

## California Semantics Meets the Great Fact

STEVEN J. WAGNER\*

*1 Introduction* I call a semantic theory extensional if it makes the meaning of a sentence a function of the referents of its semantically primitive (unstructured) parts. I provisionally assume that meanings are the objects of belief and that a 'that'-clause in a belief attribution identifies a belief by providing a sentence whose meaning is believed. Thus, 'John believes that Hesperus has risen' says that John believes the meaning of 'Hesperus has risen'. These assumptions, I think, are at best roughly accurate, but they allow a convenient statement of the main problem of extensional semantics. On the extensional viewpoint, the sole contribution of a name  $n$  to the meaning of a sentence is  $n$ 's bearer. The substitution of any coreferring  $m$  for  $n$  preserves meaning because meaning is a function of part-reference and (ordinary) names are semantically primitive. Hence, coreferring names are intersubstitutable in belief contexts, indeed in propositional attitude contexts generally. If John believes that Hesperus has risen, then he believes that Phosphorus has risen. This can seem obviously wrong, as it did to Frege. The main problem of extensional semantics is whether it is, nonetheless, true.

Russell's semantics was extensional: the meaning of a sentence was a complex of the objects, properties, and relations denoted by its terms.<sup>1</sup> But his approach fell into disfavor under the influence first of Carnap and other Fregean theorists, more recently of sententialist views of belief like Quine's and Fodor's. It is now back in fashion. One cause is simple theoretical convenience. Frege's views led him to his notoriously underspecified senses, and his followers have inherited his difficulties. If one holds that the semantic contribution of

---

\*This paper evolved from comments on an APA paper of David Austin's given in Chicago, April 1985. I am grateful for Austin's stimulation and thank him and Robert Stalnaker for showing me various manuscripts. I am very grateful to Nathan Salmon for a letter that deflated my rhetoric and forced a much more thorough presentation. (Some of the polemics remain, but I hope extensionalists will take them as friendly provocations.) Tim McCarthy's contributions are noted below.

*Received May 5, 1986*

'Hesperus' is not exhausted by its reference, one must postulate something else—call it a sense, intension, word-meaning, or what you will—that differentiates 'Hesperus' from 'Phosphorus'. If this entity is mental, one's semantics becomes involved in major problems in the philosophy of mind. If it is abstract, one must explain not only what sort of abstract entity it is, but also how it is connected to concrete linguistic practices. In any case, the extensional alternative promises to save a lot of trouble: our theory of meaning deals only with words and their referents, introducing nothing ill-understood.

The rise of possible worlds semantics also encouraged extensionalism (my name for the position). First, the possible worlds approach itself identifies the semantic contribution of a name with its bearer. In a standard possible worlds semantics, the meaning of a sentence  $s$  is the set of possible worlds in which it is true. Each possible world is in effect a model in which the truth value of  $s$  is computed according to extensional rules. Since a name is taken to have the same denotation in all worlds in which it denotes at all, its contribution to the meaning of  $s$  depends just on its bearer, as the extensionalist demands. Second, the possible worlds viewpoint accustoms one to counterintuitive results. Extensionalism permits the unrestricted substitution of coreferring names for each other in belief contexts. There are familiar ways to make this look unacceptable. But the possible worlds semanticist has to swallow much more. Since all necessary truths have the same extension across all possible worlds, believing one means believing them all. This makes us omniscient in logic and mathematics (and metaphysics?). Since true identities of the form  $n = m$  where  $n$  and  $m$  are names seem necessary (relative to the existence of the object named), it also makes us all believers in the identity of Hesperus with Phosphorus: the counterintuitive consequences of extensionalism appear as a special case. So possible worlds semantics prepares the way for extensionalism; indeed, once one has entertained the former, the intuitive difficulties of the latter may look small.

The situation semantics of Barwise and Perry [4] is now a popular variant of the possible worlds approach. Salmon [13], Soames [14], and others have revived Russell's ideas.<sup>2</sup> Although these philosophers differ on many points, the extensionalist treatment of names is common ground (at least for *standard* versions of, e.g., situation semantics). I will argue that this is enough to bring their semantic theories down. Some concept of intension for names—widely considered dispensable since [9]—will be called for. The same should then be even clearer for predicates, since names are agreed to be the extensionalist's best case. The flight from intensions will have proved illusory.

At stake here are not just semantical questions about belief contexts, of interest perhaps only to specialists. Everyday psychological talk appears to distinguish between, say, the belief that Twain is  $F$  and the belief that Clemens is  $F$ . Cases of what most people will call holding one belief but not the other are easily found. Extensionalists, on the other hand, typically argue that these beliefs are the same in spite of our strong contrary intuitions. If they are right, then fundamental confusion about our mental states is widespread. We misuse our language of propositional attitudes. This is unpalatable: although the issue is too large to be discussed here, it seems desirable to represent us as understanding our own words. In particular, we should not be systematically wrong about the logic of such an important, continuously used device as the propositional atti-

tude construction. Considerations of this kind can obviously be overridden in any given case. Still, we have reason to look for a critique of extensionalism.

Extensionalism might be considered a form of skepticism: a claim that certain common views are radically wrong. Although familiar forms of skepticism concern the external world, minds, and other traditional problems, there seems also to be common sense about language, and a view may depart from it enough to be labeled skeptical. In any case, argument against extensionalism shares a difficulty with opposition to skepticism generally: how does one refute an opponent who is already prepared to accept strongly counterintuitive results? This question sounds tendentious, since extensionalists may not agree that our intuitions about names in belief contexts are all that strong. Let us therefore consider for a moment the thesis that necessary truths are synonymous, that is, that they are identical objects of belief. Compare it to the analogous thesis for truths: that we believe every truth if we believe at least one. One claim is often taken seriously, the other, never. But if it really is clear that many truths are unknown to us, then the same holds for necessary truths. My ignorance of the number of hairs on Aristotle's head is no plainer than my ignorance of any theorem about differential equations (about which, we would ordinarily say, I know absolutely nothing). From any intuitive viewpoint, I am equally unacquainted with each. If a semantic theory can, nonetheless, credit me with one of these items of knowledge without being *ipso facto* unacceptable, why not the other? And if the theory can get away with making me mathematically omniscient, what more subtle objection is likely to have force against it? Similar questions arise for extensionalism. In a familiar kind of philosophical example, someone will appear to have as little idea of the identity of Hesperus with Phosphorus as she does of differential equations or Aristotle's anatomy. If one does claim that she knows the identity, it is far from clear that intuition is thereby violated any less than by claims of mathematical omniscience or full omniscience. The latter claims do overthrow a greater quantity of common sense, but the quality of the violation appears equal. Identities (with names) do not form a distinguished epistemic class. From the viewpoint of epistemology, there is no reason to expect them to be more easily knowable than any other propositions; we can be as far in the dark about an identity as about anything else. So it is nearly as bad to hold that we know all true identities as to hold that we know everything. Or so it seems to me. If extensionalism is indeed so counterintuitive, then it does become difficult to counter with anything better than bald denial.

Of course, these considerations alone carry little weight. The value of intuitions about names in belief contexts is debatable; and if counterintuitive results could not be right, philosophy would be extremely dull. (That was one problem with ordinary language philosophy.) Yet the following discussion will bear out these remarks. I will argue that extensionalism does — nearly — entail a version of the thesis of full omniscience. Unless the extensionalist resorts to highly unpalatable defenses, the two positions collapse into each other. There is no significant difference between them. To my mind, this is as conclusive a refutation as the subject matter will allow. But an opponent who sees nothing wrong with name-substitutions in belief contexts may be unmoved. The argument that extensionalism is as bad as the omniscience thesis might instead seem to show that the two are equally harmless.

I will use a familiar paradigm: the argument that a referentially transparent context is truth-functional if it also admits substitutions of logical equivalents. I call it the *GF* argument (for Gödel-Frege argument' or 'Great Fact argument', two of its names). The *GF* argument cannot obviously be turned against extensionalism, since the extensionalist will admit neither of its hypotheses. Belief contexts are taken to admit substitutions of coreferring names only, not of other singular terms, and logical equivalents are generally assigned distinct meanings. But I will propose an interpretation of the argument that overcomes these obstacles. Before taking that up I will try to establish two other points: that a standard defense of extensionalism and a positive argument from data about conversations both fail.

**2 Quotation** Stalnaker defended as follows the claim that we know all necessary truths [15]. Let  $T$  be some mathematical theorem that I apparently don't know—let us say that I am or appear to be investigating whether  $T$  obtains. According to Stalnaker, however, I do know  $T$ , since  $T$  is necessary. What I really don't know is the metalinguistic statement ( $M$ ) that  $T$  is (necessarily) true. In reply, several philosophers pointed out that  $M$  is a necessary consequence of things I can be assumed to know. For let  $S$  be the semantics for my mathematical language. Then the conditional  $S \rightarrow M$  is necessary—if my words have certain meanings,  $T$  must be true. So I should know the conditional. If I also know my own semantics, that is, if I know that  $S$  describes my language,  $M$  follows. (And knowing  $M$ , I can infer  $T$ .) Hence we need not even consider how well Stalnaker's defense might explain the relevant data. The move to the metalanguage reintroduces the very problem it was supposed to solve.

Stalnaker's recent variant of this defense appeals to failures of deduction ([16], pp. 72–87). His critics had assumed that we will infer  $M$  from believed premises that deductively entail it. As Stalnaker points out, however, we perform available deductive inferences only under special circumstances. There is in general no reason to expect us to believe a given  $P$  because we believe premises that entail it. Thus he denies that there is a problem about our failure to believe  $M$ . But this rebuttal is clearly inadequate. Although failures of deduction are indeed ubiquitous, this one is inexplicable. Consider, first, that  $M$  is hardly an abstruse consequence. It is the consequent of a known conditional ( $S \rightarrow M$ ) with a known antecedent, so that seeing it requires no logical ingenuity at all. Second, that in the kind of case we are imagining, I want to know specifically whether  $M$  obtains. It is not as though I could fail to make the inference because my attention, my cognitive energy, is directed elsewhere. For Stalnaker holds that if I already know  $T$ ,  $M$  is the real object of my seeming investigation into  $T$ . And I, having read Stalnaker, know that and will accordingly turn my attention to  $M$ . Third, that I recognize—I cannot be expected to overlook—the relevance of the antecedent to  $M$ . Having read Stalnaker, I also know that the desired result follows from the semantics for my mathematical language. So I will consider  $S$  and see which of the conditionals  $S \rightarrow M$  and  $S \rightarrow \neg M$  is among my beliefs. Under these circumstances, a failure to infer  $M$  would be mysterious. We generally omit deductive inferences that do not interest us. We can also fail to perform as simple a step as modus ponens if we do not recall, at a given moment, the premises that would yield our desired conclusion. But such expla-

nations cannot apply here. In fact, it would be strange enough for us somehow to overlook the step to *M* in a particular case of this kind. As a *general* account of our mathematical ignorance, Stalnaker's mechanism—a failure to detach consequences that are right before our eyes—is incredible.

What if I haven't clearly grasped *S* or don't see its relevance to *M*? Whatever the initial likelihood of this event, it certainly ought to be zero after I have read Stalnaker. Once I know the nature of my difficulty, as he analyzes it, I should have a certain cure for mathematical ignorance: given any mathematical question, I simply get clear on the relevant semantics and see which answer it entails. But that doesn't work.

The fate of Stalnaker's metalinguistic defense might have discouraged analogous moves elsewhere. Yet just such a defense is advanced in what I consider to be the best statement of the extensionalist case [14].

Consider the case of Mary and her neighbor, whom she knows as 'Samuel Clemens'. On the basis of a number of entertaining conversations, Mary believes and would assert

(1) Samuel Clemens is witty.

Now Mary likes cultural talk shows and panel discussions with major American writers. One figure she has often observed on these is a writer introduced as 'Mark Twain'. Due to the poor quality of her television, plus differences in make-up and dress, she has not recognized her neighbor, although she knows the television persona well. She has a poor impression of his wit, having found him dull compared to Hawthorne, Melville, and James. Thus, she would not assert, and would likely even deny

(2) Mark Twain is witty.

Yet Soames (whose example I've adapted) holds that Mary does believe (2); this follows from (1) for an extensionalist. What she really doesn't believe, he says, is

(3) 'Mark Twain is witty' is true.

Soames argues that Mary's disbelief (i.e., lack of belief) in (3) accounts for her apparent disbelief in (2) (as when she answers 'Is Mark Twain witty?' in the negative). Our ignorance is transferred from a statement about Twain to one about language.

Soames hardly explains how disbelief in (3) is supposed to do this job. One would have to examine the common sense psychological explanations in which disbelief in (2) seems to rationalize an agent's behavior and show that under suitable, plausible modifications, substituting (3) does as well. I would not assert in advance that this must fail; one would have to see Soames's defense worked out in detail for a range of cases in order to identify its shortcomings with any confidence. But it is hard to see how metalinguistic beliefs, or their absence, could play the role of first-order beliefs like (2) in explaining nonlinguistic behavior. Thus, any case in which someone's nonlinguistic behavior would normally be explained by a combination of belief in (1) and disbelief in (2) should pose a serious problem for the extensionalist. Whereas, if there are no such cases, the distinction between (1) and (2) becomes suspect even from the intensionalist

viewpoint. A distinction between beliefs should make a behavioral difference in suitable contexts. If no such context can be found, intensionalists would have to concede the identity of (1) and (2) by their own standards, and Soames's argument would be unnecessary. The metalinguistic defense, then, would be either hopeless or pointless.

In fact, the defense would appear pointless in any case, since (3) follows readily from (2) and the obvious

(4) 'Mark Twain is witty' is true iff Mark Twain is witty.

If Mary knows just a little semantics, what is to keep her from believing (3)?

It will be useful to state an alternative form of this objection. To take the general case, suppose that '*n*' and '*m*' are coreferring names in the vocabulary of a speaker *S*, that is,

(5)  $ref('n') = ref('m')$ ;

and that *S* believes that

(6) *n* is *F*.

According to Soames, then, *S* believes that

(7) *m* is *F*

But not that

(8) '*m* is *F*' is true.

However, it will be agreed that *S* believes

(9) '*n* is *F*' is true.

Now if we can establish that *S* believes (5), then the extensionalist loses: from (9) and (5), *S* can infer (8) if she knows just a tiny bit of semantics, contradicting Soames. Her access to (5), however, is not obvious. The extensionalist claims that we believe all true identities between names (where at least one of the names in the identity is in our vocabulary), but in (5) the identity sign links definite descriptions. And these will receive distinct contents in the extensionalist semantics, since one is built up from a name for '*n*' (namely, '*n*'), the other from a name for the distinct object '*m*'.

Disquotation comes to the rescue. Since *S*, knowing a little semantics, believes

(10)  $ref('n') = n$ ,

she believes

(11)  $ref('n') = m$  (by (4), on extensionalist principles).

Similarly, she believes

(12)  $ref('m') = m$ .

But from (11) and (12) she can infer (5) at once. (8) then follows, showing that the extensionalist has just as much trouble with our ignorance of (8) as with our ignorance of (7).

Let us review the restrictions on the argument: that both ‘*n*’ and ‘*m*’ are in *S*’s vocabulary; that *S* knows a little semantics; and that she will perform a few elementary inferences. The first is genuine. I do not see how to show that *S* believes (8) (on the extensionalist hypotheses) if she does not know the name ‘*m*’, for then (12) is not available. And this is the situation in some examples in the literature. But this does not save the metalinguistic defense, since it must also work when both names are known, as in our Twain-Clemens case. The second restriction is harmless, since the defense is clearly supposed to apply to speakers who know how to disquote. Nor could one credibly explain the phenomena that suggest absence of belief in (3) by saying that they really show ignorance of elementary semantics. Finally, I have only shown that *S* is in a position readily to infer (8), not that she actually does believe it. But that is enough. The extensionalist’s attempt to make something out of this would fail in much the same way as Stalnaker’s “failures of deduction” line.

Yet a deeper version of that line needs consideration. One might suggest that while the *sentence* (4) follows from the *sentences* (10) and (11), the *beliefs* need not stand in a corresponding inferential relation. Only the assumption that beliefs have essentially sentential structure seems to justify this correspondence. Only then does it seem obvious that one belief can be reached from the other by identity logic. Perhaps that is false; holding beliefs (11) and (12) may not always provide grounds for an immediate step to (5). Similarly, the argument for (3), depending as it does on the structures given in our sentential representations, would be blocked. The extensionalist would hold that believing (2) and (4) may put one in no position to infer (3). I see four *prima facie* difficulties for this defense.

First, it may sit uneasily with the metalinguistic defense. I do not know just how the substitution of (3) (or disbelief in it) for (2) in psychological explanation is supposed to work. But consider the extensionalist’s presumably best cases, those involving linguistic behavior. There, the question ‘Is Mark Twain witty?’ is in some sense to be taken as a request for the truth value of the sentence (2). But suppose that the content of Mary’s stored information about the world is in general nonlinguistic—representing in the first instance Mark Twain, not ‘Mark Twain’. Then the extensionalist’s model of the situation would seem to require a proper link between the nonlinguistic information and the words whose truth value is to be judged. If the nonlinguistic information is something like

(2) Mark Twain is witty

(or its denial), and if Mary is to assign truth values on its basis, then one would think that something like

(4) ‘Mark Twain is witty’ is true iff Mark Twain is witty

is needed for the transitions. And of course, it must actually be used, or the transitions will not occur. So I would conjecture that the very transitions on which my derivation of Mary’s belief in (3) depend are crucial to Soames’s own account of question-answering.

Second, the reply suits a possible worlds framework better than extensionalism. For Stalnaker, beliefs are unstructured. It is senseless to call them con-

junctive or conditional or to assign one the form of an identity. Their contents are just sets of possible worlds. The same holds on related semantic views, including situation semantics. But the extensionalist, following Russell, does give beliefs structure—a content is an ordered structure of objects and properties corresponding to the words used in its conventional expression. Because the sentential structure is preserved (up to sameness of reference for primitive elements), our reply could hardly seem *natural* to an extensionalist.

Third, the reply threatens to make belief irrelevant to psychological explanation, hence quite likely wholly irrelevant. Our accounts of practical and theoretical reasoning exploit the logical structures of beliefs or, if you like, of sentences expressing beliefs. If Mary thinks she should eat dry food, and if she believes that this is dry food, we expect her (in an appropriate context) to think that she should eat this. The success of such explanations suggests that beliefs can be related by (something like) universal instantiation. If logical relations between sentences do not induce corresponding inferential relations between beliefs, common sense psychological explanation is hopelessly crippled. (This holds even though the exact nature of the induced relations is a question for future theory.) But if a correspondence does obtain, then the derivations above should go through. The step to (3) is truth-functional, while the argument for (5) exploits simple identity logic: either the principle that if two things are identical to a third (as in (11) and (12)), they are identical to each other, or symmetry and transitivity. Since any route defeats the metalinguistic defense, all must be blocked, virtually eliminating the application of logic to belief contents.

Fourth, if one does block the application of logic to belief contents, one should save a step by doing it at once. The Clemens-Twain case seems worrisome because of problems of the following kind: if *S* does, as the extensionalist holds, believe (2), why doesn't she try Mark Twain's books for entertainment? That would, after all, be natural, given her belief plus appropriate collateral premises, e.g., that witty authors tend to write entertaining books. And why doesn't she answer 'Yes' to 'Is Mark Twain witty?'" That, too, would be natural if she believes (2) and thinks that the name in the question refers to the object of this belief. All such difficulties about the attribution of (2) are based on apparent failures to manifest the belief. But each argument for expecting a certain manifestation is easily seen to exploit logical relations between (2) and other psychological states. For example, we expect Mary to reason that Mark Twain wrote this book, and he is witty, and witty authors tend to write entertaining books, so. . . . If such inferences are not straightforward, the troublesome consequences of extensionalism will not appear anyway.

Rejoinders to each of these objections can be imagined. But they would miss the point. These are, after all, just *prima facie* challenges to a very programmatic idea. The question at the moment is whether the metalinguistic defense is worth developing. What are its likely benefits and costs? With the metalinguistic move, the extensionalist hopes to explain away the counterintuitive consequences of crediting Mary with belief in (2)—failures to buy Mark Twain's books, and so forth. These are indeed problems, but they hardly seem worse than having to explain why Mary can't perform the simplest logical steps, or why her disquotational skills are exactly those needed to make the metalinguistic account of question-answering work without allowing her to derive (3). On top



of that, recall, the metalinguistic account will need to handle nonlinguistic behavior. How could biting the bullet on (2) be worse?

These, however, are only preliminary reflections. We will shortly see that the extensionalist cannot simply choose to live with the difficulties of (2). The issues of the past few paragraphs will therefore need closer attention. But even then, the metalinguistic defense will not help.

**3 Conversation** When someone says, 'Clemens is *F*', we sometimes report her as having said that Twain is *F*, even if she did not use and would not have used those words. Similarly for reports of belief and other attitudes. Familiar examples of this kind have been offered as strong evidence for extensionalism (e.g., [12], [14]). Extensionalism, however, is a logical thesis. The question about the evidence is whether it really bears on the logic of the attitudes, not just on the rules for conversation. What is an admissible belief report would seem to vary with the context in a way that the content of our beliefs cannot.

*Example 1* *A* and *B* are arguing about the personality of Mark Twain, whom they've met at novelists' conventions. *A* finds him a great wit; *B* is not impressed. To support his own view, *A* reminds *B* of the judgment of their mutual friend Mary, a respected connoisseur of wit:

(13) Mary believes that Twain is witty.

One or both speakers may know that Mary knows Twain only as her neighbor Clemens, but neither possibility would make *A*'s remark less appropriate here. Her judgment of the man's wit is at issue.

*Example 2* Mary will host a writers' party, basing the guest list on her impressions from literary talk shows. *C* and *D* know Mary's impressions and speculate that Hawthorne, Melville, and James are likely guests. But *C*, knowing that Mary has a horror of dull parties, thinks Twain will probably be left out, since (he says),

(14) Mary does not believe that Twain is witty.

Although one of these reports is wrong, both are entirely appropriate in their contexts.<sup>3</sup> The moral of such (readily multiplied) examples is that conversational appropriateness is a poor guide to semantics. An extensionalist might of course reply that (14) is really false but a natural way to convey relevant information in a particular setting. This, however, will invite the analogous claim for (13). Once the distinction between the truth and what one ought to or might say is taken seriously, the critic gains as much ammunition as the extensionalist, and conversational data become useless.

A further limit of the argument from conversation is that it cannot, in any case, help with identity problems. If Mary has no idea who her neighbor is, I do not see how it could ever be appropriate to say that

(15) Mary believes that Clemens is Twain.

That is no disproof of extensionalism, since (15) might still be true. But it does suggest that the argument from conversation could not help with the hardest cases.

Finally, if the argument did work, it would show too much, since what goes for names also goes for definite descriptions. Let ‘the *G*’ be any description applying to Twain. Regardless of whether Mary is familiar with this description or knows that Clemens satisfies it, one can easily imagine contexts in which it is appropriate to report that

(16) Mary believes that the *G* is witty.

In fact, one can argue that this kind of report is always correct (although it will sometimes sound strange)—that one simply has to hear (16) in the right way, as carrying no import about Mary’s knowledge, about how she would express herself, etc. (cf. the attempts to establish a *de re* or “wide scope” reading of (16)). This is too familiar to need reviewing here. So it seems that if conversational data support an extensionalist view of names, they also suggest that coreferring descriptions have the same content. Yet few would claim this. (It is certainly inconsistent with the extensionalist semantics, since distinct descriptions will be structures with distinct parts.) And if it were maintained, that would pave the way for the coming application of the *GF* argument.

My views here follow Quine’s. He remarked of indirect quotation that the allowable deviation from a direct quote “depends on why we are quoting. It is a question of what traits of the quoted speaker’s remarks we want to make something of; those are the traits that must be kept straight if our indirect quotation is to count as true. Similar remarks apply to sentences of belief and other propositional attitudes” ([11], p. 218). This seems right: we represent someone’s statements or beliefs in a way that suits the conversational purposes at hand. Obviously, the appropriateness of such a representation supports no accurate conclusion about the subject’s mental state.<sup>4</sup> As Quine says, indirect quotation is “an essentially dramatic act” (p. 219). Playing the part of another speaker, we choose our lines with an eye to the audience and the intended moral of our dramatization. The argument from conversation effectively confuses theater with history.

**4 Refutation** The *GF* argument is supposed to show that if a context permits substitution of coreferential terms and logical equivalents, it is truth-functional.<sup>5</sup> But as far as I know, it always involves a major fudge. Here is a common version: We imagine that ‘*P*’ and ‘*Q*’ are true and that the substitutions take place within some context *C*:

- (A) *P*  
 (B)  $(\lambda x)x = 0 = (\lambda x)(x = 0 \ \& \ P)$  (by logical equivalence)  
 (C)  $(\lambda x)x = 0 = (\lambda x)(x = 0 \ \& \ Q)$  (substitution of coreferring terms)  
 (D) *Q* (logical equivalence)

Another version replaces (B) and (C) respectively by

(B1)  $\{x|x = 0\} = \{x|x = 0 \ \& \ P\}$

and

(C1)  $\{x|x = 0\} = \{x|x = 0 \ \& \ Q\}$ .

(The choice of 0 is conventional to the extent that any other mathematical object would do as well.) Now in neither case are the first or last two statements logically equivalent—unless one thinks that 0, or some object one might substitute for it, exists by logic. Perhaps (A) and (B), e.g., are necessarily equivalent. But if substitutability of necessary equivalents replaces substitutability of logical equivalents in the hypotheses of the *GF* argument, the result is less interesting. For example, the latter, original hypothesis is plausible for the belief contexts of an ideal reasoner, who can carry out arbitrarily complex logical derivations as she pleases. Not so for the former hypothesis unless we considerably strengthen the idealization. Thus, the weakened version of the *GF* argument will have fewer significant applications.

This weakness is no accident. The *GF* argument requires a variable-binding term operator. It requires two statements to be logically equivalent even though one contains a term, formed by the operator, that the other does not; that term must denote; and the statement without the term cannot be assumed to mention the object denoted in any way. Thus, something that exists by sheer logic is needed.

My application, however, needs nothing like that. I will recast the argument as a proof about possible inference relations: if extensionalism is true, anyone who believes one truth is in a position to infer any other truth in a few steps. The only other assumption is that the subject (*S*) truly and justifiably believe in something—0, herself, the Taj Mahal, or anything else—and have a name ('*n*') for it. An approximation to the argument runs as follows. Suppose *S* truly and justifiably believes

(A2) *P*.

Then she can infer

(B2)  $(\iota x)x = n = (\iota x)(x = n \ \& \ P)$

which follows by logic from (A2) and the existence of a bearer for '*n*'. Notice that no claim about identity of belief content is needed here. We simply observe that *S* is in a position to infer the presumably new content (B2) if she believes (A2). Now if belief contexts allowed the substitution of coreferential terms, we could conclude from (B2) that *S* believes

(C2)  $(\iota x)x = n = (\iota x)(x = n \ \& \ Q)$

where '*Q*' is any other truth. Then

(D2) *Q*

would be inferable, again by logic and the existence of the chosen object. (Again, no identity of content is needed.) But the extensionalist will reject the second step, since it involves descriptions, not names. How can we get around that?

One idea would be to stipulate rigidity for the descriptions ' $(\iota x)(x = n \ \& \ P)$ ' and ' $(\iota x)(x = n \ \& \ Q)$ '. Even if descriptions are normally nonrigid, the stipulation is clearly possible, and it would not interfere with *S*'s reasoning. The first and last steps of the argument are independent of the rigidity of the terms involved. One might then be able to argue that while the extensionalist semantics generally distinguishes between coreferential terms, coreferential *rigid* terms

should have the same content. Thus the step from (B2) to (C2) could be made. The prospects for this idea are unclear, since it turns on hard questions about rigidity and its relation to content. But a modification is more straightforward: let *S* introduce names via the descriptions, that is, by the “baptisms”

(17) Alice =  $(\lambda x)(x = n \ \& \ P)$

(18) Beth =  $(\lambda x)(x = n \ \& \ Q)$ .

In each case, it is understood that the name rigidly designates the object described. (C2) can then be reached:

(B2)  $(\lambda x)x = n = (\lambda x)(x = n \ \& \ P)$

(B21)  $(\lambda x)x = n = \text{Alice}$  (by (B2) and (17))

(B22)  $(\lambda x)x = n = \text{Beth}$  (since ‘Alice’ and ‘Beth’ corefer, the extensionalist allows their interchange in belief contexts)

(C2)  $(\lambda x)x = n = (\lambda x)(x = n \ \& \ Q)$  (by 18)

Some likely questions about these steps will be discussed. First, let us be clear on what the derivation shows if it works.

I am offering a recipe by which *S* can reason her way from any believed truth to any other truth. This shows not that if she believes one truth, she will believe them all, but rather how she can justly infer any truth she pleases.<sup>6</sup> As remarked above, the only auxiliary assumption is that she have a name for something. Granted that, let her begin with her true belief ‘*P*’. Using the existence of the named object, she infers (B2) by logic. Then she uses the description on the right hand side of this identity in a baptism of Alice. Given this introduction of ‘Alice’, *S* is justified in believing that Alice is the described object—hence (B21). Now suppose *S* wants to infer ‘*Q*’ if (and only if) ‘*Q*’ is true. Then she performs the baptism (18), adding ‘Beth’ to her vocabulary. If ‘*Q*’ is true, ‘Alice’ and ‘Beth’ are coreferring names, so that according to extensionalism, (B21) and (B22) are the same belief. Therefore *S*, believing (B21), believes (B22). Since ‘Beth’ was introduced via the description ‘ $(\lambda x)(x = n \ \& \ Q)$ ’, *S* can infer that Beth satisfies it. Hence the step to (C2). (D2) is then immediate. Of course the trick lies in the step in (B22) (which fails if ‘*Q*’ is false. Thus, *S* cannot reach a false conclusion by this route.) If *S* has no prior reason to believe ‘*Q*’, she would appear to have no reason for this inference at all. But that is just the point. For the extensionalist, no reason is needed because there is really no inference and no step, just identity of belief. To block the derivation, we should reject the extensionalist semantics. Another way of viewing it: my derivation confirms the intuitive idea, mentioned at the outset, that identities form no distinguished epistemic class. There is nothing special about them: if we credit someone with infallible knowledge of identities, even just of name-identities, we might as well suppose that she knows everything.

One might question my use of names and descriptions. Have I assumed the *synonymy* of ‘ $(\lambda x)(x = n \ \& \ P)$ ’ with ‘Alice’ (“philosophizing as though Kripke had never existed”)? No. I require only two things: that one can introduce a name via a definite description, and that having done so, one can normally infer that the descriptum bears the name and that the bearer satisfies the description. Assuming confidence that the description denotes (for the other case, see below), this looks undeniable. Philosophical controversy starts when one labels the infer-

ence, or a proposition such as (17), analytic, necessary, or a priori, but I make no such claim. Am I “giving names descriptive content”? I am not wedded to this phrase, but in a way, that is the whole issue. Identifying the content, or semantic contribution, of a name with its bearer gives ‘Twain’ the same content as ‘Clemens’, ‘Hesperus’ the same content as ‘Phosphorus’, and so on. Fregean intuitions find differences between the names in these pairs when speakers “associate different concepts” with them. The nature of a concept and of the intended association are problematic. But a description has informational content, and introducing a name with one somehow links its content to the name. The closer analysis of this situation is not necessary here, as long as we agree that the link supports the inferences I have described (e.g., from (B2) to (B21) via (17)). A coreferring name, then, can be introduced with a different description, thus associating it with information supporting other inferences. That is the essence of my version of the *GF* argument. I make no stronger claims about the meanings or contents of proper names.

I owe another objection to Tim McCarthy: the recipe is pointless if *S* already believes ‘*Q*’, but doesn’t her use of ‘Beth’ presuppose that? To believe (18)—to believe that Beth is the unique  $x$  such that  $x = n \ \& \ Q$ —one must believe that there *is* such an  $x$ . That entails believing *Q*. To take an actual example, suppose that I decide to use ‘Homer’ to name the common author of the *Iliad* and the *Odyssey*. Homer is to be the person, if there was one, who wrote both books. Now to *believe*

(19) Homer is the person who wrote both the *Iliad* and the *Odyssey*

I must believe that there was such a man—that the description with which I introduce the name denotes (uniquely). Similarly, *S* must believe that the description in (18) denotes if she is to believe (18). And that trivializes the argument.

My reply is that (18) is not a proposition *S* believes. It rather functions as a stipulation, a definition introducing the name ‘Beth’.<sup>7</sup> My presentation was therefore inaccurate. A clearer reformulation would have *S* pronouncing

(18’) Let ‘Beth’ denote the unique  $x$  such that  $x = n \ \& \ Q$ , if (and only if) there is such a thing.

Similarly, I can stipulate the following, regardless of my view on the Homeric Question:

(19’) Let ‘Homer’ denote the common author of the *Iliad* and the *Odyssey*, if there was one.

This fixes the denotation of ‘Homer’ for me. I may lay it down even if I am absolutely convinced that there was no such person. I will not then believe (19), but still, (19’) is a perfectly legitimate baptism. If I happen to be wrong—if the books really do have a common author—then (19) will state a truth for me. For example, my attempt to tell a lie with it would fail. In general, our ability to introduce names with descriptions is not limited by our beliefs about whether the descriptions denote.

My *GF* argument therefore survives McCarthy’s objection. ‘Beth’ can still be introduced and does in fact denote if ‘*Q*’ is true. Its key substitution for ‘Alice’ can thus be performed much as before. Consider the situation once *S* has

reasoned her way to (B21). If her question is whether ‘ $Q$ ’ obtains, she can introduce ‘Beth’ by (18’) without any commitment to an answer.<sup>8</sup> Since (18’) makes ‘Beth’ coreferential with ‘Alice’, she will believe (B22) if the extensionalist semantics is right. From this she can infer (C2) roughly as before. With a little help from semantics, (18’) supports  $S$ ’s reasoning as well as (18) did:

- (18’) Let ‘Beth’ denote  $(\exists x)(x = n \ \& \ Q)$ , if there is such a thing.  
 (B21)  $(\exists x)x = n = \text{Alice}$  (as before)  
 (B22)  $(\exists x)x = n = \text{Beth}$  (by extensionalism, (B21) and the stipulation for ‘Beth’)  
 (B23) ‘ $(\exists x)x = n$ ’ denotes (from  $S$ ’s beliefs about ‘ $n$ ’)  
 (B24) ‘ $(\exists x)x = n = \text{Beth}$ ’ is true iff  $(\exists x)x = n = \text{Beth}$  ( $S$  knows the semantics of her language)  
 (B25) ‘ $(\exists x)x = n = \text{Beth}$ ’ is true (by (B22) and (B24))  
 (B26) ‘Beth’ denotes (by (B23) and (B25))<sup>9</sup>  
 (B27)  $\text{Beth} = (\exists x)(x = n \ \& \ Q)$  (by (18’), (B26) and elementary semantics)  
 (C2)  $(\exists x)x = n = (\exists x)(x = n \ \& \ Q)$  (by (B22) and (B27))

I see nothing wrong with stipulating (18’) when one positively disbelieves ‘ $Q$ ’. But if that seems odd for any reason, we can restrict the argument to cases of initial agnosticism about ‘ $Q$ ’. It would still provide an infallible recipe for relieving agnosticism, which is still bad enough.

Let us consider the possible obstacles to this derivation. One might suppose that  $S$  cannot reach (B22): that *we*, knowing the coreference of ‘Alice’ and ‘Beth’, can produce (B22) by substitution, but that  $S$  herself could not thus continue the argument. But such worries can be dismissed. In assuming that the step from (B21) to (B22) requires extra knowledge, they simply assert Frege’s viewpoint. If (B21) and (B22) are the same belief, as the extensionalist holds, no “substitution” leading from one to the other is needed.

One might also object that even if  $S$  believes (B22), she does not know that she believes it; and that lacking such knowledge, she has no reason to proceed any farther than (B24). This has an intuitive point. After all, the derivation should break down somewhere, and if we must concede (B22), we should try somehow to interpret that belief in a way that prevents  $S$  from taking the remaining steps. And one may well feel that  $S$  could not fully grasp the import of her belief about Beth. The suggestion that she is unaware that she has it attempts to express this feeling. But this is no way to save extensionalism. First, we would at the very least have the result that anyone who believes one truth can infer, for any other truth  $Q$ , a statement from which  $Q$  follows by elementary logic and semantics, even if for some reason she is unable to complete the inferences. That also seems unacceptable. Second, it is unclear whether  $S$  needs to know what beliefs she has. One can imagine a smart robot carrying out the entire derivation unreflectively. When it gets to (B24), it simply continues to apply elementary logic and semantics until (D2) is reached. Self-consciousness appears irrelevant here. *Inferential access*—being in a position to perform inferences, e.g., on (B22)—is relevant, and perhaps the appeal to self-consciousness is really an attempt to raise a difficulty about that. But (B21) is a perfectly ordinary belief. (B22) is supposed to be exactly the same belief. If extensionalism must introduce an *ad hoc* theory about lack of inferential access to perfectly

ordinary beliefs in order to block our derivation, it is in trouble. (Recall the “failures of deduction” again.) Third, the identity claim for (B21) and (B22) leaves the extensionalist no room to deny *S* awareness of (B22). For she will believe

(19) I [*S*] believe that  $(\exists x)x = n = \text{Alice}$ .

But if ‘Alice’ and ‘Beth’ have the same semantic content, then the extensionalist semantics will presumably assign (19) the same content as

(20) I [*S*] believe that  $(\exists x)x = n = \text{Beth}$ .

This certainly happens in Russellian semantics, or on any construal of belief as a relation to propositions, if (B21) and (B22) express the same proposition. (It fails just on the Fregean or sentential theories of belief that extensionalism rejects.) So the attempt to distinguish between the reflective (19) and (20) runs afoul of the identity of their first-order objects.

Having exhausted these moves, we return to an issue raised earlier. The extensionalist’s best try may be to disallow (B25): one denies that the *T*-sentence (B24) plus one half yields the other half, just as (2) and (4) were not supposed to yield (3). As before, the route to the undesired conclusion ((3) or (B25)) can go via identity logic instead of a *T*-sentence, so that a general prohibition on the application of logic to belief contexts is required: agents cannot generally infer immediate implications of their beliefs, even when their logical abilities, attentional resources and such pose no obstacles. Earlier this line hardly seemed sensible. It was offered to save a metalinguistic defense that looked dubious from the start; on top of that, the consequences of the solution were not clearly better than the original difficulties. What the *GF* argument makes clear is the extent of the problem. Unless something is done, extensionalism will be stuck not just with some odd-looking belief attributions involving proper names, e.g., (2), but with the omniscience of anyone willing to go to a little trouble for it. This makes the solution — call it logical opacity — worth reconsidering.

Logical opacity looks to be an odd position for extensionalists, of all people, to embrace. A main extensionalist idea is to loosen Fregean restrictions on semantic operations within belief contexts. Structure-preserving substitutions of coextensional items are allowed, thus simplifying attitude semantics and, at the informal, intuitive level, facilitating belief attributions. Logical opacity has the opposite effects. Judicious imputations of logical ability are a main tool for belief attribution; take it away, and guessing anyone’s thoughts becomes hard. On the technical side, semantic theories explaining logical opacity could hardly be simpler than intensionalist accounts of ordinary (apparent) referential opacity, whether simplicity is measured by the amount of formal machinery or by the novelty and difficulty of the concepts employed. One wonders, then, whether the attractions of extensionalism could survive the turn to logical opacity.

Logical opacity may introduce other tensions as well. But there is no need to look for them: no one really wants opacity. If we are going to talk about belief at all, the general structure of ordinary intentional explanation is a given. Something close to our ordinary views on how reasons generate belief, desire, and action must be right. But as noted earlier, intentional explanation is impossible if we cannot expect agents to make obvious logical transitions. The exten-

sionalist must therefore combine an opacity thesis blocking the step to (B25) with a way of allowing Mary to eat her dry food and, for that matter, of letting *S* make all the other moves in the derivation of (D2). Among other things, then, there must be believers who can readily infer (B25) from (B22) and (B24); more accurately, believers whose beliefs we would represent by these sentences but who differ from *S* in a way that makes the “normal” Tarskian inference straightforward. The question is how to make sense of this position.

Although I know no developed answers, Salmon seems to have shown the form they would have to take [13]. He holds that we do not simply believe propositions, but may rather believe a certain proposition under one “guise” without believing it under another. For example, Mary does believe that Twain = Clemens, and *S* that Alice = Beth, but each of these may appear in a guise that prevents recognition. Roughly, Mary knows

(22) Twain = Clemens

but only recognizes it in certain guises, not including ordinary assertions of ‘Twain = Clemens’. It is much like failing to know a friend at a costume ball: the proposition must be presented in the right way, where the presentation may include both the linguistic expression and a variety of contextual factors. Of course this leaves it open just what a guise is, but for the moment we need only concern ourselves with the basic idea.<sup>10</sup> Broadly, one should grant that such a theory is possible, but be utterly puzzled by its claim to extensionality.

Let us focus on the inference to (B25) which, as represented above, is an instance of modus ponens (since only one direction of (B24) is used) with (B22) as the antecedent. The theory, then, is that some agents will be able to perform this step without difficulty, while others, such as *S*, cannot. Since we agree that modus ponens must be widely available to intelligent agents, and since the conditions of its availability involve guises, the extensionalist’s solution appears evident and, as far as I can tell, forced. The step to (B25) will be immediate when two guises of the proposition that (B22) expresses are the same for an agent: that under which (B22) is held, and that in which it appears in the antecedent of (B24). When, on the other hand, the guises are distinct, the agent has no reason to proceed. For *S*, the proposition that  $(\iota x)x = n = \text{Beth}$  will appear as something like, or something involving, the sentence (B21), while our sentence (B24) actually gives a correct indication of how Mary apprehends the proposition expressed there. (Recall that I am only sketching the extensionalist’s general line. Particular developments of the idea could take many forms.) Thus, guises are parameters of sorts. The application of logic to belief contents can proceed as desired, but only relative to constancy of guise for the propositions involved. This saves the indispensable patterns of intentional explanation without allowing the collapse into omniscience.

Now we know that this approach is in principle perfectly workable—we learned it from Frege. If ‘Alice = Alice’ and ‘Alice = Beth’ stand for the same object, how can they differ in cognitive value? Because they present it in different guises. If you assert the former sentence and I the latter, we have therefore expressed different cognitive states, and you may find my statement informative. Similarly, the names ‘Alice’ and ‘Beth’ are associated with different guises for the same object. Or suppose I believe ‘Alice = Alice’ (have the belief I would



thus express) and ‘Alice = Beth  $\rightarrow$  P’. Why don’t I believe P? Because the distinct guises associated with ‘Alice = Alice’ and ‘Alice = Beth’ prevent me from recognizing their shared object, hence from seeing the relevance of ‘Alice = Alice’ to my conditional belief. (For Frege, any two sentences with the same truth value would serve as well here. But his idea has no essential dependence on the identification of sentential referents with truth values, as opposed to referential structures.) Only, Frege called these guises of objects the senses of sentences and thereby produced a paradigm of intensional theorizing. Given the obvious parallels between his theory and extensionalism-plus-guises, why would anyone offer the latter as an *alternative* to Frege, or as a solution to his puzzles about names? Rather, one ends up where Frege began.

Salmon and others could well reply by emphasizing numerous differences with Frege. My presentation was of course framed so as to bring out structural analogies, and hence could be taken as a kind of caricature. But the issue here is the broad classification of semantic theories. If there is a relevant level at which guises are theoretically indistinguishable from senses, then we are justified in viewing one idea as a variation on the other. Calling tigers tabbies is, after all, also a caricature; but when someone who knows house cats claims to have found a quite novel animal in the jungle, “It’s just a big cat” is the right thing to tell him.

I can only offer a preliminary criticism of any theory of guises, since the theory itself has only been adumbrated. It would be pointless to try to anticipate every likely articulation of guise theory. For the moment, it suffices just to ask how guises could ever be more than senses under a new name.

Would senses under a new name be so bad? Since Frege never defined senses, any clarification provided by extensionalists under the guise of explaining guises would be welcome. In that respect, the guise theory would be a major advance. But if it comes, extensionalism can take no credit. The advance will in no way fall out of identifying the content of a name with its bearer, which only disguises the need for a Fregean epicycle later on. Rather, it will come from epistemological and cognitive work which could as well be done in an explicitly Fregean framework.

The upshot is that logical opacity is no defense against the *GF* argument. It does away with the very idea the extensionalist had set out to develop.<sup>11</sup> This would seem to leave no escape. However, I will comment on a new move that rejects wholesale the premises of our discussion so far ([1], [21]).

Briefly, the idea is to drop the assumed connection between semantic content and substitution in belief contexts. We have treated ‘believes that’ as a *content-functional* operator by identifying sentence meaning with content and letting identity of content for *P* and *Q* guarantee identity of the belief that *P* with the belief that *Q*. This was explicit in the opening lines of this paper and evidently presupposed in the step from (B21) to (B22). But nothing seems to force this assumption. (In fact, Almog makes it look quite groundless.) And without it, there is no way to work any *GF* argument against extensionalism. So why not leave belief contexts – which seriously challenge any theory – for another day? The extensionalist’s contribution would, then, be a notion of content for sentences without intentional (and, perhaps, other intensional) operators. It would not rule out such a theory but would rather leave it as a matter for future inves-

tigation along the lines developed for nonintentional discourse. If this is not entirely satisfactory, no other approach can claim to have achieved more. Getting away with this would be a coup: all our difficulties about substitution disappear without the need for the slightest modification in the basic theory of content. An adequate discussion of this bold move would have to investigate the entire foundation of semantic theorizing, which is not possible here. But it will be useful to raise some questions.

I agree that the relation of belief to the semantics of sentences may not be straightforward. It is very tidy to take beliefs as sentence-like objects and view belief as an agent's relation to something like a linguistic meaning. Theories as different as Frege's and Fodor's share this idea, and numerous elaborations have shown its power. But careful study reveals no good a priori argument for it and no clear empirical support (see, e.g., [8]). The idea probably gets most of its force from our use of sentences to attribute beliefs, but as noted above, belief attribution is a messy business. It is far from clear how much the form of common psychological language shows about the nature of belief. Thus, the simple picture of beliefs as attitudes toward the meanings or contents of sentences must also be questioned. For the time being, it might be prudent to isolate semantics from psychology. The semantics for natural languages, in the sense of something like truth theory, would be one issue; another would be the nature of mental states and their possible similarities to linguistic objects; a third would be the complex and perhaps quite misleading or inadequate ways in which we use ordinary language to represent minds. From this viewpoint, the comfortable assumptions with which we began cannot be maintained.

Here we have large issues that the philosophy of language has barely touched. And I quite like the initial response of putting most received ideas about belief on hold. But extensionalism can hardly welcome this situation more than any other traditional semantic theory, for it too seems to derive its motivation from the assumptions now under fire. Content has all but been defined in terms of belief. On Fregean—standard—assumptions, the force and point of identifying the content of a name with its bearer is perfectly clear. It means that 'Beth' can replace 'Alice' in belief contexts, and so on. Of course this whole idea can be rejected. The question is where that leaves the notion of content. If substitution in belief contexts isn't the issue, what can the identification mean? If we simply define agreement in content as sameness of Russellian structure, we have one notion too many. Without independent motivation, as provided by Frege, talk of content is redundant. The extensionalist might as well just say that 'Alice' has the same reference as 'Beth', which we already knew.

Let us not be too hasty. I have just spoken as though content *must* be a matter of equivalence in belief contexts, but alternatives might be proposed. The general idea is to look for a relation among names that must hold if sameness of reference holds and that guarantees something of further, independent theoretical interest. Although this latter condition is vague, its intent should be clear enough. Frege motivates a significant notion of content in one way, with reference to a conception of language and thought. What else might do?

(i) One might define content in terms of substitution in modal, rather than psychological, contexts. But as Almog notes, this is hardly a step up from Frege [1]. If the connection between belief and semantic content is in doubt, we should

find the analogous link to necessity equally dubious. In fact, assessing content via behavior in necessity contexts is probably worse. It is much clearer that people have beliefs, and that these very roughly fit our ordinary views of them, than that (“metaphysical”) necessity exists and is more or less what philosophers take it to be.

For similar reasons, a related suggestion in [2] is also insufficient. The possible worlds framework may offer a variety of appropriate notions of content; Almog’s own is too complex even to be sketched here, but the general idea is that shared referential structure may induce a significant classification of sentences, distinct from material or necessary equivalence, a priori equivalence, or any other familiar relation in modal semantics. I don’t doubt this. One might reasonably ask whether the resulting notion is really a notion of *content*. The charge that this would be a verbal quibble can backfire: after all, the extensionalist needs to convince the Fregeans that they have gotten some antecedently understood, intuitive notion of content wrong. Still, let us avoid that issue. Observe instead that Almog’s approach requires specialized assumptions. Those of us who believe in only one, our own possible world simply have no use for the (alleged) others, or for a semantic classification presupposing them. The extensionalist may reply that possible worlds only provide a convenient way of saying what can be said without them, but that remains to be shown in any particular case. (Certainly in Almog’s, where the relevant classifications depend on the stages in which semantic values are computed, first in particular worlds, then across the field of all worlds.) Further, if what is to be said without them involves metaphysical necessity, the value of saying it will remain quite unclear. Of course these remarks prove nothing. But grant that certain semantic notions, say Almog’s, are well motivated within the possible worlds framework. If we must try to cash them out by first eliminating the nonactual worlds, then finding a good philosophical interpretation of metaphysical necessity, it is hardly clear that an interesting theory of content in any nontrivial sense will survive.

(ii) One might think a “new” theory of names needs an extensionalist view of content. But what theory, and why should it have any truck with content? The idea of names as being “directly referential” might occur in this connection, but since direct referentiality is standardly explained as the identity of content with bearer, that doesn’t help. Kripkeans also argue that the bearer of a name is the object standing in a certain historical relation to its use, but whatever the virtues of this thesis, its relation to any notion of content is obscure. It concerns the determination of reference. Its formulation requires no talk of content at all. Nor is there any suggestion that if reference is determined in this way, coreferential names will share any interesting property besides their denotation. So causal-historical views on reference appear to lack even the form of views about content.

Here extensionalists may invoke the idea that names are “Milleean tags”. That idea is difficult, but I am fairly sure that no version of it will motivate an extensionalist view of content. Again it has the wrong form. (Thus the following remarks apply, *mutatis mutandis*, to causal-historical pictures.) Milleanism tries to divorce names from descriptions: perhaps by denying that any token of a name is semantically equivalent to a description, or that names have conventional, community-wide connotations. Let us grant something(s) of this sort.

Still, no useful result about coreferential names is in sight. Milleanism as just stated is, although programmatic, reasonably significant, and intelligible. And like the causal-historical picture of reference, it achieves that without mentioning content at all. Thus any step from Milleanism to claims about content would hardly be immediate. In trying to reconstruct one—a project with which the extensionalist literature helps very little—one might consider the following line of thought:

- (i) Terms (generally, often) have content, which (is something that) determines reference.
- (ii) If content isn't reference, then it is a property of kind *K*.
- (iii) But by Milleanism, names lack properties of kind *K*.
- (iv) Therefore, coreferential names have the same content.

But first, this reasoning invites the trivialization anticipated above. To find an interpretation of “content” on which it makes sense, one need look no farther than the notion of reference, which can without loss be substituted for ‘content’ in (i)–(iv). So extensionalism has still not evolved beyond the claim that referential structure fixes referential structure. Second, the extensionalist here imagines an opponent conveniently identifying content with the very property Milleanism rejects. Of course this has been done, but Fregeans need not so happily go to slaughter. If the Millean denies, say, that names have conventional connotations, Fregeans can maintain that content is not a matter of descriptions attached to names by linguistic convention. Similarly for other versions of the idea that names are semantically distinct from descriptions. The extensionalist *might* force us to grant (ii), but I know of no such argument. And even if one should appear, the trivialization would still loom. So I suggest that “new theories of reference”, however estimable they may be as theories of reference, are not theories of propositional content at all. But that is what extensionalism is all about.

Upshot: you are welcome to reject traditional ideas about semantics and belief contents. But don't do it as a favor to extensionalism.

I might add that if we do reject such ideas, the Fregean may claim a certain superiority even in defeat. For what we seem to need is a more complex, discriminating view of mental representations and their relations to our usual descriptions of them. The standard views will appear too crude. Since the Fregean view does discriminate between mental states more than extensionalism does, we can at least regard it as having been a little closer to the truth. But that is an aside. The result of our discussion is that the extensionalist can avoid the *GF* argument neither via logical opacity nor via a radical view of content. Or, more cautiously, the price of these moves is too high. Any philosophical idea involves trade-offs, which must be assessed in order to see if it is worth pursuing. I have tried to argue that the costs of certain extensionalist moves are prohibitive: if that is what it takes to be an extensionalist, then something else must be better. And I would expect some agreement on this point. Neither logical opacity nor the new view of content has until recently appealed to extensionalists. It has seemed better to allow name-substitutions in belief contexts and hope for the best—to bite the bullet on Hesperus and Phosphorus. What I have

tried to do is to close this, the principal option. The rest has been an extended rearguard action.

But we should remember our reflections on skepticism and the difficulty of refuting extensionalism. If one happily affirms a radical semantic view, what consequence will be radical enough to change one's mind? In the extreme case, even the derivation of (D2) will not suffice. Having gotten used to names, the extensionalist might find the equivalence of our *P* and *Q* no different and no worse. As I have emphasized, the *GF* argument allows that interpretation. Yet these speculations are moot. Although a determined extensionalist might not find omniscience a decisive objection, the rest of us are free to take it that way.

A notion of intension for names may now appear necessary. But since I have nothing new to offer on that score, let us briefly consider what, besides a desire to avoid intensions, underlies extensionalism.

**5 Motivation** One ground for extensionalism, or for such variants as situation semantics, is the idea that meanings ought to be truth conditions. If that is granted, the extensionalist treatment of names seems to follow, for what makes 'Twain = Clemens' true—the identity of a certain man to himself—seems no different from what makes 'Twain = Twain' true. But such considerations can bear only little weight. First, there is a difference between reasons for exploring an idea and reasons for trying to stick to it once it runs into trouble. I agree that the truth-conditional idea is initially attractive, so that it was natural to develop semantic theories based on it. But it is hardly a *compelling* intuition. We have various pretheoretic views on meaning, and among these the connection between meaning and truth conditions does not seem particularly obvious or unrevisable. Thus it would form a weak basis for the defense of a semantic theory against any serious objection. Since the objections to extensionalism and its variants are serious (whatever their ultimate merit), the intuitions in question are not very relevant. Second, these intuitions will not even take us as far as it is claimed. The notion of a truth condition is ambiguous: asked to state the conditions under which 'Twain = Clemens' is true, I could equally well say that it holds when what 'Twain' denotes is (or: bears the relation expressed by '=' to) what 'Clemens' denotes. If I then give the analogous answer for 'Twain = Twain'—that the referent of the first occurrence of 'Twain' must be identical to that of the second—then the sentences have received distinct truth conditions, and the extensionalist idea does not get started. Yet this way of responding to the demand for truth conditions is intuitively just as good as the extensionalist's. This does not, of course, refute the construal embodied in the extensionalist theory of meaning, but it does show that pretheoretic intuitions are too ambiguous to be useful here. The extensionalist must look elsewhere for support.

Another line of thought starts from versions of the 'Twain'-'Clemens' puzzle in which a speaker has two names for an object for which we only have one. The *general* idea is that they show that *no* account of belief can be intuitively right. Consider the ancient astronomer with two names for Venus. If he does not think they codesignate, rather plausible considerations may, depending on the circumstances, recommend unpalatable belief attributions, e.g., belief in

(23) Venus is rising and Venus is not rising.

Some account of the situation is needed, and the extensionalist line is said to offer advantages while violating intuition as little as possible.<sup>12</sup> Now my *GF* argument overrides any such reasoning: if I am right, extensionalism violates intuition as much as any serious semantic theory could. That aside, the argument from puzzles about names was already preempted in [7].<sup>13</sup> The idea behind the argument is that we may be in a quandry over how to represent someone's attitudes with our own sentences if his beliefs differ from ours—yet our own sentences, within our own system of belief, are all we can use. For example, the astronomer will have beliefs of the forms

(24) *a* is rising

and

(25) *b* is not rising

where '*a*' and '*b*' both name Venus. But suppose 'Venus' is *our* only name for the planet. If we are to state the astronomer's belief at all, we must apparently say that he believes that Venus is rising and that Venus is not rising, which leads to problems. Field points out, however, that simple representations of belief need not be available. Sometimes we cannot match a subject's sentence with an even roughly synonymous one of our own. Rather, we say something like, "This man has two names for Venus, but he doesn't know that they codesignate. He thinks one stands for a star seen in the morning, one for a star seen in the evening, and he doesn't know that he's seeing the same thing each time. Now he thinks that one but not the other star is rising, and he expresses this view using the names as follows [here we might give the sentences in his language, with appropriate glosses] . . ." This is (an informal version of) how we in fact represent many alien views. Statements of the '*S* believes that *P*' form are often inadequate. This too violates some intuitions about meaning (that is part of what makes it interesting) but we have to live with it in any case. So the puzzle of the astronomer is also no use to extensionalism.

Besides, one would hope to base a general approach to semantics on more than a puzzle. This is possible for the principal alternatives (e.g., Frege's or Quine's) and I think also for extensionalism. The roots of extensionalism lie in a general view of reference and meaning. They are best approached via the behaviorism and empiricism that have shaped much of our philosophy of language. The following sketch of the connections, incomplete and inadequate as it must be, may do for a start.

A behaviorist might identify the meaning of a sentence *S* with the environmental circumstances that tend to elicit *S*, or in which *S* would tend to be affirmed. Thus, the content of 'here are two cats' would be something like the set of circumstances that involve perceptual presentations of a pair of cats. One problem with this is the absence of plausible circumstances of this kind for a wide class of sentences. Affirmation of, e.g., 'some people have two cats' can hardly be associated with any particular perceptual situations.<sup>14</sup> A possible response to this difficulty would make the semantics compositional (recursive) and retain a straightforward behaviorism only for the basis of the recursion, or more likely just for part of it. One would associate 'cat' with presentations of cats, 'black' with presentations of black things, and so forth. The meanings of

complex phrases, and of sentences, would then be functions of these elementary meanings, plus those of words that are syncategorematic within this framework. The first three words in 'This is a black cat' or 'Here are two cats', for example, would not receive empirical/behavioral meanings of their own, but would produce sentence meanings by operating on the meanings of the words around them. Of course this is very vague as it stands. The nature of word meaning is barely sketched, the principles of composition are left entirely obscure. No hint of a treatment of intuitively nonobservational concepts is given. Yet the outline of an approach is recognizable here; one may hope for ways around the many obstacles. And extensionalism flows from this perspective. A semantically unstructured name would plausibly be associated with presentations of its bearer. Thus, two such names would receive the same association and the same meaning. Similarly for simple predicates expressing the same observational property: each would be associated with it and, hence, would count as making the same contribution to complex meanings. The underlying idea by no means forces this development, but it is certainly natural.

Although other reconstructions of extensionalism are certainly possible, I believe mine illuminates its present incarnation. Traditional forms of behaviorism can no longer be taken seriously; but current views of reference derive from them and their empiricist variants.<sup>15</sup> A causal theory, for example, will identify the referent of a term *t* with something in the environment bearing a certain historical relation to the use of *t*. This relation requires an appropriate causal link between speaker and referent; the causal theorist's paradigm of such a link is perceptual acquaintance. (For speakers lacking acquaintance with the referent, one tells a story about the extension of the original link through a chain of speakers.) Coreferring names are then naturally assigned the same content. For what all users of 'Mark Twain' (with its usual referent) share is the causal origin of the name they use. The particular chains linking Twain to different speakers are idiosyncratic and seem irrelevant to the name's common content. But this origin is the same for 'Samuel Clemens' as for 'Mark Twain', so these seem to share content as well as reference. These considerations are certainly no rigorous argument, but they do seem to have encouraged the widespread co-occurrence of extensionalist and causal views. And in spite of important differences, their analogy to the empiricist identification of content with a perceptual stimulus, or class of stimuli, is clear. Further, some extensionalists explicitly link their semantics to a form of this picture; e.g., [4], [14].

Indeed, the leading approach to problems of meaning and representation is even more clearly empiricist than the (closely related) causal theories. Dretske [6] has best articulated what is sometimes called "indicationalism" or "the thermometer view": roughly, that a mental state *S*, or its linguistic expression, represents whatever external state the presence of *S* most reliably indicates. Dretske casts this idea in terms of information theory, arguing plausibly that causal intuitions are thereby upheld without some of the drawbacks of causal theories. And as far as Dretske sketches his account, it closely follows the behaviorist paradigm above. He bases a theory of meaning for simple concepts or predicates, such as 'red' or 'bluebird', on the notion of informational content. A label in a system, which might as easily be a computer or simple animal as a human, means 'bluebird' if its application carries the information that a bluebird is pres-

ent: if the objective probability of a bluebird's presence, given that the label is applied, is higher than that of the alternatives in some reasonable range. This idea is developed in fair detail for cases in which the "application of the label" is something like the activation of a frog's optical bug detector. And in such cases, the behaviorist view of meaning for elementary labels is essentially Dretske's. Not surprisingly, Dretske is extensionalist as far as he goes: if label *L* indicates the presence of feature *F*, and if the coextensiveness of *F* with *G* is causally necessary, then *L* likewise indicates the presence of *G*—the relevant probabilities are identical. Similarly, two (unstructured) labels indicating the presence of the same property have the same informational value, hence content. Causally individuated properties are the referents of predicates, and we have identity of content for coreferring predicates. Like our imagined behaviorist, Dretske expresses the hope that this account of simple cases can be extended to a comprehensive theory.

Dretske's position has numerous cousins (due, e.g., to Dennett, Fodor, Loar, Putnam, and Stampe). As remarked, these together constitute the main current approach to representation. We have seen how they may suggest extensionalism. Although many indicationalists might avoid that step, the extensionalist vogue is readily comprehensible in this light. And this means that extensionalism may survive any technical refutation (such as I have attempted). As long as some form of indicationalism seems right, maybe even inevitable, the motivation to make semantics extensional will be strong. And where the motive to overcome technical problems is strong enough, technical solutions will be contrived. Those who find my objections convincing therefore have a choice. They must either sever the link between indicationalism and extensionalism or undermine indicationalism itself. The former line is undoubtedly easier. One should be able to maintain a general view about how language relates to the world without commitment to a specific form of semantics. Even so, I would, as someone convinced by my objections, concentrate on the second line [18], [20]. I think indicationalism is radically wrong; giving it up will not only remove any extensionalist temptation but also lead to better views on realism and truth. Neither option, however, can be explored here. For present purposes, it suffices to have sketched some connections between semantic theory and such broad ideologies as empiricism and behaviorism. Observe that our earlier remarks about extensionalism and skepticism fit in well. Trying to find a skeptical undercurrent in extensionalism might at first look artificial or strained. But the skeptical ties of empiricism are familiar. Once the extensionalist's empiricist affinities are made clear, it is less surprising if skeptical tendencies turn up as well. If we are generally inclined to reject skepticism, this paper suggests a reassuring fit between inclination and technical fact.

## NOTES

1. "Intensional objects", such as properties, do not make the semantics intensional. If we assign them as referents to predicates and give coreferential predicates the same meaning, then the semantics for predicates is extensional in spite of the choice of objects.



2. Extensionalism hails largely from Stanford, Westwood, and nearby places; hence “California semantics”.
3. Actually, I can imagine reconciling (13) with (14): find a hidden relativization in belief reports, perhaps to the reporter’s purposes or to some class of relevant alternative reports. ‘Believes that *P*’ would then be somewhat analogous to ‘is large’ or ‘is like Jane’, both of which can be true of Mary in one context and false in another. This idea is worth pursuing but will not help the extensionalist here. Mental states are presumably determinate; if belief reports do not describe them in absolute terms, then they are even less reliable guides to our thoughts than I am assuming in the present section—and than the extensionalist must assume to make the argument from conversation work.
4. So I interpret Quine as discussing what we will *count as true*, or accept, in a context, not which attributions of intentional states really are true—not a main concern of his, anyway.
5. [3] examines the argument. Barwise and Perry call it the “slingshot”, due to the power of this little argument to fell mighty semantic theories. Tim McCarthy suggests that if my following reasoning is sound, they should have chosen “boomerang”.
6. This is a weak omniscience thesis: God knows every truth, not just any one he pleases. But it is the appropriate thesis for computationally finite beings. As long as our memories (and other faculties) are limited, we cannot have divine omniscience anyway. My argument gives us the ability to know all the truth our heads will hold, which is all we can ask.
7. Failure to distinguish (18′) from (18) has brought difficulties for many writers, including Frege (see [17]). It is also relevant to puzzles about the (alleged) contingent a priori (e.g., [5]).
8. Strictly, *S* also needs a (17′) corresponding to (17), but since she believes ‘*P*’ anyway, there is no problem about that.
9. Whether (B26) follows from (B25) alone depends on the truth about vacuous names. But if one side of a true identity denotes, as guaranteed by (B23), then so must the other side.
10. Salmon envisions a complex view which I make no attempt to summarize. I will consider just what seem to me to be its general implications.
11. Recall that in addition, the extensionalist will be stuck with giving us all a potential omniscience inhibited only by our logical failings: if we can reach (B24), the rest follows by logic plus easy semantics. I am inclined to think this is a problem but cannot go into it here.
12. This is only a bare hint of how the extensionalist will try to exploit the case of the astronomer. [13] gives a proper account.
13. Field’s article appeared before the argument really got started (e.g., in [10]). And as Field himself notes, he is simply following Quine’s principle [11] that changes of belief bring changes of meaning.
14. [11] is the classic attempt to work out the behaviorist’s problems. The response in the text is definitely not Quine’s. [19] suggests a concise account of his difficulties.
15. I won’t try to distinguish empiricist from behaviorist theories of meaning here.

## REFERENCES

- [1] Almog, J., "Form and content," *Nous*, vol. 19 (1985), pp. 603–616.
- [2] Almog, J., "Naming without necessity," *Journal of Philosophy*, vol. 83 (1986), pp. 210–242.
- [3] Barwise, J. and J. Perry, "Semantic innocence and uncompromising situations," pp. 387–403 in *Midwest Studies in Philosophy*, eds. P. French, T. Uehling, Jr., and H. Wettstein, vol. 6, University of Minnesota Press, Morris, Minnesota, 1981.
- [4] Barwise, J. and J. Perry, *Situations and Attitudes*, M.I.T. Press, Cambridge, 1983.
- [5] Donnellan, K., "The contingent *a priori* and rigid designators," pp. 12–27 in *Midwest Studies in Philosophy*, eds. P. French, T. Uehling, Jr., and H. Wettstein, vol. 2, University of Minnesota Press, Morris, Minnesota, 1977.
- [6] Dreske, F., *Knowledge and the Flow of Information*, M.I.T. Press, 1981.
- [7] Field, H., "Logic, meaning, and conceptual role," *Journal of Philosophy*, vol. 74 (1977), pp. 379–409.
- [8] Grandy, R., "Some misconceptions about belief," in *Philosophical Grounds of Rationality: Intentions, Categories, Ends*, eds. R. Grandy and R. Warner, Oxford University Press, Oxford, 1986.
- [9] Kripke, S., *Naming and Necessity*, Harvard University Press, Cambridge, 1980.
- [10] Kripke, S., "A puzzle about belief," in *Meaning and Use*, ed., A. Margalit, Reidel, Dordrecht, 1979.
- [11] Quine, W. V., *Word & Object*, M.I.T. Press, Cambridge, 1960.
- [12] Richard, M., "Direct reference and ascriptions of belief," *Journal of Philosophical Logic*, vol. 12 (1983), pp. 425–452.
- [13] Salmon, N., *Frege's Puzzle*, M.I.T. Press, Cambridge, 1986.
- [14] Soames, S., "Direct reference and propositional attitudes," forthcoming in a volume edited by Salmon and Soames.
- [15] Stalnaker, R., "Propositions," in *Issues in the Philosophy of Language*, eds., A. F. MacKay and D. D. Merrill, Yale University Press, New Haven, Connecticut, 1976.
- [16] Stalnaker, R., *Inquiry*, M.I.T. Press, Cambridge, 1984.
- [17] Wagner, S., "Tonk," *Notre Dame Journal of Formal Logic*, vol. 22 (1981), pp. 289–300.
- [18] Wagner, S., "A philosopher looks at artificial intelligence," in *Perspectives in Cognitive Science*, eds. D. L. Farwell, S. C. Helmreich and W. D. Wallace, Linguistics Students Organization, University of Illinois, Urbana, 1983.
- [19] Wagner, S., "Quine's holism," *Analysis*, vol. 46 (1986), pp. 1–6.
- [20] Wagner, S., *Truth, Platonism, and Ultimate Theory*, in progress.
- [21] Wettstein, H., "Has semantics rested on a mistake?" *Journal of Philosophy*, vol. 83 (1986), pp. 185–209.