* referring to, when you said such-and-such?

But, second, also *words* can mean something. In the second case, what a word means depends on the *language* in question. Then, roughly, what a person usually means in using a word is what the word means in the language they *speak*. What *I* mean by ‘glory’ is what ‘glory’ means in the language *I* speak; what *you* mean by ‘glory’ is what ‘glory’ means in the language *you* speak.

A: The puzzling thing is how these words, or strings, and meanings get *connected together*.

HD: Indeed! How are strings “glued” to their meanings or referents? How does ‘cat’ get glued to cats? And ‘star’ to stars? And ‘Aristotle’ to Aristotle?
I think there are two rather different *kinds* of “glue”. I think people keep mixing these up. The first involves the language alone, and I will call it the *meaning-relation for a language*: this connects *strings* to their *meanings*. On that, I am a conventionalist. But the second connection involves a speaker or agent, and I shall call it the *cognizing relation*: this relates *minds* to the *language* they *speak*.

A: For *my* language, it’s *my mind* that does the connecting. So, I am puzzled ...

HD: The distinction is needed because there are languages that *others* speak which may have the same strings, but different meanings; and there are languages that *no one* speaks. One may study languages independently of their speakers; or their minds, if you prefer.

A: I see: languages exist independently of speakers and minds. The speaker has left the building!

HD: I want to argue that a language is an *abstract* entity: a system of syntax and, in some cases, meaning functions. A language is not a concrete thing, with a mass, or a location, or a density, or a speed. A language may have intrinsic properties quite independently of the properties of a speaker who speaks the language.

A: But there *are* speakers who, as you say, “cognize” languages?

HD: Of course!

A: I have never heard this word ‘cognizing’ before; but somehow I get what you mean ...

HD: You had never heard the word ‘magnesium’ until you studied chemistry, but you picked it up. Just as then, your idiolect has just undergone a small shift to a new one.

A: The speakers *have* to re-enter the picture somewhere ...

HD: The speaker, or their mind, is related to the language. One speaks a language by cognizing its syntax and its meaning functions. I speak a certain idiolect, *Humptese*, by cognizing its syntax and meaning functions; whereas you speak your idiolect, *Alicese*, by cognizing its syntax and meaning functions.

A: On your view, the meaning relation is intrinsic to the language itself?

HD: The languages we are discussing are *interpreted* languages. But we might consider uninterpreted, or partially, or multiply interpreted, languages too, if we like.

A: That reminds me of Professor Tarski's work. Tarski defines the set of truths for certain kinds of "formalized" languages. Tarski says that the object languages are required to be *interpreted*; and that one must have a *translation* from the object language to the metalanguage in order to formulate T-sentences, such as,

The string 'Schnee ist weiß' is true in *German* if and only if snow is white.

Tarski's object languages are simple in many respects, in order to have a precise syntactic description. But they are interpreted.

HD: And from this, one is led in a certain direction, which many authors see as counter-intuitive. In fact, some authors see it as a *reductio ad absurdum* of his work.

A: Why is defining the truths for a well-defined language problematic?

HD: Well, notice that the T-sentence for 'Schnee ist weiß' does not mention the *language users*, their *usage patterns*, their *mental states*, *intentions*, *causal factors*, and so on and so on. It states a property of an abstract system of syntax and meaning functions for an interpreted language, namely *German*.

A: Aha. I see! These authors see Tarskian semantics as missing something crucial.

HD: Yes; these authors insist that the semantic properties of a language, and in particular, the reference or meaning relation, have to be explained in terms of natural or intentional facts involving the speaker. Perhaps the *use* of the string; the *mental states* involved in speech-acts producing the string; its *role in reasoning*; its role in larger *theories*; its *causal connections* to its referent; and so on. Professor Putnam argued, in his 1985 article, 'A Comparison of Something with Something Else', that, because Tarski's approach omits facts about intentions and usage of strings, this work is not even *semantics*, and should be called a "non-semantic conception of truth". Professor Field argued, in his 1972 article, 'Tarski's Theory of Truth', that Tarski merely reduced truth to primitive denotation, which should then be reduced to a *physical* relation.

A: So, this would be a reduction along the lines of,

the string σ means M iff $R_1(\sigma, M)$
the string σ refers to x iff $R_2(\sigma, x)$

where ' $R_1(\sigma, M)$ ' and ' $R_2(\sigma, x)$ ' stand for some natural relation.

HD: Yes, but it is difficult to carry this programme through; consequently, there appears to be a deep *philosophical problem of making sense of semantic relations*. Two kinds of approach have been offered in reply. First, *naturalistic reductionism*, reducing the "real" semantic (reference or meaning) relation to relations involving

causation, or language usage, or intentional states; and, second, *deflationism*, treating ‘refers’, ‘means’ and ‘true’ in a disquotational manner.

For an example of the former, several recent authors have argued for a kind of semantic reductionism now called “reference magnetism”. This requires a certain metaphysical doctrine: only certain classes and relations of things are *natural*. These are perhaps the extensions of physical properties and relations that figure in the physical laws; or perhaps extensions whose elements are naturally similar in some way. Second, it is proposed that these special classes and relations are thought of as being the only ones that are “eligible”—that is, permitted—to be *referents* of predicates in languages. The reduction proposes that a necessary condition on an extension being a referent is its naturalness. If an extension is not natural, then it can’t be a referent of a predicate. This approach is meant to rule out certain strange, crazy or weird extensions that philosophers have come up with.

On the second, *deflationary*, approach, the strings ‘true’ and ‘refers’ are governed by primitive disquotational axioms, taken as being analytic, constituting their meaning. For example, the axioms one learns might be,

‘snow is white’ is true if and only if snow is white.
‘snow’ refers to snow.

On this deflationary view, this is all one needs to say about reference and truth. Translational indeterminacies between languages of the sort discussed by Professor Quine and others are not only unproblematic, but represent important insights into any properly scientific approach to the topic of meaning and reference.

But, at least for semantics, I disagree with both naturalistic reductionism and deflationism.

A: I am familiar with deflationism. But why do you disagree with the first approach?

HD: In a nutshell, I think it conflates *semantics* with *cognition*; semantic facts are metaphysically quite different from cognitive ones, and in particular, in their modal status. A semantic relation might be a meaning-relation or a referential string-thing relation, for an interpreted language. A cognitive relation might be the relation between an agent and the language they cognize; or the “assigning” relation between an agent, strings and meanings.

A: So you take the distinction between *semantics* and *cognition* very seriously?

HD: I think the usual philosophical problems that have been identified do not really lie in *semantic* theory. Instead, I wish to argue that the *semantic reference* relation—for example, the relation that the string ‘Schnee’ bears to snow in *German*—cannot be a natural relation, in the sense of being a relation that *contingently* holds between items.

A: Cannot?

HD: A reference relation for a language is simply any string-thing relation one likes.

A: I am not sure why you think this ...

HD: First, *semantic* relations are not *binary* relations. They are, at least, *ternary*, and involve *language relativity*. For example, consider the phrase

the string σ means M,

where M is a semantic value, such as a referent, an extension, a proposition, a concept, and so on.

As Tarski noted, this is missing a crucial parameter—a language parameter. In trying to understand the semantic properties of linguistic strings, one should include the language parameter, as follows:

the string σ means M in L .

Of course, what a string means in a language, might depend on *context* too, thus,

In context C , the string σ means M in L .

Thinking of L as determining a function, μ_L , one could write,

$$\mu_L(\sigma, C) = M.$$

The crucial point is that semantics describes *languages*, not simply strings on their own; semantics studies systems of expressions and functions like μ_L . A string doesn't "naturally" connect referentially to some meaning or referent. For semantic notions, language relativity is unavoidable. And I think the same is true for phonology, syntax and pragmatics. For example, how 'ouch' is to be pronounced, what its lexical role is, and what its governing pragmatic rule is: these are language-dependent.

A: Now I see: yes, this point is made by Tarski. The same string of symbols or phonemes might mean different things in different languages. In principle, *anything could mean anything*.

HD: We may call this *Semantic Conventionalism*. In his 1916 posthumous work *General Course in Linguistics*, Ferdinand de Saussure wrote that the sound signal is "arbitrary in relation to its signification, with which it has no natural connexion in reality". And this is all there is to the *conventionality* of language.

A: So, a language L with its meaning function μ_L , can be conventionally chosen. Much as a co-ordinate system for the sphere is, to a large degree, conventional. On this conventionalist view, no languages are "forbidden"?

HD: That sounds *very* 1968, but yes. Indeed, in 'On Sense and Reference', Gottlob Frege wrote, 'Nobody can be forbidden to use any arbitrarily producible event or object as a sign for something'. Some languages might be *socially* forbidden ... and some languages might be weird, at least from the perspective of other speakers. But our limited experience causes us to underestimate how many languages there are. There is a language in which 'gloobydooby' means Birmingham City FC, even

though no one cognizes that language. The fact that there is no person who speaks a language in which ‘gloobydooby’ means Birmingham City FC is no reason to deny the *existence* of such a language.

A: In logic, we examine infinitary languages, like $L_{\omega_1, \omega}$, with infinitely long strings. And there are languages with uncountably many constants—for example, a constant c_r for every real r . The possibilities are endless.

HD: But no one *cognizes* such languages. Presumably, for biological and psychological reasons. It’s hard to see how a human mind might *cognize* a language with infinitely long strings, or with uncountably many constants, without assuming an infinitary inflation of human cognitive capacities.

A: So, such languages are like Mr Wilde’s fox hunters: unspeakable. I rather like this idea. For languages, *existir es no ser hablado!*

HD: I do not understand ...

A: ‘To exist is not to be spoken’.

HD: Oh, like *esse est percipi*? I see. They do not *require* embodiment. We can study them, define them and prove facts about them, independently of whether anyone does or doesn’t speak them.

A: An example involves Gödel’s famous incompleteness result. This can be summarized by saying that the set of truths of the interpreted language L of arithmetic is not computably enumerable. No computer program can enumerate them all, without including some falsehoods or even inconsistencies. This property of L does not depend on speakers of L .

HD: Oh to be a logician, like you, Alice! I am just an *egg*.

A: But remarkably logical, for an *egg*.

In logic, we think of an interpreted language as a pair $L = (S_L, I_L)$, where S_L is the syntax and I_L is the interpretation function, which specifies a domain and referents or denotations for the non-logical terms and meanings for the logical terms. For example, the interpreted language L of arithmetic is (S_L, N) , where S_L is a first-order language with identity, with a constant ‘0’ and function symbols ‘ S ’, ‘+’ and ‘ \times ’; and N is the interpretation of S_L whose domain is the set $\{0, 1, 2, \dots\}$ of natural numbers, and which interprets ‘ \neg ’ as ‘not’, ‘ \wedge ’ as ‘and’, ‘ \forall ’ as ‘for all’, ‘=’ as identity, and where ‘0’ denotes 0, ‘ S ’ denotes the successor function, ‘+’ denotes addition and ‘ \times ’ denotes multiplication. An L -sentence ϕ is *true* in L iff N *satisfies* ϕ , using Tarski’s definition. Gödel’s results imply that the truth set for L is not computably enumerable.

This conception of languages is in accord with yours too: “a system of syntax and meaning functions”. So I am curious: is your view of language as an abstract entity a widely held one?

HD: Arguably, three of the founders of semantic theory, Professors Carnap, Tarski and Montague held a view like this. But the view is stated quite clearly by Professor Lewis, in his 1970 article, ‘General Semantics’:

My proposals will also not conform to the expectations of those who, in analyzing meaning, turn immediately to the psychology and sociology of language users: to intentions, sense-experience, and mental ideas, or to social rules, conventions, and regularities. I distinguish two topics: first, the description of possible languages or grammars as abstract semantic systems whereby symbols are associated with aspects of the world; and second, the description of the psychological and sociological facts whereby a particular one of these abstract semantic systems is the one used by a person or population. Only confusion comes of mixing these two topics.

And he restates this view in his 1975 article, ‘Languages and Language’:

What is a language? Something which assigns meanings to certain strings of types of sounds or of marks. It could therefore be a function, a set of ordered pairs of strings and meanings.

And a similar view was defended by Professor Soames in a 1984 article, ‘What is a Theory of Truth?’:

Although this might initially seem to be a desperate strategy, it is not. Think of a standard first-order language L as a triple (S_L, D_L, F_L) , where S_L is a family of sets representing the various categories of well-formed expressions of L ; D_L is a domain of objects; and F_L is a function that assigns objects in D_L to the names of L , subsets of the domain to one-place predicates of L , and so on. Let J be a class of such languages. Truth can now be defined in nonsemantic terms for variable ‘ L ’ in J in a straightforward Tarskian fashion.

The only significant change from before is that the notions of primitive denotation are no longer given language-specific list definitions, but rather are defined for variable ‘ L ’ using the “interpretation” functions built into the languages. In particular,

a name n refers to an object o in a language L iff $F_L(n) = o$.

And Professor Soames adds,

This way of looking at things was suggested to me from two sources: David Lewis’s ‘Languages and Language’, and one of Saul Kripke’s seminars on truth, Princeton, 1982.

A: But it is still very counter-intuitive to understand language *not* as a sociological or psychological phenomenon, involving language speakers, their usage patterns, and so on. Suppose we call this the *social-psychological view* ...

HD: Well, Professor Lewis *also* describes this more intuitive view too:

What is language? A social phenomenon which is part of the natural history of human beings; a sphere of human action, wherein people utter strings of vocal

sounds, or inscribe strings of marks, and wherein people respond by thought and action to the sounds or marks which they have observed to have been so produced.

Lewis tries to do justice to both the abstract view and the social-psychological view by examining what it is for a population to *use* a language. Perhaps I can return to this later?

A: Professor Field's main objection to Tarski, if I remember right, is that the *definiens* for a semantic concept is simply a "list". For example,

x refers in *German* to y iff ($x = \text{'Schnee'}$ and $y = \text{snow}$) or ($x = \text{'Wasser'}$ and $y = \text{water}$) or ...

HD: My response is simply that this is precisely what one expects on the abstract view: definitions of semantic relations are *listlike* because semantic notions are *language relative* and *languages* are like *functions*. Notice that the name of the language disappears in the *definiens*. The resolution of the conundrum is that reference is not some peculiar physical or natural relation; rather, an interpreted language is—well, more exactly, determines—a *function*. So, for example, *German* is a language L such that $\rho_L(\text{'Schnee'}) = \text{snow}$, and so on and so forth. Consequently, if one looks for something else that explains this referential "glue", one will not find it. Referential glue is quite arbitrary and promiscuous.

A: Perhaps the following is an analogy. Consider a function F from alphabet $\{\text{'a'}, \text{'b'}, \text{'c'}\}$ to, say, the set $\{0, 1, 2\}$. Such a function is a set of ordered pairs. For example, suppose F is $\{(\text{'a'}, 1), (\text{'b'}, 2), (\text{'c'}, 0)\}$. Then we may give the following definition of F :

$y = F(x)$ iff ($x = \text{'a'}$ and $y = 1$) or ($x = \text{'b'}$ and $y = 2$) or ($x = \text{'c'}$ and $y = 0$).

Since F might be any function, there is nothing wrong with this. There is no special intrinsic connection between the arguments and values, no special glue, other than that they appear in the given ordered pairs. There is no such thing as the "natural" function from the alphabet $\{\text{'a'}, \text{'b'}, \text{'c'}\}$ to the set $\{0, 1, 2\}$ of referents.

Still, if I look at the Tarskian definition of truth or reference for a language, I see no *semantic* terms in the *definiens*. So, Tarski *does* seem to give a reduction of semantic terms ...

HD: Tarski's definitions of semantic notions replace semantic notions with *logical* and *mathematical* ones along with notions from the *object language*. The definitions of semantic notions do *not*, however, explain ' x refers in L to y ' in terms of some more basic natural, physicalistic or intentional binary relation that is language-independent or language-transcendent. Semantic relations just aren't what these authors think they are.

I think that, at bottom, these authors are making precisely the mistake that Professor Lewis alluded to, and 'only confusion comes of mixing these two topics': *semantic* relations, which connect linguistic strings to their meanings, are being mixed up with

cognitive relations, involving agents' minds. The *cognitive* glue is quite different from the *semantic* glue.

A: Is the following an objection to your view? I surely *could* have spoken a different language, *L* say, whose meaning function maps the string 'cat' to dogs. And if I had spoken that language *L*, then I would *not* have been speaking the language that I do in fact speak. For 'cat' means cats in *Alicese*, but means dogs in *L*.

HD: Yes ...

A: So, if 'cat' means dogs in *L*, then *L* cannot be *Alicese*. So, the *semantic* fact that 'cat' means cats in *Alicese* *couldn't have been otherwise!*

HD: Quite so! Indeed, as Professor Soames notes:

On this conception, languages are abstract objects, which can be thought of as bearing their semantic properties essentially. There is no possibility that expressions of a language might have denoted something other than what they do denote; or that the sentences of a language might have had different truth conditions. Any variation in semantic properties (across worlds) is a variation in languages. Thus, semantic properties aren't contingent on anything, let alone speaker behavior.

A: So, *Alicese* has an essential semantic property: the property of being a language in which 'cat' means cats.

HD: This point has appeared many times before, as the *modal objection* to the semantic conception of truth. For example, because semantic axioms,

'Schnee ist weiß' is true in *German* if and only if snow is white

'Schnee' refers in *German* to snow

behave as *metaphysical necessities*, Professors Field and Putnam took this as an objection.

A: I see ... but, *prima facie*, that seems reasonable.

HD: Similarly, Professor David, in his 1994 book *Correspondence and Disquotation*, argues that, on a correspondence theory of truth, semantic facts—such as what state of affairs a sentence corresponds to in language *L*—should be *contingent*. But it seems to me perfectly compatible with a correspondence theory of truth that a semantic fact of the kind,

the string 'Schnorbl este fnoffle' represents snow's-being-white in language *L*

should be *necessary*. It may be contingent whether the state of affairs of snow's-being-white *obtains* or not; it may be contingent that anyone does or doesn't *speak L*. But it is *not* contingent that snow's-being-white *is* the state of affairs that *L* assigns to the string 'Schnorbl este fnoffle'. In fact, one simply says that the value of the string 'Schnorbl este fnoffle' under *L* is snow's-being-white. With a distinct value, we just get a distinct language.

A: But, in the case of semantics, does not the *extension* of ‘cat’ in *English* change from world to world?

HD: Yes. But when we consider possible worlds semantics, such a language L specifies a reference or denotation function on the *worlds*, say ρ_L . Given a predicate P , and world w , each extension $\rho_L(P, w)$ is a set or relation at that world w , the *values* of this function ρ_L . But values of a function should not be confused with the *function* itself. In different worlds, the extensions of the same word or string may be different. However, the *reference function* ρ_L doesn’t itself change or vary in some further sense. The contingency lies in the *cognizing* relation; the necessity resides in the *reference* relation.

A: So, your position is this: an interpreted language L determines a semantic reference or meaning *function*, μ_L . Defining μ_L is separate from defining any *other* function, say μ_{L^*} for some other L^* . Every *language* is different from every other, much as every *function* is different from every other. If we compare a definition of the reference relation $\rho_{English}$ for *English* with a definition of the reference relation ρ_{French} for *French*, there is nothing that $\rho_{English}$ and ρ_{French} have “in common”. They’re functions from strings to things. To suppose otherwise would violate Semantic Conventionalism.

HD: Yes.

A: It is initially counter-intuitive to recognize that semantic facts are necessities. But on reflection, I agree with you.

HD: And this yields the crucial conclusion: semantic facts cannot be reduced to use facts. For necessities cannot be reduced to contingencies.

A: In a sense, the *semantic* “gluing” of strings to their meanings by languages is *mathematical*; and consequently not reducible to a physical relation at all. Whereas *cognitive* “gluing”, of a mind to the strings that the mind uses and therefore to some language, is not arbitrary in the same way.

HD: Yes, precisely.

A: Can nominalism make sense of your account of *languages* and *language cognition*? For languages are mathematical entities, and language cognition is now understood as a form *mathematical* cognition!

HD: The languages that an agent cognizes are *mixed* or *impure* mathematical entities, to use the jargon from the philosophy of mathematics. An example would be a set of concrete things. The set is an *impure* set. Similarly, when an agent cognizes a language L , this capacity must, I think, involve a cognitive capacity to recognize and discriminate symbol *tokens*, such as letter or phoneme *tokens*: these are concrete things. Even so, the syntactical entities of any interesting language L are, normally, *expression types*: *arbitrarily long finite sequences* of symbols. Though *all* the individual symbols will have tokens, not all the sequences of these symbols will: there are some, of course, but infinitely many sequences are too long to be tokened. Motivated by nominalistic concerns, one might try and make do with just the tokens

instead. Professors Quine and Goodman tried, in a 1947 article ‘Steps Toward a Constructive Nominalism’, to formulate syntax without assuming types. But Quine, recognizing insurmountable difficulties, abandoned the proposed view rather quickly. For example, how many strings are there in the language of Professor Field’s reformulation of Newtonian gravitational physics? There are \aleph_0 strings in this language! So one ends up denying that there are *numbers*, while asserting that there are *strings*, indeed \aleph_0 -many; but these strings are *finite sequences*, which are themselves abstract.

A: I agree that this is odd, to say the least. Even so, isn’t linguistics also an empirical science, concerned with the psychology of speakers? But on the view you defend, the semantic facts for a language are necessities; and these cannot be explained by facts about how strings are used, for such facts are contingent; while semantic facts are necessary. So, *semantic facts cannot be explained by usage!*

HD: I think the same goes for phonology, syntax and pragmatics too.

A: But this is where it becomes quite puzzling. For we want linguistics to be an empirical science. Do you think the semantic axiom that ‘Schnee’ refers in *German* to snow is *empirically* testable?

HD: I think a confusion arises here because *German* is a natural language, and is thought of in terms of a *language community* of speakers. So, to check facts about *German* I might, for example, travel to Munich as an amateur field linguist, and perform some bio-culinary experiments on Bavarians, observing their behaviour around cats and doughnuts, and so on; upon hearing utterances of ‘Das ist eine Katze!’, ‘Zwei Krapfen, bitte!’, I accumulate evidence that, in *German*, ‘Katze’ means cat and ‘Krapfen’ means doughnut.

A: That seems correct! Why is there any confusion?

HD: Let us call the usage facts U-facts. Let us call the semantic facts S-facts. But let us also introduce a category of C-facts, of the form,

agent *A* cognizes language *L*.

We might add certain parameters too, such as ‘cognizes at time *t*, at world *w*’. But let us simplify. To further simplify our analysis, suppose we accumulate U-facts for a *single agent*, *A*.

I believe that these U-facts may form *evidence* in favour of a more complicated claim, involving a C-fact *and* an S-fact, joined by a quantifier. Namely the contingent claim that,

There is some language *L* such that *A* cognizes *L* and ‘Katze’ means-in-*L* cat and ‘Krapfen’ means-in-*L* doughnut.

Such a hypothesis helps to *explain* *A*’s word usage. But I still maintain that the semantic subclause expresses some *essential* property of whatever *L* is asserted to exist.

If I accumulate U-facts for another agent, *B*, they might form evidence for a contingent claim that,

There is some language *L* such that *B* cognizes *L* and ‘Katze’ means-in-*L* cat and ‘Krapfen’ means-in-*L* doughnut.

From these two, we can reasonably begin to conjecture that *A* and *B* cognize *similar idiolects*.

A: So, the *logical role* of U-facts—word usage, associated intentions, inferential practices and so on—in linguistic theory is to *constrain* the C-facts, what *language* an agent cognizes?

HD: That is the idea. If an agent *uses* the word ‘Katze’ in a certain way, then they might *cognize* a language *L* such that ‘Katze’ refers in *L* to dogs.

A: So, it seems that, properly speaking, the various agents in a speech community do not, usually, cognize *exactly the same language*.

HD: Precisely.

A: But meaning in the sense of what an *individual speaker* means is related to *their use* of strings and their mental representational states?

HD: Yes. Though it needn’t be *public* or *conscious*. On the contrary, the mental state of cognizing a language is largely *private* and *unconscious*. I think this is the norm. I have even heard people disparage this view as *Humpty Dumptyism*, but I am not upset by that!

A: I suppose that Professor Wittgenstein is the nemesis of any Humpty Dumptyist?

HD: Opposition to private languages is found in Professor Dewey’s writings before it was adopted by Wittgenstein many years later. For both, the arguments were behaviouristic or naturalistic, as Quine was to emphasize, for example, in his 1968 John Dewey Lectures, “Ontological Relativity”. There, Quine himself repudiates as “uncritical semantics”, the ‘myth of a museum in which the exhibits are meanings and the words are labels’.

A: I find it hard to adjudicate ... where would modern linguistics stand on this?

HD: Well, this is a thorny topic. What I think—and I think it is the correct view within scientific descriptive linguistics—is that agents speak an *idiolect*, which may be changing all the time. Maybe it can change in a nanosecond. Maybe you shift into different idiolects as the context changes. Why not? Perhaps one speaks more than one, perhaps many, idiolects. Also, I cannot see why one must insist that idiolects be “public” or “manifestable”. I have no objection to *innate* ideas, to *private* grasping of concepts, to *intentional realism*, to there being very *peculiar psychological states* shared by no one, and so on. Linguists are bound to take a variety of views on such matters, though.

A: But there *do* seem to be public languages. All the natural languages which have names are usually thought of as *public* languages. They involve speech communities. For example, what is *English*?

HD: The abstract view naturally fits with the notion that *idiolects* are basic, *because* there is fluctuation and variation across communities and over time. Since these idiolects are *different*, they must be individually cognized by speakers. Consequently, public languages, allegedly common to a speech community, are not theoretically basic for linguistic theory, except perhaps for branches like sociolinguistics. For example, in his 1975 article, “Knowledge of Language”, Professor Chomsky writes,

Customarily, one speaks of a child as having partial knowledge of a language, taking as the external standard the grammar of the adult language, of the speech community to which he belongs. In another sense, the child has, by definition, a perfect knowledge of his own language. There is no contradiction here; rather, there are two notions that must be distinguished. The study of language should be conceived, in the first place, with the speaker’s perfect knowledge of his own language. The notion of ‘language’ as a common property of a speech community is a construct, perfectly legitimate, but a higher-order construct. In the real world, there are no homogeneous speech communities, and no doubt every speaker controls several grammars, in the strict sense in which a grammar is a formal system meeting certain fixed conditions.

A: So linguistics should be methodologically individualistic?

HD: Yes. The language cognized is an *idiolect*. Although one may cognize many at a time. Bilingual speakers usually cognize distinct idiolects—for example, varieties of *English* and *Punjabi*—that are not easily mutually interpretable. And individuals can modify the idiolect they cognize in various ways: for example, phonetically. If they cognize two languages, they can perform interesting operations on them, such as mixing them and fusing them. For I can say, ‘Would you like a Krapfen?’ or ‘Ich bin very very glücklich!’. The point is that cognizing an idiolect, or several, is the *mental state* of an *individual* cognitive system. There is no role for a mental state of cognizing, or partially cognizing, a “common shared language”.

More recently, Chomsky talks of “I-languages”, for *internal* languages, as being the proper subject of scientific linguistics. As I understand the terminology, the I-languages are *idiolects*. And he talks also of “E-languages”, for *external* languages, each of which corresponds roughly to a single public, shared, language spoken by a speech community. There is an important difference, however, in metaphysics between myself and Chomsky, who claims that I-languages are *mental objects*, whereas I claim that they are *abstract objects*, carrying their intrinsic syntactic, semantic, etc., properties essentially, for the reasons we discussed earlier.

A: So, *English* is an E-language? And *Alicese* is an I-language?

HD: You cognize *Alicese*. I cognize *Humptyese*. This conversation is witness to the great overlap of these idiolects, which are, to a large degree, mutually intelligible. The problem is that we must recognize that there is a large, heterogeneous speech community; there is no single language that *all* speakers cognize. The idiolects spoken overlap a good deal, but they inevitably vary somewhat. Not everyone in a speech community has the same lexicon, the same meanings for words, the same pronunciations, the same pragmatic rules, and so on. If they vary, even a tiny bit, they are different.

A: The tiniest change produces a new idiolect or language?

HD: To use Professor Quine's terminology, languages are *individuated* very finely: by exact sameness of phonology, syntax, semantics and pragmatics.

A: So, ... *English* isn't really a language at all in your sense?

HD: Although we were using *English* in our examples earlier, that is, in fact, somewhat misleading, yes. I am not sure what to say *English* is. Often, we can easily switch to talking of the idiolect we call '*A-English*', where *A* is the speaker in question. Or, we may talk of the *variety* of *English* spoken by *A*.

A: Perhaps we should take these "varieties" seriously, and think of each "natural" language not as a particular language but as a *set* of languages? So, when we say that I speak *English*, really what we mean is that I speak a *variety* of *English*—Alice cognizes some idiolect that is an *element* of *English*?

HD: Alternatively, we might try and define a unique language, *English*, which functions as an external normative standard for the language community, but which no individual speaker cognizes, of course. This would require smoothing out the discrepancies between idiolects, and invoking a powerful version of collective normativity: for example, counting those who assign the concept UNINTERESTED to the string 'disinterested' as being *mistaken* as to what they *ought* to mean, which is the concept NOT HAVING A PARTICULAR INTEREST.

A: On this view, there is a common shared E-language, which is an external standard, and from which idiolectic deviations are *mistakes*. Whereas on the I-language view, the idiolectic variations are not *mistakes*.

HD: *This* is very controversial: *prescriptivism*. I am more sympathetic to what linguists call 'descriptivism', but prescriptivism seems to be very popular with philosophers.

A: Maybe the following is a case to illustrate. I used to pronounce "schedule" as "shedule", but increasingly I say "skedule".

HD: Well, if you make a *performance* error relative to the idiolect you speak, then this is a violation of the intrinsic rules of your idiolect, the language you cognize. You will normally recognize this yourself. There is a kind of *normativity* that arises from the properties of your own idiolect; let's call it *idiolectic* normativity.

A: So this would include things like stammerings, ums and errs, mispronunciations ...

HD: A performance error occurs when your speech performance violates, or is out of kilter with, the properties of your own idiolect. But, in your ‘schedule’ case, it does not seem to be a performance error. It seems that your idiolect has a string pronounced as ‘shedule’ at time t_1 and a string pronounced as ‘skedule’ at later time t_2 . So, your idiolect is shifting. Or perhaps you are, in an important sense, cognizing several idiolects, and one is “fading out”, as it were.

A: Yes, idiolects can shift, for various reasons ...

HD: Of course! In fact, *very quickly*. I can “baptise” something, introducing a new name in my *idiolect*, tagging the object. As in the case of the examples discussed by Professors Marcus and Kripke. My idiolect has been modified: *extended*. Also, I can *borrow* meanings, somehow. I do not quite know the details of the cognitive mechanism; but clearly we can copy or borrow meanings from others.

With language acquisition, when children acquire a language, they do not just acquire any old language from syntactico-semantic heaven. They gradually acquire an idiolect which *resembles* the idiolects of those around them. Their idiolect—more exactly, the sequence of idiolects—slowly converges towards an approximate standard in their speech community.

A: To apply it to my case, I can see a kind of reason why it’s better for me say ‘skedule’ instead of ‘shedule’: namely, *social co-ordination*. The speakers around me tend to say ‘skedule’ instead of ‘shedule’: my idiolect shifts and it becomes easier for me to process incoming messages from others and for them to process messages from me. Mutual “decoding” becomes easier. It isn’t conscious, though. But I think that social co-ordination has normative value. And I may defer to expertise too in order to improve social co-ordination.

HD: It is rational for humans to improve co-ordination and to share expertise!

A: Any rational *creature*, talking eggs included.

HD: Of course.

A: So, let’s assume that this is so. Then while you are entitled to *your* meaning for ‘glory’, meaning what it does in *your* idiolect, and likewise I am entitled to *my* meaning for ‘glory’, meaning what it does in *my* idiolect, nonetheless there is a broader language community relative to which your meaning is non-standard, whereas my meaning is standard. By the social co-ordination principle, I’m therefore entitled to request that you improve social co-ordination by using a *shared* meaning. And there’s a nice knock-down argument for you!

At this, Humpty Dumpty sat and thought for a while.

HD: Do you recall Professor Quine’s famous puzzle? A human rushes past The Mad Hatter, who utters ‘gavagai’. The Mad Hatter always tends to behave like this. Let us

define *Madhatterese* to be the idiolect that The Mad Hatter cognizes. Then, given his speech behaviour, what does the *Madhatterese* word ‘gavagai’ mean?

A: Yes, Quine’s argument for indeterminacy of translation ...

HD: Quine suggests that it is not semantically determinate whether ‘gavagai’ should be translated as ‘Lo, a human!’, or ‘Lo, humanhood instantiated!’, or ‘Lo, an undetached human part!’, and so on.

A: I heard something like this from my set theory teacher once. The example she discussed involved Skolem’s Paradox.

Let me now see how to reformulate the ‘gavagai’ example on your conception of languages. First, we define two languages L and L^* , with reference functions ρ_L and ρ_{L^*} such that,

$$\begin{aligned}\rho_L(\text{‘gavagai’}) &= \text{humans,} \\ \rho_{L^*}(\text{‘gavagai’}) &= \text{undetached human parts.}\end{aligned}$$

Then there is no question as to what the *strings* mean! It is a matter of definition that ‘gavagai’ refers to humans in L and refers to undetached human parts in L^* . So, it really isn’t a puzzle about *semantics*.

HD: Precisely.

A: But, on the other hand, we still have not settled what The Mad Hatter *himself* means in using these strings, given only the U-fact that The Mad Hatter utters ‘gavagai’ when there are humans nearby. If there are humans nearby, then there are undetached human parts nearby, and *vice versa*. We do not know if The Mad Hatter cognizes L or cognizes L^* . Or perhaps neither. If we cannot pin this down, then *the U-facts do not uniquely determine the C-facts*.

HD: It seems to me that all of these puzzles allegedly about semantics, for which Wittgenstein, Quine, Putnam and Kripke are rightly famous, can be reformulated in the same manner. Suppose the language L assigns electrons to ‘electron’, while L^* assigns protons to ‘electron’. Do scientists *cognize* L or L^* ? Or, suppose that the symbol ‘plus’ means addition in L , and means quaddition in L^* . Do we *cognize* L or L^* ? These are not problems in semantics at all! They are problems in *cognitive science*.

A: As I said a moment ago, Professor Skolem has a similar argument, in connection with non-standard models. In your formulation, it is like this. Suppose L maps the string ‘number’ to the standard numbers 0, 1, 2, ..., while L^* maps ‘number’ to elements of a non-standard model of Peano arithmetic. Then do number theorists *cognize* L or L^* ?

HD: Problems previously thought of as *semantic* indeterminacies have now become indeterminacies of language cognition: a certain *mental cognizing state* of the speaker is undetermined or perhaps even indeterminate.

A: So, if I understand the overall argument correctly, there are three main steps.

First, through a thought-experiment concerning a certain hypothetical scenario and a certain body of U-facts, we obtain a dilemma: *either* there is no fact of the matter as to which language *L* is cognized, *or* some underlying U-facts have been omitted. Second, we insist that *all* the U-facts have been specified. This yields Quine's conclusion.

HD: Consequently, to dispute Quine's conclusion, we need to clarify Quine's conception of what a *full* specification of the U-facts amounts to. For Quine, this involves *physicalism about facts* and *empiricism about meaning assignment*.

A: Which rules out certain certain kinds of "mysterious" *assignment of meaning*? For example, a primitive grasping of universals, or forms, or concepts or propositions?

HD: Yes, a certain kind of *physicalistic empiricism* implies that it is *indeterminate* what *language* is *cognized*. Regarding the relation of indeterminacy and the empiricist theory of meaning, Quine comments in his 1969 article "Epistemology Naturalized",

For the uncritical mentalist, no such indeterminacy threatens. Every term and every sentence is a label attached to an idea, simple or complex, which is stored in the mind. When, on the other hand, we take a verification theory of meaning seriously, the indeterminacy would appear to be inescapable. The Vienna Circle espoused a verification theory of meaning, but did not take it seriously enough. If we recognize with Peirce that the meaning of a sentence turns purely on what would count as evidence for its truth, and if we recognize with Duhem that theoretical sentences have their evidence not as single sentences but only as larger blocks of theory, then the indeterminacy of translation is the natural conclusion.

Should the unwelcomeness of the conclusion persuade us to abandon the verification theory of meaning? Certainly not. The sort of meaning that is basic to translation, and to the learning of one's own language, is necessarily empirical meaning and nothing more. A child learns his first words and sentences by hearing and using them in the presence of appropriate stimuli. These must be external stimuli, for they act on both the child and on the speaker from whom he is learning. Language is socially inculcated and controlled; the inculcation and control turn strictly on the keying of sentences to shared stimulation. Internal factors may vary *ad libitum* without prejudice to communication as long as the keying of language to external stimuli is undisturbed. *Surely one has no choice but to be an empiricist so far as one's theory of linguistic meaning is concerned.*

Well, I suppose the question is: do we have no choice?

HD: In 'Languages and Language', Professor Lewis introduces an indirect proposal for analysing 'A cognizes *L*', but this is based on the notion of a *population* speaking a *common, or shared, language L*.

My proposal is that the convention whereby a population *P* uses a language *L* is a convention of truthfulness and trust in *L*. To be truthful in *L* is to act in a certain way: to try never to utter sentences of *L* which are not true in *L*. ... To be trusting in *L* is to form beliefs in a certain way: to impute truthfulness in *L*

to others, and thus to tend to respond to another's utterance of any sentence of L by coming to believe that the uttered sentence is true in L .

Professor Lewis then goes on to embed this account of a population "using" a common language within his more general account of *convention*. It is here that my view diverges. His account assumes that languages are E-languages. As I said earlier, I think this is the wrong place to *begin*, because cognizing a specific language, an idiolect, is an *individual mental state*.

A: So, in each case, we are looking for some *extra* U-facts about the agents themselves which tell in favour of their cognizing L and not L^* .

HD: Yes, this is the central problem identified by Quine. *What are the kinds of meanings that a mind can grasp and can assign to strings?* How are we to understand notions such as:

A assigns meaning M to string σ ,
A uses string σ to mean M,
A uses string σ to refer to x ?

For example, it seems to me that the debates within the "theory of reference" for proper names are crucially about the assignment notion '*A uses the name n to refer to x* '. We might call these *semantico-cognitive assignment relations*.

If we simply assume some such notions, we can *define* '*A cognizes the semantics of L* ' by saying simply that *A assigns meanings to strings just as the language L does*. The agent and the language are in harmony. So, suppose we write ' $\mu_L(\sigma, C)$ ' for the meaning of σ in context C in L . Then the definition, more formally, is:

(COG1) *A cognizes the semantics of L iff, for any lexical L -string σ , context C , A assigns $\mu_L(\sigma, C)$ to σ in C .*

In one of the "Objections" in "Languages and Language", Lewis cites a definition like (COG1) and replies,

Yet there is no such thing as an action of bestowing a meaning (except for an irrelevant sort of action that is performed not by language users but by creators of language) so we cannot suppose that language-using populations have conventions to perform such actions. Neither does bestowal of meaning consist in forming some belief.

But my reply to this is that an agent's cognizing a language is not really a *convention* to perform such an action, because I do not think the speakers all follow the *same* conventions. The conventionality of language is encapsulated by Semantic Conventionalism, that we discussed before: the *meaning relation* is conventional, much as a co-ordinate system is.

Furthermore, the language user *does* assign meanings to strings, *does* bestow meanings, *does* assign referents, *does* baptise new referents, *does* borrow meanings,

and so on. It seems to me that these *are* actions performed by language users, frequently unconsciously. So, here my view diverges from Lewis's.

A: Is the real problem one of cognitive "access" or "contact"? We must look for some *meaning entity* that minds are in *contact* with. For example, it is difficult to see how the mind might "assign" the appropriate referents to strings like 'the electromagnetic field', which requires a physical field as referent, or 'the Church-Kleene ordinal', which requires an abstract referent. And, on the other hand, we do seem to be in *immediate contact* with sense data or stimuli. So, it's easy to agree with Quine that, given our broadly scientific picture of how the mind relates to the world, "one has no choice but to be an empiricist" for linguistic meaning.

HD: The age-old epistemological conflict between Empiricism and Rationalism resurfaces as a conflict concerning language cognition, concerning what kinds of entities can be assigned as meanings of strings and how this might work. For Quine there are determinate meanings of a sort: *stimulus meanings*. But such stimulus meanings are nothing like *propositional content*. Let M be the stimulus meaning of 'gavagai' for The Mad Hatter. This is a set of retinal and auditory stimuli types. By language inculcation, i.e., conditioning, The Mad Hatter's cognitive system has assigned M to the string 'gavagai'. Behaviourists called this "association". And, the consequence is that the Mad Hatter has a disposition to assert the string 'gavagai' whenever the environment causes a *token* of any stimulus type in M. Such a token might be caused by a cleverly disguised dummy, an experimental setup in a neuro-psychology laboratory, or by a Cartesian demon.

A: Perhaps Quine's own approach, with its behaviourist overtones, is a bit too limited? Must all empiricists work within such limited framework? For example, can the empiricist not appeal to *inference*? Associating sentences, or mental representations, with *each other*, by finitary rules? Can the empiricist also not appeal to *theories* as implicitly determining meanings? You mentioned earlier that perhaps one might appeal to *natural* properties and kinds?

HD: Empiricists vary quite a bit on how much metaphysics they can stomach. But an extended empiricist theory of meaning assignment might involve *sensory meanings*, *inferential connections*, *conceptual role in accepted theory*, *causal connections*, and *natural kinds/natural similarity relations*. And although many authors have contributed in various ways, probably the main architect of this is Professor Lewis himself. We might call it the Lewisian Standard Model for Metasemantics.

A: And does it succeed?

HD: I am unsure. But, as Professor Lewis mentions, in one case where it doesn't fix a *unique* language cognized, one might conclude that, here at least, this is what one expects. This is for *vagueness*; for terms like 'bald' and 'tall', it is indeterminate what *precise* language one *cognizes*. Perhaps the physicist can *define* these precise languages. Vagueness turns out to be neither semantic, nor epistemic, nor ontic; it is cognitive.

A: I am still bothered by the examples we discussed above. Consider our earlier puzzle about whether Professor Gowers cognizes a language *L* in which 'number'

denotes the standard numbers or a language L^* in which ‘number’ denotes the non-standard numbers. Although, for any sentence in both languages, the truth values are the same, it seems reasonable to say that the *propositions* expressed usually differ. For example,

‘every natural number has a unique prime factorization’ expresses in L the proposition that *every natural number* has a unique prime factorization,

while,

‘every natural number has a unique prime factorization’ expresses in L^* the proposition that *every natural number and non-standard number* has a unique prime factorization.

And presumably, when asserting the string ‘every natural number has a unique prime factorization’, Professor Gowers thereby expresses a belief in the former proposition, and not the latter. So, it is *determinate* that Professor Gowers cognizes L and not L^* . This C-fact is determined by a certain kind of U-fact involving appeal to intentional states, with *propositional content*. That is, the Quinian claim to have specified *all* the U-facts is not correct.

HD: This suggests that we take the cognitive assignment of a *proposition* to a string as basic.

A: Suppose that, for the languages in question, we have a function π_L which maps any sentential string σ of L , any context C , to a *proposition* $\pi_L(\sigma, C)$. For example,

π_L (‘Samuel möchte einen Brezel’, C) is the proposition that Samuel would like a pretzel at t_C ,

where t_C is the time of the context C .

Then, taking ‘agent A assigns proposition p to string σ in context C ’ as primitive, we try the following necessary condition for ‘ A cognizes L ’:

(COG2) A cognizes L only if, for all L -sentences σ , contexts C , A assigns $\pi_L(\sigma, C)$ to σ in C .

HD: This is plausible, though it may violate the limits of empiricism. How might it be connected to verbal behaviour? How is it connected to U-facts?

A: Perhaps assigning a proposition to a string correlates with being disposed to assert it when one believes it:

A assigns p to σ iff (A is disposed to assert σ in context C only if A believes p in context C).

HD: But people sometimes *lie*. A person might assign p to σ , and be disposed to assert σ in context C , even though believing $\neg p$ in context C . *Speech acts* cannot be

tied so closely to propositional content. Rather they are tied to a network of interconnected beliefs and desires.

A: Yes, that is too quick. Instead, try:

A assigns p to σ iff A accepts σ in context C only if A believes p in context C.

One might allow then for contexts where lying is justified, and then what one asserts may differ from, and even contradict, what one accepts. And instead we can adopt a somewhat hedged claim about speech acts:

(SA) Other things being equal: *A* is disposed to assert σ in context C only if *A* believes in context C the proposition that *A* assigns to σ .

HD: Is the claim about speech acts meant to be an analysis, or a contingent claim?

A: I am unsure; but, in any case, I think it is roughly right. If so, then I can see roughly how this theory logically relates C-facts to U-facts. Suppose that our evidence includes,

(U) Other things being equal, Naomi is disposed to assert ‘Das ist eine Katze’ in context C only when she believes that the ostended entity in C is a cat.

We can present a proposed *explanation* of that U-fact based on the theoretical assumptions:

- (COG2) *A* cognizes *L* only if, for all *L*-sentences σ , contexts C, *A* assigns $\pi_L(\sigma, C)$ to σ in C.
- (SEM) $\pi_{German}(\text{‘Das ist eine Katze’}, C) =$ the proposition that the ostended entity in C is a cat.
- (SA) Other things being equal: *A* is disposed to assert σ in context C only if *A* believes (in C) the proposition that *A* assigns to σ (in C).

along with the empirical claim,

(E) Naomi cognizes *German*.

For suppose (E): Naomi cognizes *German*. So, by (COG2), Naomi assigns $\pi_{German}(\text{‘Das ist eine Katze’}, C)$ to ‘Das ist eine Katze’ in context C. So, by (SEM), Naomi assigns the proposition that the ostended entity in C is a cat to ‘Das ist eine Katze’ in context C. So, by (SA), other things being equal: Naomi is disposed to assert ‘Das ist eine Katze’ in context C only if Naomi believes in context C the proposition that Naomi assigns to ‘Das ist eine Katze’ in context C. So, other things being equal: Naomi is disposed to assert ‘Das ist eine Katze’ in context C only if Naomi believes that the ostended entity in C is a cat. QED!

HD: Very good! Dispositions to verbal behaviour of an agent can be explained by what *language* the agent *cognizes*. So, as we concluded earlier, the U-facts do *not* determine the S-facts. Rather, the S-facts and C-facts together explain the U-facts.

A: I have a rather strong intuition that, unless one simply accepts the indeterminacies, the indeterminacy puzzles can't be resolved without requiring that speakers assign genuine *propositional content* to sentential strings; and perhaps one should explore meaning assignment to *subsential* strings too, using the more basic "semantico-cognitive notions", as you call them, such as,

agent *A* assigns the *intension* *I* to string σ in context *C*
agent *A* assigns the *extension* *E* to string σ in context *C*

suitably qualified, depending on which lexical category σ belongs to.
And we need not insist that these cognitive assignment notions be reducible or definable in more basic physicalistic terms.

HD: So, when one uses a word, it means just what one *assigns* it to mean—neither more nor less.