Why Truth is Not Important in Philosophy

1. John Austin said that any philosophical claim consists of the bit where you say it and the bit where you take it back. I am going to explicate the claim made in my inflammatory title in the reverse order.

The claim is not that *truthfulness* is not an important virtue. For it is a species of *honesty* or *sincerity*: honesty or sincerity in expressing one's beliefs. A truthful person does not assert what she does not believe. We *can* describe this virtue in terms of truth—"Don't put forward *as* true what you don't believe *is* true,"—but the appeal to truth in expressing it is optional, not essential, as my previous formulation shows. So I am not saying that philosophers can forget about telling the truth (as Plato thought poets do).

The claim is not that *epistemic conscientiousness* is not an important virtue. We should not be heedless, reckless, or careless in forming or adopting our beliefs. Once again, we *can* put this point in terms of truth: We should work hard to see to it that our beliefs are true. But once again, expressing this point in terms of truth is optional. For what doing that consists in is paying critical attention to our *evidence*, to the *justification* we have for endorsing various claims that we consider. What is incumbent on us as conscientious believers is not to be *credulous*, that is, not to acquire beliefs on the basis of insufficient evidence. It is not to be *prejudiced* or biased, that is, not to allow our preferences or desires—how we would *like* things to be—to suborn our assessment of the reasons there are to think that things actually are that way. It is to be *critical*,

that is, actively to seek out and honestly to assess possibly countervailing reasons: carefully to consider what justifications there might be for claims incompatible with the one we are assessing. So long as we pay sufficiently close attention to the *reasons* that can be offered for and against various claims, their *truth* will take care of itself—or at least, we will have done everything we can do about it.

The claim is not that *knowledge* is not important. Since Plato we have been told that knowledge is not just *belief*, but belief that can be *justified*; and not just justified belief, but *true* justified belief. And that is a perfectly good way of picking out a centrally important cognitive status.[1] But I think it is a fundamental mistake to think that what is important is the possession by beliefs of a certain metaphysically weighty *property*: being *true*. I think the beginning of wisdom in assessing the significance of the justified true belief analysis of knowledge is to think about what one is *doing* when one *attributes* knowledge to another, or *assesses* the credentials of another as a knower. For me to take you to know, for instance, that the Washington Monument is 555 feet tall, I must do three things. First, I must *attribute* to you a *belief* that the Washington Monument is 555 feet tall. You can't know what you don't believe. (Notice that we *could* say "believe to be *true*," but that doing so adds nothing.) Second, I must take it that you are *entitled* to that belief or commitment, that you have *reasons* for it, that you can *justify* it. An accidentally acquired belief is not yet knowledge. If you just picked some number out of the air, even if your lucky guess was right, you don't *know* that it is. Third, I must *myself* endorse the belief, that is, believe that the Washington Monument is 555 feet tall. That is, besides *attributing* to you both a *commitment* (corresponding to the *belief* condition), and an *entitlement* to that commitment

[1] Edmund Gettier famously argues that *more* is required for our ordinary conception of knowledge. But his arguments don't tell against the claim that justified true belief is still an important cognitive status.

(corresponding to the *justification* condition), I must myself *undertake* the corresponding commitment. That is what corresponds to the *truth* condition on knowledge. But *all* that condition is doing is marking the coincidence of belief across social perspectives: I only count as knowledge beliefs that I *share*.

Suppose you are standing in a darkened room, and seem to see a candle ten feet in front of you. I attribute to you the belief that there is a candle ten feet in front of you. And so long as you have no reason to think anything funny is going on, I take you to be justified in that belief, since you can see it. So I take you to be committed to there being a candle ten feet in front of you, and entitled to that commitment: to have a justified belief. Nonetheless, I will not take it that you *know* that there is a candle ten feet in front of you if *I* don't believe that—if, for instance, I, but not you, can see that there is an angled mirror five feet in front of you, and that the candle you see is actually quite close to you, hidden from you by a curtain. My assessment that your justified belief is not *true* is a way of expressing the fact that under the circumstances described, I am not willing myself to *undertake* commitment to the claim I *attribute* to you. Assigning some belief the honorific status of *knowledge* is important, because in doing that I am classifying it as being of the kind that I think *everyone* should employ as premises in their own inferences, should appeal to in their own reasoning. These are the beliefs that I take to be eligible to serve as reasons on the basis of which to form further beliefs. For I take it both that any good inference in which they figure as premises is one whose conclusions I should endorse, and I take it that good reasons can be given to believe them, in turn. Thus, these are the beliefs that I take it deserve to spread.

2. So far, in "the bit where I take it back", I've said that my claim that truth is not important in philosophy should *not* be understood as denying the importance of truthfulness, epistemic conscientiousness, or assessments of knowledge. But I've also said that in each of those cases, though we may if we like talk about the phenomena in question in terms of truth, we need not do so, and lose nothing essential if we do not. That last claim points towards a more positive characterization of the view I want to defend: Truth is not a concept that has an important *explanatory* role to play in philosophy. Appearances to the contrary are the result of *mis*understanding its distinctive *expressive* role. The word 'true' does indeed let us *say* things that in many cases we could not say without it. But when we understand *what* it lets us say, and *how* it does that, we will see that the very features that make it *expressively* useful make it completely unsuitable to do the sort of theoretical *explanatory* work for which philosophers have typically enlisted its aid.

The expression "…is true" looks like a predicate that ascribes a property. If it were, it would be a very special kind of immediately and unconditionally normatively significant property: a kind of "to-be-believed-ness" property. No wonder metaphysicians, ethicists, and especially epistemologists have regarded it with fascination. Nor is its normative weight exclusively of an abstract, disinterested, ethical sort—a high ideal that is a suitable object of selfless commitment by those of good character, lofty aspiration, and sufficient leisure. For, we are assured by the philosophical tradition, the *truth* of our beliefs is the touchstone and sole possible guarantor of the *success* of our practical endeavors—including the lowest and most narrowly self-interested. Having *beliefs* with the special, desirable property of being *true* is the only reliable way to get what you want—to imbue your *desires* with the most important and

desirable property *they* can aspire to: being *satisfied.* So truth is of supreme *practical* importance.

Besides its central significance for both the most ethereal principles and the most egoistical practices, truth has also seemed to hold the key to our inmost, ownmost nature. For (as Chapter Five emphasized) we are not merely *sentient* creatures, but also *sapient* ones. That is, in addition to consciousness in the sense of having *feelings* and *sensations*—awareness in the sense that underwrites a distinction between being *awake* and *asleep*—as our mammalian cousins such as cats do, we have states with conceptually articulated contents that can be expressed in *sentences*. We can believe *that* the international monetary system needs to be reformed and desire *that* it be reformed. These are the *propositional attitudes* that can constitute *knowledge*. And the standard way to understand the *propositional contents* that distinguish these states from the images and raw feels that are the contents of merely sentient states is that they can be assessed as to their *truth*. The *meaning* of a declarative sentence, expressing the *content* of a possible belief (or desire, or intention), consists in the circumstances under which it would be *true*. To grasp or understand that meaning or content just is to know its truth conditions: how the world would have to be for it to be true. So the sort of mindedness that distinguishes us from the beasts of the field—the sapience that gives our species its very name—consists in the relations we stand in to the very special property of truth: that we can think things that could be true, desire and intend that they be true. Take away that relationship to truth and you take away our sapience, relegate us to the cognitive torpor of mere sentience. This sapience- constituting directedness at *truth* is the essence and the motor of our ascent out of that primeval sea into the

broad highlands of thought. Philosophical concern with us, our nature and our spirit, *is* philosophical concern with truth.

3. This familiar philosophical scene, with truth at center-stage and in the leading role, is no doubt uplifting and inspiring. But I think it is deeply confused and almost totally wrong. Consider to begin with the idea that truth is the property of beliefs that conduces to the *success* of practical projects based on those beliefs. This thought is so deeply entrenched that some pragmatists have even sought to *define* truth as the success-producing property of beliefs.[2] But even those not inclined to endorse such an order of definition have felt free to appeal to the intimate connection between the truth of beliefs and the satisfaction of desires for other philosophical projects—for instance when scientific realists argue that the at least approximate truth of our scientific theories is the only possible explanation for the practical success of our technologies: the extent to which they provide powerful instruments for getting what we want (at least, for some kinds of things we want).

The idea is that it is the truth of my belief that there are cookies in the cupboard that explains the fulfillment of my desire for cookies. This is an intuitively compelling thought, but we need to be careful with it. The truth of that belief will not lead to satisfaction of my desire in the context of the collateral *false* belief that the cupboard is in the kitchen, rather than in the pantry. And, to vary the example, the *false* belief that one can tan leather by boiling it with birch-bark will result in practical success if it is combined with the false collateral belief that the

[2] For a sophisticated contemporary example, see J.T. Whyte's "Success Semantics", Analysis, 50.3 (June 1990) 149-157; followed up by "The Normal Rewards of Success", Analysis, 51.2 (March 1991) 65-73. I present a more detailed version of the argument that follows in my commentary: "Unsuccessful Semantics", Analysis Vol. 54 No. 3 (July 1994) pp. 175–8.

oak in front of me is really a birch. So the practical utility of a belief's being true is wholly hostage to the truth or falsity of the collateral beliefs with which it is combined.

Well then, perhaps one should only talk about the truth of a whole *set* of beliefs—indeed, of *all* one's beliefs. The requirement that we banish *all error* from our beliefs is a tall order, and probably not very realistic. But surely *that* would reliably produce successful, desire-satisfying actions? Not really. For the effects of collateral *ignorance* are just as bad as those of collateral *error*. If I am unaware that wet weather has swelled the cupboard door so that it cannot be opened, all my true beliefs about the location of the cookies and of the cupboard will be of no practical avail. But banishing *ignorance* as well as *error* seems over the top: is truth really only of practical use to the omniscient?

At this point one might be tempted to assimilate ignorance to error, by claiming that in the case where the swelling shut of the cupboard thwarted the good practical effects of my believing the cookies were in there, I really did believe that the door had *not* swelled and stuck, and so was in fact done in by my false belief. But there is not just an infinite, but, worse, an *indefinite* number of ways in which my plan could go awry through circumstances of which I am unaware. Is it at all plausible that I have beliefs about them all? That the cupboard has not been nailed shut, moved, mined, infested with voracious beetles, encased in glass, shrunk to microscopic size, and so on? As a last, desperate attempt, one might stipulate that I had a collateral belief that there were no impediments to my opening the cupboard and eating the cookies. Insisting that *that* collateral belief be present and true will indeed ensure that my true belief that there are cookies in the cupboard will lead to practical success in satisfying my cookie-desire. But it does so only

by trivializing the original claim. For it is, in effect, the demand that the original true belief be accompanied by a *true* belief that it will lead to practical success. And *that* would hold also for my *false* belief that boiling with birch-bark will tan leather, in all those cases in which I have a *true* collateral belief that there are no impediments to my tanning leather by boiling it with birch-bark. I think that when one looks closely at the claim, one finds that there is no way to make sense of the idea that true beliefs produce successful actions so that it is both true and non-trivial. This sort of explanatory appeal to truth, at any rate, collapses when weight is put upon it.

4. What about the role of truth in *semantic* explanation, via a definition of propositional content in terms of truth conditions? We certainly do use 'true' to *say* what the content of a claim is. If you don't understand the sentence "The surgeon performed a cholecystectomy," I can explain it to you by telling you that it is true just in case the surgeon removed the patient's gall bladder. And we can say more general things, such as: Any claim of the form ~*p* is true just in case *p* is not true. But it would be a mistake to infer from this sort of appeal to truth conditions to *express* propositional contents that one can *explain* what propositional contents *are* by appeal to the conditions under which sentences are true. That would be a possible order of explanation *only* if one can make sense of the notion of truth prior to and independently of making sense of the notion of propositional content. And there is good reason to think that that cannot be done.

What is the expressive role of "true"? Let us start by drawing a lesson from our discussion of the JTB account of knowledge. What is one *doing* when one *says* that some claim is true? The answer offered there is that one is *endorsing* it, undertaking or acknowledging a

commitment with that same content. So here is a first try: Saying or believing that "Snow is white" is true—that it is true that snow is white—is just saying or believing that snow is white. This, however, by itself, won't do. For it only covers some of the cases. When I say "If it is true that *p*, then *q*," I have *not* said that *p* is true, and have not asserted *p*. We need something more general, something that applies to free-standing, assertional uses of 'true', but also to embedded, non-assertional uses of 'true', in which its use just contributes to the *content* of what is said, without being itself invested with assertional *force*.

So here is a try. "It is true that *p*," is in *all* cases, *both* free-standing, force-bearing, *and* embedded uses, equivalent to *p*.[3] This is, roughly, Tarski's famous "Convention T". It defines 'true' by the principle that the quoted-or-named sentence ⌜*p*⌝ is true just in case p, where 'p' is any claim that has the same *content* as is expressed by the expression ⌜p⌝. This can be called a "content-redundancy" theory of truth.[4] This is plausible. But notice that it *uses* the notion content, or at least same-content (co-contentful) to *define* 'true'. If *that* is how we define truth—Tarski-wise—then truth cannot be appealed to in explaining the notion of content, since that notion is rather defined by appealed to the notion of content, and accordingly presupposes it.

I think that the theory that best captures the insight behind the intuitive thought that "It is true that *p*," just *means* the same (expresses the same content as) ⌜*p*⌝, is the *prosentential* account.[5] The basic idea behind this theory is to exploit the analogy between identity of content

[3] For a sophisticated exposition and defense of this view, see Paul Horwich *Truth* [Oxford University Press, 1998].
[4] It is usually associated with Frank P. Ramsey, on the basis of his article "Facts and Propositions (1927), reprinted in *F. P. Ramsey*: *Philosophical Papers* edited by D. H. Mellor (Cambridge University Press, 1990) pp. 34-51.
[5] First put forward in Dorothy Grover, Joseph Camp, and Nuel Belnap, "A Prosentential Theory of Truth", *Philosophical Studies,* 27 (1975): 73-4.

within the category of *sentences* (which is what the content-redundancy theory appeals to in defining 'true') and identity of content within the category of *singular terms*. For in that case, there are *two* ways that one can commit oneself to the identity of content of two term-tokenings. One can explicitly assert an *identity*. So in the context of the claim "Benjamin Franklin was a printer," the explicit identity statement "Benjamin Franklin was (=) the inventor of the lightning rod," licenses the conclusion: "The inventor of the lightning rod was a printer." But one also might use a *pronoun* to achieve the same effect, by implicitly *stipulating* the relevant content-identity. Thus one can derive the very *same* conclusion from the claim "Benjamin Franklin was a printer," by appending the auxiliary hypothesis "And *he* was the inventor of the lightning rod." For the fact that the anaphoric *antecedent* of the pronoun 'he' is the term "Benjamin Franklin" licenses the very same substitution as the explicit identity. The prosentential theory of truth is what you get if you transform the content-redundancy view by assimilating it to the second paradigm, rather than the first. On this view, when Hegel says "History is the progress of the consciousness of freedom," and I say "That is true," I have made myself a same-sayer with Hegel in exactly the same way that I have managed to refer to the same person when he says "Hamann was a great philosopher," and I say "But he was an incredibly strange human being." In both cases the content of my remark inherits its content—respectively what I say and who I am talking about—from the antecedent provided by Hegel.

The anaphoric account of the identity (redundancy) of content of "It is true that *p*," and $\lceil p \rceil$ has many technical advantages. In the most sophisticated version of the theory, "…is true" is understood as a prosentence-forming operator. It applies to a singular term, which specifies the sentence that is the anaphoric antecedent of the prosentence, from which, accordingly, it inherits

its content, in the same way a pronoun like 'he' or 'it' inherits its content from its antecedent. There are as many ways of using 'true' as there are ways of specifying sentences that can serve as the anaphoric antecedents. There are quote-names, such as "Snow is white," yielding prosentences such as "'Snow is white,'is true," which inherit their content from the named-by-being-quoted sentence. But there are also descriptions of sentences, such as "Fermat's last theorem", and "The first full sentence on page 37 of Quine's *Word and Object*", and prosentences that inherit their content from the sentences so indicated, such as "Fermat's last theorem is true," and "The first full sentence on page 37 of Quine's *Word and Object* is true." The rule for applying sentential operators such as 'not' and 'probably' in prosentences such as "What the policeman said is not true," and "Tarski's favorite sentence is probably true," is that the resulting prosentence is equivalent to the result of applying the operator to the sentence (or sentences) picked out as an antecedent by the phrase the prosentence-forming operator applies to. So if the policeman said "The suspect fell down," the prosentence user is claiming something equivalent to "The suspect did not fall down," and in the other case, something equivalent to "Snow is probably white." And notice that all of these equivalences work equally well for embedded uses. If I say "If what Hegel said is not true, then I don't know what history is," I've committed myself to the conditional claim that if history is not the progress of the consciousness of freedom, then I don't know what it is.

I assert, though I certainly cannot pretend to be able adequately to justify the claim here[6], that the prosentential theory offers a complete, adequate, and satisfactory account of the use of

[6] I try to do that in Chapter Five of *Making It Explicit* (Harvard University Press, 1994) and "Expressive vs. Explanatory Deflationism About Truth" in *What Is Truth?* Richard Schantz (ed.) Hawthorne de Gruyter, Berlin & N.Y, 2002, pp. 103-119

the term ‘true’. If that is so, two consequences are worthy of note. First of all, one can use ‘true’ to *express* conceptual or propositional contents, since claims formed using ‘true’ inherit their contents from their antecedents. But the pronominal or prosentential relation is one of content-*inheritance*: one must *already* have a good grip on the notion of the content that is inherited in order to understand it. Second, on this account there is no more a *property* expressed by ‘true’ than there is a kind of *object* picked out by ‘it’. After all, “that is so” (which I used just above) is *also* a prosentence, and one ought not (though some have been, are, and will be) be tempted to look for a property of so-ness that claims could have. And surely, even if one succumbed to that temptation, one would have to admit that one needed an *antecedent* grasp of the content that was a *candidate* for being ˢsoˢ. There is no go at all to the idea that one could *explain* those contents (the things that could be so) in terms of a *prior* notion of so-ness or being so (what, exactly, being so?). If we are going to explain one in terms of the other, surely the right way to go about things is to make sense first of the notion of a proposition—of something that *can* be true (as, for instance, the doorknob on my office cannot).

Prosentences, like pronouns, do let us say things we could not say otherwise. Besides the lazy or redundant uses of pronouns, there are quantificational ones. Using pronouns, I can say not only “If John thinks that, then he is being foolish,” but also “If anyone beats his donkey, then he is mean.” And using prosentences, I can say not only “What John said is true,” but “Everything the policeman said is true.” Our language would be expressively impoverished without pronouns, and it would be expressively impoverished without prosentences. But that important fact has no bearing on the suitability of the ˢconceptˢ it to serve to explain the notion of

object[7], nor, correspondingly, on the suitability of the ˢconceptˢ so, or true to explain the notion of a proposition.

5. So far, my aim has been relentlessly critical. I've claimed that one should no more think of truth as a philosophically weighty property that can be wheeled in to perform various sorts of crucial explanatory labor than one should think of no-one as a particularly spooky kind of person: the one who is in your room even when *no-one* is in your room. (Think here of crafty Ulysses, telling the Cyclops that his name was "Noman", so that when he escaped, blinding the Cyclops, the Cyclops' neighbors responded to his cries "Noman is killing me!" by saying: "If no man is killing you, shut up and let us sleep!") But it is fair enough to ask: if truth is not important in philosophy, what is? What *does* articulate our sapience and so constitute the difference that makes a difference in distinguishing us from the beasts of the field? What *does* matter in epistemology, in semantics, and in the philosophy of mind? It is no use deflating truth unless there is some candidate available to replace it. And after all, the anaphoric, prosentential account, it was insisted, *presupposes* a notion of the propositional contents that truth-claims inherit from their antecedents. How *are* we to understand and explain propositional content, if not in terms of truth conditions? And what *is* the overarching cognitive goal we are supposed to be pursuing, if not truth?

[7] Not such a quixotic enterprise that it has never been attempted. Think of Quine's slogan: "To be is to be the value of a variable." But Quine uncritically took for granted the availability of a *metalanguage* in which one *could* pick out the domain elements over which his variable (a kind of pronoun) range. Though that is unobjectionable for certain *formal* purposes, it does not support a corresponding order of *philosophical* explanation.

The slogan that expresses the view I want to recommend as in answer to these reasonable questions is: from *knowledge* to *understanding*; from *truth* to *inference*. On the first point, we may paraphrase T. S. Eliot: "Where is the understanding we have lost in knowledge? Where is the knowledge we have lost in information?" But it is by thinking about propositional and conceptual content—and thereby our sapience—in terms of *inference* rather than *truth* that we will find a way through and forward from the confusions I have been diagnosing and criticizing.

Frege, the founder of modern logic, explicitly codified a semantic principle relating truth to inference: good inferences never take premises that are true into conclusions that are not true. The traditional way of exploiting this principle is to use it to underwrite an explanation of the goodness of inference in terms of truth: any inference is good, so long as it does not have true premises and a false conclusion. This is a very weak notion of <u>good inference</u>: the inference from "Snow is white," to "Grass is green," is a good one in this sense, even though the premise has nothing to do with the conclusion. It is, of course, a good feature of an inference that it does not have true premises and a false conclusion—at least, it would be a bad feature if it did. Surprisingly, Frege was able to show that this weak notion of good inference suffices for many purposes in mathematics. But we could exploit Frege's semantic principle to underwrite an explanation going the other way around: Truth is what is preserved by good inferences. Of course, to do that, we would need an antecedent, independent grip on the notion of <u>good inference</u>.

In the very same work in which he puts forward his semantic principle—his seminal *Begriffsschrift* of 1879—Frege suggests how we could think about propositional content in terms

of inference. His avowed aim is to explicate the notion of "conceptual content" [begriffliche Inhalt]. The qualification "conceptual" is explicitly construed in inferential terms:

> ...there are two ways in which the content of two judgments may differ; it may, or it may not, be the case that all inferences that can be drawn from the first judgment when combined with certain other ones can always also be drawn from the second when combined with the same other judgments. The two propositions 'the Greeks defeated the Persians at Plataea' and 'the Persians were defeated by the Greeks at Plataea' differ in the former way; even if a slight difference of sense is discernible, the agreement in sense is preponderant. Now I call that part of the content that is the same in both the conceptual content. Only this has significance for our concept language [Begriffsschrift]... In my formalized language [BGS]...only that part of judgments which affects the possible inferences is taken into consideration. Whatever is needed for a correct ['richtig', usually misleadingly translated as 'valid'] inference is fully expressed; what is not needed is...not.[8]

Two claims have the same conceptual content iff they have the same inferential role: a good inference is never turned into a bad one by substituting one for the other.

Frege's idea is that we think of conceptual content in terms of role in reasoning. To be *propositionally* contentful is to be able to play the role both of *premise* and of *conclusion* in an inference. (We have to insist on both, since one common way of thinking about *practical* reasoning is in terms of inferences whose conclusions are *doings* rather than *sayings*.) Frege's semantic principle relating truth to inference tells us that anything that can play the role both of premise and conclusion in an inference will be the right sort of thing to be evaluated as to its *truth*. So we can pick out declarative *sentences*, as the linguistic items that express *propositions*, in terms of their role in *inference*.

[8] Frege, *Begriffsschrift* (hereafter *BGS*), section 3.

An order of semantic explanation that approaches the notion of conceptual content, paradigmatically, *propositional* content, through that of *inference* rather than through that of truth has the advantage of establishing at the outset of the enterprise a connection between the *contents* possessed or expressed by sentences, on the one hand, and what language users actually *do*, on the other. For *inferring* is something that thinkers and speakers *do*. It is, in Wittgenstein's phrase, making a *move* in the language game. So, of course, is asserting. Asserting *can* be thought of as putting a sentence forward *as true*. But then we need to ask what doing that consists in: putting it forward as *what*, exactly? Following out Frege's strategy of focusing on inference suggests that putting something forward as true just is putting it forward as an appropriate *premise* from which to make *inferences*. That is putting it forward as something that is appropriately responded to practically by *doing* something: making an inference. It is not possible to make sense of the notion of inference apart from that of assertion, since assertions are the termini of inferential moves. But it may be possible to make sense of these two kinds of doing together, without having to appeal to the notion of truth at the outset, and then get a grip on the notion of truth as what is both preserved by good inferences and what one is putting something forward *as* when one asserts it.

6. Consider how the notion of understanding looks from the two different semantic perspectives being considered. From the point of view of *truth*, understanding a sentence, associating a propositional content with it, is knowing its truth conditions: knowing how the world would have to be for it to be true. Here understanding is explained in terms of a special kind of knowledge. But what does this kind of knowledge consist in? Is it a justified true belief? If so, *what* belief? I understand the content of the sentence "Snow is white." I know that it is

true just in case snow is white. Is *that* the justified true belief that my knowledge of the truth conditions of "Snow is white," consists in? If so, how do I understand *that* content? We seem to be moving in a very small circle, and not really explaining the content or what understanding it consists in. From the point of view of *inference*, understanding a sentence, associating a propositional content with it, is having practical mastery of its inferential role: knowing what follows from it, and what would be evidence for or against it. The talk of 'knowledge' here is very different from that involved in knowledge of truth conditions. For it is a kind of knowing *how* rather than knowing *that*: knowing how to *do* something, namely distinguish in practice between good inference and bad inference in which the sentence appears as a premise or conclusion, rather than knowing *that* the truth conditions are such-and-such. Understanding shows up on this account as a practical ability, a kind of skill: sorting possible inferences into good ones and bad ones, endorsing or being disposed to make some of them, and rejecting or being disposed not to make some others.

Of course, one might try to understand the knowledge of truth conditions that understanding consists in on the traditional approach also as a kind of practical know-how, rather than as a kind of theoretical knowing that—that is, as not consisting in an explicit (hence itself propositionally contentful true and justified *belief*, which is what gets us into circularity problems). But what kind of know-how or practical ability *is* knowing the truth conditions of a sentence? The standard answer—that of Michael Dummett, for instance—is that it is the capacity to *recognize* the truth of the sentence: to sort situations into those in which it is true and those in which it is not. This is a practical *classificatory* ability: the ability reliably to respond differentially to situations in which the claim is true. The case where this proposal has the most

intuitive appeal is that of *observational* concepts. One who understands the concept red can respond differentially to the visible presence of red things: paradigmatically, by saying "That's red." But we need to be careful here. After all, a parrot could learn to do that. Do we want to say that a parrot that reliably responds to red things by uttering the noise "Rrawk—that's red!" shows that it grasps or understands the concept red? What about a photocell hooked up to a tape-recorder? A chunk of iron reliably responds differentially to wet environments by rusting. Should we therefore credit it with understanding the concept wet, and treat its rusting as the utterance of a sentence whose truth conditions are that it is true just in case something is wet?

Something has gone wrong here. The parrot, the photocell, and the chunk of iron can serve as *instruments* for the detection of red things or wet things, because they respond differentially to them. But those responses are *not* claims *that* things are red or wet, precisely because they do *not understand* those responses *as* having that meaning or content. By contrast, when *you* respond to red things or wet things by saying "That's red," or "That's wet," you *do* understand what you are saying, you *do* grasp the content, and you *are* applying the *concepts* red and wet. What is the difference that makes the difference here? What practical know-how have you got that the parrot, the photocell, and the chunk of iron do not? I think the answer is that *you*, but *not* they, can use your response as the premise in *inferences*. For *you*, but *not* for them, your reponse is situated in a network of connections to other sentences, connections that underwrite inferential *moves* to it and from it. You are disposed to accept the inference from "That's red," to "That's colored,", to reject the move to "That's green,", and to accept the move to it from "That's a stoplight." You are willing to make the move from "It's wet," to "There is water about," to infer it from "It is raining," to take it as ruling out the claim "We are in a

desert," and so on. Because you have the practical ability to sort inferences in which it appears as a premise or conclusion into good ones and bad ones, *your* response "That's red," or "It's wet," is the making of a *move* in a language game, the staking of a claim, the taking of a stand that commits you to other such claims, precludes some others, and that could be justified by still others. Having practical mastery of that inferentially articulated space—what Wilfrid Sellars calls "the space of reasons"—is what *understanding* the concepts red and wet consists in. The responsive, merely classificatory, *non*-inferential ability to respond differentially to red and wet things is at most a necessary condition of exercising that understanding, not a sufficient one.[9]

We can come to the same conclusion from another direction. We were considering a semantic theorist who claims that understanding the propositional content expressed by a sentence consists in knowledge of its truth conditions. This claim faces the objection that if that knowledge is understood as knowing *that* things are thus-and-so, that is, as having some justified true *belief*, that the account is going to be objectionably circular or fatally incomplete, since we will not have an account of what understanding the content of *that* belief consists in. The response was then that one might instead construe knowledge of truth conditions as a kind of know-how, as practical mastery of inferential roles is. In particular, it could be thought of as *recognitional* capacity. We've seen that this will work in the central cases that motivate the idea, namely *observational* concepts, concepts that have *non*-inferential, responsive uses, *only* if we supplement it with an appeal to practical mastery of *inferences*. But what about the contents of other kinds of sentences, sentences that do *not* have observational uses? How am I supposed to manifest the practical ability to recognize the truth conditions of sentences such as "The

[9] This line of thought is pursued further in Chapter Seven.

Washington monument is 555 feet tall," or "The international monetary system needs to be reformed,"? It is hard to avoid the conclusion that this practical ability would itself have to be understood in *inferential* terms. I could manifest my capacity to recognize the truth conditions of the first sentence for instance by endorsing the inference to the claim about the height of the Washington monument from a picture of it next to a giant ruler, or from a measurement of the length of its shadow and a simultaneous measurement of the length of the shadow of a stick of known length. And I could manifest my practical capacity to recognize the truth conditions of the second sentence by accepting an argument to that conclusion from a claim about the tendency of currency fluctuations to further impoverish the poorest of the poor countries, while rejecting an argument to that conclusion from the claim that international trade is growing by great leaps and bounds.

So it looks as though, whether we think about observational concepts or theoretical ones, the practical abilities associated with grasp of truth conditions turn out to be inferential ones. And there seems to be a large number of other philosophically interesting kinds of concepts whose contents are *only* going to be intelligible in terms of their roles in reasoning: for instance, modal, normative, and probabilistic concepts, and those having to do with the past or the future. At the very least, approaching these potentially puzzling kinds of concepts by asking about their role in reasoning promises to avoid the sort of spooky metaphysical questions that arise were we to ask instead about the truth-makers of claims formulated in those vocabularies: about the metaphysical status of *past* (hence no longer existent) facts, or merely *possible* fat men stuck in the doorway, or of that famous 10% chance of rain on a day when it does not in fact rain. So, for instance, the difference between merely accidental regularities, such as "all the coins in my

pocket are copper," and genuinely lawlike regularities, such as "the melting-point of copper is 1083.4° C," is that the latter, but not the former support *counterfactual* inferences. It is does *not* follow from the first claim that if this nickel *were* in my pocket, it *would be* copper, but it *does* follow from the second that if I *were* to heat this penny to 1083.4° C, it *would* melt. Again, you will think much more usefully about the content of claims such as "Freedom is better than slavery," and "Blake is a better poet of the imagination than Milton," by thinking in terms of their inferential roles—what follows from such claims, what do they preclude, what would be evidence for or against them?—than if you try to think in terms of their truth conditions.

7. Thinking of conceptual content in terms of inferential role, and of understanding correspondingly as practical mastery of such an inferential role—as the ability to sort into good and bad the inferences in which the concept appears in the premises or conclusions—has other advantages as well. It is a powerful corrective to the philosophically unilluminating and pedagogically damaging cartesian picture of the achievement of understanding as the turning on of some kind of inner light, which permits one then to see clearly. This is what the elementary-school kid thinks happens in math class when the girl next to him "gets it", and he doesn't. He is waiting for the light to go on in his head, too, so that he'll understand fractions. In fact, he's just got to practice making the moves, and distinguishing which ones are OK and which ones are not, until he masters the practical inferential abilities in question. It is not unusual for teachers of technical material to have students who can do all the problem-sets and proofs, can tell what does and doesn't follow from some situation described using the concepts being taught, but still think that they don't understand those concepts. A feeling of familiarity and confidence in knowing one's way around in an inferential network often lags one's actual mastery of it. The

important thing is to realize that the understanding *is* that practical mastery, and the feeling (the cartesian light) is at best an indicator of it—often an unreliable one.

We professors tell our students that it is important to think and write *clearly*. No doubt it is. But this can be frustrating advice to receive. After all, presumably no students think *that* fuzzy thinking and fuzzy writing are better than the alternative. The hard thing is to tell the difference. What, exactly, is one supposed to *do* in order think or write more clearly? Thinking about meaning and understanding in terms of inference provides some more definite guidance in this area. Thinking clearly is a matter of knowing, for each claim that you make, what else you are committing yourself to by making it, what you are ruling out, and what would be evidence for or against it. You can test the clarity of your thinking by rehearsing sample inferences, so as to test your practical mastery of the inferential vicinity of your thoughts. Of course, you may be mistaken about what really does follow from your claims. But that is *just* a mistake. So long as you are sure what you *take* to follow from and be evidence for your claim, your mistaken thought is at least *clear*. And writing clearly is giving your reader enough clues that *she* can tell what you *mean* to be committing yourself to by the claims you make, what *you* would take to be evidence for or against them, what follows from them, and what they preclude. And once again, this is something you can check for yourself when writing, by asking yourself, for each important consequence you take to follow from one of your claims, how your reader is supposed to *know* that you take it to be a consequence: what clues have you given to that effect?

The prosentential approach to truth-talk not only supports the claim that the use of 'true' presupposes a notion of <u>semantic</u> <u>content</u>, and so cannot be the basis of an *explanation* of that

notion, but also shows how to build an account of truth-talk out of an account of semantic content. Since we have in view the possibility of explaining semantic content in terms of role in reasoning, the prosentential theory accordingly offers an account of the important *expressive* (not explanatory) role that 'true' plays. And that account will underwrite Frege's semantic principle relating truth and goodness of inference.

8. One might ask whether the inferentialist approach does not require overestimating the extent to which we are rational. Are we really very good at telling what is a reason for what? How often do we act for reasons—and in particular, for *good* reasons? The question betrays a misunderstanding. We are rational creatures in the sense that our claims and aims are always liable to *assessment* as to our reasons for them. How good we are at satisfying those demands doesn't change our status as rational. And it must be kept in mind that on this way of thinking about the nature of semantic content, it makes no sense to think of us *first* having a bunch of sentences expressing definite propositions, which accordingly stand in inferential relations to one another, and only *then* having there be a question about how many of those inferences we *get right*. For it is our practices of treating what is expressed by some noises as *reasons for* what is expressed by other noises that makes those noises express conceptual contents in the first place. Once the enterprise is up and running, we can certainly make mistakes about what follows from the commitments we have undertaken, and what would justify them. But there is no possibility of us *massively* or *globally* getting the inferences wrong (for very much the same Quinean reasons that Davidson has emphasized).

I have been arguing that it is better to think in terms of understanding than knowledge, and better to think of meaning-and-understanding (which on this approach are two sides of one coin) in terms of inference than in terms of truth. So far, I have approached this issue largely from the point of view of semantics and the philosophy of language. But there is more at stake here. For this way of thinking about semantic content goes to the heart of the question of what it is to be sapient—to be the kind of creature we most fundamentally are. It says that we are beings that live, and move, and have our being in the space of *reasons*. We are, at base, creatures who give and ask for reasons—who are sensitive to that "force of the better reason", persuasive rather than coercive, which so mystified and fascinated the ancient Greek philosophers. Crossing that all-important line from mere sentience to sapience is participating in practices of giving and asking for reasons: practices in which some performances have the pragmatic significance of *claims* or assertions, which accordingly, as both standing in need of reasons and capable of serving as reasons (that is, of playing the role both of conclusion and as premise in inference) count as expressing *propositional* semantic content.

This *semantic rationalism*—which goes with thinking of content in the first instance in terms of *inference* rather than *reference*, *reason* rather than *truth*—flies in the face of many famous movements in 20th century philosophical thought. The American pragmatists, above all, John Dewey, used the possibility of explaining knowing *that* in terms of knowing *how* not only to assimilate our sapient intellectual activity to the skillful doings of merely sapient animals, but at the same time to blur the sharp, bright line I am trying to draw between sapience and sentience. Wittgenstein famously said that language does not have a 'downtown': a core set of practices on which the rest depend, and around which they are arrayed, like suburbs. But

inferentialism says that practices of giving and asking for *reasons are* the 'downtown' of language. For it is only by incorporating such practices that practices put in play propositional and other conceptual contents at all—and hence count as *discursive* practices, practices in which it is possible to *say* anything. The first 'Sprachspiel', language game, Wittgenstein introduces in the *Philosophical Investigations* has a builder issuing ˢordersˢ to an assistant. When he says 'Slab!' the assistant has been trained to respond by bringing a slab. When he says 'Block!' the assistant has been trained to respond by bringing a block. From the inferentialist point of view, this does not qualify as a *Sprach*spiel at all; it is a *vocal*, but not a *verbal* game. For the assistant is just a practical version of the parrot I considered earlier: he has been trained reliably to respond differentially to stimuli. But he grasps no concepts, and if this is the whole game, the builder expresses none. An *order* or command is not just any signal that is appropriately responded to in one way rather than another. It is something that determines *what is* an appropriate response by *saying* what one is to do, by *describing* it, specifying what *concepts* are to apply to a doing in order for it to count as *obeying* the order. Derrida's crusade against what he calls the 'logocentrism' of the Western philosophical tradition has brilliantly and inventively emphasized all the *other* things one can do with language, besides *arguing, inferring, explaining, theorizing,* and *asserting*. Thus we get the playful essays in which the key to his reading of Hegel is that his name in French rhymes with 'eagle', his reading of Nietzsche that turns on what Derrida claims is the most important of his philosophical writings (a slip of paper that turned up in his belongings after his death, reading only "I have forgotten my umbrella,"), and the unforgettable meditation on the significance of the width of the margins of the page for the meaning of the text printed there. But if inferentialism is the right way to think about contentfulness, then the game of giving and asking for reasons *is* privileged among the games we

play with words. For it is the one in virtue of which they mean anything at all—the one presupposed and built upon by all the other uses we can then put those meanings to, once they are available. Again, the master-idea of Foucault's critique of modernity is that *reason* is just one more historically conditioned form of *power*, in principle no better (and in its pervasive institutionalization, in many ways worse) than any other form of oppression. But if giving and asking for reasons is the practice that institutes *meanings* in the first place, then it is does *not* belong in a box with violence and intimidation, which show up rather in the contrast class precisely insofar as they *constrain* what we do by something *other* than reasons.

9. My initial claim was that truth is not important in philosophy. I offered some reasons for that critical assessment, culminating in the more specific claim that when we properly understand the *expressive* role of 'true', we see that that role precisely disqualifies it from doing the sort of important *explanatory* work in various areas of philosophy, beginning with semantics, for which it has often been pressed into service. But rather than leaving it at that—content, as it were, with the bit where one takes it back—in the second half of this talk I have sketched an alternative constructive account of some of the most important phenomena that one *might* have thought one needed an oomphy concept of <u>truth</u> to address. On this view, we are seekers and speakers of *truth because* we are makers and takers of *reasons*. I think this way of thinking about ourselves can be as edifying and inspiring as the one it seeks to supplant.