

A PROBLEM FOR INFORMATION THEORETIC SEMANTICS*

ABSTRACT. Information theoretic semantics proposes to construe predicate reference in terms of nomological relations between distal properties and properties of representational mental events. Research on the model has largely concentrated on the problem of choosing the nomological relation in terms of which distal properties are to be singled out. I argue that, in addition to this, an information theoretic account has to provide a specification of which properties of representational mental events will play a role in determining reference, *qua* bearers of nomological relations. I contend that this task poses a serious additional challenge to the viability of the model.

1. INTRODUCTION

When I judge that my car is dirty, my judgment may or may not be satisfied by my car. Which of the two is the case depends on whether my car fulfils a certain condition that my dirtiness-judgment imposes on objects. It is natural to think that what plays the role of this condition is a property, whose presence in, or absence from my car, or any other object, will determine whether my judgment is satisfied by it. Let me refer to these judgments in which a property is ascribed to an object as *predicative judgments*. And let me speak of the property (or properties)¹ whose instantiation conditions determine the satisfaction conditions of a predicative judgment as the property (or properties) *expressed*² by it, and of the relation between each judgment and the property (or properties) that it expresses as the *expression relation*.

Jerry Fodor has chosen this relation as the point at which to break into the circle of semantic notions.³ The relation between a predicative judgment and the property that it expresses is no doubt a semantic relation. But for a naturalist like Fodor, if it is a real relation, it cannot be irreducibly semantic. It has to be possible to specify in non-semantic, non-intentional terms how a predicative judgment would have to be related to a property in order for the former to express the latter.⁴ I shall refer to this kind of account of the expression relation as a *naturalization* of the notion. Fodor's recent work on the theory of content has been largely devoted to providing this kind of account.⁵

Fodor's approach to the task is shaped by two main features. The first feature of his approach to the task that I want to highlight is the involvement of the Language of Thought in his account of the expression relation. In Fodor's picture, predicative judgments are tokens, of a certain kind,⁶ of predicates of an inner language. These predicates play a crucial role in his account of the expression relation. It is to Mentalese predicates, not to predicative judgments, that distal properties are primarily assigned.⁷ The expression relation is supposed to emerge from this assignment: a predicative judgment j expresses a property P just in case P has been assigned to the Mentalese predicate tokened by j . This means that the question whether a predicative judgment expresses a property P has to be answered in two steps. First, we need to determine which of the properties that j instantiates is the property of tokening a Mentalese predicate. Then we need to determine whether this property bears a certain relation to P . An account of the expression relation along these lines would be an instance of the following schema:

- (1) For every predicative judgment j and property P , j expresses P iff $R(f(j), P)$,

where f is the function that assigns to each predicative judgment the Mentalese predicate that it tokens, and R is a relation specified in non-semantic, non-intentional terms.

This feature of Fodor's approach can be analysed into two different claims. First, we have the claim that the expression relation has to be construed in terms of a relation between distal properties and a certain kind of properties of predicative judgments. Second, we have the claim that each of the latter properties is to be thought of as the property of tokening a predicate of the Language of Thought.⁸ I want to leave aside the second claim and concentrate on the first one. We can characterise the first claim as the idea that an account of the expression relation would have to take the following form:

- (2) For every predicative judgment j and property P , j expresses P iff there is a property Q such that $Q(j)$, $S(Q)$, and $R(Q, P)$,

where S is a second order property that singles out the properties of each predicative judgment with respect to which distal properties will be assigned to it, and R is a relation between S -properties and distal properties. In the Language of Thought version of (2), S would be the property of being a Mentalese predicate. In this version, a predicative judgment j that tokens Mentalese predicate M would express a property P just in case P is R -related to the property of tokening M .

In general, we can think of S as a criterion for selecting which of the properties of a predicative judgment play a role in singling out the distal property that it expresses. Hence, I shall refer to S as the *semantic relevance criterion*, and to the S-properties of a predicative judgment as its *semantically relevant properties*. In the Language of Thought version of the model, each predicative judgment has exactly one semantically relevant property: the property of tokening a certain Mentalese predicate.

Needless to say, if an instance of (2) is to provide a naturalization of the expression relation, both the semantic relevance criterion S and the relation R that assigns distal properties to S-properties will have to be specified in non-semantic, non-intentional terms. I shall refer to the claim that a naturalization of the expression relation will have to be an instance of (2) as the *semantic relevance (SR) claim*.

The second feature of Fodor's approach on which I wish to focus is his commitment to information theoretic semantics.⁹ The basic tenet of information theoretic semantics is the idea that the property expressed by each predicative judgment will be singled out in terms of a nomological relation. The following considerations make this thought compelling. A predicative judgment is an event in the physical world, presumably in the subject's nervous system.¹⁰ In some cases, the event consisting in a predicative judgment is caused by an event consisting in the perceptual (e.g. visual) presentation of the object of predication, e.g., when looking at my car prompts me to judge that it is dirty. It is natural to think of this scenario as in some sense fundamental.¹¹ In these cases, there will be a causal chain of events connecting the event consisting in the presentation of the object with the event consisting in the predicative judgment, via the effect of the presentation of the object on the subject's sensory transducers. These causal relations will be governed by laws. The basic intuition of the information theoretic approach is that the property expressed by each predicative judgment can be singled out in terms of these laws. I shall refer to the claim that the expression relation will have to be specified in this manner as the *information theoretic (IT) claim*.

For someone who holds SR, IT would have to be understood as imposing a restriction on the kind of relation that can be expected to play the role of R in (2): it will have to be a nomological relation. Hence, for someone who holds both claims, a predicative judgment j will express a property P just in case P bears a certain nomological relation to an S-property of j. This is indeed the shape that Fodor's proposals have taken. In each of them, the property expressed by each predicative judgment has been singled out as the bearer of a nomic relation to the Mentalese predicate that it tokens.¹²

SR and IT may seem *prima facie* independent of one another. Thus it seems possible to think, on the one hand, that the expression relation should arise from a relation R between distal properties and a certain kind of properties of predicative judgments, but that R need not be a nomic relation. On the other hand, it seems that one could try to single out the property expressed by each predicative judgment as the bearer of a nomic relation, without having to single out some of the properties of each predicative judgment as its semantically relevant properties. However, I am going to argue that the latter is not true. I shall contend that the only viable way of specifying the expression relation along the lines of information theoretic semantics is by means of an instance of schema (2) in which a nomic relation is substituted for R.

Then I shall focus on a difficulty for information theoretic semantics that derives from this fact. Notice that finding an adequate instance of schema (2) would involve two tasks. First, one would have to specify the relation R that assigns distal properties to S-properties. Second, one would have to specify the semantic relevance criterion – the second-order property S that singles out the semantically relevant properties of each predicative judgment.

Most work in information theoretic semantics has been devoted to the first of these two tasks – to finding the nomological relation in terms of which the property expressed by each predicative judgment can be singled out. In this paper, I am going to focus on the second task. I shall contend that an acceptable semantic relevance criterion may prove hard to find, and hence that finding the right nomic relation is not the only significant hurdle that an information theoretic account of the expression relation would have to overcome.

The problem arises from the fact that, as I shall argue in Sections 4 and 5, a plausible semantic relevance criterion must have the resources for ascribing to predicative judgments properties that are disjunctive from a neurological point of view. This forces the proponent of the model to provide an account of which disjunctions are to play this role in each case. The most promising approach to this task is to invoke the functional roles played by neural properties. But, as I shall argue in Section 6, neural properties can be expected to realize a number of functional roles. Hence construing the semantic relevance criterion in terms of functional roles would involve providing an account of which functional roles are to be treated as semantically relevant. In Section 7 I consider some of the difficulties that may make this account hard to achieve. But before we turn to the problems faced by the task of construing the semantic relevance criterion, we need to spend some time making sure that the task is unavoidable

– that an information theoretic account of the expression relation cannot succeed unless it imposes a restriction on which properties of each predicative judgment are to be treated as semantically relevant. This will be the burden of Sections 2 and 3.

2. INFORMATION THEORETIC SEMANTICS AND THE DISJUNCTION PROBLEM

The most immediate challenge for the proponent of an information theoretic approach to the naturalization of the expression relation is to find the specific law that will single out the properties expressed by each predicative judgment. Fodor presents the problems confronted by this task by considering the difficulties faced by a crude proposal as to which nomological relation will serve this purpose.¹³ In his presentation of the crude proposal, he presupposes that predicative judgments are tokens, of a certain kind, of Mentalese predicates. Let f be the function that assigns to each predicative judgment the Mentalese predicate that it tokens. Then, according to the crude proposal, the naturalization of the expression relation would have to take this form:

- (3) For every predicative judgment j and property P , j expresses P iff it is a law that P s cause tokenings of $f(j)$.

Fodor argues convincingly that this proposal is unsatisfactory. The reason is what has come to be known as the *disjunction problem*. Consider the Mentalese predicate that is tokened whenever I judge of an object that it is a horse. A satisfactory account of the expression relation would have to treat the property *horse*¹⁴ as the (only) property expressed by a predicative judgment of mine that tokens Mentalese “horse”. But according to (3), these judgments express whatever property is nomically sufficient to cause Mentalese “horse” to be tokened. Hence (3) would yield an acceptable account of what properties are expressed by these judgments just in case *horse* is the only property that is nomically sufficient to cause tokenings of Mentalese “horse”. But this is not so. By hypothesis, “horse” is tokened whenever I judge of something that it is a horse. But I occasionally judge *wrongly* that something is a horse. I may judge, e.g., of a cow, that it is a horse, and hence “horse” would be tokened as an effect of a presentation of a cow, if the angle, the lighting conditions, my lack of attention, or any other factor makes me take the cow for a horse. But these deviant tokenings of “horse” will also be governed by laws. Thus, if cows that look like horses would cause “horse” to be tokened, the property *horsy-looking-cow* is also nomically sufficient to cause “horse” to be tokened. Hence, we would have to conclude that my predicative judgments that

token “horse” express, not just the property *horse*, but also *horsy-looking-cow*, or rather the disjunctive property *horse or horsy-looking-cow*.

The problem for the crude proposal seems to arise from the fact that the criteria by which it singles out distal properties aren’t restrictive enough. According to (3), the tokenings of a Mentalese predicate *M* express every property that is nomically sufficient to cause *M* to be tokened. But we have seen that many properties that satisfy this criterion could not be expressed by tokenings of *M*. Thus what the proponent of the information theoretic approach needs is a nomic relation that Mentalese predicates bear to fewer properties – only to those properties that their tokenings express. In our example, we need that relation to link “horse” to *horse*, but not to *horsy-looking-cow*.

Fodor presents a number of influential proposals as to how to circumvent this difficulty as following a common pattern.¹⁵ The idea is to bring into the picture a set of conditions *C*, and to treat as the property expressed by tokenings of a Mentalese predicate *M*, not every property that is nomically sufficient to cause *M* to be tokened, but only those that are nomically sufficient to cause *M* to be tokened *under C-conditions*. Thus the idea would be to replace (3) by the following schema:

- (4) For every predicative judgment *j* and property *P*, *j* expresses *P* iff it is a law that *Ps* cause tokenings of *f(j)* under *C-conditions*.

C-conditions would have to be chosen in such a way as to filter out unwanted distal properties. Thus, in our example, we would want *horse*, but not *horsy-looking-cow*, to be nomically sufficient to cause “horse” to be tokened under *C-conditions*.

Fodor claims that his most recent information theoretic account of the expression relation does not follow this pattern.¹⁶ Be this as it may, in what follows I shall focus on proposals that invoke a nomic relation of this form. However, my main points will not obviously depend on this stipulation – and they may turn out to be more widely applicable.

3. INFORMATION THEORETIC SEMANTICS AND THE SEMANTIC RELEVANCE CRITERION

My goal in this section is to argue that an information theoretic account of the expression relation that invokes this kind of nomological relation (it is a law that *Ps* cause *Qs* under *C-conditions*) would also have to invoke a semantic relevance criterion. In other words, I am going to argue that such an account would have to take the following form:

- (5) For every predicative judgment j and property P , j expresses P iff there is a property Q such that $Q(j)$, $S(Q)$, and it is a law that P s cause Q s under C -conditions.

Notice that (4) is an instance of (5), in which the unique S -property of each predicative judgment is the property of tokening a certain Mentalese predicate.

That properties of predicative judgments have to be involved in an information theoretic account of the expression relation should be beyond any doubt. One might be tempted to replace (5) by

- (6) For every predicative judgment j and property P , j expresses P iff it is a law that P s cause j under C -conditions.

But this, as it stands, does not constitute an intelligible proposal. Properties, not particulars, are the proper bearers of nomological relations. A causal law can connect particular events only insofar as they instantiate the properties that the law connects. The only way to take (6) as an intelligible proposal would be to read it as

- (7) For every predicative judgment j and property P , j expresses P iff there is a property Q such that $Q(j)$ and it is a law that P s cause Q s under C -conditions.

Thus, according to this approach, a predicative judgment j will express any property that bears the right nomic relation to one of the properties that j instantiates. (7), unlike (6), singles out the properties expressed by a predicative judgment j as the bearers of a nomic relation to *properties* of j . But, unlike (5), it imposes no restriction on which of j 's properties can play this role. (7) construes the expression relation without invoking a semantic relevance criterion. According to (7), the semantically relevant properties of a predicative judgment are all the properties that it instantiates.

I want to argue that this feature renders any instance of (7) unsuitable as an account of the expression relation. The problem has to do with the fact that a predicative judgment, as any other particular, can be expected to instantiate a large number of properties.¹⁷ Each of the properties that a predicative judgment instantiates will bear various nomological relations to other properties. In particular, for every property Q that a predicative judgment instantiates, there may be a distal property that is nomically sufficient to cause Q to be instantiated. Bringing in C -conditions won't alter this situation. Predicative judgments that occur under C -conditions can also be expected to instantiate many properties, and we cannot rule out the possibility that many of them will have distal properties that are nomically sufficient to cause them to be instantiated.

According to (7), a predicative judgment j would express every property that is nomically sufficient to cause one of the properties of j to be instantiated under C -conditions. Hence, (7) doesn't make room for the possibility that a predicative judgment j has a property Q such that some distal property P is nomically sufficient to cause Q to be instantiated under C -conditions, but j does not express P . (7) cannot accommodate the possibility that some of the properties of a predicative judgment are semantically irrelevant, that they are mere noise from a semantic point of view. I am going to argue that this is unacceptable. I shall contend that a plausible account of the expression relation has to make room for semantically irrelevant properties of predicative judgments.

Let's suppose that the properties of predicative judgments that should be treated as semantically relevant are as salient as they could possibly be. Suppose that things inside someone's head are literally the way Fodor, following Schiffer, describes them metaphorically.¹⁸ This subject has inside her head a tiny box marked "Belief Box". When she makes a predicative judgment, an inscription of a predicate is thrown into this box. Let's waive questions concerning which object is the object of predication in each case. Let's suppose also that the inscriptions that go into the box are English predicates, with subscripts as needed to remove ambiguities.

According to information theoretic semantics, one of these inscriptions would express a property P just in case there is a property Q of the inscription such that P is nomically sufficient to cause an inscription with property Q to go into the belief box under C -conditions. The question that we need to ask is whether we have to impose any restrictions on which properties of each inscription can play this role in singling out the properties that it expresses. I want to suggest that if no restrictions are imposed, the resulting account of the expression relation would be utterly unsatisfactory.

Suppose that our subject judges of an object that she is observing that it is a horse. Her making this judgment consists in an inscription of the word "horse" going into the belief box. Suppose also that the property *horse* is nomically sufficient to cause a "horse"-inscription to go into the belief box under C -conditions. Then, if being a "horse"-inscription is a semantically relevant property of the inscription, it will express the property *horse*, as desired. But her regimented Mental English also contains the predicate "hornet", and inscriptions of this predicate go into the belief box whenever she judges of something that it is a hornet. If the information theoretic account is going to come out right for inscriptions of this word also, the property *hornet* will have to be nomically sufficient to cause a "hornet"-inscription to go into the belief box under C -conditions.

But notice that “horse”-inscriptions and “hornet”-inscriptions have at least one property in common: both are inscriptions starting with the sequence “h”^“o”^“r”. Moreover, it follows from what we have said that both the property *horse* and the property *hornet* are nomically sufficient to cause an inscription starting with the sequence “h”^“o”^“r” to go into the belief box. Hence, if every property of an inscription is semantically relevant, we would not be able to say that “horse”-inscriptions express *horse*, but not *hornet*, and that “hornet”-inscriptions express *hornet*, but not *horse*. We would be forced to conclude that each “horse”-inscription and each “hornet”-inscription expresses the disjunctive property *horse or hornet*.

There are some structural similarities between this problem and the disjunction problem. In both cases, the difficulty takes the form of unwanted distal properties (*horsy-looking-cow*, or *hornet*) being assigned to a predicative judgment by an inadequate account of the expression relation. But the source of the difficulty is different in each case. The disjunction problem is created by a nomological relation that links the property of tokening “horse” with a distal property (*horsy-looking-cow*) that isn’t expressed by tokens of “horse”. The present problem, by contrast, arises even if we have succeeded in identifying a nomological relation that connects the property of tokening “horse” only to the property *horse*. The problem arises from the fact that this nomological relation connects another property of tokens of “horse” (the property of starting with the sequence “h”^“o”^“r”) with a property (*hornet*) that isn’t expressed by them. The problem is created, not by our choice of a nomological relation, but by the lack of restrictions on which properties of a predicative judgment are to be treated as singling out, *qua* bearers of a nomological relation, the distal properties that the judgment expresses.

We have to conclude that in order to have a plausible information theoretic account of the expression relation, even in these unrealistically favourable circumstances, we need to impose some restrictions on which properties of each predicative judgment should count as semantically relevant. In addition to finding the nomological relation that will do the job, an information theoretic account of the expression relation would have to find a semantic relevance criterion. It would have to tell us, in non-semantic, non-intentional terms, which properties of each predicative judgment are fit for semantic duties. In other words, an instance of (7) could not provide a plausible account of the expression relation. Only a (non-trivial)¹⁹ instance of (5) could do the job.

Of course, in the situation that we have envisaged, formulating a semantic relevance criterion wouldn’t be very hard. The semantically relevant

property of each inscription would be the property of containing exactly a certain sequence of letters. But needless to say, things are bound to be much messier in real life. Predicative judgments are unlikely to be inscriptions of English predicates going into a box. More likely, they will be neural events of some sort, and deciding which properties of each of them should be counted as semantically relevant is bound to prove trickier. In the rest of this paper I am going to argue that the task will be very difficult indeed. I shall look in some detail at the standard strategy for accomplishing this task, and argue that it faces serious obstacles.

4. NEURAL PROPERTIES

When I judge that my car is dirty, a number of neural events take place in my brain. On the information theoretic account, my judgment will be identical with one of these neural events, and the property that it expresses will be singled out as the bearer of a nomological relation to some of the properties of this neural event. Suppose that we have determined which neural event is identical with my predicative judgment. As we have seen in the previous section, if we count all the properties of this event as semantically relevant, we are bound to end up with an implausible account of the expression relation. Hence, we need to find a criterion for deciding which of the properties of this event are to be counted as semantically relevant.

My judgment being a neural event, it is natural to think that its semantically relevant properties will be found among its neural properties. According to this proposal, the semantic relevance criterion will appoint as the semantically relevant properties of each predicative judgment some of its neural properties. In this section, I want to argue that such an account of the semantic relevance criterion is unlikely to result in an acceptable naturalization of the expression relation.

The question that I propose to consider is whether we can expect that any of the neural events that take place when I make a predicative judgment will instantiate a neural property that bears the right sort of nomological relation to the property expressed by the judgment. Notice that an affirmative answer to this question will not suffice for establishing the viability of the approach that takes neural properties as semantically relevant. It's not enough that the right properties be there. The viability of the approach would also require that we have a non-semantic, non-intentional specification of which neural properties are to be treated as semantically relevant in each case.

However, a negative answer to the question would suffice for establishing that this approach cannot succeed. If we cannot expect that each predicative judgment will involve the instantiation of a neural property that bears the right nomological relation to the property expressed by the judgment, this approach is bound to fail. I am going to argue that this is in fact the case. I shall contend that we have no good reasons for assuming that every predicative judgment will involve the instantiation of a neural property suitable for the job.

I shall start by describing a situation in which no neural property with the desired features seems to be present. Let's focus again on the predicative judgments that make the information theoretic approach most compelling: predicative judgments that are caused by a visual presentation of the object of predication. Among these, I propose to further restrict our attention to judgments that cause (sincere) utterances of English predicates.²⁰ Thus a visual presentation of my car might prompt me to judge it to be dirty, and I may express this judgment with an utterance of "dirty". In this case, we can expect that there will be a causal chain of events leading from the visual presentation of my car to my utterance of "dirty", via a retinal stimulation and a sequence of neural events. In these cases, my judgment will have to be one of the neural events along this causal chain, and our search for semantically relevant properties can be restricted to the neural properties instantiated by these events.

Consider Mary, a fairly standard English speaker, whose use of English words isn't noticeably different from that of other speakers of the language. In particular, there is nothing strange about her use of the predicate "red": she applies it to the objects to which the rest of us, in her situation, would apply it, and her use of the predicate is equally standard in all other respects. So much so that no one has ever doubted that she means *red* by "red" – that each of her utterances of "red" expresses a judgment that is satisfied precisely by red things.

But one day Mary's nervous system is examined by a team of super-neurologists, who make the following surprising discovery. It turns out that the perception of dark red objects and the perception of light red objects have totally different effects on Mary's nervous system. The neural processes triggered by dark red objects are fairly uniform, and so are the neural processes triggered by light red objects. But the former processes have absolutely nothing to do with the latter. In fact, the difference between these two kinds of process is so huge that no single neural property is instantiated by both – no property, that is, that isn't instantiated also by many other kinds of process.

Such huge differences are to be found, in particular, between the processes that lead from the presentation of a dark red object to an utterance of “red”, and those that lead from the presentation of a light red object to an utterance of the same predicate type. Hence no single neural property will be instantiated in the causal ancestry of all of Mary’s utterances of “red”. Many neural properties will be instantiated in the causal ancestry of all her utterances of “red” caused by dark red objects, and many others by all her utterances of the predicate caused by light red objects. What we won’t find is a neural property instantiated in the causal ancestry of both kinds of utterance of the predicate.

This means, of course, that if we look for semantically relevant properties among the neural properties in the causal ancestry of each predicate utterance, we won’t be able to assign the same semantically relevant property to all the judgments that Mary expresses with utterances of “red” – to her “red”-judgments, as I shall call them. At best, we could assign the same semantically relevant property to all of her “red”-judgments caused by dark red objects, and another one to all of her “red”-judgments caused by light red objects.

One could argue, correctly, that this result does not establish the invariability of this approach. Judgments that are assigned different semantically relevant properties could still turn out to express the same property, since different neural properties could bear the right nomological relation to the same distal property. Thus even if we could not assign the same semantically relevant properties to all of Mary’s “red”-judgments, it is still possible that they will all be assigned *red* as the property that they express. This would be so if *red* turns out to be nomically sufficient to cause the semantically relevant neural properties of all these judgments to be instantiated under C-conditions.

But this remark doesn’t offer much solace to the proponent of this approach. The reason is that we cannot expect that any of the neural properties instantiated in the causal ancestry of Mary’s utterances of “red” will stand in the right nomic relation to *red*. Take the neural processes that lead to Mary’s utterances of “red” from the presentation of dark red objects. I want to suggest that there is no reason to expect that any of the properties instantiated in these processes will stand in the right nomic relation to the property *red*. For any plausible account of C-conditions, if Mary were presented with a light red object under C-conditions, none of those properties would be instantiated. Only the presentation of a dark red object would have this effect. Hence it is with *dark red*, and not with *red*, that we should expect these properties to be nomically connected.

Thus we have to conclude that if we let the semantic relevance criterion assign to Mary's "red"-judgments neural properties, we are likely to end up with an account of the expression relation according to which "red", as meant by Mary, is an ambiguous predicate. Some of her utterances express judgments that are satisfied by dark red things and others express judgments that are satisfied by light red things. None of them would express a judgment that is satisfied by both dark and light red things.

Let's now consider how this result bears on the plausibility of this approach to the definition of the semantic relevance criterion. The question that we need to ask is whether we can accept an account of the expression relation according to which someone with Mary's neural make-up cannot express, with her utterances of "red", judgments that are satisfied precisely by red things, even if her use of the predicate is in all respects indistinguishable from ours. I want to suggest that this question has to be answered in the negative. The reason is not that our pre-theoretical ascriptions of satisfaction conditions couldn't ever be overridden by a scientific account of the notion. The reason is rather that we would want the notion of judgment satisfaction to enjoy a degree of independence from speakers' hardwiring that this approach wouldn't allow for.

There are two different ways in which one may try to defend the neural approach to defining the semantic relevance criterion in the face of cases like Mary's. One possible strategy would be to argue that the case is impossible, because someone with that kind of neural constitution wouldn't be capable of using "red" the way it is used by those who mean *red* by the predicate. In particular, the claim would go, such a person wouldn't be capable of applying "red" (under normal circumstances) precisely to red objects.

This claim is plainly wrong. It is true that the way in which an organism is hardwired limits the ways in which it can learn to attune its behaviour to features of its environment. On the one hand, an organism could not learn to react to the presence in the environment of a feature that its perceptual devices cannot detect. Thus consider an organism with black and white vision, i.e., one on whose cognitive devices reflected light of a given intensity has always the same effect, no matter what colour it has. This organism could not learn to display a certain behaviour precisely when a red object is present. In particular, it couldn't learn to utter "red" in such circumstances.

On the other hand, an organism's powers of discrimination could be so fine-grained, that they make it impossible for it to learn to attune its behavioural responses to comparatively coarse features of the environment. Thus imagine an organism with a remarkable sensitivity to colour. Each

of 75 shades of red affects it in a substantially different way, and these 75 kinds of process don't have anything in common setting them apart from the many other neural processes that non-red colour shades trigger. One would expect that this organism would have great difficulty in learning to behave in a certain way – e.g., producing an utterance of “red” – whenever a red object is present.

These remarks are undoubtedly correct. It is true that not every neural configuration could give rise to the ability to utter “red” in the presence of red objects. But this is not to say that there cannot be several quite different types of configuration that could give rise to this ability. It doesn't establish, in particular, that Mary's neural configuration could not give rise to it. We can concede that it may be a little bit harder for Mary to learn to use “red” properly than for speakers on whose nervous systems the perception of red objects causes uniform processes. But the difficulty would certainly not be insurmountable. All of us can, after all, perform the feat. We can easily learn to use a predicate that is applicable, say, to things that are red or square, even though, presumably, we would face the same difficulty with respect to this predicate that Mary faces with respect to “red”.

The second way in which one could try to defend the neural approach is to bite the bullet and argue that someone with Mary's neural constitution cannot make judgments that are satisfied precisely by red things. In particular, the judgments that she expresses with utterances of “red” cannot have these satisfaction conditions, even if her use of the predicate is indistinguishable from that of speakers who regularly express such judgments with their utterances of it. In other words, according to this approach, we should accept the idea that “red”, as meant by Mary, is ambiguous between dark red and light red.

I want to suggest that the account of meaning and judgment satisfaction that would arise from this response is at odds with our intuitions concerning the relationship between semantics and the neurological ‘infrastructure’. We think that all of a speaker's utterances of “red” could express judgments with the same satisfaction conditions even if there are several different neurological processes that would lead from the perception of an object to an utterance of the predicate. In fact, there seem to be empirical reasons for thinking that this is not an uncommon situation. I am thinking of results to the effect that “a given type of psychological process is in fact often associated with a variety of distinct neural structures”.²¹ But it is hard to see why the situation would be substantially different if the various neural processes associated with a predicate failed to be nomically related to a unique property.

However, an account of the semantic relevance criterion that treats neural properties as semantically relevant cannot make room for this possibility. On this kind of account, predicate utterances caused by different neural processes that don't bear the same nomic relations to the environment couldn't express judgments with the same satisfaction conditions. Hence we have to conclude that this approach to the definition of the semantic relevance criterion cannot result in an adequate account of the expression relation.

5. DISJUNCTIVE PROPERTIES

The outcome of the previous section still leaves the proponent of the information theoretic model some room for manoeuvre. We have concluded that neural properties cannot be expected to play the role of semantically relevant properties in each case. For in cases like Mary's "red"-judgments, this would result in incorrect assignments of satisfaction conditions. But even if neural properties cannot in general serve as semantically relevant, they can still play some role in singling out the expression relation. The reasoning would go as follows. Let S_N be a second order property satisfied only by neural properties. Consider the relation E_N that we can specify in terms of S_N and a nomic relation:

- (8) For every predicative judgment j and property P , $E_N(j, P)$ iff there is a property Q such that $Q(j)$, $S_N(Q)$ and it is a law that P s cause Q s under C -conditions.

The outcome of the previous section is that if we treat E_N as the expression relation, we would render some predicative judgments satisfied by *fewer* objects than we should. For some predicative judgments are satisfied by objects that don't instantiate any of the properties to which they are E_N -related. Thus, some of Mary's "red"-judgments can only be E_N -related to *dark red*, but we want all red objects, dark and light, to fall under their extension.

Notice, however, that as far as the outcome of the previous section goes, treating E_N as the expression relation may render predicative judgments satisfied only by objects by which they are, as a matter of fact, satisfied. None of Mary's "red"-judgments is satisfied *only* by dark red objects, but they are satisfied by *all* dark red objects. The extension that E_N would assign to them will not coincide with their actual extension, but the former may be a subset of the latter. In other words, E_N cannot be expected to fulfil the following constraint:

- (9) For every predicative judgment j , and every object x , if j is satisfied by x , then x instantiates a property to which j is E_N -related,

but it might fulfil its converse:

- (10) For every predicative judgment j , and every object x , if x instantiates a property to which j is E_N -related, then j is satisfied by x .

This may suggest a way of using neural properties in the specification of the expression relation that circumvents the difficulty raised in the previous section. Neural properties cannot be expected, in general, to play the role of semantically relevant properties. But in cases in which they fail, *disjunctions* of neural properties may do the job. Thus I have argued that none of the neural properties instantiated by Mary's "red"-judgments can act as their semantically relevant properties. Suppose, however, that each of them instantiates one of the neural properties R_{DARK} and R_{LIGHT} , such that *dark red* is nomically sufficient to cause R_{DARK} to be instantiated under C-conditions, and *light red* is nomically sufficient to cause R_{LIGHT} to be instantiated under C-conditions. Then the 'nomological pedigree' of R_{DARK} and R_{LIGHT} will include the nomological connections depicted in Figure 1, in which arrows represent causal laws linking the properties on both ends.




Fig. 1

Then, if we treat the disjunctive property R_{DARK} -or- R_{LIGHT} as the semantically relevant property of each of Mary's "red"-judgments, we will assign to them the right extension.

Pursuing this strategy would involve two tasks. First, we would need an account of which neural properties of each predicative judgment can be used in this manner. Second, we would need an account of which disjunctions of these neural properties will do the job. Let's grant, for the sake of the argument, that the first task can be accomplished. That is, suppose that we have specified a class of neural properties such that the semantically relevant properties of each predicative judgment are to be found either among these properties or among disjunctions thereof. I want

to concentrate on the second task – that of providing an account of which disjunctions of neural properties are to be treated as semantically relevant in each case.

The task is not a trivial one. For if our ultimate goal is a naturalization of the expression relation, this account would have to be constructed in non-semantic, non-intentional terms. In the case of Mary’s “red”-judgments, we have seen that treating $R_{\text{DARK-or-LIGHT}}$ as the semantically relevant property of each of them would result in a correct ascription of satisfaction conditions. But the naturalist needs an account of why this is the right disjunction that doesn’t make use of the fact that all of Mary’s “red”-judgments have the same satisfaction conditions.

6. FUNCTIONAL ROLES

The most promising approach consists in appealing to the *functional roles* of neural properties. On this account, what enables us to choose $R_{\text{DARK-or-LIGHT}}$ over the other possible candidates is the fact that R_{DARK} and R_{LIGHT} realize a functional role that other neural properties don’t realize. Stephen Schiffer has provided a very accurate characterisation of the notion of functional role. He writes:

*A functional role is simply any second-level property of first-level state-types possession of which entails that the state-type possessing it is causally or counterfactually related in a certain way to other state-types, to outputs, to inputs, or to distal objects and their properties.*²²

Let’s consider how functional roles would help us to define the semantic relevance criterion. Predicative judgments can be expected to enter in various kinds of causal relations. As we have seen, they are often caused by perceptual presentations of objects. They may also be caused by other predicative judgments, as when my judging of Fido that it is an animal is caused by my judging that it is a dog. They also cause other events, including other predicative judgments and certain sequences of behaviour, verbal and otherwise.

The neural properties of predicative judgments can be expected to figure in some of the causal laws that cover these causal transactions. Thus, in Mary’s case, *dark red* is nomically sufficient to cause R_{DARK} to be instantiated under C-conditions, and R_{DARK} is, in turn, nomically sufficient to cause “red” to be uttered under certain conditions. We may suppose also that R_{DARK} bears nomic relations to other neural properties. The functional roles that might enable us to specify the right disjunctions are features of the nomological pedigree of neural properties. The proposal is to specify the

relevant disjunctions of neural properties in terms of similarities between their nomological pedigrees.

Notice, however, that simply invoking functional roles does not provide us with a full account of how the task is to be accomplished. The problem is that each neural property realizes many functional roles. Suppose, e.g., that in Mary's case the nomological pedigree of R_{DARK} is the one represented by Figure 2.




Fig. 2

Then R_{DARK} will realize the functional role expressed by the open sentence

- (11) x is such that a presentation of a dark red object is nomically sufficient to cause x to be instantiated under C-conditions, A is nomically sufficient to cause x to be instantiated (under certain conditions), x is nomically sufficient to cause D to be instantiated (under certain conditions), and x is nomically sufficient to cause "red" to be uttered (under certain conditions).

But it will also realize the functional role expressed by

- (12) x is such that a presentation of a dark red object is nomically sufficient to cause x to be instantiated under C-conditions, B is nomically sufficient to cause x to be instantiated (under certain conditions), x is nomically sufficient to cause E to be instantiated (under certain conditions), and x is nomically sufficient to cause "red" to be uttered (under certain conditions),

as well as many others.

Our goal is to decide which disjunction of neural properties should be treated as the semantically relevant property of instances of R_{DARK} , and the present proposal is to specify the relevant disjuncts as the ones that realize the same functional role as R_{DARK} . But which disjunction we end up with is going to depend on which of the functional roles that R_{DARK} realizes we use for this purpose. For we can expect, e.g., that the functional roles expressed by (11) and (12) will be realized by different classes of neural properties. Someone who wants to decide which disjunctions of neural

properties are semantically relevant by reference to the functional roles that they realize would have to provide a specification of which functional roles are to be used for this goal – of which functional roles of neural properties are semantically relevant. And, needless to say, if the ultimate goal is a naturalization of the expression relation, this account will have to be provided in non-semantic, non-intentional terms.

We can take the task to consist in defining an equivalence relation on neural properties – “realizes the same semantically relevant functional role as” – that satisfies the following constraint:

- (13) Neural properties φ and ψ realize the same semantically relevant functional role iff for every predicative judgment j that instantiates φ or ψ , and every distal property δ that is nomically sufficient to cause either φ or ψ to be tokened under C-conditions, j is satisfied by every instance of δ .²³

Notice that R_{DARK} and R_{LIGHT} satisfy this criterion. *Dark red* is nomically sufficient to cause R_{DARK} to be instantiated under C-conditions, and *light red* is nomically sufficient to cause R_{LIGHT} to be instantiated under C-conditions. And every instance of R_{DARK} or R_{LIGHT} (i.e., each of Mary’s “red”-judgments) is satisfied by every instance of *dark red or light red* (i.e., by all red objects). I shall refer to the equivalence relation that satisfies this constraint as the *same role* relation. My goal in the remainder of this paper is to argue that the task of defining the same role relation may prove quite hard to accomplish. In particular, I am going to suggest that the task may face a difficulty that would invalidate the whole information theoretic approach.

7. THE PREEMPTION PROBLEM

Let’s consider once more the position that an account of the same role relation would occupy in the overall information theoretic project. The ultimate goal of the model is to provide an account of which properties are expressed by each predicative judgment. The basic idea of the information theoretic approach is to construct this account in terms of the nomological relations between properties of predicative judgments and distal properties. But I have argued that a plausible account along these lines would have to introduce some restrictions on which properties of each predicative judgment are to be used in this manner. This is the goal that an account of the same role relation is expected to attain. The properties expressed by each predicative judgment will be singled out by the nomological relations of the disjunctive properties defined by the same role relation.

The fact that the same role relation is to be put to this use imposes an important constraint on a satisfactory account of it. Suppose that the task of specifying the same role relation were *just as difficult* as the task of specifying the expression relation. That is, suppose that the facts that we need to invoke in the former task sufficed by themselves for determining which distal properties are expressed by each predicative judgment. In this situation, we would have to conclude that the information theoretic approach cannot succeed. For nomological relations cannot be used for specifying the expression relation until the same role relation has been construed. But if the facts that we need to invoke for specifying the same role relation sufficed for specifying the expression relation as well, the use of nomological relations would be preempted. We would only be able to specify the expression relation in terms of nomological relations if we already had a specification of it. The construal of the expression relation in terms of nomological relations would have to take as input what it is expected to deliver as output.

Hence the viability of the information theoretic model rests on an important assumption. The task of specifying the same role relation has to be *essentially easier* than the task of specifying the expression relation, in the following sense: the facts that determine whether two neural properties φ and ψ realize the same semantically relevant functional role cannot suffice for determining as well which properties are expressed by the instances of φ and ψ . I shall refer to the challenge posed by this assumption as the *preemption problem*. I am going to consider the difficulties faced by the task of specifying the same role relation, and suggest that the preemption problem may pose a serious obstacle to the viability of the information theoretic model.

In order to present the difficulties encountered by this task, I propose to look first at a crude account of the same role relation, and consider how it would have to be modified in order to become a plausible proposal. The basic thought of the crude account is that *any* difference between the nomic profiles of two neural properties will make them realize different semantically relevant functional roles. Thus, the crude account can be formulated as follows:

- (14) Neural properties φ, ψ , realize the same semantically relevant functional role iff φ and ψ bear the same nomic relations to distal properties, to behavioural properties and to other neural properties.

I want to argue that the crude account fails on all counts. A plausible account, I shall contend, will have to treat some neural properties as real-

izing the same semantically relevant functional role even if they bear different nomic relations to distal properties, to behavioural properties or to other neural properties. Let's consider each of these in turn.

7.1. *Nomic Relations to Distal Properties*

Mary's case illustrates that a plausible account of the same role relation will have to count some neural properties that bear different nomic relations to distal properties as realizing the same semantically relevant functional role. As we saw, *dark red*, but not *light red*, is nomically sufficient to cause R_{DARK} to be instantiated under C-conditions, and *light red*, but not *dark red*, is nomically sufficient to cause R_{LIGHT} to be instantiated under C-conditions. But I argued that a plausible account of the expression relation must have the resources for counting the disjunctive property $R_{\text{DARK-OR-LIGHT}}$ as the semantically relevant property of instances of each of its disjuncts. Hence a plausible account of the same role relation would have to treat R_{DARK} and R_{LIGHT} as realizing the same semantically relevant functional role.

This is not to say that none of the nomic relations that R_{DARK} and R_{LIGHT} bear to distal properties could enter in the definition of the semantically relevant functional role that they realize. The nomic relations that each of them bears to distal properties have many common features. And one may hope to find a non-semantic, non-intentional specification of which of these features should figure in the definition of the semantically relevant functional role realized by R_{DARK} and R_{LIGHT} . Notice, in particular, that the nomic relations that R_{DARK} and R_{LIGHT} bear to distal properties have the following common feature: both are such that a shade of red is nomically sufficient to cause them to be instantiated under C-conditions. Presumably, this feature could enter into a successful definition of the semantically relevant functional role realized by R_{DARK} and R_{LIGHT} . In general, if P_φ is the property expressed by instances of neural property φ , the following feature of φ 's nomological pedigree could play this role:

- (15) a certain kind of P_φ is nomically sufficient to cause φ to be instantiated under C-conditions.

Hence it may seem that if we have a non-semantic, non-intentional way of specifying this feature in each case, we could use some of the nomic relations between neural properties and distal properties in the definition of the same role relation.

Nevertheless, the possibility of using these nomic relations in the definition of the same role relation is not good news for the proponent of the information theoretic model. The suggestion is that we could use some

nomic relations between neural properties and distal properties if we had a suitable specification, for every neural property φ , of the property P_φ in (15). The problem is that providing this specification is precisely the ultimate goal that the information theoretic model is meant to attain. For P_φ is the property that instances of φ express. A specification of the semantically relevant functional role realized by φ that invoked the feature expressed by (15) would fall prey to the preemption problem. Performing one of the subtasks that the information theoretic model involves would require having already the account that the model is meant to provide. Of course, it may turn out that the feature expressed by (15) is not the only feature of the nomic relations between φ and distal properties that could figure in a successful definition of the semantically relevant functional role realized by φ . But if it were, then we would have to conclude that we can't use the nomic relations between neural properties and distal properties in the definition of the same role relation.

7.2. *Nomic Relations to other Neural Properties*

Nomic relations between neural properties pose a similar difficulty. I want to suggest that a plausible account of the same role relation would have to make room for the possibility that neural properties that bear different nomic relations to other neural properties realize the same semantically relevant functional role. Imagine, e.g., that in Mary's case there are two neural properties, A, B, that bear the same nomic relations to distal properties, to behaviour and to other neural properties, except for the following difference: A is nomically sufficient to cause R_{DARK} , but not R_{LIGHT} , to be instantiated (under certain conditions), and B is nomically sufficient to cause R_{LIGHT} , but not R_{DARK} , to be instantiated (under those same conditions). I want to suggest that this difference between the nomic pedigrees of A and B should be treated as compatible with the possibility that A and B realize the same semantically relevant functional role. We want it to be possible for A-or-B to be the semantically relevant property of the predicative judgments that instantiate each of the disjuncts. This might be the right thing to say if, e.g., A and B have the nomological pedigree depicted in Figure 3, i.e., if both are nomically sufficient to cause "saturated red" to be uttered by Mary, and *saturated red* is nomically sufficient to cause both A and B to be instantiated under C-conditions.




Fig. 3

The same possibility would have to be available if A and B also bore different nomological relations to distal properties. Thus suppose that only *dark saturated red* is nomologically sufficient to cause A to be instantiated under C-conditions, and only *light saturated red* is nomologically sufficient to cause B to be instantiated under C-conditions:




Fig. 4

The same reasons that I adduced in support of the hypothesis that Mary's "red"-judgments express the property *red* would also lend support to the hypothesis that her "saturated red"-judgments express the property *saturated red*. But that would seem to require treating A-or-B as the semantically relevant property of the instances of each of the disjuncts. If this is correct, it follows that some differences in the nomic relations between neural properties should not be counted as differences in their semantically relevant functional roles, even if combined with differences in their nomological relations to distal properties. We would want to say this whenever two neural properties bear otherwise identical nomic relations to two neural properties that realize the same semantically relevant functional role.

7.3. *Nomic Relations to Behaviour*

Let's now turn to nomic relations to behavioural properties. According to the crude account of the same role relation, two neural properties would be counted as realizing the same semantically relevant functional role only if they bore exactly the same nomic relations to behavioural properties. I am going to argue that this restriction is bound to result in unacceptable assignments of satisfaction conditions. The reason is that we cannot rule

out the possibility of nomic connections to behavioural properties that are mere noise from the semantic point of view.

Consider again Mary's case. We said that a correct account of the semantic relevance criterion should treat the disjunctive property R_{DARK} -or- R_{LIGHT} as the semantically relevant property of each of her "red"-judgments. In order for the crude account of the same role relation to yield this assignment, R_{DARK} and R_{LIGHT} would have to stand in exactly the same nomic connections to behavioural properties. But it is perfectly possible that this is not so. The behavioural sequences in which R_{DARK} and R_{LIGHT} issue cannot exhibit differences of a certain kind – of the kind that we consider relevant for semantic purposes. Otherwise we would presumably take Mary's "red"-judgments as having different satisfaction conditions. But that they don't exhibit differences of that kind does not mean that they exhibit no differences at all. Let's imagine, e.g., that those of her utterances of "red" that are caused by instances of R_{DARK} are consistently louder than those that are caused by instances of R_{LIGHT} :




Fig. 5

We would want to count this difference as compatible with the hypothesis that all of the judgments that she expresses with these utterances have the same satisfaction conditions. But the crude account of the same role relation would leave no room for this possibility. It would result in an account of the expression relation that assigns different satisfaction conditions to each of them depending on which of the two neural properties they instantiate. Hence a plausible account of the same role relation must have the resources for counting some differences between the nomic relations that two neural properties bear to behavioural properties as compatible with the hypothesis that they realize the same semantically relevant functional role.

Nevertheless, some of the nomic relations that a neural property bears to behavioural properties could figure in a plausible definition of the semantically relevant functional role that it realizes. In Mary's case, we could use the property of being nomically sufficient to cause utterances of "red" under certain conditions in our definition of the semantically relevant functional role of R_{DARK} and R_{LIGHT} . This approach would seem to yield the

right result in this case, if, as we are supposing, R_{DARK} and R_{LIGHT} are nomically sufficient to cause “red” to be uttered under certain conditions, and no other neural property bears this nomic relation to “red”.

This may seem to suggest a promising strategy for deciding which nomic relations to behaviour should enter into the definition of the same role relation. The idea would be to say that two neural properties realize the same semantically relevant functional role whenever they are nomically sufficient to cause the same predicate type to be uttered under certain conditions. But this strategy will not, in general, yield the right results. The proposal would fail in the case of ambiguous predicates. Thus consider the predicative judgments that I would express with utterances of “bank”. Some of them express the property *money-lending institution* and others the property *river side*. We may suppose that judgments of the first kind instantiate a neural property, B_{MONEY} , that bears the right nomic relation to the property *money-lending institution*, and that judgments of the second kind instantiate a neural property, B_{RIVER} , that bears that nomic relation to the property *river side*.




Fig. 6

If we treat two neural properties as realizing the same semantically relevant functional role whenever they are nomically sufficient to cause the same predicate type to be uttered (under certain conditions), $B_{\text{MONEY-OR-RIVER}}$ will turn out to be the semantically relevant property of each of my “bank”-judgments. This would result in an inadequate account, since each of my “bank”-judgments would end up expressing the disjunctive property *money-lending institution or river side*.

The strategy would only work if we defined predicate types in such a way that phonologically identical predicates that express different properties are treated as type-distinct. But the naturalist cannot achieve this goal by stipulation. The criterion according to which different tokens of “bank” are counted as type distinct cannot appeal to the fact that they express judgments with different satisfaction conditions. The naturalist would need a criterion in which only non-semantic, non-intentional facts are invoked. Providing an account of which behavioural types should be counted as semantically relevant is by no means a trivial task in the case of linguistic

behaviour. But the task would be even more complicated in the case of non-linguistic behaviour. The proponent of the information theoretic model would have to provide an account of which properties of a subject's bodily motions are to be counted as relevant for semantic purposes. And this account would have to steer clear of the preemption problem. It could not invoke facts that would suffice for specifying which properties are expressed by the judgment that causes each behavioural sequence.

8. CONCLUSION

Let's consider where this leaves us. I have argued that in order to obtain an information theoretic account of the expression relation, one needs to provide first a semantic relevance criterion – an account of which properties of each predicative judgment are to be counted as semantically relevant. Then I have argued that their neural properties cannot be expected to do the job. A plausible account of the semantic relevance criterion must have the resources for treating as semantically relevant properties that are disjunctive from a neurological point of view. Hence the proponent of the information-theoretic model has to provide an account of which disjunctions of neural properties are to be treated as semantically relevant in each case. I have suggested that the most promising strategy for providing this account is to appeal to the functional roles realized by neural properties. But pursuing this strategy requires providing an account, in non-semantic, non-intentional terms, of which functional roles are semantically relevant. Finally I have argued that providing this account can be expected to be a non-trivial task. For not all the features of the nomic pedigrees of neural properties can be taken into account for this purpose. The information theoretic project could only succeed if there were a principled way of telling apart the features of nomic pedigrees that should be taken into account from those that shouldn't. We could help ourselves to semantic information for deciding which criterion will do the job, but it has to be possible to specify the criterion in non-semantic, non-intentional terms.

Furthermore, the project would fail if specifying this criterion turned out to be *just as difficult* as deciding which property is expressed by each predicative judgment. If the information required for accomplishing the former task sufficed for accomplishing the latter, the information theoretic approach to the naturalization of the expression relation would have to be abandoned. For, in that case, providing an information theoretic account of the expression relation would require having already an account of the notion at our disposal. I have suggested that this may turn out to be the case.

Be this as it may, my main goal has been to argue that finding a suitable nomic relation is by no means the only substantial challenge that confronts the proponents of information theoretic semantics. In addition to that, they need to provide an account of which neural properties of each predicative judgment are to be counted for semantic purposes, and an account of which of these realize the same semantically relevant functional role. These tasks may turn out to be feasible, but I suspect that they add up to a heavier burden than some proponents of the information theoretic model had bargained for.

NOTES

* I am indebted to Paul Boghossian, Allan Gibbard, Eric Lormand, Peter Railton, Gideon Rosen, Crispin Wright, and Steve Yablo for their comments on drafts of this paper.

¹ When several properties P_1, \dots, P_n play this role with respect to a predicative judgment, it is satisfied by an object just in case the object instantiates one (or more) of these properties. Hence we can also speak of the satisfaction conditions of the judgment as being determined by the instantiation conditions of the disjunctive property $P_1 \vee \dots \vee P_n$.

² This is the term used by Fodor for the relation between a Mentalese predicate and the property that determines the satisfaction conditions of its tokens. Cf., e.g., *Psychosemantics*, p. 99, and 'A Theory of Content, I', p. 52.

³ See, e.g., 'A Theory of Content, I', p. 52:

[...] recent developments in "informational" semantics suggest the possibility of a naturalistic atomistic theory of the relation that holds between a predicate and the property that it expresses. Such a theory would, of course, amount to a good deal less than a complete understanding of intentionality. But it would serve to draw the skeptic's fangs since his line is that irreducibility and holism are *intrinsic* to intentionality and semantic evaluability. Given any suitably atomistic, suitably naturalistic break in the intentional circle, it would be reasonable to claim that the main *philosophical* problem about intentionality had been solved.

⁴ Cf., e.g., *Psychosemantics*, p. 97:

It is hard to see [...] how one can be a Realist about intentionality without also being, to some extent or other, a Reductionist. If the semantic and the intentional are real properties of things, it must be in virtue of their identity with (or maybe of their supervenience on?) properties that are themselves *neither* intentional *nor* semantic. If aboutness is real, it must be really something else.

Fodor has recently toned down his commitment to this kind of reductionism. Cf. his reply to Boghossian in 'Replies', pp. 271–72.

⁵ See his 'Psychosemantics, or: Where Do Truth Conditions Come From?', for an early proposal. In *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, he launched a new proposal that is further developed in "A Theory of Content, II".

⁶ It is a central aspect of Fodor's picture that beliefs, desires, and other propositional attitudes are to be seen as different ways of tokening the same mental (type) expressions. Thus, my believing of my car that it is dirty, and my wishing that it were dirty would

involve two different ways of tokening the same mental predicate. Cf. *Psychosemantics*, p. 17.

⁷ Cf., e.g., *Psychosemantics*, p. 98.

⁸ Cf. the Appendix to *Psychosemantics* for a discussion of what this second claim comes to.

⁹ For an early defence of the information theoretic model, see D. Stampe, 'Toward a Causal Theory of Linguistic Representation'. Another milestone in the information theoretic program is Fred Dretske's *Knowledge and the Flow of Information*.

¹⁰ Throughout this paper, I shall assume that token physicalism is correct. In particular, I shall assume that every predicative judgment is (token-)identical with an event in the subject's nervous system.

¹¹ See Fodor, *Psychosemantics*, pp. 112–14, for a defence of the fundamental status of these cases.

¹² Or rather, as Fodor puts it, to the property of tokening that predicate (Cf. *Psychosemantics*, p. 99). This is the shape both of his early proposal in 'Psychosemantics' and of the more recent one in *Psychosemantics* and 'A Theory of Content, II'.

¹³ Cf. *Psychosemantics*, pp. 99–102.

¹⁴ I shall follow Fodor's orthographic convention and use italics to name properties. Cf. Fodor, *Psychosemantics*, p. 160, fn. 5.

¹⁵ Cf. 'A Theory of Content, I', p. 60.

¹⁶ Cf. 'A Theory of Content, II', p. 90. Paul Boghossian has contested this claim. Cf. his 'Naturalizing Content', pp. 71–73.

¹⁷ Throughout this paper, I am assuming an account of events according to which they are not individuated in terms of properties. Donald Davidson has defended such an account. Cf. his 'The Individuation of Events'. For the opposing view, see, e.g., Jaewon Kim, 'Events as Property Exemplifications'. My argument could be reformulated in terms of a conception of events as property exemplifications. Then the problem would not be to identify the properties of a predicative judgment that are to be treated as semantically relevant, but rather to identify the event that each predicative judgment is identical with.

¹⁸ Cf. *Psychosemantics*, p. 17.

¹⁹ (7) is a trivial instance of (5).

²⁰ Of course, the satisfaction conditions of a predicative judgment need not coincide with the applicability conditions of the English predicate with which the speaker attempts to express it, as malapropisms illustrate. However, it seems natural to treat the notion of the satisfaction conditions of the judgment as identical with the notion of the applicability conditions of the predicate *as meant by the speaker at the time of utterance*.

²¹ N. J. Block and J. A. Fodor, 'What Psychological States Are Not', p. 160.

²² S. Schiffer, *Remnants of Meaning*, p. 21. As Schiffer goes on to suggest, "each functional role determines a unique *functional property*, viz., the property of having some property which has that functional role" (Ibid.).

²³ In fact, a somewhat weaker constraint would be more accurate. The idea would be to keep the 'only if' part of (13):

- (13*.1) Neural properties φ and ψ realize the same semantically relevant functional role only if for every predicative judgment j that instantiates φ or ψ , and every distal property δ that is nomically sufficient to cause either φ or ψ to be tokened under C-conditions, j is satisfied by every instance of δ ,

and to replace the 'if' part by:

- (13*.2) For every neural property φ , and every object x that satisfies the instances of φ , there is a property δ , such that $\delta(x)$, and for some neural property ψ that realizes the same semantically relevant functional role as φ , δ is nomically sufficient to cause ψ to be tokened under C-conditions.

REFERENCES

- Block, N. J. and J. A. Fodor: 1972, 'What Psychological States Are Not', *Philosophical Review* **81**, 159–81.
- Boghossian, P. A.: 1991, 'Naturalizing Content', in G. Rey and B. Loewer (eds.), *Meaning in Mind: Fodor and His Critics*, Blackwell, Cambridge, Mass., pp. 65–86.
- Davidson, D.: 1969, 'The Individuation of Events', in N. Rescher et al. (eds.), *Essays in Honor of C. G. Hempel*, Reidel, Dordrecht, pp. 216–34.
- Dretske, F.: 1981, *Knowledge and the Flow of Information*, The MIT Press, Cambridge, Mass.
- Fodor, J. A.: 1987, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, MIT Press, Cambridge, Mass.
- Fodor, J. A.: 1990, 'Psychosemantics, or: Where Do Truth Conditions Come From?', in William G. Lycan (ed.), *Mind and Cognition*, Blackwell, Cambridge, Mass., pp. 312–37.
- Fodor, J. A.: 1990a, 'A Theory of Content, I: The Problem', in *A Theory of Content and Other Essays*, The MIT Press, Cambridge, Mass., pp. 51–87.
- Fodor, J. A.: 1990b, 'A Theory of Content, II: The Theory', in *A Theory of Content and Other Essays*, The MIT Press, Cambridge, Mass., pp. 89–136.
- Fodor, J. A.: 1991, 'Replies', in G. Rey and B. Loewer (eds.), *Meaning in Mind: Fodor and His Critics*, Blackwell, Cambridge, Mass., pp. 255–319.
- Kim, J.: 1976, 'Events as Property Exemplifications', in M. Brand and D. Walton (eds.), *Action Theory*, Reidel, Dordrecht, pp. 159–177.
- Schiffer, S. R.: 1987, *Remnants of Meaning*, the MIT Press, Cambridge, Mass.
- Stampe, D. W.: 1977, 'Toward a Causal Theory of Linguistic Representation', in P. French et al. (eds.), *Midwest Studies in Philosophy*, vol. 2, University of Minnesota Press, Minneapolis, pp. 42–63.

Department of Philosophy,
University of Birmingham,
Edgbaston,
Birmingham B15 2TT,
England

E-mail: J.L.Zalabardo@bham.ac.uk