

Hesperus and Phosphorus: Sense, Pretense, and Reference

Mark Crimmins

In “On Sense and Reference,” surrounding his discussion of how we describe what people say and think, identity is Frege’s first stop and his last. We will follow Frege’s plan here, but we will stop also in the land of make-believe.

Identity challenges reflection, says Frege. By *identity* Frege intends whatever it is we attribute when we say, for instance, that Hesperus is *the same thing* as Phosphorus. It can seem odd, he says, to think of identity as a relation between *things*; for, what interest could attach to a thing’s being *itself*? It is odd, too, to think of identity as the relation that holds between different *names* of a thing, for then the statement that Hesperus *is* Phosphorus would be held to concern neither Hesperus nor Phosphorus, but only the name ‘Hesperus’ and the name ‘Phosphorus’ (and this would alienate the identity statement from such apparent kin as the statement that Hesperus is the same *size* as Phosphorus). Frege’s resolution of the dilemma is well known: intermediate between a name and the thing named is a *mode of presentation* of the thing; the name expresses *that* way of thinking about the thing. A true identity statement like ‘Hesperus is Phosphorus’ is not *trivial* because the entity that in the statement is said to be identical to itself is conceived differently under the two names: the names ‘Hesperus’ and ‘Phosphorus’ express different modes of presentation of a single object.

Russell, too, finds his reflection challenged by identity. In “The Philosophy of Logical Atomism” he writes:

Identity is a rather puzzling thing at first sight. When you say ‘Scott is the author of *Waverly*’, you are half-tempted to think there are two

This paper has benefited from very generous help. Thanks to participants in my seminar at the University of Michigan, and to audiences at MIT, Wayne State University, New York University, and Stanford. I am particularly grateful to Kent Bach, David Braun, Alex Byrne, Ray Elugardo, Avrom Faderman, Graeme Forbes, Ned Hall, David Hills, Jim Joyce, Michael McKinsey, Joe Moore, Larry Powers, François Recanati, Ian Rumfitt, Stephen Schiffer, Laura Schroeter, Dan Sperber, Loretta Torrago, Ken Walton, Steve Yablo, Ed Zalta, and the editors of the *Philosophical Review*.

people, one of whom is Scott and the other the author of *Waverly*, and they happen to be the same. That is obviously absurd, but that is the sort of way one is always tempted to deal with identity.¹

Russell instead deals with identity by holding that nontrivial, true identity statements like 'Hesperus is Phosphorus' do not attribute relations at all. Rather, they assert complex claims about what *sorts* of things exist: for instance, 'Hesperus is Phosphorus' might assert that there exists a thing that is both the first star visible in the evening sky (in one season) and the last star visible in the morning sky (in another season).

I want to underline, however, Russell's phenomenological confession: *when we say* 'Scott is the author of *Waverly*', we are half-tempted to see what we are doing as stating of *two things* (Scott and the author of *Waverly*), that these two things 'happen to be the same'. Russell is right about the semantic phenomenology. At least, it is like that for me, too, and even more strongly in other cases (Hesperus and Phosphorus, Superman and Clark, the child and the adult). I agree, as well, that it would be absurd to think that we seriously mean to be claiming of *two things* that they are the same thing. This peculiar feeling of talking about two things cannot reveal what we are seriously doing. But before we dismiss this peculiar feeling, perhaps we should consider whether it is clear that in making a serious identity statement, everything we do is done seriously.

1. Make-Believe and Truth

Statements (and more generally uses of sentences) that rely on make-believe can be used to express genuine claims, and can be candidates for genuine truth and falsehood. Let me support this with an example in which the role of make-believe is more blatant.

Suppose that I am trying to describe to you Ann's cleverness and modesty, and that I cannot find words to express neatly and perspicuously just the traits I want to characterize. I might compare her to some of our mutual acquaintances. I might, that is, say:

(1) Ann is as clever as X and more modest than Y.

¹ *The Philosophy of Logical Atomism* (La Salle, Illinois: Open Court, 1985), 115.

But let us suppose that none of our mutual acquaintances will do. The traits I have in mind, however, are related in a relatively direct way to traits that Conan Doyle portrays his characters as having. So I might say:

- (2) The degree of cleverness and the degree of modesty that actually are such that in the Sherlock Holmes stories there is portrayed there being a person named 'Holmes' with that degree of cleverness and there being a person named 'Watson' with that degree of modesty, are such that Ann's degree of cleverness is comparable to the former, and her degree of modesty is greater than the latter.

While perspicuous, that takes a long time to say, it is not easy to follow, and one needs considerable conceptual sophistication to formulate or to understand it. What I actually say, of course, is:

- (3) Ann is as clever as Holmes and more modest than Watson.

In saying this I am not really comparing Ann to two other people. But it is important to see that I am *making as if* to do exactly that. My use of the sentence (3) is framed from within a shallow pretense: a shared, conspiratorial make-believe that we can refer with the names 'Holmes' and 'Watson' to people who are as described in Conan Doyle's stories. It is this pretense that explains my choice of these words in this particular arrangement. Sentence (3) is a sentence for comparing a person to two other people, and, viewed from inside the context of our shallow pretense, the sentence indeed expresses a comparison of Ann to two other people. Certainly, the kind of "making as if" this exemplifies is not the sort of pretense that draws us into imaginative play; that is why I call it a shallow pretense. But it is nonetheless a form of pretense: to understand me, you have to see that I intend my utterance to be *as if* a statement comparing Ann to two other people, and *only as if* that. You have, that is, to distinguish what's so from what's pretend-so.

The point of pretending, shallowly, with you that I am saying something about three people is that, in so doing, I genuinely say what I would have said in the laborious (2). The crucial distinction is between the sort of claim I make as if to express (I make as if to express a claim about three people), and the sort of claim I genuinely, seriously express (I genuinely express a claim about just

one person, namely Ann). I speak within the pretense, pretendedly about the world as we pretend it to be. But I also speak *through* the pretense, about the world as it really is. Somehow, the context of pretending allows me to generate with a pretend assertion of one sort of claim a genuine, serious assertion of a different sort of claim.

But how does this work? What is the connection between what goes on within the pretense and what is really, seriously going on?

Answering that question satisfactorily means confronting a web of issues about the nature of assertion and content. But until we have made a bit of headway in our main tack (we are headed, I remind you, to a consideration of statements about what people think and say), the relevant subtleties are likely to seem more subtle than relevant.

Let's focus on the question of the truth condition of my utterance—that is, on the question what the world really must be like if my utterance is in fact to be a true utterance.² The answer I favor will require elaboration, but it is straightforward: the utterance is really true just in case it is true within the pretense. In elaborating this, I will borrow heavily from Kendall Walton's work on make-believe.³

If we make believe as I am inviting us to with statement (3), then we are making believe that with 'Holmes' and 'Watson' I can refer to persons with certain levels of cleverness and modesty. So *within the pretense*, what it takes for (3) to be true is that Ann's cleverness approximate Holmes's and that her modesty exceed Watson's. Thus, for the statement to be correctly attributed truth within the pretense (for it to be *fictionally true*), what must *really* be the case is that Ann's cleverness and modesty compare suitably to what is

²Speaking of *the* truth condition of an utterance is in general dangerous, since one needs to decide to hold fixed certain features of the utterance in order to ask what *more* is needed if it is to be a true utterance. To see this, notice that what's wrong with the following view is not *falsity*: every utterance has the same truth condition: that the utterance be a true one. It is only when certain aspects of the utterance are held fixed (perhaps including its syntax and some "lexical" features of the constituent expressions, contextual facts and speaker intentions), that a more specific truth condition can be isolated. But for the moment let's go with the flow.

³*Mimesis as Make-Believe* (Cambridge: Harvard University Press, 1990), and "Metaphor and Prop Oriented Make-Believe," *European Journal of Philosophy* 1 (1993): 39–56.

attributed to the characters of Conan Doyle's stories. That is to say, (2) specifies the condition of the fictional truth of (3). And this condition, I propose, is also the genuine truth condition of (3), the condition that must be met if my utterance is to be seriously true.

To lay this out a bit more systematically, I need to develop a distinction between two ways in which propositions can come to be fictionally true: propositions can be *expressly made-believe*, or their being fictionally true can be *generated from reality*.⁴ If we expressly make believe that a certain hill is Mount Olympus and that you and I are the gods, then those propositions are directly *forced* to be fictionally true; they are in a sense stipulative, foundational truths of the make-believe. But consider the proposition that a god soon will tumble down the mountain. In our context, this proposition might well be fictionally true, because you or I might well soon tumble; so, if you say 'a god soon will tumble down the mountain', you might turn out to have spoken fictionally truly. In this case, whether it is fictionally true that a god will tumble depends not only on what we are expressly making believe, but also on whether one of us in fact will soon tumble down the hill. Understanding our make-believe involves not only grasping what facts are expressly made-believe to obtain, but also grasping how further pretend-facts are generated from anticipated or unanticipated real facts. As Walton puts it, the *principles of generation* governing our make-believe determine how real facts and fictionally true propositions generate further fictionally true propositions. In this example, we might naturally and spontaneously adopt principles of generation according to which actual tumbles down the hill generate fictional truths about tumbles down the mountain. The important point here is that the fictional truth of some claims depends not only on the imposed parameters of the make-believe (which include the claims that are expressly made-believe and the principles of generation), but also on real-world facts that are in no sense constitutive of the make-believe. I think that statement (3) is such a claim, and that it is made precisely to express a serious

⁴Walton's distinction between directly and indirectly generated (or primary and implied), fictionally true propositions seems to cross-cut this one. See *Mimesis as Make-Believe*, 140–44.

commitment to the real world's being such as to make the statement fictionally true.

Here, then, is what I take to be going on in (3):

- There is a transient context of make-believe, in which certain propositions are expressly made-believe; for instance: that there are two persons to whom we can refer with the names 'Holmes' and 'Watson', that the former was a brilliant detective, and that the latter was his faithful sidekick.
- The fictional truth of certain other propositions is generated from reality. That is to say, these propositions are fictionally true in consequence of various real-world facts and of the parameters of the make-believe. Among the fictionally true propositions, depending on whether Ann indeed is sufficiently clever and modest, might be the proposition that Ann is as clever as the detective we can refer to as 'Holmes' and more modest than the physician we can refer to as 'Watson'.⁵
- It is fictionally true that I have expressed a proposition with my utterance of (3)—more particularly, it is fictionally true that there are people to whom we can refer as 'Holmes' and as 'Watson' such that I have asserted the proposition that Ann is as clever as the former and more modest than the latter. Note that since there really are no such people, there really is no such proposition; however, since it is fictionally true that there are such people, it is fictionally true that there is such a proposition.
- Given the parameters of the make-believe, for it to be fictionally true that the proposition I have expressed with (3) is true, is for (2) to be true. In short, (2) expresses the fictional-truth condition of my utterance.
- For my utterance to be genuinely, seriously true it is necessary and sufficient that it be fictionally true. Thus, (2) expresses the truth condition of my utterance.⁶

⁵Why mention these metalinguistic propositions rather than, say, 'the proposition that Ann is as clever as Holmes and more modest than Watson'? Since 'Holmes' and 'Watson' do not really refer to anything, I do not assume that there is any such proposition. Indeed I think that there is no such proposition, and so that there is no such proposition that can be even fictionally true. In explaining *what* is fictionally so, I want to identify propositions that are fictionally true, and so of course I restrict myself to all the propositions that there are.

⁶It is compatible with this account that there is a "fictionally operator"

This account of (3) involves a postulation of *semantic pretense*. This puts it in league with quite a few other semantic accounts, including for example certain theories of metaphor and of talk apparently about numbers. Accounts that postulate semantic pretense break in two the semantic explanation of a statement. First, we explain the semantic properties (such as the logical form, truth conditions, and modal content) of the sort of statement the speaker makes *as if* to make, and second, we explain how the facts about the pretense fix the genuine semantic properties of the serious statement.

The point of postulating semantic pretense is to answer the question, what is a sentence like that doing in a speech act like this? We start with strong opinions about the serious semantic properties of the speech act and strong opinions about the semantic properties for which the sentence is really suited, but these opinions clash: the sentence doesn't seem suited to expressing what it actually is used to express.

Part of postulating semantic pretense is to claim that the apparent structure of a sentence is misleading as to the claims made in uses of it. But we do not simply claim this; we explain the use of *that* sentence by its being structurally suited to perform a certain task *within a pretense*, and we explain how its performance of that task suits it for use in genuinely expressing what it otherwise has no business expressing. The aim is to provide satisfying accounts both of what we say and of how we can manage to say it *like that*.

Semantic pretense might find use in a language for any number of reasons. A key one is to let us express using ready, tidy linguistic resources claims that, perspicuously stated, would require cumbersome formulations or unfamiliar terminology.

Semantic pretense involves a special path from the semantic

governing the statement, but I don't see any point in postulating hidden syntax in this sort of case. I prefer to see this statement as employing a trope, akin to metaphor, characterized by a distinctive path connecting utterances to their truth conditions.

It perhaps is preferable to view this case not as involving a pretend (and also serious) *assertion*, but as involving two pretend (and also serious) *predications*—namely, of 'is as clever as Holmes' and of 'is more modest than Watson'. We then would inquire after the (fictional and serious) conditions of applicability, analogously to how we have sought truth conditions. There is no reason to restrict this sort of analysis to entire speech acts, nor even to utterances of whole sentences.

properties of words to the semantic properties of utterances. In that sense it is a kind of figure of speech, and I expect that some philosophers would prefer to say that in such cases the notions of genuine *saying* and of utterance *truth* are out of place. These philosophers would hold that statement (3) could be *true* only if the names ‘Holmes’ and ‘Watson’ really referred. But, for one thing, ordinary intuitions about truth and saying are on my side: I have *said that* Ann is very clever and modest, and, supposing she is, my utterance is true. Moreover, the behavior of the sentence (3) when embedded in larger sentences (for instance, in a denial or as the antecedent of a conditional) shows that we linguistically treat this use of the sentence as a genuinely expressive use. Certainly there is a legitimate notion of a statement’s *perspicuously* portraying a state of affairs, on which (3) does not perspicuously portray the state of affairs that I take it actually to portray. But it would be unfortunate to assume that every genuine linguistic portrayal of a state of affairs must be a perspicuous portrayal. Obviously, one can opt for a narrow concept of utterance truth by staunchly refusing to admit truth of all but perspicuous sayings. The only complaint against that concept of utterance truth is its irrelevance. The narrow concept of utterance truth does not in general capture the kind of utterance correctness that we normally use ‘true’ to mark, and it does not track the features of utterances that are central in their status as assertions, denials, and so on. I will adopt here the ordinary concept of utterance truth that allows imperspicuous sayings such as (3) to be true, and I hope that what follows will constitute a partial defense of the utility of this concept.

2. Back to Hesperus and Phosphorus

Let’s assume for now that Frege is correct in taking the names ‘Hesperus’ and ‘Phosphorus’ to somehow signal distinct ways of thinking of Venus. I am not asking you to assume that there are abstract entities suited to being called modes of presentation, nor to assume that ways of thinking have any particular connection to descriptions or features of the object thought about, nor that those who share a mode of presentation are in any intuitive sense cognitively similar, nor that a way of thinking is what a competent user of a name grasps. Indeed, all I ask you to assume with me is that we in fact use the names to effect *some* way of classifying actual and

possible thoughts about Venus, some as “Hesperus thoughts,” and others as “Phosphorus thoughts.” The classification may be obscure, artificial, vague, and certainly it may vary from context to context, but I will assume that in fact we do make such a classification. I hope that this is uncontroversial. I will speak of Hesperus and Phosphorus thoughts as involving “the Hesperus-mode (of presentation)” and “the Phosphorus-mode,” which unfortunately gives the impression that modes of presentation are assumed to be constituents of thoughts, but I want to be understood simply as meaning that the thoughts are on one side or the other of the distinction that we in fact use the names to signal. Given this, it is trivially true that the thoughts we classify as “Hesperus thoughts” all involve exactly the same mode of presentation—but of course it is open to us (and indeed it is plausible) to hold that this single mode of presentation encompasses very different cases. After placing more cards on the table, I will return briefly to questions about how classifications of thoughts by modes of presentation might be explained. But in this paper I am primarily concerned not with what distinctions we make among thoughts when we distinguish by ways of thinking, but rather with how the language that we use to report those distinctions manages to do that.

I am going to explore the idea that our talk about what people say or think often involves semantic pretense. As a way of sneaking up on that idea I propose first simply to plant a pretense in order to examine its fruit.

Let’s make believe that Hesperus and Phosphorus are two things, that thinking of a thing using the Hesperus-mode constitutes thinking of one of them, and that thinking of a thing using the Phosphorus-mode constitutes thinking of the other. Thus, we are making believe different things about the thoughts we in fact classify as “Hesperus thoughts” and as “Phosphorus thoughts.” This make-believe is best described by giving its principles of generation:

- (4) It is fictionally true that there are two things to which we can refer as “Hesperus” and “Phosphorus”; when and only when a thought involves the Hesperus-mode, it is fictionally true that the thought concerns the one thing, and when and only when a thought involves the Phosphorus-mode, it is fictionally true that the thought concerns the other thing.

Real facts about people thinking of Venus in the two different ways make it fictionally true that they are thinking of two different things. Thoughts are correctly classified within the pretense as being “about Hesperus” just when they in fact involve the Hesperus-mode, and they are correctly classified as “about Phosphorus” just when they in fact involve the Phosphorus-mode. In a sense, we are pretending-apart Venus, by pretending of two kinds of thoughts about it that they are thoughts about different things.

This leaves a great deal open: which features of Venus are we to make believe that the pretended two things possess—just how are we to pretend-apart all that we believe about Venus? Suppose for the moment that, apart from what I have already laid down, I don’t say; suppose that I leave our make-believe in this extremely fragmentary state. Remember, I am not suggesting guidelines for imaginative play here, but only establishing a shallow pretense for semantic purposes; we are instituting a manner of speaking, not painting a fantasy world. As I will clarify, we should think of the pretense as a license for representing distinctions about modes of presentation *as* distinctions among the objects presented.

Having imposed the shallow pretense, I want to show how to exploit it as a tool in saying things. I want to show that there are various claims about the real person Hammurabi that our pretense makes much easier to state, since, given what we are now making believe, we can simply report how Hammurabi is related to the *things* Hesperus and Phosphorus.

Let’s start with:

- (5) Hesperus is visible in the evening, and Hammurabi correctly attributed this feature to it.

There is a strong intuition about the semantic *structure* of this statement that might be expressed as follows: Hammurabi is said here to attribute a certain *feature* to a certain *thing*; his way of thinking of the thing is not brought into play. This is what the sentence *feels like*, semantically; it does not feel as though there are modes of presentation mentioned explicitly, nor even brought in by ellipsis or by tacit proviso. In the tradition of philosophical semantics, this would be deemed a paradigmatically *de re* belief sentence: designed for relating Hammurabi to Hesperus without entailing anything special about how he thinks of the thing. I call statements of that kind *notionally open*—they classify the agent’s thought by reference

only, and leave open the question how the agent thinks of the thing.

Accordingly, in the context of our make-believe, let's regard (5) as entirely notionally open: as expressing a claim simply about how Hammurabi is related to a particular object, namely the one we call "Hesperus," without mention of modes of presentation. But remember that, given what we are pretending, it is fictionally true that someone has a belief about the thing we call "Hesperus" exactly when they *really* have a belief involving the Hesperus-mode. So this statement is fictionally true just in case Hammurabi in reality had a belief attributing evening visibility to a thing, in which belief he employed the Hesperus-mode. I will adopt a semi-formal notation that lets me express such claims without horrible awkwardness. Where m_H is the mode of presentation in question, the fictional truth turns on the claim diagrammed in (6):

(6) Hammurabi believed: [m_H] is visible in the evening.

This bracket notation is a cognitive twist on the sort of concatenated use-and-mention quotation used in:

(7) Chip says that we are "politically correct."

In (6), we are using the bracket notation in a formula that, in describing a belief, partly describes what is allegedly believed to be so, and also describes the agent's alleged way of thinking of the subject matter of the belief. The formula (6) portrays a state of affairs that obtains just in case Hammurabi had a belief ascribing evening visibility, and this belief involved the mode of presentation m_H in the "subject-position." (It does not entail that Hammurabi's belief is *about* the mode of presentation.) At the cost of a more complex notation, this state of affairs can be diagrammed more perspicuously:⁷

(8) Hammurabi's belief: $\langle\langle$ evening-visible, $\overset{m_H}{\downarrow}$?? $\rangle\rangle$

The point I want to dwell on is that (6), which is the (nonfictional) condition for the fictional truth of (5) is what we might call "Fre-

⁷See my *Talk About Beliefs*, chapter 4. The need for complexity in a perspicuous notation is an important motivation for semantic pretense.

gean” or *de dicto*, or, to choose a term with less baggage, *notionally loaded*—it characterizes Hammurabi’s belief by mode of presentation. This shows that by situating ourselves within a shallow pretense, we can exploit a sentence that is *intrinsically* notionally open to characterize beliefs in a notionally loaded way, at least in the sense that the condition of the statement’s *fictional* truth is a notionally loaded condition. But I ask you to treat the utterance not merely as a pretend-statement, but also as a genuine, serious statement—one that, like our earlier example about Holmes and Watson, has as its genuine truth condition the condition of its fictional truth. Voilà: a notionally open sentence determines a notionally loaded truth condition. (More revealingly, a notionally open *truth condition* within the pretense generates a notionally loaded serious truth condition.)

This strategy pays off too in:

- (9) Hesperus, but not Phosphorus, was thought by Hammurabi to be visible in the evening.

Supposing that the statement, viewed from within the pretense, is notionally open (and once again that is how the sentence structurally *feels*),⁸ it is fictionally true just in case Hammurabi attributed evening visibility to a thing under the Hesperus-mode, but not under the Phosphorus-mode. The condition of fictional truth, then, and hence the genuine truth condition, is the proposition given in (10):

- (10) Hammurabi believed: [m_H] is visible in the evening, but
Hammurabi did not believe: [m_P] is visible in the evening.

Once again, our simple make-believe bends an intrinsically notionally open sentence to notionally loaded purposes.

Many philosophers have followed Frege in holding that sentences like

- (11) Hammurabi believed that Hesperus was brighter than Phosphorus

⁸The reader may feel, too, that this sentence would naturally be used in making a notionally loaded claim. I agree, and what I hope to make clear is just how this sentence can do that. But the first step is to notice how the intuition about what the sentence can be used to *claim* conflicts with an intuition about how it is semantically structured. I will show how to reconcile these intuitions.

are, in virtue of the semantic conventions of our language, equipped for ascribing modes of presentation—they are intrinsically notionally loaded sentences. Others have insisted that, despite appearances, such sentences express claims that are entirely notionally open. Now, our make-believe has set up a middle ground. Suppose we consider (11) to be intrinsically notionally open, and suppose I use it within our pretense, in order to make a serious statement. Then, just as before, the pretend-statement is notionally open, but the requirement for its fictional truth—and hence its serious truth condition—is notionally loaded. The statement is genuinely true just in case Hammurabi believed of a thing thought of under the Hesperus-mode that it was brighter than a thing thought of under the Phosphorus-mode—that is, just in case:

(12) Hammurabi believed: $[m_H]$ is brighter than $[m_P]$.

3. The Pretense Account

Maybe the manner in which our stipulated make-believe has enabled these notionally loaded claims is revealing of how ordinary notionally loaded attitude talk really works. Maybe, that is to say, attitude-ascribing sentences really are intrinsically notionally open, and yet standard utterances of them have notionally loaded truth conditions, owing to our standard use of something relevantly like the pretense I have suggested. This is the “pretense account” that I want to pursue. There will be obstacles. We need, for one thing, to address the issue of what *modal contents* are expressed by attitude reports (as contrasted with their truth conditions); and we need to confront a variety of other interesting issues and examples.

But the obstacle that must first be addressed is stark implausibility. For one thing, the pretense account about attitude reports may seem theoretically far-fetched—attitude reports, after all, form a large, important, and relentlessly scrutinized class of serious statements, and the pretense account has at the very least not often struck theorists as an obvious line of explanation. To be sure, there are many partly related ideas in the literature: Quine deems attitude reporting a “dramatic idiom,” several philosophers treat the talk as describing unreal notional or belief “worlds” or “theories,” and cognitive simulation theorists have suggested that reporting another’s attitudes is really a matter of voicing the attitudes one

takes on in simulating the other. Nonetheless, nothing very like the present hypothesis has gained currency.⁹

More importantly, it would surprise and confuse an ordinary speaker to suggest that in discussing Hammurabi's beliefs about Venus, she herself is pretending that Hesperus and Phosphorus are two things rather than one. She might insist that she knows how to pretend that, and that she is quite sure that that is not what she is doing. Normally, we use notionally loaded attitude talk without engaging in the kind of imaginative exercise that would be instigated by the instructions 'Let's pretend that Hesperus and Phosphorus are two things'.

But the pretense account does not require this kind of imaginative pretense. I have said that all it requires is a *shallow* (limited, provisional) pretense that really amounts to the institution of a manner of speaking. Let me clarify what I think the pretense account should be committed to. According to the pretense account, in distinguishing thoughts by mode of presentation we talk as if we are distinguishing among the things thought about—we use our linguistic tools that are designed for specifying and distinguishing among the things thought about. In particular, when we are distinguishing among multiple modes of presentation of a single thing we standardly talk as if we are referring to different things (for example, Hesperus and Phosphorus). It seems to me that this much should be granted plausibility by every ordinary speaker. According to the pretense account, the way this sort of talk accomplishes the task is illuminated by an account in terms of principles of generation. This requires there being a distinction between what is *as-if-so* and what is really so. What is *as-if-so* needn't be what we are licensed to imagine (as in more familiar cases of pretense); for

⁹See W. V. O. Quine, *Word and Object* (Cambridge: MIT, 1960), esp. 219; Daniel Dennett, "Beyond Belief," in *The Intentional Stance* (Cambridge: MIT, 1987); François Recanati, "Domains of Discourse," *Linguistics and Philosophy* 19 (1996): 445–75; Walter Edelberg, "Intentional Identity and the Attitudes," *Linguistics and Philosophy* 15 (1992): 561–96, and "A Perspectivalist Semantics for the Attitudes," *Noûs* 29 (1995): 316–42; Robert Gordon, "Folk Psychology as Simulation," *Mind and Language* 1 (1986): 158–71; Alan Leslie, "Some Implications of Pretense for Mechanisms Underlying the Child's Theory of Mind," in *Developing Theories of Mind*, ed. J. Astington, P. Harris, and D. Olson (Cambridge: Cambridge University Press, 1988), 19–46. See also Fred Landman, *Towards a Theory of Information* (Dordrecht: Foris, 1986).

the correctness of the pretense account, it need only be that what is as-if-so governs what is correct to say when we are speaking as if we are referring to different things. If we do quite standardly talk as if we are linguistically manipulating things when really we are manipulating modes of presentation, and if the pretense account gives a plausible story about how thing-language manages to do this work, then it matters not at all that we do not ordinarily think of ourselves as pretending that Hesperus and Phosphorus are two things.

So the objection from ordinary-speaker-bafflement can be met by distinguishing, in the class of pretendings, shallow speaking-as-if from imaginative play. Nothing depends on whether the word 'pretense' is appropriate to speaking-as-if. The important commitments of the pretense account are that in ordinary notionally loaded talk there is a distinction between what is as-if-so and what is so, and that the principles of generation proposed by the account are correct.

That the pretense account has not emerged as a contender among theorists (which anyway is not a great objection), can be explained by too much focus on the compositional model on which what a sentence is used to say is built up largely from the meanings of its components. The mechanism of semantic pretense (which surely is sufficiently systematic as not to raise special worries about how finite minds can grasp it) allows dramatic shifts from component-meanings to serious statement-content.

Thus, we should be careful not to overstate the degree of immediate implausibility we face here. There are also several considerations boosting the pretense account's plausibility. Principally, the account manages to answer nicely to key truth-conditional and phenomenological intuitions, as might best be seen in contrast with the Russellian and Fregean accounts of attitude statements.

Synthesizing Russellianism and Fregeanism

The Russellian account¹⁰ has it that all attitude reports are notionally open. To say that Hammurabi believes that Hesperus is bright-

¹⁰See for instance Nathan Salmon, *Frege's Puzzle* (Cambridge: MIT, 1986), and Scott Soames, "Direct Reference, Propositional Attitudes and Semantic Content," *Philosophical Topics* 15 (1987): 44–87.

er than Phosphorus is not to require anything whatever of the modes of presentation Hammurabi employs in the attributed belief. He simply must believe of the *thing* Hesperus (and how he thinks of that thing is irrelevant), that it is brighter than the *thing* Phosphorus (and, again, how he thinks of that thing is irrelevant). Famously, this has implausible consequences: if Hammurabi believes that Hesperus is brighter than Phosphorus, then he necessarily also believes that Phosphorus is brighter than Hesperus (and, even, that Hesperus is brighter than Hesperus). This has led to the stigmatization of Russellianism as a view that sacrifices all truth-conditional plausibility to the dubious idol Logical Purity. But attention to the bizarre inferences licensed by Russellianism obscures a deep source of plausibility for the account. When I say ‘Hammurabi believes that Hesperus is brighter than Phosphorus’, it *feels like* I am simply talking about Hammurabi, believing, Hesperus, brightness, and Phosphorus—it feels like I am simply portraying a state of affairs in which an individual (Hammurabi) believes that a thing (Hesperus) is brighter than a thing (Phosphorus). Davidson expresses a related point when he writes that if we could regain our pre-Fregean semantic innocence, we would find it incredible to think that names in that-clauses stand for anything different from what they stand for in other contexts.¹¹ That is surely right, but even more is true: when we focus on the phenomenology of attitude reporting, keeping at arm’s length the truth-conditional intuitions informed by worries about substitutivity, it can seem that Russellianism *has to be right*, for what it feels like semantically is that there are no references to modes of presentation, utterances, syntactic structures, or the like. Russellianism is truth-conditionally absurd, but it is the phenomenologically natural semantics for attitude sentences; we have to be led *away* from it by the substitution puzzles.

Frege is led by those puzzles to postulate an interesting—to my mind *too* interesting—hypothesis about our language: in the context of attitude sentences, he proposes, names stand for modes of presentation rather than for their bearers. The effect of this device is that I use the names in (11) to refer *not* to Hesperus and Phosphorus, but to the Hesperus-mode and the Phosphorus-mode. Put-

¹¹“On Saying That,” in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984).

ting aside qualms about whether modes of presentation really are things to be referred to, I believe that Frege provides an apparatus that issues in rather plausible judgments about the truth conditions and logical properties of a wide class of attitude reports. We do use attitude sentences to ascribe modes of presentation, and Frege's proposed device would explain this neatly, were it not for a major problem. As Davidson remarks, it seems frankly incredible that our language actually employs such a device: when we use names in attitude ascriptions it feels like we are using them *as names*, and not as names of modes of presentation. Fregeanism accommodates notionally loaded truth conditions, but it is phenomenologically absurd.

The pretense account we are exploring, it seems to me, keeps the good bits of both Russellianism and Fregeanism, while avoiding their worst problems.

According to the pretense account, there really are notionally loaded uses of attitude ascriptions. In using an attitude ascription to attribute a mode of presentation, the speaker adopts a pretense that allows her to talk about thoughts involving that mode of presentation as if she were talking about thoughts *about* a distinctive object. So she uses language that is designed for talking about objects rather than about modes of presentation.

On this account, it is no surprise that the phenomenology of using names in attitude sentences (just like the use of 'Holmes' in (3)) is that of using names to refer to things, since, within the pretense, this is exactly what the speaker is doing. Indeed, within the pretense, this is *all* that she is doing. The logical form of the pretend claim is entirely Russellian: the agent's belief is classified simply by what things it is about, and modes of presentation do not come into it. In one phenomenological respect the pretense account even has the advantage over Russellianism. Russellianism captures the feeling that the names 'Hesperus' and 'Phosphorus', in (11), are simply used to refer. But, and this is related to Russell's observation about identity, even when one knows that Hesperus is Phosphorus, in using (11) one feels as though one is referring to different things with those names. The phenomenology of (11) is precisely that of expressing a state of affairs relating Hammurabi to *two* things. It feels like we are distinguishing *Hesperus*, and Hammurabi's attitudes about *it*, from *Phosphorus*, and Hammurabi's attitudes about *it*. Sober-minded philosophers with robust senses of

reality notice this, realize that it cannot be taken seriously, and decide, too hastily, that it cannot be taken *at all*.

The pretense account allows essentially Fregean truth conditions for attitude ascriptions, and so it garners the plausibility that comes with that. But there is no need for the Fregean rule that names take on special sorts of referents in such contexts. As regards names, the pretense account gets by with the straightforward principle that they are devices for talking about their bearers. We reconcile these intuitions: that some attitude reports are notionally loaded, that the names in them function as names normally do, that the names are in some sense not treated as names of the same thing, and that the sentence wears its logical form on its sleeve—there is no tacit or explicit component of the sentence that stands for a mode of presentation.

De Re/De Dicto Unmarked

Another merit of the account is that it allows a *de re* / *de dicto* distinction that is not syntactically marked. Most theorists who have expressed a view on the matter believe that many attitude reports are notionally open, and of course it is compatible with the pretense account that sometimes attitude ascriptions are used independently of any pretense and that they are in these cases notionally open. For reasons that I will not pursue here, I am not convinced that there are any ordinary notionally open reports.¹² However, even if I were convinced that there are lots of them, I would remain unconvinced that the distinction between notionally open and notionally loaded reports is syntactically marked. For I very much doubt that a report in which names “take wide scope” with respect to the attitude verb must be notionally open. Consider again:

- (9) Hesperus, but not Phosphorus, was thought by Hammurabi to be visible in the evening.

Here, the names take wide scope (meaning roughly that they are not logically contained in the attitude verb phrase), and yet the overwhelmingly natural interpretation of this statement is the notionally loaded one. The idea that wide scope forces notional open-

¹²See my “Notional Specificity,” *Mind and Language* 10 (1995): 464–77.

ness, it seems to me, comes from assuming that the question whether a statement is notionally loaded must be settled by the logical form of the sentence being used. But the pretense account shows that this assumption is dispensable.

Empty Names No Worry

The pretense account accounts particularly smoothly for the function of nonreferring names in attitude reports such as:

- (13) Elijah believes that Santa is overworked.

Just as we talk as if there are distinctive objects corresponding to the Hesperus and Phosphorus concepts, we talk as if there is a distinctive object corresponding to the Santa-mode. Hence, the statement's truth requires that Elijah have an overwork-attributing belief involving that mode of presentation. We will see below how this dovetails with a nice account, due to Walton and Gareth Evans, of statements about existence and nonexistence.

Birds of a Feather

The pretense account suggests a tight connection between notionally loaded propositional attitude reports and certain other statements.¹³ Consider first:

- (14) Hammurabi was fond of Hesperus but not of Phosphorus.

If we employ the pretense account exactly as before, then for this statement to be true is for Hammurabi in reality to direct fondness toward a thing thought of under the Hesperus-mode but not under the Phosphorus-mode. Evidently, objectual attitudes like fondness, fearing, and remembering present no new challenge to the pretense account. Consider also:

- (15) It was Hesperus, not Phosphorus, that Hammurabi often saw in the evening sky.

¹³There recently has been a flurry of interest in apparent failures of substitution in other than the "classic" propositional attitude contexts. See for instance Anne Bezuidenhout, "Pragmatics and Singular Reference," *Mind and Language* 11 (1996): 133–59; Jennifer Saul, "Substitution and Simple Sentences," *Analysis* 57 (1997): 102–8; Graeme Forbes, "How Much Substitutivity?" *Analysis* 57 (1997): 109–13.

- (16) Hammurabi saw Hesperus more often than he saw Phosphorus.
- (17) Hammurabi's way was lit more often by Hesperus than by Phosphorus.

Each of these sentences can be used naturally by a speaker aware of the identity, in each case the speaker would naturally take the truth of her utterance *not* to allow switching the names, and yet in each case it seems very awkward to hold that there are special opaque senses (for example, of 'lit by') or that names are functioning other than as names of their bearers. It seems *prima facie* plausible that what is at work in these cases is at least approximately what is at work in notionally loaded attitude reports. And, independently, it seems *prima facie* very plausible that in these cases we are making as if to refer to different things with the names 'Hesperus' and 'Phosphorus'. This in itself is support for the pretense account: it groups like with like.

Here is how I would explain these statements. In these cases our make-believe about the modes of presentation helps to generate further fictional truths about the pretended two objects. To square our make-believe with the fact that sightings of Venus in different situations are connected with the Hesperus-mode and the Phosphorus-mode, we make believe that, in certain situations, it is Hesperus and not Phosphorus that is seen (indeed, it is Hesperus and not Phosphorus that is *there*), and in others vice versa. Our make-believe does not, of course, resolve all the questions one might have as to just how this is a coherent fiction. But for the purposes of the statements in question, it need not do so, so long as it is clear what to count, within the make-believe, as manifestations of Hesperus versus Phosphorus. And this is pretty clear: we count evening manifestations of Venus as manifestations of Hesperus, and morning manifestations as manifestations of Phosphorus. Thus, for instance, the truth condition of (16)—that is, the real-world requirement for its fictional truth—is perhaps that Hammurabi saw Venus in the evening more often than he saw it in the morning.

Suppose that Hammurabi awakens abruptly and is wrong about what time it is (and confused about what season it is). It might be correct (the pretense account predicts) to report, "Hammurabi believes that Hesperus is in the sky; but he is wrong—really it's

Phosphorus.” This indeed seems a natural thing to say, and the pretense account shows how it is entirely compatible with holding that Hammurabi’s belief is *really* true (its falsity being merely fictional).

No doubt, special features of the Hesperus/Phosphorus example account for there being a natural way to pretend-apart manifestations of Hesperus and of Phosphorus. There is no similarly natural way to pretend-apart manifestations of Cicero and of Tully (except in the sense of pretending-apart cases of *hearing about* Cicero and cases of *hearing about* Tully, as well, of course, as *believing about* Cicero versus Tully). Where such pretendings-apart are natural, the present account would predict, statements analogous to (15), (16), and (17) will be available.

The Usual Case

One respect in which the Hesperus/Phosphorus example is exceptional is that in it, distinctions about modes of presentation are undeniably live issues. Few deny that in the case of ‘Hammurabi believed that Hesperus is brighter than Phosphorus’, the speaker is at some communicative level or other distinguishing among (what might be called) modes of presentation. This is exceptional. In most ordinary attitude reports, modes of presentation seem to be the farthest things from our minds. If I say,

(18) Buchanan thinks that Gingrich is a communist

it is not nearly so obviously natural to regard me as talking about modes of presentation.

It would be compatible with the pretense account to hold that the statement (18) is notionally open (for the account is an account only of notionally loaded statements). However, I am convinced that the statement is in fact notionally loaded. If (18) were notionally open, then for its truth it would suffice that Buchanan recalled once seeing a man adrift in a pro-communist parade (who in fact was Gingrich, though Buchanan did not recognize him), and took him to be a communist. But surely that in fact does not suffice for the truth of my statement. For my statement to be true, Buchanan must attribute communism to Gingrich thinking of him, as we might put it, relevantly normally. Given the thin understanding of a mode of presentation that we have been employing, we

can as well say that Buchanan must employ the relevantly normal mode of presentation of Gingrich (it is unimportant whether that strikes you as just the right characterization of how Buchanan must think of Gingrich so long as you grant that the report is notionally loaded).

On the face of it, these two observations are in tension. Intuitively, modes of presentation are far from our minds in the statement, and yet what is claimed is notionally loaded. But there is a synthesis: the possibility of other modes of presentation of Gingrich is *truth-conditionally ignored* in the statement.

There are at least two senses in which a possibility might be said to be ignored in a statement. First, it might be *presuppositionally ignored*—it might be presupposed that the possibility is not actual; but this surely is not necessary here—nothing requires presupposing that Buchanan really has no irrelevant mode of presentation. The crucial sense in which a possibility can be ignored is that it can be truth-conditionally ignored—it can be tacitly settled that the possibility is not to be regarded as relevant to the truth of the statement. In a statement of

(19) There is milk in the refrigerator

the possibility that there is a dried drip of milk on the shelf might be ignored—it might be tacitly settled that this is not to be *counted* a case of there being milk in the refrigerator despite its really *being* a case of there being milk in the refrigerator. I believe that this is aptly regarded as a form of pretense, since it is a (tacit) determination not to count as a P what really is a P (the proper explanation of such cases¹⁴ will advert to principles of generation). Returning to our example, in context it is clear that only thoughts involving the relevantly normal mode of presentation *are to be counted thoughts about Gingrich*. More precisely, we are engaging in a make-believe with the following principle of generation:

(20) Thoughts involving the relevantly normal mode of presentation of Gingrich (and only those) generate fictional truths that one is thinking *about Gingrich*.

¹⁴Such cases are enormously common. See John Searle, "Literal Meaning," in *Expression and Meaning* (Cambridge: Cambridge University Press, 1979).

The parallel with the pretense operative in the Hesperus/Phosphorus case is this: in both cases thoughts involving a certain mode of presentation (and only such thoughts) generate fictional truths that the thoughts concern a distinctive object. In the Hesperus/Phosphorus case, this requires pretending Venus apart into two *merely* fictional objects; in the case of Gingrich, no pretending apart is required—Gingrich himself can serve within the pretense as the needed distinctive object. By ignoring other modes of presentation in this way, our talk simply of beliefs *about Gingrich* acquires notionally loaded truth conditions. Like the pretense about milk in (19), the pretense about thoughts about Gingrich normally doesn't draw notice until truth-conditional intuitions are canvassed. This resolves the observed tension and brings the usual case under the wing of the pretense account.

It is absolutely essential to the plausibility of this application of the pretense account that we keep in mind the shallow nature of the claimed pretense. The claim is simply that we use one categorization of thoughts (being about Gingrich) to effect a different categorization (involving the normal mode of presentation of Gingrich). We talk as if we are discussing any thoughts about Gingrich whatsoever in order really to discuss only thoughts that involve the normal mode of presentation of him. The proposed principle of generation embodies that idea.

If a pretense really plays this role in (18) then surely a similar pretense operates in just about every ordinary attitude report, for the puzzling tension we observed about that statement is rather the rule than the exception: modes of presentation are not often on our minds in attitude reporting, but it is usually counterintuitive to view reports as notionally open. In the unexceptional cases, some mode of presentation of the entity has an obvious claim to relevance; so it is plausible that, as in the Gingrich case, we quite standardly truth-conditionally ignore all but the obviously relevant mode of presentation of an entity.¹⁵ In the exceptional cases, such

¹⁵Given our thin conception of mode of presentation, remember, there is no distinction between specifying a single mode of presentation and specifying a type or class of modes of presentation. Considerations like those supporting the notional loadedness of the Gingrich case seem to me to support taking most ordinary reports to be notionally loaded as regards not only particulars (like Gingrich) but also universals (like being a communist), but I ignore this in the present discussion.

as the Hesperus/Phosphorus case and many others familiar in the philosophical literature, it is still more plausible to postulate pretenses involving modes of presentation, for modes of presentation are quite plausibly “in the air” in such statements, and the pretense account offers an account (which I hope by now has begun to seem reasonably plausible) of how they enter into the truth conditions of the statements.

Beginning at Home

If the pretense account is correct, then it seems that speaking about modes of presentation as if we were speaking about the objects presented is something that we do spontaneously and naturally. But why should that be? Let me distinguish this question from an easier one. It is not surprising to find that our natural ways of thinking about the thoughts of others involve pretense in one form or another. It is certainly plausible—and recent work in philosophy and cognitive psychology has only made it more so—that we often make sense of others by figuring out how to see things their way. So it would come as no surprise that attitude reporting involves pretense in some way. But the pretense account postulates a specific sort of pretense that is not a matter of pretending to be someone else, or to be in their shoes. Why should it be natural to make as if one is distinguishing among objects when one is really distinguishing modes of presentation?

There may be a deep explanation of the naturalness of this pretense based on features of normal first-person access to one’s attitudes. Consider how absurd it would be for Hammurabi, in thought, to self-ascribe attitudes in a way that makes the attributed modes of presentation explicit:

- (21) I believe that Hesperus, thought of under my Hesperus-mode, is brighter than Phosphorus, thought of under my Phosphorus-mode.
- (22) When I see Hesperus tonight, I will imagine holding Hesperus, thought of under my Hesperus-mode, in my hand.

The absurdity here is something beyond the mere awkwardness and terminological unfamiliarity of the third-person case. It has to do with the fact that the modes of presentation explicitly self-attributed in these cases are *already there*, in that they are *employed* in

the ascriptions to represent the individuals that the attitudes are about. Normally, when one has reason to consider a self-ascription of an attitude about a thing, in the context of ascription one thinks of the thing using the very mode of presentation that one has reason to ascribe to oneself.¹⁶ Thus, in the first-person case, the strategy of distinguishing among modes of presentation as if among objects presented is an utterly natural and elegant cognitive shorthand.

A similar consideration extends to the third-person case to whatever extent it is typical that when we represent to ourselves *which thing* another's attitude is about, we represent it *using* the very mode of presentation we ascribe to the other. Perhaps that indeed is typical, but the question may be controversial. But even if it is not typical, the familiarity of the pretense strategy from the first-person case might make it thoroughly natural in the third-person case as well (especially if third-person ascription in some way involves simulated or proxy first-person ascription). It is important to see, in any case, that the pretense hypothesis is not obviated by a simulation account of third-person ascription, since it is needed even for understanding first-person ascriptions.

4. Modes of Presentation

The pretense account employs the notion of a mode of presentation, and nothing could be more important in assessing the account than determining the prospects for satisfying explanations of what modes of presentation are, and of the access to them that speakers and hearers rely on in producing and understanding attitude reports.

I have little novel to say about these matters here.¹⁷ The use to which modes of presentation are put in the pretense account does require that modes of presentation group together actual and possible cases of thinking about (or thinking as if about) a thing. It is not required that they be natural kinds of mental representation,

¹⁶For an abnormal case in which this is not so, see my "I Falsely Believe that P," *Analysis* 52 (1992): 191.

¹⁷For useful recent work on these questions, see Jennifer Saul, "The Pragmatics of Attitude Ascription," forthcoming. See also my *Talk About Beliefs* (Cambridge: MIT, 1992), especially chapter 5, and "Notional Specificity."

nor that they be stably associated with natural language expressions, nor that to engage in discourse *about* them one need to *grasp* them.

For a given mode of presentation, it might be that knowing what it is (knowing what it is to employ it) is prior to one's capacity to participate in the kind of pretense that the pretense account describes. But, however natural that seems, it is not necessary. I am attracted to the view that our fundamental access to certain modes of presentation comes from our capacities to participate in pretenses of just this kind. In such cases, we in the first instance have the capacity to understand what it takes to be, fictionally, thinking *about* this or that object, and only derivatively on this do we have any conception of the modes of presentation—of the requirements for its being fictionally true that someone is thinking of these objects.¹⁸

5. Modal Content and Innocence

So far we have explored a proposal restricted to the question of truth conditions; we have not decided what the pretense account should say about the modal contents expressed by attitude reports.

In addition to having truth conditions, I will assume, utterances of sentences can express modal contents. My chief assumption about modal content is this: the modal content of an utterance answers the question, what must a possible situation be like for it to be accurately portrayed by this (actual) utterance?¹⁹ In many semantic theories, utterances of sentences are said to express propositions, or portray states of affairs; these semantic values (according to many such theories) determine the modal content of the utterance.

¹⁸This is connected to certain suggestions Evans makes in his discussion of existence statements in *Varieties of Reference* (New York: Oxford University Press, 1982). Relevant also is one way in which, according to Walton, metaphors can be *essential*; see “Metaphor and Prop Oriented Make-Believe.”

¹⁹A conception of content built around this assumption is familiar from Kripke's *Naming and Necessity* (Cambridge: Harvard University Press, 1981). The theoretical centrality of such a notion of content (but not its semantic reality) is questioned by Michael Dummett in appendix 3 of *The Interpretation of Frege's Philosophy* (Cambridge: Harvard University Press, 1981), and by Evans in “Reference and Contingency,” in *Collected Papers* (New York: Oxford University Press, 1985).

So conceived, the modal content of an utterance and the utterance's truth condition have to agree about the utterance's truth *value* (they have to "agree at the actual world"), but the truth condition of an utterance need not be equivalent to its modal content. For instance, an utterance of the sentence:

(23) This utterance does not exist

has an impossible truth condition (it cannot occur truly) but a contingently false modal content (the utterance correctly describes just those possible situations in which it does not exist). An utterance's truth condition and its modal content are the answers to different questions. The truth condition answers the question, what must be so if this is to be a true utterance? The modal content answers the question, what must a possible situation be like for it to be accurately described by this (actual) utterance? The answers to these questions can come apart, sometimes dramatically. Indeed, shared truth *value* is the only necessary link.²⁰ A semantic account of truth conditions, then, seriously underdetermines an account of modal content.

Now, it is open to us in developing the pretense account simply to identify the modal contents of attitude ascriptions with their truth conditions. However, this would be a mistake. To see why, let's return to the idea of semantic innocence—the idea that a name in an attitude report stands for just its ordinary referent.

The pretense account as we have developed it thus far seems two-faced about semantic innocence. The account entails that a name in an ascription behaves completely innocently *within* the pretense; we pretend to relate the agent simply to the bearer of the name. But even if the name names a real thing, the ascription's truth condition is not a proposition that contains the thing as a constituent. This may appear to be a problematic loss of semantic innocence. Consider the argument from (24) and (25) to (26):

²⁰Given any two true sentences p and q , whatever their truth conditions and modal contents, we can construct (using the 'Actually' operator of philosophical logic) a sentence whose truth condition is necessarily equivalent to that of p and yet whose modal content is necessarily equivalent to that of q , namely, 'Actually p , and if actually q then q '. And for any two false sentences p and q , the same feature is possessed by 'Actually p , or if not actually q then q '.

- (24) Hammurabi believed that Hesperus is visible in the evening.
- (25) Hesperus is a planet.
- (26) So, there is a real thing such that Hammurabi had a belief attributing evening visibility to it.

Surely this is a valid argument. One might worry that if the pretense account were true, the inference would be invalid, since the truth condition of (24) would be a proposition not about Hesperus but about a mode of presentation—

- (27) Hammurabi believed: [m_{ij}] is visible in the evening

—which has no *logical* connection to the thing Hesperus itself.

A first response to this worry is to observe that on the pretense account there indeed is an entailment here: so long as the premises are true, so is the conclusion. This is because the mode of presentation that our make-believe concerns is always a mode of presentation of the name's bearer, if it has one.²¹ Thus, if (24) is true, then Hammurabi's belief is about what our use of the name 'Hesperus' names, if anything. If (25) is true, then our use of the name 'Hesperus' really does name something, and so the conclusion (26) must be true.

But demonstrating a truth-conditional entailment may not be enough. There is an intuition that not only is there a truth-conditional entailment, there is also an entailment in modal content: the modal contents of the premises necessitate the modal content of the conclusion. Any possible situation correctly described by the premises is correctly described by the conclusion as well. There is a natural line of reasoning that might lead one to doubt the entailment in modal content. I will reject this line, but it will be useful to sketch it. The line starts with an argument that the modal content of (24) does not *by itself* entail that of (26), because the ar-

²¹This needs to be unpacked carefully if names are individuated in the ordinary way, such that different things can have the same name. We must understand the use of the name in the ascription as borrowing from one of one's personal traditions of univocally using the name to refer (one has different 'Aristotle' traditions for what one at least suspects might be different Aristotles). The mode of presentation in question must concern not just any bearer of the name, but what the name in fact names within that tradition. The source of this constraint will become clear later in the present section.

gument from just (24) to (26) is relevantly like the clearly invalid argument from (28) to (29):

- (28) Elijah believes that Santa is overworked.
- (29) So, there is a real thing such that Elijah has a belief attributing overwork to it.

So *if* there is an entailment in modal content from (24) and (25) to (26), the line goes, the premise (25) is indispensable. But if the modal content of (24) does not entail the modal content of (26) by itself, then it is absurd to suppose that it might do so when conjoined with the information *that Venus is a planet* (which is the modal content of (25)). For if (24) does not by itself entail (26), then it does not entail that Hammurabi has a belief about *Venus*. The modal content of (25) could be relevant to the entailment only if the modal content of (24) were itself about Venus; but if it were, surely premise (25) would not be required for the entailment. So, while there is a logical entailment, one might well doubt that there is an entailment in modal content.²²

But I cannot accept this line of reasoning. If there is not an entailment in modal content from (24) on its own to (26), it must be that there are possible situations in which Hammurabi believes that Hesperus is visible in the evening, but in which Hammurabi's belief is not about Venus. But how could that be? Might a belief correctly so described be about a different thing? Then we face the unattractive prospect that Hammurabi might believe, and believe *truly*, that Hesperus is visible in the evening, in a situation in which Hesperus is not visible in the evening. It seems hardly more palatable to admit a situation in which Hammurabi believes that Hesperus is visible in the evening, but in which his belief is about nothing whatever. This constitutes a strong case for the claim that in any situation in which Hammurabi believes that Hesperus is visible in the evening, his belief is about Venus—and so a strong

²²I should note the possibility of considering both premises to be framed from within the pretense. We can interpret (25) as meaning 'Hesperus is *really* a planet' (in analogy to 'Hesperus (really) exists'. As will be clearer after our discussion of existence and identity, this might give (25) the truth condition that the Hesperus-mode denotes a planet. If (25) is given this interpretation, there can be an entailment in modal content even if the modal content expressed by (24) has nothing essential to do with Venus.

case for the entailment in modal content from (24) on its own to (26).

Fortunately, the pretense account can admit an entailment in modal content. Because of the gap between truth condition and modal content, and because the Hesperus-mode actually presents Venus, it is consistent with the pretense account that (24) expresses a modal content true of just those possible worlds in which Hammurabi attributes evening-visibility to *Venus* using the Hesperus-mode (for if so, modal content and truth condition would agree as to truth value). Of course, it is one thing for a view to be coherent and another for it to be defensible.

One way to defend it would be to hold that the Hesperus-mode denotes Venus essentially, so that it follows from Hammurabi's having a belief employing it that the belief is about Venus. But I will set this possibility aside, since a far less contentious motivation is available.

We need to answer the question, what possible situations are truly described by the utterance of (24)? If this question is capable of an answer, it must be possible to view the utterance not simply as determining a real-world condition for its actual fictional truth, but also as determining a condition that must be met by a possible world if the utterance is to be fictionally true in relation to that world. But what can it mean for an utterance to be fictionally true in relation to another really possible world? I suggest that we need the notion of genuine possibilities generating fictional possibilities. This notion is needed anyway, since even ordinary pretenses like our gods on the hill game extend modally. It is fictionally true that I *might* have pushed a god and made him fall, and this fact about fictional possibility is generated by the real possibility that I should have pushed you and made you fall. Just as real-world facts about the Hesperus-mode make it fictionally true that there are facts "about Hesperus," so certain genuine possibilities make it fictionally true that there are certain sorts of possibilities "for Hesperus." The attitude report correctly describes just those possible worlds that generate fictional possibilities that, fictionally, are described truly by the utterance.

The key question regarding semantic innocence, we can now see, is: which real possibilities generate fictional possibilities in which there are attitudes "about Hesperus" and "about Phosphorus"? The argument from (24) to (26) fails to exhibit an entail-

ment in modal content if we answer this question in a particular way: by deciding that *all* possibilities in which there are attitudes involving the Hesperus-mode generate fictional possibilities involving attitudes “about Hesperus.” If this is the right answer, then the possible worlds (if there are any) in which Hammurabi uses the Hesperus-mode to think not about Venus but about Mercury generate fictional possibilities in which Hammurabi is thinking “about Hesperus.” But this clearly cannot be the right answer. Instead, the contours of our pretense seem normally to be such that only possibilities about attitudes that both involve the Hesperus-mode *and are about Venus* generate fictional possibilities concerning attitudes “about Hesperus.”²³ Plausibly, this is a reflection within the pretense of the rigidity of names.²⁴

If this is right, then (24) gives a true description only of possible situations in which Hammurabi’s belief is about Venus, and so there is an entailment from that statement to (26). There is no similar entailment from (28) to (29) because the principles of generation that generate fictional thoughts (and fictionally possible thoughts) “about Santa” are not similarly tied to any real individual. Indeed, plausibly *no* genuine possibilities in which the Santa-mode succeeds in denoting generate fictionally possible cases of thinking “about Santa” (another reflection within the pretense of the rigidity of names).

These considerations suggest a general semantic account of attitude reports. The modal contents of attitude reports (typically) concern the real objects, if any, that the ascribed attitude allegedly concerns, as well as the required modes of presentation. For instance, we can express the modal content of ‘Hammurabi believes that Hesperus is brighter than Phosphorus’ as follows:

- (30) Hammurabi believed: Venus[m_H] is brighter than Venus[m_P].

²³If ‘the Hesperus-mode’ owes its nature to the contours of this pretense, then it seems to follow that it is necessary that thoughts involving it concern Venus.

²⁴Certain examples suggest that this principle of generation is defeasible: ‘Had Hesperus not been Phosphorus, I would have believed that Hesperus was not Phosphorus’. Related examples are considered by Michael Dummett in *Frege: Philosophy of Language* (Cambridge: Harvard University Press, 1973), 113; Landman in *Towards a Theory of Information*, and Bezuidenhout in “Pragmatics and Singular Reference.”

That is, Hammurabi believed that Venus (thought of under the Hesperus-mode) is brighter than Venus (thought of under the Phosphorus-mode). This seems to me just the right account of modal content. It is not an unfamiliar account: Stephen Schiffer's "hidden indexical theory," John Perry's and my "unarticulated constituents" account, Mark Richard's "Russellian Annotated Matrix" account, and François Recanati's "quasi-singular proposition" account offer roughly the same story.²⁵ The key problem for this semantic account has been that despite its plausible verdicts concerning which statements are true about which possible situations, the various developments of it have seemed ad hoc and false to the phenomenology of ordinary speech. The pretense account offers a deep and phenomenologically plausible explanation for the semantics. Perhaps surprisingly, what grounds the semantics is not any technical riff from philosophical logic, but an understanding of indirectly serious discourse that has been made possible by Walton's work in aesthetics.

6. Existence and Identity

Now the promised return to identity. To set the stage, I will make two observations: one about a curious pattern in our talk about identity, the other about a prima facie difficulty for the pretense account of attitude reporting.

Here, like it or not, is how we (and not only unjustly rebuked undergraduate metaphysicians!) talk about identity:

- (31) When *two things* are identical, the one thing has the same properties as *the other thing*.

As for the prima facie problem for the pretense account, consider:

- (32) Galileo knew that Hesperus is Phosphorus.

According to the pretense account, it is fictionally true that the names 'Hesperus' and 'Phosphorus' name two objects, and so the

²⁵Schiffer, "Naming and Knowing," in *Midwest Studies in Philosophy* 2, ed. P. French, T. Uehling, and H. Wettstein (Minneapolis: University of Minnesota Press, 1977); Crimmins and Perry, "The Prince and the Phone Booth," *Journal of Philosophy* 86 (1989): 685–711; Crimmins, *Talk About Beliefs*; Richard, *Propositional Attitudes* (Cambridge: Cambridge University Press, 1990); Recanati, *Direct Reference* (Oxford: Blackwell, 1993).

speaker here fictionally claims that Galileo knew of *two* things that they are numerically identical. But no two things are identical, and so surely the statement is fictionally false, *whatever* Galileo believed. This would be a devastating problem for the pretense account, since surely (32) is relevantly like the other attitude reports we have considered, and in fact it is true.

In addressing these issues, I will draw on an account due to Walton and Evans of statements attributing existence and nonexistence. The details of their presentations differ, but both hold that in saying

(33) Santa does not exist

one is pretending (or at least alluding to a pretense) that one can refer to a thing with a certain kind of use of the name 'Santa', in order to *disavow* such uses; one's statement is true just in case such uses do not refer. One very attractive feature of these accounts, as Walton and Evans emphasize, is that they capture the semantic phenomenology of such uses of empty names: they *feel* just like ordinary uses of referring names. This is respected by taking the use of the name, within the pretense, to *be* an ordinary use of a referring name.

Walton suggests, somewhat tentatively, that in talking about what does and does not exist we might be engaging in a pretense to the effect that 'exists' expresses a discriminating property—a property that not everything has. This makes marvelous sense of a lot of puzzling talk about existence:

(34) I've discovered that there are some things which don't really exist but which I had been duped into thinking exist. Among them are Santa, the tooth fairy, and Barney.

This talk makes perfect sense if we view it as framed within a pretense that there is a thing for every mode of presentation, except that some of these things do not have a certain property—one that we can refer to (within the pretense) as 'existence'. The things of which it is fictionally true that they have this property are exactly *everything*—everything in reality, that is. We pretend that there are more things than just those, but that everything else doesn't have this property. More precisely, while it is fictionally true of all modes of presentation that they denote things, it is fictionally true that a mode of presentation denotes a thing *that has the property expressed*

by 'exists' just when that mode of presentation actually denotes a thing. When the speaker says 'There are some things that do not exist', I would say, she is relying both on the genuine universal property of existence (which informs 'There are') and the pretended discriminating property (which informs 'do not exist'). Within the pretense, this might be described as the distinction between *being* and *really existing*. Of course, no such ontological distinction is seriously employed by the account, nor does the account attribute such an ontological distinction to us—our talk really concerns only real things and real modes of presentation.

The considerations about modal content in the previous section may suggest a contribution to this account. So far, I have sketched only an account of the truth conditions of statements about existence, as contrasted with their modal content. Evans does not discuss the modal content of these statements. Walton suggests that in a nonexistence statement "what one asserts is simply that to attempt to refer in a certain way is to fail."²⁶ But this is problematic. Consider: it does not seem that possible worlds in which the Santa-mode denotes someone normally are correctly described as worlds in which Santa exists. And it does not seem right that possible worlds in which Venus does not exist but in which the Hesperus-mode denotes something other than Venus are normally correctly described as worlds in which Hesperus exists. Further, it does seem right to describe a world in which Venus exists as one in which Hesperus exists, even if no attempts to refer happen in it or in any nearby worlds. Probably these observations do not constitute a knock-down argument. But it seems to me that our attention is misdirected if, in assessing whether a statement like (33) correctly describes a possible situation, we focus on (potential) attempted acts of reference *in that situation*. It seems to me that the critical question must be what an attempted reference *from here* picks out *there*. More precisely, we need to ask which real possibilities make it fictionally true that there are possibilities in which "Santa exists." I believe that the following ideas are defensible: normally, (a) no possible world makes it fictionally true that there is a possible world in which "Santa exists" (so that (33) normally expresses a neces-

²⁶ *Mimesis as Make-Believe*, 426. He also considers the possibility that one seriously asserts that it is fictionally true that one speaks truly in making the statement.

sary content despite having a contingent truth condition), and (b) all and only possible worlds in which Venus exists make it fictionally true that there are possible worlds in which “Hesperus exists.”²⁷ It seems to me that this suggestion has the merit of construing the modal content of existence claims as ontological rather than meta-representational: they are really about what things there are, and not about our referential access to things.

In the case of identity, the observation about our odd identity-talk and the difficulty for the pretense account about (32) cry out for a parallel explanation. In talking about identity we standardly talk as if we think that things can be identical to *other* things. We talk as if identity, which in fact is as celibate as a relation can be, is promiscuous. More precisely, we pretend that with certain of our linguistic devices that normally express identity, we can express a relation that can hold between distinct objects. It is fictionally true that this relation holds between two objects when these fictional objects result from pretending-apart a single object (as in the case of Hesperus and Phosphorus).²⁸ In our pretense, for instance, Hesperus and Phosphorus are *two* objects that fictionally bear to each other the pretended promiscuous identity-relation: the one thing is, fictionally, the “same as” the other thing.

One pay-off for the pretense account about attitude ascription is that statements like (32) are accommodated neatly. Within the pretense, the speaker of (32) is claiming that Galileo knew of two things that they bear the promiscuous identity-relation to each other. And it is fictionally true that the two things named by ‘Hesperus’ and ‘Phosphorus’ *do* bear that relation to each other. The real-world requirement for the fictional truth of (32) is the (true) claim (35):

(35) Galileo knowledgeably believed: [m_H] is [m_P].

(This proposition entails that the two modes of presentation in fact co-denote.) And, using our general apparatus, we see that the serious modal content of the statement is equivalent to:

²⁷However, as noted earlier, such constraints may be defeasible, offering resources to explain such sentences as ‘Had Santa existed, I would have gotten more toys’ and ‘Had Hesperus not been Phosphorus . . .’.

²⁸We might allow that there are other ways for it to become fictionally true that this relation holds, to allow, for instance, for the fictional truth of ‘Zeus is Jupiter’.

(36) Galileo knowledgeably believed: Venus[m_H] is Venus[m_P].

What of identity claims themselves?

Frege's account, remember, allows even informative and true identity claims simply to attribute the relation of strict identity between the referents of referring expressions. The *distinction* between the two sides of the equation that explains the substantial nature of the statements is simply the difference in the modes of presentation tied to the terms. This seems to me a satisfying explanation both of the truth conditions of these claims and of their substantiality. Frege did not address the question of modal content, but I believe that there is also a satisfying account of modal content open to him, namely Kripke's: identity claims (at least those in which the singular terms are referring devices and indeed refer) have necessary or impossible contents, depending on whether they are in fact true. That Frege's account of truth conditions coheres with Kripke's account of modal content is another manifestation of the gap between truth conditions and modal content.

I am happy to allow that Frege's explanation of identity statements is often exactly right. But it does not offer a fully satisfying explanation of certain identity statements, in particular, those of which Russell's phenomenological confession rings most true: it feels as though we are saying of one thing that it is the same thing as *another thing*. In these cases, the pretense account offers a promising alternative picture. Consider, as framed within our running pretense, the statement:

(37) Hesperus is Phosphorus.

Suppose that for this statement to be true is for it to be fictionally true. For it to be fictionally true is for it to be fictionally true that the two objects denoted by these names stand in the pretended promiscuous identity relation. For that to be fictionally true is for it (really) to be the case that the object denoted by the Hesperus-mode is the very object denoted by the Phosphorus-mode. Thus, the truth conditions offered by the pretense account are entirely Fregean. (Actually, that is only half true, since identity statements with nonreferring singular terms are deemed by the Fregean account not to possess truth values, whereas sometimes they are deemed *false* on the pretense account.²⁹ This offers accurate pre-

²⁹This too is only approximately true, depending on how the contours

dictions; consider: ‘He is Santa Claus’.) Furthermore, the Kripkean account of modal content open to Frege is open as well to the pretense account. The extra apparatus of the pretense account retains approximately Fregean truth conditions and Kripkean modal content, but it supplies a satisfying explanation of Russell’s phenomenological observation—one, moreover, that coheres with an enhanced understanding of how we ascribe modes of presentation in reporting thought and talk.

7. Appendix

We will run through a number of examples and issues to exercise the pretense account.

7.1 *Speakers Mistaken About Identities*

Here is a potential problem. Suppose the speaker of

- (11) Hammurabi believed that Hesperus was brighter than Phosphorus

herself believes that the two names name different objects. Surely, one might object, she is not *making believe* that the names name different objects, and yet the report means just what it means in our mouths—so the pretense account does not seem to apply correctly.³⁰

Recall our distinction between what is expressly made believe and what is fictionally true. Fictional truths that are not themselves expressly made believe are generated from what is made believe with help from reality. In the statement

- (3) Ann is as clever as Holmes and more modest than Watson

we do not *make believe* that Ann really is so related to individuals bearing those names. Whether this statement is fictionally true depends on Ann’s actual cleverness and modesty. It may in fact be fictionally true, and we may believe that it is fictionally true, but

of the pretended promiscuous “identity” relation are to be spelled out. In particular, we might well allow ‘Santa is Santa’ and ‘Zeus is Jupiter’ to be counted true—in the pretense, “identity” extends to “things that don’t exist.”

³⁰Steve Yablo first called such examples to my attention.

even if so it is not something we are expressly making believe to be true, since it could be fictionally false consistently with everything we are making believe. While it may be only in odd cases that someone is mistaken about what she is expressly making believe, it is quite common to be mistaken about what is fictionally true; we might well be mistaken about whether (3) is fictionally true.

The pretense account requires that for each of the two modes of presentation ascribed in (11), real facts about attitudes involving that mode of presentation generate fictional truths (fictionally) about a distinctive object. It is not entailed by this that the speaker must make believe that there are two objects, for she might assume that the two required objects are supplied by reality. Recall that we took the speaker of

(18) Buchanan thinks that Gingrich is a communist

to be making believe of a certain mode of presentation that it is *the way* to think of Gingrich. For the same reasons, it is plausible to take the speaker of (11) to be making believe, of the two modes of presentation, that each is the only way to think of the thing thus thought of. Since the modes of presentation in fact denote just one object, in fact her make-believe generates a certain *merely* fictional truth—a claim that is fictionally true and not really true—namely, the claim that the modes of presentation denote two objects. The speaker of (11) knows that it is *fictionally true* that the modes of presentation denote two objects. But since she is unaware that Hesperus is Phosphorus, she thinks that the *reason* that it is fictionally true that there are two objects so denoted is that it is *true* that there are two objects so denoted. She thinks something is genuinely true (and only derivatively fictionally true) when really it is merely fictionally true. She rightly assumes that it is correct to talk as if there are two objects, but she is mistaken about why it is correct to talk that way. The response to the objection, then, is that, true, the claim that there are two objects rather than one is not something she is making believe; nonetheless, this is something that is fictionally true for her, and that is all that is entailed by the pretense account. This is a case of *resolute fictional truth*, in which a believed proposition counts as fictionally true even if the belief is false.³¹

³¹In a reenactment of Waterloo, one might believe that it is in fact Napoleon's sword in one's hand, and yet the fictional truth of this claim would not depend on its truth.

7.2 *Intentional Identity*

Consider Geach's example of what he calls *intentional identity*:³²

- (38) Hob thinks a witch has blighted Bob's mare, and Nob wonders whether she killed Cob's sow.

Let us assume that the speaker and hearers do not believe in witches, and assume that no person has harmed the mare and sow.

This sentence might admit of importantly different readings. Perhaps on one reading (which is promoted if one inserts 'in addition' after 'Nob'), it is asserted that Hob believes the claim that exactly one witch blighted Bob's mare, Nob is said to believe that too, and to wonder whether it is true that there is just one witch who blighted Bob's mare, who also killed Cob's sow. Let us call this the in-addition reading. It can be explained by taking there to be a tacit claim that Nob too believes that a witch has blighted Bob's mare, and by taking 'she' to be, in Geach's terminology, a pronoun of laziness for the description 'the witch who has blighted Bob's mare' (and the description can be treated in Russell's way, with narrow scope):

- (39) Hob thinks that just one witch has blighted Bob's mare, Nob thinks that just one witch has blighted Bob's mare, and Nob wonders whether there is someone who is the one witch who has blighted Bob's mare and who killed Cob's sow.

There is also a reading that involves something more aptly called intentional identity (and it is this reading that Geach has in mind). This reading is promoted if one replaces 'a witch' with 'a certain witch', and replaces 'she' with 'that witch'. It is difficult to explain this reading without portraying the statement as entailing that *there is a witch*. For now let us yield to this difficulty, and paraphrase the reading in question as:

³²"Intentional Identity" and "The Perils of Pauline," in *Logic Matters* (Berkeley: University of California Press, 1972). In thinking about intentional identity I have profited from Edelberg's "A New Puzzle about Intentional Identity," *Journal of Philosophical Logic* 15 (1986): 1–25, "Intentional Identity and the Attitudes," and "A Perspectivalist Semantics for the Attitudes," as well as Michael McKinsey's "Mental Anaphora," *Synthese* 66 (1986): 159–75.

- (40) A certain witch is such that Hob thinks that she has blighted Bob's mare, and Nob wonders whether she killed Cob's sow.

This reading tends to be dismissed by the sober; after all, it is stipulated that the speaker does not believe in witches. But that may be to take the statement too seriously. In a pretense in which a mode of presentation generates fictional facts about a distinctive object, it might well be that attitudes involving a certain mode of presentation make it fictionally true that there are attitudes about a certain witch (a witch who fictionally does not "exist," supposing the mode of presentation fails to refer). Let us call such a mode of presentation a "witch-generating" one. On the pretense account, the statement is fictionally true just if it is fictionally true that there is a witch, about whom Hob and Nob have such-and-such beliefs. Hence, the real-world requirement for fictional truth, and also the truth condition of the serious statement, would be this:

- (41) There is a witch-generating mode of presentation *m* such that Hob believes: [*m*] has blighted Bob's mare, and Nob wonders whether: [*m*] killed Cob's sow.³³

Let us call this the fictional-object reading of (38). Notice that it does not require that Nob share Hob's belief. This is appropriate, for (as Geach points out) statement (38) might be correctly used in the following sort of case: the community is agreed that there is a witch that has been damaging livestock; Hob says 'That witch has blighted Bob's mare' and Nob says (elsewhere at the same time) 'That witch may have killed Cob's sow'.

It is possible, too, that there is a reading that combines the in-addition reading with the fictional-object reading, in the obvious way (adjust (41) by requiring Nob to have the same belief that Hob is said to have).

A related phenomenon also is nicely described by Geach:³⁴

Although I see no reason to doubt that:

³³This need not entail, as we would put it colloquially, that Nob assumes that the witch exists; it does plausibly entail that he does not assume that she does not exist. This seems right: we can unproblematically append to (38) 'but is unsure whether she even exists'.

³⁴"Intentional Identity," 151 (sentence labels altered).

(42) The witch who blighted Bob's mare killed Cob's sow
is analyzable as:

(43) Just one witch blighted Bob's mare and she killed Cob's sow
it seems doubtful whether these two are mutually replaceable *salva
veritate* in a context like "Nob wonders whether". If we prefix 'Nob
wonders whether' to (42), the result seems to be analyzable, not as:

(44) Nob wonders whether (the following is the case): just one
witch blighted Bob's mare and she killed Cob's sow

but rather in some such way as this:

(45) Nob assumes that just one witch blighted Bob's mare, and Nob
wonders whether she (that same witch) killed Cob's sow.³⁵

The pretense account might offer an explanation as follows. The
use in

(46) Nob wonders whether the witch who blighted Bob's mare
killed Cob's sow

of the description 'the witch who blighted Bob's mare' is, pretend-
edly, a wide-scoped description—a specification of *the thing* Nob's
wonder is said to be about—which means that really it serves to
specify the mode of presentation that Nob's wonder is said to in-
volve. This allows, correctly, that (46) need not imply that Nob
really assumes that there is such a witch (the ascribed wonder
seems consistent with serious doubts on that issue). It explains, too,
why it is admissible to add 'of course, that witch doesn't exist'. And
it explains why Geach's strategy of paraphrase will not work for:

(47) Nob wonders whether the witch who blighted Bob's mare
exists.

Another very perplexing puzzle about intentional identity is due
to Walter Edelberg.³⁶ On natural readings, neither of (48) and (49)
entails the other:

(48) Arsky thinks someone murdered Smith, and Barsky thinks
he murdered Jones.

³⁵As Geach notes, the example is not completely compelling. A better
case (McKinsey's) is: Nob wishes that he caught the fish that got away
(where the speaker doubts there really is such a fish).

³⁶"A New Puzzle about Intentional Identity."

- (49) Barsky thinks someone murdered Jones, and Arsky thinks he murdered Smith.

Now, the in-addition readings of these statements do not entail one another. So we are home free if the in-addition readings are in question. But the most puzzling version of this case defeats those readings. Suppose the circumstances known to the speaker and hearer are as follows: No one really has murdered or even has wounded Smith or Jones, but both have suffered serious wounds. Barsky and Arsky are aware of these wounds. Barsky has no suspects, but he believes that just one person wounded both men, and he believes that Jones, but not Smith, has died from his wounds. Arsky has no suspects, he does not believe that just one person wounded both men, and he believes that Smith, but not Jones, has died from his wounds. In these circumstances, Edelberg holds, there are natural readings on which (48) seems true and (49) seems false.

I am not sure he is right about this; our ability to interpret these statements in the described circumstances may involve some non-trivial charitable creative reinterpretation. But let's assume that there are the needed readings, and explore what they might be. The circumstances block in-addition readings, because Barsky does not believe that Smith has been murdered and Arsky does not believe that Jones has been murdered. Fictional-object readings may be possible, however. Here is a try at expressing such readings:

- (50) The one Arsky thinks murdered Smith, Barsky thinks murdered Jones.
(51) The one Barsky thinks murdered Jones, Arsky thinks murdered Smith.

Now, these seem to me on the right track: to whatever extent (48) and (49) have the readings that Edelberg needs, so do (50) and (51). One apparent difficulty with these readings is that these sentences seem to entail that Arsky and Barsky have particular suspects in the attacks. But really all that is required by the fictional object treatment is that they have modes of presentation of the attackers (as in the circumstances plausibly they do) that fictionally generate objects of thought (as again is plausible). If you can get yourself in the frame of mind wherein (50) sounds true of the described situation, ask yourself: *who* is it that Arsky thinks murdered Smith

and that Barsky thinks murdered Jones. I think the answer will be: *Smith's attacker*. Plausibly, then, it is fictionally true that there is an object denotable as 'Smith's attacker', who makes (50) true. Now ask yourself, in that same frame of mind, but now focusing on the false-seeming (51), who is it that Barsky thinks murdered Jones (and that Arsky doesn't think murdered Smith)? I think the answer will be: *Jones's attacker*. What plausibly is going on, then, is this: fictionally there are two individuals, denotable as 'Smith's attacker' and 'Jones's attacker'; it is fictionally true that Barsky believes both that these individuals are identical and that since Jones's attacker murdered Jones, so did Smith's attacker; nonetheless, when we say 'the one whom Barsky thinks murdered Jones', it is fictionally Jones's attacker rather than Smith's attacker that is denoted—perhaps because, while he thinks that Smith's attacker is guilty of the same murder, it is only guilt by association ("identity"!) with the primary suspect (Jones's attacker).

7.3 *Embedded Ascriptions*

Consider:

(52) Sarah believes that Ray suspects that Laurie is pregnant.

Applying the pretense account, we face a new issue: what does it take for it to be fictionally true that a given object of thought (in this case, Laurie) is being ascribed by one person (Sarah) as the object of another person's (Ray's) attitudes? The pretense account has it that for Ray's suspicion fictionally to be about Laurie, it must involve a certain mode of presentation, call it m_L . For Sarah's belief fictionally to attribute to Ray a suspicion about Laurie, then, should require Sarah really to have a belief whose truth requires Ray to think of Laurie using m_L .

Now, here we face a rather abstruse complication. For the truth of Sarah's belief to turn on a fact about m_L (namely the question whether Ray has a suspicion involving it) does not rule out that she has a very odd, irrelevant way of thinking *about* m_L itself—perhaps one she would be likely to acquire only in a philosophy class devoted to the alleged Fregean hierarchy problems. But surely that possibility *is* ruled out by (52)—the statement, then, is *second-order notionally loaded*; it entails that Sarah has a certain kind of grasp on the way she takes Ray to think of Laurie. To capture this,

it is plausible to hold that Sarah counts fictionally as ascribing an attitude *about Laurie* just if she really has the proper grasp on the mode of presentation required for an attitude fictionally to be about Laurie. The truth condition of (52) might be represented in this way:

(53) Sarah believes: [m_R] suspects: [[m_{mL}]] is pregnant.

The double brackets, I stipulate (and it really is an additional stipulation about the notation), indicate that the mode of presentation specified within is, according to this proposition, Sarah's second-order mode of presentation—that is, her mode of presentation of the mode of presentation of Laurie that she attributes to Ray. This likely seems complicated, but in my view the “output” is just what we should hope for: (52) is treated as a notionally loaded ascription of a notionally loaded ascription of an attitude.³⁷

7.4 Anaphora and De Se Reports

Consider:

(54) When Fred came in, Doris doubted that he was Fred.

This notionally loaded ascription seems to present the following difficulty for the pretense account: If the account applies to it, then in the statement it must be fictionally true that the uses of ‘he’ and of ‘Fred’ in the complement clause refer to two objects. However, the use of ‘he’ is anaphoric on the earlier use of ‘Fred’, so surely it is not fictionally true that the use of ‘he’ refers to anything but what the use of ‘Fred’ refers to.

The analogous example (55) is if anything even more troublesome:

(55) Sometimes, when I see an important philosopher, I don't realize that she *is* that philosopher [or: that she *is she*].

The pretense account would predict that by the time we get to the complement clause, it is fictionally true that ‘she’ and ‘that philos-

³⁷I continue to find persuasive the argument in *Talk About Beliefs* (193–94) for the claim that modes of presentation of modes of presentation are ascribed in embedded reports. The hierarchy has no end, but this is not a *problem*: thought is nowhere bare.

opher' (or the second '*she*') refer to two things (and no doubt that these things are 'identical'). But *each* of these two expressions seems to be anaphoric on 'an important philosopher'.

One response is that while anaphora requires co-reference, co-reference in the context of the sort of pretense we are exploring is just a bit liberal: expressions co-refer only when they refer to things that are identical, but remember that it can be fictionally true that *two* things are identical. The response, then, is that the anaphoric connections in these examples are consistent with its being fictionally true that the expressions in the complement clauses refer to two things. In each case, that is, it is clear that the speaker is pretending that there are two "numerically identical" things, and it is clear to which of them the various names and pronominal phrases refer.³⁸

A different response to the problem about (54) would be to treat some of the pronominal phrases here as "lazy for" descriptions (such as 'the person who came in'). In spelling out this response, there is a decision to be made about the scope of these descriptions. If they take narrow scope with respect to the attitude verbs (if they are descriptions that the agents allegedly would use), then,

³⁸A possible reason for discomfort with this response is that it seems to predict that (56) would be an apt surrogate for (54), when really it is at best awkward:

(56) ? When Fred came in, Doris doubted that he was the same person as Fred.

In contrast, (57) does seem an apt surrogate:

(57) When someone who in fact was Fred came in, Doris doubted that he was the same person as Fred.

This is *prima facie* evidence that the response is not on the right track, since if (56) did involve the predicted sort of pretense, as (57) clearly does, surely it would be no more awkward than (57).

On the other hand, perhaps (56) is not really so very awkward, and perhaps the awkwardness it has can be explained consistently with taking it to be semantically equivalent to (54). For instance, it might be that the crucial difference between (54) and (56) is a garden-path phenomenon—that it owes to a cognitive effect of the increased "parsing distance" between 'he' and 'Fred'. The hearer discovers that she must, so to speak, pretend-apart Fred only when she encounters the second occurrence of 'Fred'. In (56), perhaps, this realization is more surprising, since it comes only after the hearer has been lulled into being unprepared for it. So it seems that this response can be defended.

for instance, the claim made in (54) is just the claim that would be made in (58):

- (58) When Fred came in, Doris doubted that the person who came in was Fred.

But this does not require, what should be required, that Doris be said to have a doubt that a certain *person* is Fred. If the descriptions take wide scope (if they are descriptions that the speakers are using to describe things), then the pretense account enters into the response more centrally: in statement (54), that is to say, it would be fictionally true that the speaker's tacit use of 'the person who came in' denotes a thing thinkable only in a certain way. Statement (54) would express, within the pretense, what is expressed in (59):

- (59) When Fred came in, the person who came in was such that Doris doubted that he was Fred.

Similar considerations apply to "*de se*" cases. Suppose that Lingens has amnesia and is lost in a library,³⁹ and that you are reading on a computer screen a report of his progress. The report presents evidence that the lost man is, unbeknownst to himself, Lingens. You say:

- (60) That man doesn't think that he is Lingens.
(61) I think that he is Lingens.

The explanation of (60) is that you are pretending-apart Lingens into two people, and (pretendedly) claiming that the man doesn't believe that these two people ("he" and Lingens) are identical. For that man fictionally to have a belief about the first of these two people, the belief must involve a certain mode of presentation (call it m_1), which is such that the man thinks of someone that way when he thinks of them first-personally. For him fictionally to have a belief about the second person, his belief must involve a certain mode of presentation (call it m_2) which requires thinking of Lingens as being called "Lingens." In (61), similarly, the speaker is pretending apart "he" and Lingens, and (pretendedly) claiming that he believes that the two are identical. The two fictional indi-

³⁹See John Perry, "The Problem of the Essential Indexical," in *The Problem of the Essential Indexical* (New York: Oxford University Press, 1993).

viduals again correspond to two ascribed modes of presentation m_3 (corresponding to 'he') and m_4 (corresponding to 'Lingens').

Plausibly, m_2 and m_4 are the same mode of presentation. What about m_1 and m_3 ? Using m_1 requires of Lingens that he think of himself first-personally, while using m_3 requires of you that you think of Lingens as "that man in the report." Now, this is not obviously incompatible with m_1 and m_3 being the same mode of presentation. Remember that we are not assuming that modes of presentation reflect cognitive similarities between the agents who employ them; here we are taking them to be simply whatever classification of thoughts we in fact make in notionally loaded attitude reports. For the pretense account, the question of the identity of m_1 and m_3 amounts to the question of the fictional identity of the individuals ("he" and "he") that the reports fictionally concern. Is it fictional that you have claimed of just one person both that you believe him to be Lingens and that Lingens does not? I am not sure, but suppose that you send an electronic message to Lingens's terminal in the library:

(62) Unlike you, I think that you are Lingens.

Here, surely, the pretense account sees a single fictional individual being used to ascribe a first-personal belief to Lingens and a "that man in the report" belief to you. We might say that an interpersonally heterogeneous mode of presentation is ascribed. This is perfectly coherent, and thus we have seen how (60), (61), and (62) can all be true.

Exercise for the reader: explore the consequences for this account of (62) of your receiving the error message, 'User cannot send message to own terminal'.

University of Michigan, Ann Arbor