

Reasoning and Representing

MARK ELI KALDERON

m.kalderon@ucl.ac.uk

Up to now, in constructing a language, the procedure has usually been, first to assign a meaning to the fundamental mathematico-logical symbols, and then to consider what sentences and inferences are seen to be correct in accordance with this meaning...The connection [between meaning and inference] will only become clear when approached from the opposite direction: let any postulates and any rules of inference be chosen arbitrarily; then this choice, whatever it may be, will determine what meaning is to be assigned to the fundamental logical symbols.

Rudolph Carnap

The Logical Syntax of Language

0 Introduction

According to an inferential role semantics, the content of a logical expression is determined by (some aspect of) its inferential role. The *inferential role* of an expression is a competent speaker's disposition to accept the validity of a distinguished set of inferences involving that expression. The insight, if it is one, is that in a speaker's reasoning in accordance with the validity of certain inferences, the logical expression comes to have whatever content is required for their validity.¹

There are two ways in which inferential role semantics is an unfortunate label. The term *inference* is ambiguous. In one sense of the term it refers to an abstract pattern of argument that is either valid or invalid. Specifically, to say that an inference is valid is to say that (the proposition expressed by) the conclusion is a logical consequence of (the propositions expressed by) the premises. There is another sense of the term according to which an inference is something we *do* as opposed to an abstract relation among sentences (and the propositions that they express). It is natural, in this sense, to describe the process

1. Contemporary advocates of inferential role semantics include, *inter alia*, Block (1986,1993), Boghossian (1993), Brandom (1994), Field (1977), Harman (1986a,1986b), and Peacocke (1976,1987,1992). In what follows, I shall not distinguish between meaning and content—none of the circumstances that requires one to distinguish these notions, such as context sensitivity, obtain with respect to logical discourse. **Boldface** will be used throughout as a device of quasi-quotation.

of belief fixation as performing an inference.² In what follows, I shall be using *inference* in the former but not the latter sense. Another way in which the label is unfortunate is that *semantics* suggests that it is a theory that pairs expressions with semantic values, but an inferential role semantics is an account of what determines this pairing. Though it is not the best of names, the use of inferential role semantics is now well entrenched, and I shall follow that usage here.

I shall argue that given the best explanation of logical understanding, the content of a logical expression must be determined by its inferential role. This might strike some as surprising. Why believe logical understanding to be explicable at all? Why isn't logical understanding *primitive*, *i.e.*, not subject to any further explanation? If logical understanding were indeed primitive, then this line of argument would be foreclosed. There are two possible sources for the conviction that logical understanding is primitive. The first, and least reputable, arises from a general skepticism about the ability of philosophy to explain anything. *Quietism*, as such skepticism has come to be known, would preclude an explanation of logical understanding, for there would be no explanation of anything of philosophical interest. For present purposes, we may dispense with such skepticism. The second and more interesting source of the conviction that logical understanding is primitive involves an initially plausible assumption concerning the *form* any such explanation must take. This assumption is that logical understanding should be explained in terms of propositional knowledge. With the primitivist, I maintain and shall argue that logical understanding *is* inexplicable in terms of propositional knowledge. But against the primitivist, I shall argue that this is not the only available form of explanation. Logical understanding is explicable not in terms of propositional knowledge but, rather, in terms of practical knowledge—not knowledge *that*, but knowledge *how*. I shall argue that once logical understanding is conceived as a species of practical knowledge, then the truth of some version of inferential role semantics is a consequence.

I am not arguing for any particular version of inferential role semantics, but rather, for the class of such theories. If my argument is cogent, then some version of inferential role semantics must be correct. But I do not take a stand on which version that is. Moreover, I am not arguing that inferential role semantics is a correct account of content generally. I am only arguing that some version of inferential role semantics must be the correct account of the

2. For this reason Harman (1986) has reserved the term *implication* for the relation of logical consequence among sentences (and the propositions that they express)—though a similar difficulty faces Harman's own choice of terminology. There is a corresponding sense of *implication* according to which implying that S is something that a speaker does. What's important, however, is that the relevant distinction be clearly marked, not what labels we attach to it.

content of logical vocabulary.³

1 Content and Inferential Role

What makes it the case that a logical expression has the content that it does? What, for example, determines the content of $\exists x$? Certainly not the orthographic properties of the expression—after all there is nothing semantically significant about a backwards E (unless, of course, it's being used as an existential quantifier). So what determines the content of existential quantification?

According to the inferentialist, the content-determining feature of a logical expression is (some aspect of) its *inferential role*, i.e., the disposition of competent speakers to accept the validity of certain inferences involving that expression.⁴ Thus, for example, the content of $\exists x$ is at least partly determined by the disposition to accept as valid inferences of the form:

$$S(\dots n \dots) \quad \exists x S(\dots x \dots)$$

I say “at least partly” since the disposition to accept the validity of other forms of inference might also be content-determining. Similarly, the inferentialist might claim that \vee functions semantically as disjunction partly in virtue the disposition to accept as valid inferences of the form:

$$S \quad S \vee S^*$$

Let M be the distinguished set of inferences involving a logical expression L that

3. The discussion will be restricted to the logical vocabulary of *complete* logics. I believe that inferential role semantics can be extended to, say, *full* first- and second-order logic and modal logic but special difficulties face an inferential role semantics for these. Thus, for example, Skolem's paradox can be transformed into an argument that an inferential role semantics for second-order logic is committed to the semantic indeterminacy of our second-order idiom. (Suppose that the content of second-order quantification is fixed by the axioms of second-order logic and the rules of inference. Very roughly, the problem is that this account of the content of second-order quantification is consistent with their being subject to not only the standard interpretation but also to a non-standard interpretation.) And Field (1989, pp. 278–9) has argued that a similar semantic indeterminacy affects our modal discourse. As interesting as these issues are, they are beyond the scope of the present essay.

4. Harman (1986a, 1986b) has argued that the disposition to accept the validity of certain inferences is insufficient to determine logical content. The content of negation, say, is determined as well by competent speakers' disposition to accept the inconsistency of certain sets of sentences. There are thus two components to an expression's inferential role. However, as nothing in the present discussion depends on this complexity, I shall ignore it.

competent speakers are disposed to accept as valid (where *M* may or may not be a proper subset of valid inferences involving *L*). The inferentialist believes that in a speaker's accepting the validity of the *M*-inferences *L* comes to have whatever content is required for their validity.

The inferentialist is committed to, at a minimum, the following supervenience thesis:

The Supervenience Thesis: The content of a logical expression *L* supervenes on its inferential role. Specifically, for any metaphysically possible worlds *w* and *w** and any logical expressions *L* and *L** that are tokened in *w* and *w** respectively, if *L* and *L** occur in the same distinguished set *M* of inferences that speakers competent with these expressions are disposed to accept as valid (in their respective worlds), then *L* has the same content in *w* as *L** in *w**.⁵

The supervenience thesis is not by itself a species of reductionism, at least not in the strong sense of proposing a substantive identification. It claims that logical content supervenes on inferential role, not that logical content and inferential role should be identified. It is consistent with the supervenience thesis that the content of a logical expression and the disposition to accept the validity of the *M*-inferences be distinct, if necessarily coextensive, properties.

The supervenience thesis is fairly weak—it is merely the claim that there could be no independent variation of logical content and inferential role. I have claimed that, at a minimum, the inferentialist must accept the supervenience thesis. However, the inferentialist also claims *more than this*. The core commitment of inferential role semantics is not just that the content of a logical expression supervenes on its inferential role, but that it has its content *because* of the inferential dispositions of speakers competent with that expression. Inferential role semantics essentially involves a determination thesis that has the supervenience thesis as a consequence:

IRS: The content of a logical expression is determined by its inferential role. Specifically, the fact that a speaker accepts the validity of a distinguished set *M* of

5. Let me explain what I mean by “the same distinguished set *M* of inferences.” For present purposes let inferences be finite sequences of sentences. The initial segment of a sequence consists of the premises of the inference and the last member of the sequence is its conclusion. Two distinct expressions occurring in inferences that competent speakers are disposed to accept as valid cannot participate in *identical* sets of inferences. Since the expressions are distinct, so are the sentences of which they are a part, and thus the inferential sequences in which they occur will themselves be distinct. That *L* and *L** occur in *the same set M* cannot mean that they occur in identical sets of distinguished inferences, rather only that speaker's competent with these expressions are disposed to accept as valid the *same forms* of inference.

inferences involving L, makes it the case that L has the content required for the validity of the M-inferences.

Inferential role semantics is usefully thought of as a variant of *implicit definition*—with validity instead of truth constraining the content of the *definiendum*. Consider the introduction of a novel expression *e* by means of implicit definition (as normally understood). Let A be a distinguished set of sentences involving *e*. The A-sentences are stipulated to be true and, as a result of such stipulation, competent speakers become disposed to regard them as true. In being disposed to hold true the A-sentences, *e* comes to have the content required for the truth of these sentences. In claiming that the disposition to hold true the A-sentences implicitly defines *e*'s content, the advocate of implicit definition needn't further claim that the disposition makes the A-sentences true. He can consistently claim that while the disposition to hold true the A-sentences plays a content-determining role, their truthmakers constitute an independent domain of fact. Suppose this procedure did in fact secure a content for *e*.⁶ Notice *e*'s content would then supervene on the disposition of competent speakers to hold true the A-sentences. Had they come to regard a distinct set of sentences involving *e* as true, *e* would have had a distinct content—the content required by the truth of those sentences. It is important to emphasize that the advocate of implicit definition isn't necessarily committed to a historical claim about how *e* was actually introduced into the language. Typically, they are only committed to the weaker claim that *e* in fact has the content it would have if it were introduced by means of implicit definition.⁷

IRS differs from paradigm cases of implicit definition in that what determines the content of the *definiendum* is the *validity* of a set of inferences, rather than the *truth* of a set

6. There is no general guarantee that such a procedure will in fact secure a content for *e*. The content of the *definiendum* is supposed to be the one and only one content such that the A-sentences come out true. There may be more than one such content, or there may be no way to make the A-sentences jointly true.

7. Thus Lewis (1972/1980, pp. 212–3), having claimed that psychological idiom is implicitly defined by its role in folk psychology, writes as follows:

Imagine our ancestors first speaking of external things, stimuli, and responses—and perhaps producing what we, but not they, may call *Ausserungen* of mental states—until some genius invented the theory of mental states, with its newly introduced T-terms, to explain the regularities among stimuli and responses. But this did not happen. Our common sense psychology was never a newly invented term-introducing scientific theory—not even in prehistoric folk-science. The story that mental terms were introduced as theoretical terms is a myth. It is, in fact, Sellars' myth of our Rylean ancestors. And though it is a myth, it may be a good myth or a bad myth. *It is a good myth if our names of mental states do in fact mean just what they would mean if the myth were true.* [my emphasis]

of sentences. The fact that a speaker accepts the validity of the M-inferences, makes it the case that L has the content required for their validity. The inferentialist, in claiming that logical content is determined by the disposition to accept the validity of the M-inferences, is not claiming that the disposition to accept the M-inferences *makes them valid*. The inferentialist can consistently claim that while a speaker's disposition to accept the validity of certain inferences plays a content-determining role, their validity is an independently constituted property. Just as e's content supervenes on competent speakers' disposition to hold true certain sentences, so the inferentialist maintains that L's content supervenes on competent speakers' disposition to accept the validity of certain inferences. Had speakers competent with L been disposed to accept the validity of a distinct set of inferences, L would have had a distinct content—the content required by the validity of those inferences.⁸ Moreover, just as the advocate of implicit definition isn't (necessarily) committed to any historical claim about how the *definiendum* was in fact introduced into the language, the inferentialist needn't be committed to any similar historical claim.

There is a further parallel between inferential role semantics and implicit definition as traditionally conceived. The advocate of implicit definition merely claims that the content of the *definiendum* is determined by the disposition to hold true certain sentences. He needn't make the further claim that the content of the *definiendum* is identical to the linguistic disposition that implicitly defines it. There is no obvious inconsistency in believing, say, that an expression's content is implicitly defined and that contents are mind- and language-independent *abstracta*. Notice if contents are mind- and language-independent *abstracta*, they are necessarily distinct from linguistic dispositions that implicitly define them. The claim that an expression is implicitly defined is a claim about what determines its content and not a claim about what contents are. Similarly, the inferentialist merely claims that the content of a logical expression is determined by its inferential role. He needn't make the further claim that logical contents are identical to inferential roles.⁹

8. IRS and the supervenience thesis are thus distinct claims. One can coherently believe the supervenience thesis while maintaining that logical expressions have their inferential roles *because* of their content. Belief in supervenience is consistent even with the denial that there is an explanatory asymmetry between content and inferential role. Thus while IRS entails the supervenience thesis, the converse entailment fails.

9. The analogy with implicit definition, though useful, is imperfect—for it obscures two ways in which an inferential role semantics might be developed. The first conforms closely to the model of implicit definition. Just as an implicitly defined expression has the descriptive content assigned to it by its implicit definer, the inferentialist might also claim that logical expressions have descriptive contents. Roughly speaking, a logical expression L would have the descriptive content: the logical object that makes the M-inferences valid (where *logical objects* are truth functions, quantifiers, and the like). Thus logical expressions with distinct inferential roles would have distinct descriptive contents even if one and the same logical object satisfies the descriptive

The inferentialist claims the content of a logical expression is determined by and hence supervenes on the inferential dispositions of speakers who are competent with that expression. These claims are only as clear as our understanding of the relevant subvenient disposition. While I am inclined to believe that the relevant disposition is analytically primitive, it will be useful to consider what the relevant disposition is not.

I doubt that the disposition to accept the validity of an inference is analytically reducible to any *non-intentional* property. Consider an inference that a speaker is disposed to accept as valid. Let **P** be the conjunction of its premises, and let **C** be its conclusion. The speaker's disposition to accept the validity of the inference is not, for example, the verbal disposition to utter **C** if the speaker is disposed to utter **P**. The possession of this non-intentional disposition is neither necessary nor sufficient for being disposed to accept the validity of the inference. It is not necessary since a competent speaker may accept the

conditions. Conversely, logical expressions with identical inferential roles would have identical contents. According to this Fregean variant of inferential role semantics, inferential role *individuates* logical content.

There is, however, a Millian variant of inferential role semantics that disavows this commitment. According to the Millian variant, a logical expression **L** would not have the descriptive content: the logical object that makes the **M**-inferences valid; rather, **L**'s content is exhausted by the logical object that satisfies this description. The descriptive conditions determined by an expression's inferential role are not *synonymy-supplying*, but *reference-fixing*. The inferentialist who accepts the Millian variant of inferential role semantics is not in any way committed to the claim that inferential role individuates logical content. Expressions with distinct inferential roles can have the same content if the descriptive conditions determined by their inferential roles are satisfied by one and the same logical object. Moreover expressions with identical inferential roles need not have identical contents. Suppose the following pair of expressions have identical inferential roles:

&
 $\neg((\neg\dots) \vee (\neg\dots))$

Suppose further that the inferential role of **&** determines a descriptive condition that picks out the truth function associated with conjunction, **CONJ**. Suppose as well that the inferential roles of \neg and \vee determine descriptive conditions that pick out the truth functions associated with negation, **NEG**, and disjunction, **DISJ**, respectively. Thus whereas the content of **&** would be the semantically unstructured two-place truth-function:

CONJ(.....),

given compositionality, the content of $\neg((\neg\dots) \vee (\neg\dots))$ would be the semantically structured two-place truth-function:

NEG(DISJ(NEG(...),NEG(...))).

These expressions would thus have distinct, if logically equivalent, contents despite having identical inferential roles.

Despite its implausibility, the Fregean variant is not without its adherents (see for example Peacocke, 1987). Much of the discussion in the literature, however, seems to unwittingly vacillate between these distinct positions.

validity of the inference and even be disposed to assertively utter its premises without wishing to advertise his conclusion. Nor is it sufficient. A speaker may possess the relevant verbal disposition while meaning nothing by **P** and **C**. Under such circumstances, it would be wrong, solely on this basis, to interpret the speaker as accepting the validity of the given inference. (Even a *Furby* could possess the relevant verbal disposition.) I suspect that similar difficulties will attend any attempt to analytically reduce the relevant disposition to a non-intentional property.

Not only is the disposition to accept the validity of an inference not the verbal disposition to utter **C** if one is disposed to utter **P**, neither is it the disposition to infer that **C** if one believes that **P**. (Notice that a speaker's *inferring that C* and *believing that P* are propositional attitudes and hence that this disposition, unlike the verbal disposition, is not specified in non-intentional terms.) A speaker may believe that **P** and accept the validity of the inference and yet not come to believe that **C**. The proposition that **C** may, after all, be unreasonable given the speaker's background beliefs, in which case the speaker may be disposed to disavow the belief that **P** (or, at the very least, suspend judgement concerning it).¹⁰

Notice, whether or not the speaker comes to believe the conclusion or disavow (or suspend judgement concerning) one of the premises, his disposition to accept the validity of the inference is manifest in the fixation of the relevant propositional attitudes. The disposition to accept the validity of the inference mediated, if not determined, the rational fixation of belief. While the speaker's acceptance of the validity of the inference did not *rationaly oblige* him to believe the conclusion if he believed the premises, it did *rationaly prohibit* prohibit him from believing the premise and denying the conclusion. Being disposed to accept the validity of an inference consists in the ability to reason in accordance with its validity. Given the (albeit, limited) role it plays in the rational fixation of belief, the relevant subvening disposition should be understood as an intentional and indeed normative property (which is not, of course, to say that intentional/normative properties are not themselves *metaphysically* reducible to the non-intentional.)

Notice that maintaining the subvenient disposition to be an intentional property in no way trivializes IRS nor renders it circular as an account of the determinants of logical content. Although the disposition to accept the validity of an inference is an intentional property, the account is still non-circular, so long as the disposition can be specified without

10. The present example highlights the way in which the two senses of *inference* can come apart.

employing the content property.¹¹

Not only is the disposition to accept the validity of an inference not the disposition to infer the conclusion if the premises are believed, neither is it the metalinguistic disposition to believe the inference to be valid. A competent speaker may be disposed to accept the validity of an inference without being disposed to believe it to be valid for the simple reason that he fails to grasp the content of validity and so is incapable of forming the relevant belief.

Just as I suspect there to be no analytic reduction of the subvenient disposition framed in non-intentional idiom, I also suspect that there is no analytic reduction of the notion in intentional terms. For present purposes, I shall regard the disposition to accept the validity of an inference as an analytically primitive, intentional property. I have no objection to a metaphysical reduction of this notion if such a reduction is indeed possible. However, I strongly suspect that no such reduction is forthcoming.

2 Inferential Role and Objectual Truth-Conditions

Call someone who claims, *contra* the inferentialist, that content determines inferential role, a *representationalist*. The representationalist claims that it is *because* a logical expression has the content that it does, that speakers competent with that expression are disposed to accept the validity of certain inferences. Against the inferentialist he maintains that what is constitutive of existential quantification is *not* the fact that competent speakers are disposed to accept the validity of certain inferences; rather, what is essential to existential quantification is the objectual truth-conditions of sentences employing such quantifiers. Specifically, the occurrence of $\exists x$ in the sentence, $\exists xS(\dots x\dots)$, functions semantically as an existential quantifier because its content requires the extension of the sentential matrix to be non-empty.

Notice the inferentialist needn't deny that existential sentences have objectual truth-conditions—he is merely committed to the weaker claim that *if* existential sentences have objectual truth-conditions, they have them *in virtue of* the inferential role of the constituent

11. Moreover, since the disposition to accept the validity of an inference is an intentional/normative property, it is at least unobvious that IRS is subject to Kripke's (1982) normativity objection to dispositional reductions of content. Kripke's Wittgenstein argues that since our non-intentional dispositions to apply an expression fail to determine that we would be justified in so applying it, facts about its content cannot be identified with any facts about non-intentional dispositions. For a sympathetic reconstruction of the argument see Boghossian (1989). Some philosophers have argued that the argument establishes, at best, an epistemological conclusion: that non-intentional facts fail to determine the intentional facts only in the sense that facts about content are not *a priori* consequences of non-intentional facts. For discussion see Byrne (1993, ch. 4), Horwich (1995), and Soames (1997).

existential quantifiers. I suspect that the temptation to suppose otherwise is do to not clearly distinguishing the determination thesis—that logical content is determined by inferential role—from the stronger reductive thesis—that logical content is identical to inferential role. If logical contents were inferential dispositions, then they would be distinct from objectual truth-conditions. Truth-conditions are constraints on the way the world must be if a sentence or proposition is true. Since truth-conditions essentially involve relations to extra-linguistic reality, they cannot be identified with anything purely intralinguistic. Thus one might conclude that inferentialist is committed to denying that existential sentences having objectual truth-conditions. However, as I have emphasized, inferential role doesn't have *to be* logical content in order *to determine* it. As such, the inferentialist needn't (for this reason) deny that existential sentences have objectual truth-conditions.

The representationalist might have a further worry. The representationalist might doubt that inferential role suffices to determine objectual truth-conditions. Consider some logical operator **L**. If the truth of **LxS(...x...)** doesn't require the extension of **S(...x...)** to be non-empty, then **L** is not an existential quantifier—even if **L** has precisely the same inferential role as legitimate existential idiom. If there could be an operator that had the same inferential role as an existential quantifier but differed in content, then the explanatory pretensions of inferential role semantics are empty.¹²

More explicitly, the representationalist might argue by *reductio ad absurdum*:

- (1) **L** functions semantically as an existential quantifier (by hypothesis).
- (2) *The Supervenience Thesis*: The content of a logical expression **L** supervenes on its inferential role. Specifically, for any metaphysically possible worlds *w* and *w** and any logical expressions **L** and **L*** that are tokened in *w* and *w**

12. Certain forms of ontological reduction are committed to the envisioned possibility. Consider the existential sentences of mathematics, say, **There is at least one prime number less than 17**. It is plausible that the apparent existential sentences of mathematics have the same inferential role as genuine existential sentences. Just as competent speakers are disposed to accept as valid the inference from **Paul is a skeptic** to **There is at least one skeptic**, so too are they disposed to accept as valid the inference from **3 is a prime number less than 17** to **There is at least one prime number less than 17**. A speaker who failed to accept the validity of this form of inference would exhibit his incompetence with the apparent existential idiom of mathematics just as he would in the non-mathematical case. Certain nominalists claim that, despite having the same inferential role as legitimate existential idiom, the seeming existential assertions of mathematics are not ontologically committed to the existence of mathematical entities—their truth-conditions require the existence nothing. If this were correct, then the apparent existential sentences of mathematics would lack the objectual truth-conditions of genuine existential sentences despite the identity of their inferential roles. Anyone who accepts that apparent existential quantification over mathematical objects has the same inferential role as genuine existential quantification, but maintains that mathematical truth requires the existence of nothing are committed to the relevant possibility. For further related discussion see Kalderon (1996), Rosen (1993), and Wright (1983).

respectively, if **L** and **L*** occur in the same distinguished set **M** of inferences that speakers competent with these expressions are disposed to accept as valid (in their respective worlds), then **L** has the same content in *w* as **L*** in *w**.

- (3) If **L** and **L*** occur in the same set **M** of inferences that a competent speaker is disposed to accept as valid, then **L** and **L*** must be synonymous or share the same content (by (2)).
- (4) *The Standard Explanation*: **L** functions semantically as an existential quantifier iff: **LxS(...x...)** is true iff there is an individual in the range of objectual variable **x** that is in the extension of **S(...x...)**.
- (5) There could be a logical operator **L*** such that (i) **L** and **L*** would occur in the same set **M** of inferences that competent speakers are disposed to accept as valid but (ii) **L*** would lack the objectual truth-conditions described by the standard explanation.
- (6) Since, by (5), **L** and **L*** would occur in the same set **M** of inferences that competent speakers are disposed to accept as valid, then, by (3), **L*** and **L** would be synonymous or share the same content.
- (7) **L*** would function semantically as an existential quantifier (by (1) and (6)).
- (8) **L*** would not function semantically as an existential quantifier (by the standard explanation and (5)).

We have a contradiction, and the representationalist recommends that we avoid it by rejecting the supervenience thesis. And since the supervenience thesis is a consequence of IRS, we should reject IRS as well.

It is initially tempting to suppose that the difficulty for the inferentialist lies with premise (4)—the insistence that legitimate existential sentences have objectual truth-conditions. However, as I observed earlier, the inferentialist is not committed to denying that existential sentences have objectual truth-conditions, rather, only to the weaker claim that if existential sentences have objectual truth-conditions, they have them in virtue of the inferential role of the ingredient existential idiom. I believe, moreover, that the inferentialist cannot coherently deny (4). As Kripke (1976, pp. 379–80) has emphasized, the standard explanation of existential quantification is ultimately couched in terms of the existential idiom of natural language. As such, *there could be no question* as to whether the existential sentences of natural language possess objectual truth-conditions. Objectual truth-conditions just are the truth-conditions determined by the content of natural language quantifiers. And if IRS is the correct account of their content, then inferential role could not fail to determine objectual truth-conditions.

The real problem for the inferentialist doesn't consist in the insistence that existential sentences are subject to the standard explanation and so possess objectual truth-conditions. Rather, the problem consists in the distinctive modal claim, (5)—that there *could be* a logical operator that behaved inferentially just like an existential quantifier but differed in content. Notice, the envisioned possibility just is the denial of IRS. According to IRS *to be* an existential quantifier is *to behave* like one (where the behavior in question is inferential behavior). There could be no such operator, at least if IRS is correct. To repeat, it is less the insistence that existential sentences have objectual truth-conditions, than the insistence that logical content and inferential role could come apart that is the difficulty posed by the *reductio*.

3 The Regress

Why couldn't there be a logical operator that had the same inferential role as legitimate existential idiom but wasn't itself an existential quantifier? Why are logical content and inferential role necessarily coextensive? Perhaps surprisingly, the answer to this question will depend on an answer to a prior question: In what does a speaker's understanding of existential idiom consist?

3.1 Logical Understanding and Knowledge of Truth-Conditions

The representationalist claims that the content of existential idiom is captured by the standard explanation and the content described therein determines their inferential role. If asked in what does a speaker's understanding of existential idiom consist, the representationalist might claim that it consists in knowing the standard explanation.¹³

I suspect that the apparent plausibility of this response is due, in no small part, to a *conflation*. There is an ambiguity in our initial question: In what does a speaker's understanding of existential idiom consist? Two distinct issues are at stake. The question might be: How does a speaker attach an expression to a logical content that he can antecedently entertain? Or the question might be: How can a speaker come to entertain a logical content that he antecedently could not? So consider a speaker that grasps the content of existential quantification and so is capable of entertaining existential propositions. (In what follows I shall be using *grasping* and *entertaining* interchangeably.) The first question asks what facts about the world establish that he is capable of expressing these

13. Though of course he needn't. The consequence of dropping this assumption is discussed in section 5.

propositions by uttering certain sentences. The second question asks what facts about the world establish that he is capable of entertaining existential propositions in the first place. The problem of logical understanding, in the sense that shall concern us here, is to determine the fact in virtue of which a speaker is capable of entertaining a logical content, not in determining how the speaker comes to express such contents. While knowing the standard explanation is plausible (at least initially) as an account of how a speaker associates an expression with the content of existential quantification, it is hopeless as an account of what it is to grasp that content and so entertain existential propositions—or so is the burden of this section to show. The temptation to suppose otherwise is due to not clearly distinguishing these two questions.

We are considering the suggestion that a speaker's understanding of existential idiom consists in knowing the standard explanation where understanding, here, is being able to grasp the content of existential quantification and so entertain existential propositions. Specifically, grasping the content of $\exists x$ consists in knowing every instance of the standard explanation: $\exists xS(\dots x\dots)$ is true iff there is an individual in the range of the objectual variable x that is in the extension of $S(\dots x\dots)$. This, however, is clearly circular. The representationalist cannot coherently claim that a speaker could come to entertain the content of existential quantification by knowing the standard explanation since the standard explanation presupposes the antecedent intelligibility of existential quantification. On the standard explanation, an existential sentence is true just in case *there is* an individual that is in the extension of the sentential matrix. In order to know the proposition expressed by the standard explanation the speaker must already be able to entertain the content of existential quantification (since one can only know propositions that one can entertain).

One might claim that knowing the standard explanation is not circular as an account of logical understanding if the quantifier used in the explanation is a distinct expression. Just as it is legitimate to explain the content of **lapin** by using a synonymous expression that one antecedently understands, e.g., **rabbit**, so it is legitimate to explain the content of object-level existential quantification in terms of unexplained metalinguistic quantification. But knowledge of the standard explanation presupposes the antecedent intelligibility of the metalinguistic quantifier, and, we may fairly ask, in virtue of what is the speaker capable of entertaining the content of metalinguistic quantification? If, moving up the hierarchy of metalanguages, the representationalist now claims that the speaker's understanding of the metalinguistic quantifier consists in knowing a version of the standard explanation framed in terms of meta-metalinguistic quantification, then the representationalist faces an infinite regress, for in what does the speaker's grasp of meta-metalinguistic quantification consist?

One might claim that knowing the standard explanation is, nevertheless, a good explanation of the speaker's understanding of logical vocabulary understood as expressions

of public language. Understanding the English expression **something**, for example, is simply a matter of knowing the standard explanation. It is in virtue of the standard explanation being (part of) the content of a speaker's belief or semantic intention that the speaker understands the public language expression. Of course, the logical contents of the relevant belief or semantic intention cannot themselves be entertained in virtue of having other beliefs or semantic intentions on pain of infinite regress. But such an account provides at least a partial reduction of the problem of logical understanding.

Even this restricted claim is unwarranted. It, in effect, changes the subject. Recall *logical understanding*, in the sense that concerns us here, is a matter of being able to entertain certain logical contents, not in attaching these contents to expressions of public language. We don't want an account that presupposes the speaker's ability to entertain the relevant logical content; we want an account of the presupposed ability.

In making this observation I am not thereby committed to the Gricean contention that linguistic content depends on, and is derived from, mental content. Essentially the same remark can be made by an advocate of the Sellarsian contention that mental content depends on, and is derived from, linguistic content, or by the Davidsonian who claims that linguistic and mental content are *interdependent*. The general point is that logical understanding is not a matter of being able to explain the content of some logical expression (whether a public language expression or some mental expression) in terms of a distinct expression whose content the speaker can antecedently entertain. Rather, an account of logical understanding, in the sense that concerns us here, is an account of the *antecedent* grasp of logical content presupposed by such explanations.

The regress is perfectly general in that it relies on no feature specific to existential quantification. Consider, for example, *the standard explanation of conjunction*:

S & S* is true iff **S** is true and **S*** is true.

If someone were to claim that entertaining the content of conjunction consists in knowing the standard explanation of conjunction, then such an account would itself be circular or subject to an infinite regress. The regress provides us with a general method for arguing against any proposal to the effect that our ability to entertain a logical content consists in knowledge of truth-conditions.

One might be tempted to avoid the regress by making the following observation. It is inessential that the standard explanation of existential quantification be formulated in terms of metalinguistic existential quantification. One can reformulate the standard explanation in terms of negation and universal quantification:

$\exists xS(\dots x\dots)$ is true iff it is not the case that for all individuals i in the range of the objectual variable x i is not a member of the extension of $S(\dots x\dots)$.

Similarly, the standard explanation of conjunction can itself be reformulated in terms of negation and disjunction:

$S \ \& \ S^*$ is true iff it is not the case that (S is false or S^* is false).

The problem is that this merely postpones the inevitable. For whatever choice of logical primitives the question will always arise: In virtue of what is the speaker capable of entertaining *their* content?

With respect to determining what it is to entertain logical content, knowledge of the standard explanation is hopeless as an account. Specifically, it faces the following dilemma: It is either circular (since such knowledge presupposes the antecedent ability to entertain the relevant contents) or it is subject to an infinite regress. Indeed it should be emphasized that the temptation to avoid circularity and thereby embrace the regress is due to not clearly distinguishing the ambitious project of determining the fact in virtue of which a subject can entertain a logical content from the more modest project of determining how a subject can attach an expression to a logical content that he can antecedently entertain.¹⁴

3.2 Logical Understanding and Knowledge of Content

It is worth emphasizing that the problem for the representationalist is not due to an overly narrow focus on truth-conditions. Knowledge of content, as opposed to mere truth-conditions, fares no better as an account of what it is to entertain a logical content. Consider the following revision of the standard explanation:

14. The representationalist might maintain that knowledge of the standard explanation is nevertheless an adequate account of how speakers associate expressions with logical contents that they can antecedently entertain. I believe even this more modest claim is false. As Foster (1976) and Soames (1989) have argued, knowledge of truth-conditions is not sufficient for understanding in the relevant sense since such knowledge is consistent with false beliefs about content. Someone could know that **Il y a des lapins** is true iff there are individuals in the extension of **des lapins** without understanding the French sentence. Indeed it is consistent with such knowledge that the subject believes that **Il y a des lapins** expresses the proposition that there are rabbits and arithmetic is incomplete. Since knowledge of truth-conditions is compatible with false beliefs about content, such knowledge is insufficient for understanding even in the modest sense.

Soames (1989) has further argued that knowledge of truth-conditions is not necessary either. It is plausible that a child can come to understand sentences of his native language before he can entertain the content of a truth predicate, in which case he would be incapable of such knowledge. For an independent argument that one must be able to entertain propositional contents prior to grasping the content of a truth predicate see Kalderon (1997).

The Revised Standard Explanation: $\exists x$ functions semantically as an existential quantifier iff $\exists xS(\dots x\dots)$ expresses the proposition that $xS(\dots x\dots)$.

The current suggestion on behalf of the representationalist is that a speaker's ability to grasp the content of existential quantification and so entertain existential propositions consists in knowing every instance of the revised standard explanation: $\exists xS(\dots x\dots)$ expresses the proposition that $xS(\dots x\dots)$. Notice, however, that such knowledge presupposes the speaker's ability to entertain the content of existential quantification. The problem is that in order to entertain the content of a **that**-clause one must be able to entertain the proposition expressed by the embedded sentence. Thus, in order to know an instance of the revised standard explanation the speaker must already be capable of entertaining existential propositions (since one can only know propositions that one can entertain). Any such account would clearly be circular.

Nor would it help to observe that a speaker can designate propositions in ways that do not require the capacity to entertain them by, for example, *naming* them. Consider a speaker that knows no French. Suppose further that due to some educational impoverishment he has yet to grasp the content of **rabbit**. Since he cannot entertain this content neither can he entertain propositions about rabbits. Nevertheless we can teach him to designate the proposition that there are rabbits in the following fashion: We stipulate that **Gavagai** refers to the proposition that competent French speakers normally express by assertively uttering **Il y a des lapins** (where the stipulation is reference-fixing, not synonymy supplying). Given his competence with **Gavagai**, the speaker could, for example, entertain the proposition expressed by **The truth of Gavagai surprises and delights Bernice** despite his inability to entertain the named proposition. What this example shows is that it is possible to designate propositions that one cannot entertain. It is thus possible to entertain a proposition with a propositional constituent that is not itself entertainable.¹⁵

Can the representationalist exploit this fact to give a non-circular reformulation of the revised standard explanation? The revised standard explanation might be reformulated as follows:

$\exists x$ functions semantically as an existential quantifier iff $\exists xS(\dots x\dots)$ expresses p

(where p is schematic letter whose permissible substituends are names for existential

15. In addition to *naming* propositions, one can designate propositions that one cannot entertain by *quantifying over* or *describing* them. For further discussion see Soames (1989) and Kalderon (1997).

propositions). Suppose now, not that the speaker is incapable of entertaining propositions about rabbits, but rather that he is incapable of entertaining existential propositions. The representationalist might maintain that the speaker could come to grasp the content of existential quantification and so entertain existential propositions by knowing, among other things, that **there are rabbits** expresses Gavagai. While clearly avoiding the charge of circularity, it ought also to be clear that circularity is avoided at the cost of providing no account whatsoever of a speaker's grasp of existential quantification. *Ex hypothesi*, knowing that **there are rabbits** expresses Gavagai doesn't require that one grasp the content of existential quantification. Reformulating the revised standard explanation by designating propositions in ways that do not require that one can entertain them has the consequence that knowledge of instances of the revised standard explanation is insufficient for grasping the content of existential quantification.

The difficulty for the representationalist, then, is not just the relatively parochial problem that a speaker's grasp of the content of existential quantification (and his consequent ability to entertain existential propositions) could not consist in an explicit representation of objectual truth-conditions—*neither* could it consist in an explicit representation of the content of existential quantification. The problem is perfectly general: Either such a representation presupposes the antecedent ability to entertain the logical content and is thus circular, or it makes an essential reference to a further representation and is thus subject to an infinite regress.

The present considerations suggest a stronger and more general conclusion. *Grasping the content of an expression is not a propositional attitude*. Propositional attitudes are, of course, relations to propositions, but one can only adopt an attitude toward a proposition that one can entertain. Grasping a content could not itself be a propositional attitude; for insofar as propositional attitudes are representational, a speaker's having a propositional attitude already involves a grasp of its content. If grasping is not a propositional attitude, then any explanation of it in terms of propositional knowledge is no explanation at all.¹⁶

16. The regress argument is not without historical antecedents. It is a variant of Quine's 1935 argument in "Truth by Convention." There Quine argues that linguistic conventions cannot determine the content of our logical vocabulary. Roughly the problem is that in order to explicitly state the convention one must deploy the very vocabulary whose content the convention was supposed to determine. And if the logical vocabulary used in the statement of the convention are distinct expressions whose contents are themselves determined by further conventions, then we have a beginning of an infinite regress.

It has, perhaps, been insufficiently appreciated how Quine's argument in "Truth by Convention" anticipates Wittgenstein's regress of interpretation argument. Wittgenstein argues that following a rule cannot consist in possessing an explicit representation of the dictates of that rule. For such a representation must itself be rule-governed, and if following a further rule is a matter of possessing yet another representation, then we are off on an infinite regress. And thus Wittgenstein writes (1945/1968, §201): "What this shews is that there is a way of grasping a rule

4 From Logical Understanding to Logical Content

The ability to grasp the content of existential quantification (and so entertain existential propositions) could not consist in knowledge of objectual truth-conditions. Nor could it consist in knowledge of the content of existential quantification. Nor, would it seem, could it consist in a competent speaker's having any other propositional attitude. In what then does logical understanding consist?

These difficulties might encourage the conviction that logical understanding is *primitive*, i.e., not subject to further explanation. But does the regress in fact justify belief in primitive logical understanding? Only if we assume that logical understanding, if explicable at all, should be explained in terms of propositional knowledge—for it is a consequence of the regress that no such explanation is tenable. But this is not the only form such an explanation can take. I believe the reason speakers competent with existential idiom, say, are disposed to accept the validity of existential generalization is that their competence consists precisely in their having this (and perhaps other) inferential dispositions. A speaker's grasp of logical content doesn't consist in any propositional knowledge but rather in his practical knowledge—grasping the content of logical expression is a matter of knowing how to reason in accordance with the validity of the M-inferences. If an explanation of logical understanding in terms of practical knowledge were to survive the regress, then belief in primitive logical understanding would be rationally unmotivated.

An inferentialist account of logical understanding is committed, at a minimum, to the following supervenience thesis:

For any metaphysically possible worlds w and w^* , and any speakers A and B that exist in w and w^* respectively, if there are expressions L_A and L_B such that A and B are disposed to accept the validity of the same set M of inferences involving L_A and L_B , then A and B grasp the same logical content.

The supervenience thesis is fairly weak—it is merely the claim that there could be no independent variation of logical understanding and inferential role. The inferentialist claims more than this. According to the inferentialist the relevant inferential dispositions *determine*

which is not an interpretation, but which is exhibited in what we call 'obeying the rule' and 'going against it' in actual cases." The regress of interpretation argument is itself a generalization of Quine's argument against logical conventionalism.

Notable as well is Ryle's (1949, ch. 2) argument that any attempt to reduce practical knowledge to propositional knowledge is subject to an infinite regress.

a speaker's grasp of logical content.

What is the nature of the determination relation? Recall the inferentialist offered a semantic explanation of the supervenience of logical content on inferential role. Logical content supervenes on inferential role because, roughly speaking, the disposition to accept the validity of the M-inferences implicitly defines L's content. Unfortunately, no corresponding semantic explanation is available for the supervenience of logical understanding on inferential role. Here, unlike the case of inferential role semantics, I am inclined to believe the stronger reductive thesis. That is, while I believe logical content to be distinct from the relevant inferential dispositions, I am inclined to believe that logical understanding *just is* the possession of the relevant inferential dispositions:

Inferential Competence: A speaker's ability to grasp a logical content just is his disposition to accept the validity of the M-inferences involving a logical expression L (where L is either a linguistic or mental expression).¹⁷

Notice that the capacity to reason in accordance with the validity of an inference, though an intentional property, does not itself involve the content the grasp of which it determines. According to Inferential Competence, a speaker's ability to entertain a logical content does not require him to represent that content. Inferential Competence avoids the regress precisely in virtue of this denial.

Insofar as the relevant inferential dispositions do not themselves involve the target content property, Inferential Competence is not circular. But why think that it is an account of logical understanding at all? What is it about our knowing how to reason in accordance with the validity of certain inferences that determines our grasp of logical content?¹⁸ The

17. It is worth pointing out that on the Millian variant of inferential role semantics there will be more than one way to grasp a logical content. Recall on the Millian variant, the inferential role of an expression L determines a descriptive condition, and L's content is the logical object that satisfies the description. While distinct inferential roles will determine distinct descriptions, it is possible that expressions with distinct inferential roles have identical contents since the descriptions determined by inferential role are reference-fixing, not synonymy-supplying. But if being disposed to accept the validity of the M-inferences is what it is for a speaker to grasp L's content, then there will be more than one way to grasp that content. While something of a heresy in Fregean semantics, the claim that there is more than one way to grasp a content is not implausible. For a useful related discussion see Fodor (1998, pp. 16–7).

18. Recall the account of logical understanding in terms of knowledge of the revised standard explanation. The claim was that a speaker grasps the content of existential quantification and is thereby capable of entertaining existential propositions by knowing every instance of the revised standard explanation: $\exists xS(\dots x\dots)$ expresses the proposition that $xS(\dots x\dots)$. The account is, of course circular, if existential propositions are denoted by *that*-clauses, but not if they are instead *named*—for naming propositions allows one to designate propositions that one cannot entertain.

inferentialist is well placed to answer this question for he maintains that the inferential dispositions of competent speakers do double duty—they at once constitute the competence of such speakers and determine the contents which they entertain. A speaker grasps the content of a logical expression by being disposed to accept the validity of certain inferences because this disposition implicitly defines the expression's content. Knowing how to reason in accordance with the validity of certain inferences imposes constraints on a logical expression's content sufficient to determine that content. It is for this reason that such practical knowledge is a good candidate for being that in which logical understanding consists.

According to the inferentialist, a speaker's ability to entertain a logical content consists in his disposition to accept the validity of certain inferences involving a logical expression (whether of public language or some mental expression). Inferential Competence survives the regress by denying that grasping a logical content involves the possession of an explicit representation of that content. Moreover, it would seem that one cannot reasonably believe that a speaker's ability to entertain a logical content consists in his disposition to accept the validity of the M-inferences without also believing that this disposition determines the very content that the speaker can entertain. If that's right, then a commitment to Inferential Competence carries with it a commitment to IRS. Not only do I believe that Inferential Competence survives the regress argument, but I also believe that it is the *only* fully adequate account of logical understanding that can. The case here should be understood as a *challenge*: Give some other account of logical understanding that survives the regress, or accept Inferential Competence.

If Inferential Competence is forced on us (and, with it, IRS), then (5) must be false and the *reductio* unsound. Recall the representationalist attempted to undermine IRS by envisioning an operator that participated in the same set M of inferences that speakers competent with existential idiom are disposed to accept as valid but lacked the appropriate truth-conditions. If, however, a speaker's ability to entertain the content of existential quantification consists in his disposition to accept the validity of the M-inferences, and this disposition is content-determining, then (5) must be false—there could be no operator that behaved inferentially just like an existential quantifier but differed in content.

5 Could Logical Understanding Consist in an Implicit Grasp of Truth-Conditions?

Once the schema is so reformulated, knowledge of its instances does not require the ability to entertain existential propositions. The proposal avoided the regress at the cost of providing no account whatsoever of logical understanding. The present worry is of a piece. Inferential Competence is not circular as an account and so too avoids the regress, but perhaps at the cost of providing no account of logical understanding.

The representationalist might feel up to the challenge posed above. Specifically he might object as follows: If Inferential Competence survives the regress by claiming that a speaker's grasp of a logical content is determined by some *implicit* feature of his inferential usage, why couldn't the representationalist make a similar claim on his own behalf? He can agree that the regress establishes that logical understanding could not consist in an *explicit* representation of truth-conditions. But he could argue that logical understanding consists in an *implicit* grasp of truth-conditions.

In assessing this objection, let's begin by examining what this implicit grasp of truth-conditions could consist in. The representationalist might claim that it consists in a speaker's disposition to unhesitatingly accept certain elementary logical truths involving the target logical expression. The disposition to hold true certain sentences counts as an implicit grasp of truth-conditions because this disposition *implicitly defines* the content of the logical expression.

Corresponding to every inference in the set M is a conditional whose antecedent consists of the conjunction of its premises and whose consequent consists in its conclusion. Thus corresponding to instances of existential generalization are conditionals of the form:

If S(...n...), then $\exists xS(...x...)$

Given this observation, the representationalist might develop the present suggestion as follows: A speaker's implicit grasp of objectual truth-conditions consists, at least in part, in his disposition to unhesitatingly accept all conditionals of this form. Logical understanding consists in an implicit grasp of truth-conditions, and this implicit grasp, in turn, consists in a speaker's disposition to hold true conditionals corresponding to valid inferences involving the target logical operator. This disposition counts as an implicit grasp of truth-conditions since it *implicitly defines* the content of the logical expression. In holding true certain distinguished sentences, the logical expression comes to have whatever content is required for their truth. On this proposal, the speaker's logical understanding consists in his disposition to *hold true* certain sentences, not in the speaker's disposition to *accept* certain inferences *as valid*. The representationalist can consistently claim that while the acceptance of certain sentences plays a role in determining the content of these sentences, it is independently constituted facts which makes these sentences true (given that they have those contents).

It is worth emphasizing one commitment of any such account. The representationalist who adopts such an account is committed to denying that there could be a logical operator L^* that participates in the same set M of inferences as L but differed in content. If an implicit grasp of truth-conditions consists in holding true conditionals corresponding to

valid inferences in which L occurs, then there could be no logical operator L^* that participates in the same set M of inferences but differed in content. A speaker's implicit grasp of the content of L^* would consist in his disposition to hold true precisely the same conditionals that he is disposed to hold true of L . But if holding true certain sentences is content-determining, then L and L^* must share the same content. The representationalist can only avoid the regress by rendering the *reductio* unsound.

Recall, the *raison d'être* of representationalism consists in the claim that content determines inferential role. The representationalist claims that it is *because* a logical expression has the content that it does, that speakers competent with that expression are disposed to accept the validity of the M -inferences. How would logical content determine inferential role if logical content is itself determined by the holding true of certain elementary logical truths? Suppose that one of the speaker's content-determining dispositions with respect to the conditional, \rightarrow , is his disposition to hold true conditionals of the form:

If $S \rightarrow S^*$ and S , then S^*

How in terms of this are we to explain the speaker's disposition to accept the validity of the following inference:

$S \rightarrow S^*, S \vdash S^*$?

One might be tempted to claim that given the premises of the inference:

$S \rightarrow S^*$
 S

and the conditional:

If $S \rightarrow S^*$ and S , then S^*

the conclusion of the inference follows by *modus ponens*. But notice any such explanation itself presupposes the speaker's disposition to accept the validity of the inference—indeed the very disposition it was supposed to explain.¹⁹ The explanation relies on being *antecedently* disposed to accept the validity of the inference and is thus circular.

19. This is, in effect, the problem posed to Achilles by the Tortoise in Lewis Carroll's (1895) "What the Tortoise said to Achilles." (Though Carroll himself seems not to be able to see clear of this problem.) For a useful discussion see Thomson (1960).

It might be objected that I am ignoring the resources that are available to the representationalist. If holding true certain sentences genuinely constitutes an implicit grasp of truth-conditions, then valid inference can be explained *à la* Tarski in terms of the machinery of model-theoretic semantics. Given a description of the truth-conditions of $S \rightarrow S^*$, S and S^* one can establish that there is no model M and assignment function s relative to which $S \rightarrow S^*$ and S are true and S^* is false.

Consider then a simple language consisting of names, predicates, and truth-functional connectives. A *model* is a specification of non-empty set plus an assignment of denotations from that set to the non-logical vocabulary. Specifically, the denotations of names will be elements of the set, and the denotations of n -place predicates will be n -tuples of elements of the set. Given the notion of a model we can inductively define the notion of truth in a model as follows:

- (a) A name n denotes an individual i relative to a model M and assignment function s iff M and s assign i to n .
- (b) An n -place predicate P applies to an n -tuple of individuals relative to a model M and assignment function s iff the n -tuple is in the set M and s assign to P .
- (c) An atomic sentence $P(n_1, \dots, n_n)$ consisting of an n -place predicate P and n names n_1, \dots, n_n is true relative to a model M and assignment function s iff P applies to an n -tuple of individuals $\langle i_1, \dots, i_n \rangle$ that are denoted by n_1, \dots, n_n respectively relative to M and s .
- (d) A sentence $\neg S$ is true relative to a model M and assignment function s iff S is not true relative to M and s .
- (e) A sentence $S \rightarrow S^*$ is true relative to a model M and assignment function s iff it is not the case that (S is true relative to M and s) and (S^* is false relative to M and s).

Given the notion of truth in a model one can define the notion of logical consequence as follows:

A sentence S is a logical consequence of a set of sentences P (the inference from the sentences in P to S is valid) iff there is no model and assignment function relative to which every sentence in P is true and S is false.

We are supposing that in virtue of the fact that a speaker is disposed to hold true

certain distinguished sentences, \rightarrow has whatever content is required for the truth of these sentences. We are further supposing that the implicitly defined content of \rightarrow is such that clause (e) is an adequate representation of that expression's truth-conditional content. The representationalist claims that we can explain the speaker's disposition to accept the validity of the inference from $\{S \rightarrow S^*, S\}$ to S^* in terms of such a content. What form would such an explanation take?

It is tempting to suppose that the explanation would involve a metalinguistic proof that S^* is a logical consequence of $\{S \rightarrow S^*, S\}$. According to the definition of logical consequence, S^* is a logical consequence of $\{S \rightarrow S^*, S\}$ just in case there is no model and assignment function relative to which $S \rightarrow S^*$ and S is true and S^* is false. Consider all the models relative to which S is true and S^* is false. Given clause (e), the conditional $S \rightarrow S^*$ is false relative to such models since these are models in which the antecedent is true and the consequent false. It follows that there are no models relative to which $S \rightarrow S^*$ and S is true and S^* is false, and hence that S^* is a logical consequence of $\{S \rightarrow S^*, S\}$.

To complete the explanation the representationalist would have to make the following epistemological assumptions. A competent speaker knows that: $S \rightarrow S^*$ is true iff it is not the case that S is true and S^* is false. Given this knowledge, he would be disposed to accept the model-theoretic representation of this fact. Specifically, if he knows what a model is, then given his knowledge of truth-conditions, he could know that a sentence $S \rightarrow S^*$ is true relative to a model M and assignment function s iff it is not the case that S is true relative to M and s and S^* is false relative to M and s . Suppose further that the speaker knows the model-theoretic definition of logical consequence. That is, the speaker knows that a sentence S is a logical consequence of a set of sentences P iff there is no model and assignment function relative to which every sentence in P is true and S is false. Given this knowledge, the speaker could give the metalinguistic proof sketched above, and thereby come to know that S^* is a logical consequence of $\{S \rightarrow S^*, S\}$. If the speaker comes to know this fact about logical consequence, then the speaker would consequently be disposed to accept the validity of *modus ponens*.

I believe that skepticism is warranted concerning this overly rationalistic explanation of a speaker's inferential dispositions. But it suffices to observe that the explanation is only cogent if the speaker is disposed to accept the validity of the metalinguistic inference. The explanation of validity in terms of truth in a model itself presupposes the speaker's acceptance of the validity of certain metalinguistic inferences, and, we may fairly ask, in virtue of what is the speaker disposed to accept *their* validity? The appeal to the model-theoretic representation of validity is less an advance than a detour.

I suspect that the representationalist may have been misled by failing to clearly distinguish between validity and the acceptance of validity. Validity is, of course, definable

in terms of logical content. The key idea of Tarski's definition of logical consequence is that an inference is valid iff *holding the content of the logical vocabulary fixed*, there is no way to interpret the non-logical vocabulary such that the premises are true and the conclusion is false. This may encourage the thought that the acceptance of validity is itself ultimately explainable in terms of logical content. But validity and the acceptance of validity are distinct properties—it is possible for someone to accept the validity of an inference that is in fact invalid. And if the considerations of the present section are to be believed, the acceptance of validity is inexplicable in terms of logical content.

In general, while one apparently avoids the regress by claiming that logical understanding consists in an implicit grasp of truth-conditions, one's success is *only* apparent. For the regress will return when it comes to explaining how logical content, so conceived, determines inferential role.

6 Conclusion

I have argued that if a speaker's ability to entertain a logical content is constituted by his disposition to accept the validity of the M-inferences involving some logical expression (whether it be a linguistic or mental expression), then the content of that expression is itself determined by this inferential disposition. I have further argued that this account of logical understanding survives the regress described in §3 since grasping a logical content is not a matter of possessing an explicit representation of its content. More than that, I have suggested that this is the only fully adequate account of logical understanding that can. The case here should be understood as a challenge: Give some other account of a speaker's ability to grasp the content of a logical expression or accept the account offered here. If the challenge cannot be met, then the advertised conclusion stands—given the nature of our understanding, some version of inferential role semantics must be the correct account of the determinants of logical content.²⁰

20. I would like to thank Paul Benacerraf and the participants of the University of California, Riverside faculty seminar for comments on an earlier version of this material. Special thanks are due to Fiona Cowie, James Pryor, and Mike Thau for their many helpful suggestions.

References

- Block, Ned, 1986: "Advertisement for a Semantics for Psychology." In *Midwest Studies in Philosophy*, 10, pp. 615–28.
- _____, 1993: "Holism, Hyper-analyticity and Hyper-compositionality." *Mind & Language*, 8, pp. 1–26.
- Boghossian, Paul, 1989: "The Rule-Following Considerations". *Mind*, 98, pp. 507–49.
- _____, 1993: "Does and Inferential Role Semantics Rest Upon a Mistake?" *Mind & Language*, 8, pp. 26–.
- Brandom, Robert, 1994: *Making It Explicit*. Cambridge, MA: Harvard University Press.
- Byrne, Alex, 1993: *The Emergent Mind*. Ph.D. thesis, Princeton University.
- Carnap, Rudolph, 1937: *The Logical Syntax of Language*. London: Kegan Paul, Trench, Trubner & Co. Ltd.
- Carroll, Lewis, 1895: "What the Tortoise Said to Achilles." *Mind*, 4, pp. 278–80.
- Field, Hartry, "Logic, Meaning and Conceptual Role." *Journal of Philosophy*, 74, pp. 379–408.
- _____, 1989: "Realism, Mathematics and Modality." In *Realism, Mathematics and Modality*. Oxford: Basil Blackwell.
- Fodor, Jerry, 1998: *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Foster, J.A.: "Meaning and Truth Theory." In Gareth Evans and John McDowell (eds.), *Truth and Meaning*. Oxford: Oxford University Press, pp. 1–32.
- Harman, Gilbert, 1986a: *Change in View*. Cambridge, MA: MIT Press.
- _____, 1986b: "The Meaning of Logical Constants." In Ernest LePore (ed.), *Truth and Interpretation*. Oxford: Basil Blackwell, pp. 125–34.
- Horwich, Paul, 1995: "Meaning, Use and Truth." *Mind*, 104, pp. 355–68.
- Kalderon, Mark Eli, 1996: "What Numbers Could Be (And, Hence, Necessarily Are)." *Philosophia Mathematica*, 3, vol. 4, pp. 238–55.
- _____, 1997: "The Transparency of Truth." *Mind*, 106, pp. 475–97.
- Kripke, Saul, 1976: "Is there a Problem about Substitutional Quantification?" In Gareth Evans and John McDowell (eds.), *Truth and Meaning*. Oxford: Oxford University Press, pp. 325–419.
- _____, 1982: *Wittgenstein on Rule-Following and Private Language*. Cambridge, MA: Harvard University Press.
- Lewis, David, 1972/1980: "Psychophysical and Theoretical Identifications." Reprinted in Ned Block (ed.), *Readings in Philosophy of Psychology, Volume One*. Cambridge, MA: Harvard University Press, pp. 207–15.

- Peacocke, Christopher, 1976: "What is a Logical Constant?" *Journal of Philosophy*, lxxiii, pp. 221–40.
- _____, 1987: "Understanding Logical Constants: A Realist's Account." *Proceedings of the British Academy*, lxxiii, 153–200.
- _____, 1992: *A Study of Concepts*. Cambridge, MA: MIT Press.
- Quine, Willard Van Orman, 1935/1979: "Truth by Convention." Reprinted in *The Ways of Paradox*. Cambridge, MA: Harvard University Press, pp. 77–106.
- Rosen, Gideon, 1993: "A Refutation of Nominalism(?)." *Philosophical Topics*, 21, pp. 149–86.
- Ryle, Gilbert, 1949: *The Concept of Mind*. New York: Barnes & Noble, Inc.
- Soames, Scott, 1997: "Skepticism about Meaning: Indeterminacy, Normativity, and the Rule-Following Paradox." In A. Kazmi (ed.) *Meaning, Reference and Truth*, *Canadian Journal of Philosophy*, supplementary volume.
- Tarski, Alfred, 1935/1983: "On the Concept of Logical Consequence," J.H. Woodger (tr.). Reprinted in *Logic, Semantics, Metamathematics*, (ed.) John Corcoran. Indianapolis: Hackett Publishing Co., Inc., pp. 409–20.
- Thomson, J.F., 1960: "What Achilles Should Have Said to the Tortoise." *Ratio*, 3, pp. 95–105.
- Wittgenstein, Ludwig, 1945/1968: *Philosophical Investigations*, G.E.M. Anscombe (tr.). New York: MacMillan Publishing Co., Inc.
- Wright, Crispin, 1983: *Frege's Conception of Numbers as Objects*. Aberdeen: Aberdeen University Press.

