PHILIP PETTIT

DESCRIPTIVISM, RIGIDIFIED AND ANCHORED*

ABSTRACT. Stalnaker argues that, while the two-dimensional framework can be used to give expression to the claims associated with rigidified descriptivism, it cannot be used to support that position. He also puts forward some objections to rigidified descriptivism. I agree that rigidified descriptivism cannot be supported by appeal to the two-dimensional framework. But I think that Stalnaker's objections can be avoided under a descriptivism that introduces a causal as well as a descriptive element – a descriptivism in which the relevant descriptions are allowed to be, not only rigidified, but anchored in causal exposure to referents.

In 'Conceptual Truth and Metaphysical Necessity' (2003) and 'Assertion Revisited' (this volume), Robert Stalnaker sets out his views on the use to which the two-dimensional framework of modal thought can be put in thinking about issues of necessity and contingency, the *a priori* and the *a posteriori*, and about questions in the theory of representation that those issues presuppose. In this brief comment, I give a somewhat different casting to the matters involved, I characterise the main lines of his argument, and I try to rebut one charge that he levels against opponents.

Stalnaker's discussion assumes that what goes for representational language will go for representational thought and I should mention at the outset that I shall abstract in what follows from this assumption. In discussing the representational properties of thought, it is hard to distinguish the sort of thought that a non-linguistic animal like a dog may be able to instantiate and the sort that occurs in creatures like us. Presumably a dog has thoughts that are representational, if indeterminately representational at certain margins, so that a story has to be told about how its mental states get to have contents. And presumably what holds of the dog will hold at a certain level of us, however far we may be transformed by our linguistic and communicative competence. But the story to be told is certainly complex and I do not want to try to cover thought as

Philosophical Studies **118:** 323–338, 2004.

well as language in my discussion (Pettit, 1993). I shall concentrate, therefore, only on what Stalnaker says about language.

1. RIGIDIFIED DESCRIPTIVISM AND SOME COROLLARIES

According to standard descriptivism, a community which uses a name 'N' will associate a definite description with it – one, we may presume, that is cast in entirely general terms – and will intend to refer to whatever it is that satisfies that description. Such a term may fail of reference, as when nothing satisfies the description, or when too many things do. But normally there will be something unique, or uniquely selected by context, that answers to the description and in that case the term will refer to that entity.

While descriptivism has been criticised on many grounds in the last few decades, the most distinctive problem raised for the doctrine is associated with the work of Saul Kripke (1980). The Kripke problem is that the candidates that it identifies for various names in common usage do not intuitively fit. Suppose that 'Aristotle' is used to refer to 'the last great philosopher of antiquity'. This descriptive assignment – and any similar proposal – will be counter-intuitive, so far as it is readily conceivable that Aristotle, the very man we refer to by that name, might not have been a philosopher at all. Given that we can coherently entertain that possibility, how can the use of the term in our mouths and minds be governed by the description 'the last great philosopher of antiquity'?

This problem is distinctive among the other difficulties raised for descriptivism, because it has prompted the appearance of a new variant of the doctrine, sometimes cast as rigidified descriptivism. The idea in this approach is to allow that at least some names and referring terms are governed by rigidified descriptions: that is, by descriptions that refer to the same thing under various counterfactual imaginings. One version would say that the sort of description associated with the name 'Aristotle' is not 'the last great philosopher of antiquity' but rather 'the actual last great philosopher of antiquity'. We can imagine that Aristotle, so understood, might not have been a philosopher; that will involve contemplating the possible world where the person who in the actual world was the last great philosopher of antiquity is not a philosopher at all. We will think of him

in every possible world as the person who in the actual world is the last great philosopher of antiquity – the name will rigidly refer to that individual – and we will be able coherently to inquire after how things go in this or that possible world where he is not a philosopher.

Rigidified descriptivism, and some allied approaches, are particularly interesting, because they lead us naturally towards a two-dimensional way of thinking about possible worlds. For if Aristotle, the actual last great philosopher of antiquity, is not a philosopher in a certain possible world, w_1, that invites us to think about what we would say were w_1 to play the role of the actual world: to be the world where we are speaking, and where 'Aristotle' gets its semantic value determined. With our world playing the actual, interpretative role, 'Aristotle' is given such a semantic value that we say that in w_1 Aristotle – that very person – is not a philosopher; we evaluate the sentence 'Aristotle is the last great philosopher of antiquity' as false. But how would things stand were w_1 the interpretative world?

If we are to imagine ourselves or our counterparts using the term 'Aristotle' in w_1 in a way that sustains the descriptive presumptions that we by hypothesis employ – that is, referring to the last great philosopher of antiquity – then we must say that the name would refer to the different person in that world considered as actual, assuming there is one such person, who counts there as the last great philosopher of antiquity. In every world considered as actual – this is a phrase from Davies and Humberstone's path-breaking paper (1980) – Aristotle will refer, and refer rigidly, to the last philosopher of antiquity in that world; we restrict attention to those worlds where speakers remain compliant with the hypothesised, descriptive associations. There will be a connection between being the referent of 'Aristotle' and being the last great philosopher of antiquity, then, that any compliant user of the term will grasp. But despite its intimacy, that connection will still allow that no matter which world is considered as actual, the sentence 'Aristotle is the last great philosopher of antiquity' – assuming that 'Aristotle' has a reference – will come out there as true in that world and false in other possible worlds.

Rigidified descriptivism is readily generalised, with the rigidification move being introduced for more and more descriptions, so that they build in a variable actuality-index. Under that approach,

which Stalnaker describes as the generalised Kaplan program, 'practically all descriptive expressions of the language will have a variable character' (Stalnaker, 2003, p. 208; see also this volume, p. 309). That is, they will serve to pick out one referent when the 'actually' is interpreted relative to one world, another when it is interpreted relative to a different world, and so on.

If we go along with this rigidifying approach, then we may be willing to go a little further and describe a sentence like 'Aristotle is the last great philosopher of antiquity' as one that, however contingent, is *a priori* true in some sense. What certainly holds is that any speaker who remains party to the descriptive connection hypothesised will be in a position to recognise that it is true. The sentence will be *a priori* true modulo that hypothesis about what guides users of the name. It will be true in every world considered as actual, where speakers sustain the descriptive presumption in question.

If we adopt this way of speaking, then we may in turn be willing to go on and acknowledge the possibility of sentences that are only *a posteriori* but still necessarily true. They will only be *a posteriori* true so far as they are not true in every world considered as actual, even ones where speakers sustain our descriptive presumptions; they will be necessarily true so far as they are true in every possible world. Take the sentence 'Water is H_2O', where 'water' is governed by descriptive associations between being water and being colourless, odourless, potable and so on. The actual stuff that 'water' refers to here is indeed H_2O and, if water is picked out by speakers as the actual stuff that is colourless, odourless, potable, etc, then the word refers in every possible world to that stuff with those properties here and so to H_2O; thus the sentence 'Water is H_2O' is necessarily true. But there are some possible worlds where the stuff that satisfies the relevant descriptive associations is not H_2O – it may be, say, XYZ – and as we consider those worlds in the role of actual world, assuming that speakers there continue to use 'water' under the descriptive presumptions we employ, we can see that it will not refer in that usage to H_2O but to XYZ. Thus the sentence 'Water is H_2O' will be necessarily but only *a posteriori* true. The sentence 'Water is colourless, odourless, potable etc', by contrast, will be like the sentence 'Aristotle is the last great philosopher of antiquity'. It will be *a priori* but only contingently true.

But reader beware. The sense of '*a priori*' employed in these comments is available only so far as we assume that the community of speakers is indeed guided in its usage of 'Aristotle' or 'water' by the descriptive linkages hypothesised: only so far as we assume, in effect, that a certain sort of descriptivism is sound. No one who rejects that sort of theory will be impressed by the use of the two-dimensional framework to explicate notions like those of the contingent *a priori* and the necessary *a posteriori*. The two-dimensional framework may be useful, then, in spelling out the implications of rigidifying descriptivism but it does not offer any support for that move.

2. STALNAKER'S OBJECTIONS

Robert Stalnaker is not hospitable to the rigidified descriptivism sketched, which he finds elaborated in the work of Frank Jackson (1998) and David Chalmers (1996). He rejects any suggestion that the position derives support from the two-dimensional framework for modal thought, as we have just done, and he offers some objections to the position itself.

Rigidified Descriptivism is Not Supported by the Framework

Stalnaker locates the basic semantic role of the two-dimensional framework in the fact that it provides us with some important distinctions among the intensions that may be assigned as semantic values to terms and sentences. For example, it gives us the distinction between primary and secondary intensions, in David Chalmers's phrasing.

Take the name 'Aristotle' and consider the fact that in different worlds considered as actual – whether or not they are worlds that sustain the descriptivist presumptions linking the bearer to the last great philosopher of antiquity – it will be associated with a different referential condition and will have a different referential value. And take any sentence that uses the name 'Aristotle' and consider the fact that in different worlds considered as actual it will be associated with a different truth condition and will have a different truth value.

The secondary intension of such a term or such a sentence at any world is the reference condition or the truth condition that it

has at that world. This is the function that determines the referential value – the referent – of the term and the truth value of the sentence, both at that world and at other worlds that are possible relative to that world. What then is the primary intension? It is the yet more abstract function which for any world considered as actual fixes or yields – in a purely formal sense – the secondary intension of the term or sentence at that world; and in fixing the secondary intension fixes also the referent of the term and the truth-value of the sentence at that world.[1]

According to Stalnaker, the two-dimensional framework is useful in enabling us to make distinctions like that between primary and secondary intensions among the candidates that we survey as we ask about the semantic values of different terms. In that sense, it underpins the enterprise of assigning worldly items to linguistic terms. But he denies that the framework gives any particular support to rigidified descriptivism.

Rigidified descriptivism holds that terms like 'Aristotle' and 'water' have both sorts of intensions, primary and secondary, as their semantic values. It suggests that with such a term, we speakers have an immediate grasp of the primary intension: the primary intensions 'represent the cognitive values of our thoughts' (Stalnaker, 2003, p. 213; see also this volume, p. 309). And it holds, then, that so far as we intend to refer to whatever satisfies that primary intension in the actual world – to whoever is actually the last great philosopher of antiquity, to whatever substance is colourless, odourless, potable etc – we will succeed, indirectly as it were, in fixing on the secondary intension of the term.

What goes for terms carries over to sentences. Each sentence that employs a rigid term will have a primary intension that we grasp as speakers and it will enable us to pick out a secondary intension that that primary intension determines in our world. *A priori* true sentences will have primary intensions that are true at all relevantly constrained worlds considered as actual;[2] necessarily true sentences will have secondary intensions that are true at all possible worlds from the interpreting perspective of this, the actual world.

Stalnaker agrees that the two-dimensional framework can be used to give expression to the claims associated with rigidified descriptivism. But he wants to insist, as we have done, that it cannot

be used to support those claims. He does this by showing how the framework can also be used to express the claims associated with other views.

The view which he favours himself would assign as the semantic value of a term like 'Aristotle' or 'water' only a secondary intension: the intension relevant in the actual world. And equally it would assign as the semantic value of a sentence that uses such a term only the secondary intension of the sentence: only its truth condition at the actual world. What do we learn on recognising that had a different world been actual, the term would have had a different reference condition and a different referential value? We learn only about the different semantic assignments that would have held 'if the world had been different in certain ways' (2003, p. 209). We do not learn anything about the semantic assignment that the sentence has here in this world.

How do we know, then, that the secondary intension of the term would be such and such in the event of this or that world playing the actual-world role? Not on the basis of access to the primary intension that would be guaranteed by our competence as speakers. Rather on the basis of a mixture of things that we know, some perhaps derivable from habits associated with competence, but others of a decidedly empirical character (Stalnaker, 2001, pp. 155–156).[3]

Under Stalnaker's view, there will be 'no account of conceptual truth – truths knowable a priori on the basis of the semantics of the language': 'The framework, on this interpretation, gives us no account of *a priori* truth' (2003, p. 213). This view and rigidified descriptivism, then, will 'tell very different stories about necessary *a posteriori* truths' (*ibid*., p. 211). Under the descriptivist story, the necessary *a posteriori* sentence has a necessary secondary intension and a contingent primary intension as its semantic values. Under Stalnaker's alternative view, it only has a secondary intension as its semantic value – a necessarily true one – and the contrast between it and other necessary utterances and thoughts comes out in the fact that it would have had a necessarily false secondary intension as its semantic value had the actual world been different in certain ways. 'To say that a primary proposition associated with a sentence was necessary would be to say that the sentence would express a

truth whatever it meant, and that notion, of course, will have no application' (Stalnaker, 2001, p. 155).

Stalnaker's claim that the two-dimensional framework does not give any support to the rigidified descriptivist position is entirely persuasive, reinforcing the line of argument at the end of the last section. There can be no quarrel with his illustration of how the framework might be used to spell out the implications of his alternative view just as well as being used to spell out the implications of the descriptivist alternative. The two approaches are, as he says, 'two interpretations of the abstract two-dimensional framework' (2003, p. 202). I don't think that Chalmers or Jackson need disagree, even though Stalnaker sometimes suggests that they do.

Rigidified Descriptivism is Independently Unappealing

But Stalnaker claims more than that rigidified descriptivism is not particularly supported by the two-dimensional framework. He also wants to argue that that approach is inherently unattractive.

His argument begins with an important distinction between two forms of semantic enterprise. The one, which he calls descriptive semantics, is concerned with which items in the world should be assigned as semantic values to which terms in the language under discussion. The other, foundational semantics, deals with the question of why the items assigned have a claim to be assigned to those terms: it seeks to explain the connections that the other enterprise describes.

The two views discussed so far are exercises in descriptive semantics only. The one, as we saw, argues that many terms and sentences have both primary and secondary intensions as their semantic values, the other that they have secondary only. But what is to be said on the foundational side in favour of those different patterns of assignment? And does it argue for one over the other?

Stalnaker takes the sort of assignment – the descriptive semantics – he prefers to be based on a straightforward, causal-informational story. Why does this term refer to that item; why is it paired off with that item in the assignment given? Because that's the item that figures in the causal ancestry – and perhaps in the continuing causal control – of the term's usage. The story can hold, not just for names, but for terms of all kinds; it 'takes the kind of reference fixing

illustrated with proper names as an account of how expressions in general – predicates as well as names – get their reference' (2003, p. 212).

What of the foundational semantics that would serve in defence of the rigidified descriptivism? The natural, descriptivist thing to say is that certain names like 'Aristotle' and 'water' are semantically associated with certain primary and secondary intensions so far as users intend to use the names according to those primary intensions – intensions they can express in general terms – and the secondary intensions are those intensions that happen to satisfy the primary intensions. But this, Stalnaker thinks, would be to pass the 'buck from names to general terms', on the grounds that 'one still needs an account of how the general terms get their meaning' (*ibid*., p. 211). The idea is that if general terms are capable of having their semantic values fixed independently of descriptions, then there is no reason why this should not also be true of names and singular terms.

He argues, in view of that consideration, that if a foundational semantics is to give robust support to the sort of position outlined earlier, then it had better be globally descriptivist, applying to all the terms in the language. Under global descriptivism, we think of the terms of a language being networked with one another in descriptive sentences, with speakers intending to refer to corresponding items. We imagine that confronted with the world as a whole, the terms so networked will each have a determinate semantic value fixed by the way the world is, assuming that they are intended to have values which make the networking sentences come out as true. 'Roughly, the terms refer to whatever things, properties and relations do the best job of making the theory, or at least as much of it as one can, true' (*ibid*., p. 212).

There is a well-known permutation problem that arises for global decriptivism, so characterised (Putnam, 1982; Lewis, 1984). This is that consistently with making the networking sentences come out as true the terms might be mapped onto items in the world, now according to one assignment, now according to a permuted version of it. If the original assignment has 'big' refer to bigger things, for example, 'small' refer to smaller things, then the permutation might reverse this and, provided it introduced compensating changes elsewhere, still manage to make the networking sentences come out as

generally true. The possibilities of such permutation are endless and so it appears that global descriptivism is not an attractive option in foundational semantics.

While remaining 'sceptical', Stalnaker concedes that this problem may be overcome if global descriptivism is constrained, as for example David Lewis wants to constrain it (Lewis, 1982). Lewis's approach would require, not just that the items assigned to various terms make as many as possible of the relevant sentences true, but also that they are in a certain sense natural in kind (for a variant approach see Pettit and Stoljar, forthcoming). But Stalnaker thinks that no matter how it is constrained, such a position will fall foul of two objections.

The first is that it would give rise to a troublesome holism and solipsism, with speakers each referring to the satisfiers of the different sets of descriptions that they happen to endorse. This objection is sourced more or less directly in a Quinean resistance to the idea that there is a divide, robust and community-wide, between the sentences that guide us in how we use certain terms – sentences the truth of which is presupposed to that usage – and sentences that do not play such a role. If descriptivism is true then our usage of terms is guided by the sentences we adopt. And so, in the absence of a distinction between privileged and unprivileged sentences, we must conclude that each of us is likely to be guided by any of the sentences we hold true – hence holism – and that we are guided in a different way from others: hence solipsism.

Stalnaker's second complaint is more original and more pointed. The objection is that on the approach in question we speakers don't ever get into direct touch with items in the world (2003, p. 213; see also this volume p. 317). The contents that we immediately grasp are primary intensions; 'it is the primary intensions that represent the cognitive values of our thoughts – the things that we understand when we know what we are saying and thinking.' As for the reference conditions and the truth conditions that those primary intensions pick out in the world as it actually is, we only know them indirectly: 'they are uniquely determined, as a function of the facts, but not something to which we have cognitive access'.

This indirectness objection begins from a traditional distinction between two ways of knowing something: by acquaintance – by

being directly exposed to it, as it were – and by description. The idea is that we don't know any individual thing or property in itself but only as that thing or property, whatever it may be in itself, that satisfies a certain description. As he puts it (2003, p. 212), 'we don't refer to and think about particular individuals, as contrasted with whoever it is that fits some description. We don't describe things in terms of ground-level properties and relations, but only in terms of whatever properties and relations are the ones that best fit the abstract structure given by our uninterpreted theory.'

The complaint extends from things and properties to truths. We will assent to the proposition that is expressed by a sentence like 'Water is found on other planets' only in the way in which we may assent to what we know as 'the content of the first assertion that Napoleon made to Josephine after he was crowned emperor' (*ibid*., p. 213). The proposition will not be there, at our direct disposal, as something we can assert or believe in itself. It will be identified only as something of which we can assert that it is true, without in an important sense knowing what truth we are asserting.

3. RESPONDING TO THE OBJECTIONS

I think that the objections can be avoided under a global descriptivism that is constrained in a different manner from Lewis's. The theory may not deserve to be characterised as descriptivism, since it introduces a causal as well as a descriptive element. It amounts to a sort of descriptivism in which the relevant descriptions are allowed to be, not only rigidified, but anchored in causal exposure to referents.

This theory is presented as a way of dealing with the permutation problem in joint work with Daniel Stoljar (Pettit and Stoljar, forthcoming). I shall try to illustrate the approach here by telling a story as to how I as a speaker may have learnt the use of the term 'blue'. The story should suggest a response to the Quinean worry that no single set of sentences can ever be selected as those that play the role, in a robust and community-wide way, of guiding speakers in the use of their terms. And, more particularly, it should silence the concern that any terms that are guided in that way will allow us only an indirect sort of contact with their referents.

I set out the story in ten stages. It is in need of further elaboration and defence, of course, but it may at least serve to indicate a line of response to Stalnaker that those in the position of Chalmers and Jackson may find congenial.

1. Imagine that I am trained to use 'blue' under the causal impact, now in this instance, now in that, of the blueness property. My parents or teachers point out blue things to me, set them in contrast to things of other colours, and make the colour blue into the salient property to ascribe with the predicate 'blue'.

2. This training, if it is to be successful, requires that on being presented with blue things in the course of my tutelage, I take them under the relevant aspect: as things, to use a language I will not yet have mastered, that look blue or that look blue, at least, under conditions that I am given no independent reason to treat as abnormal. (On how to interpret 'looks', see Pettit, 2003.)

3. This means in turn that the causal linkage establishing the blueness property as the referent of the term will have an effect that is epistemically resonant; its effects will not be confined to the subpersonal arena. The linkage will tune me to believe now of this object that I take to be blue, now of that object that I take to be blue, that it looks blue, though I may not be in a position to express the belief in such words. Otherwise put, it will lead me to believe *in sensu diviso* – in a case-by-case way – the general proposition linking the blueness of an object with its looking blue.

4. If I have the requisite concepts available, I will therefore be in a position to recognise by reflecting on my practice – without recourse to further empirical inquiry – that blueness is linked in this way with looking blue. And I will be in a position to recognise that as the things I am disposed by training to call 'blue' all look blue (at least in normal conditions), so training disposes me – by a sort of stop clause – to call nothing else blue. I will be able to conclude, then, that blueness is a property of things that is associated with their looking blue (at least under normal conditions).

5. Will I be able to see, beyond this, that blueness is that property – presumptively, that unique property – which lies at the causal origin of the use of 'blue' on my part, and indeed that of my fellows? I think so. Assume, as seems right, that without yet having a word for 'cause' an intentional subject may yet have causal beliefs. The

exercise of learning the referent of 'blue' by causal exposure to blueness will have sensitised me, not just to the looks-blue aspect of blue things, but to their being things under the causal impact of which I was inducted in the use of 'blue'; the exercise will have alerted me to their ostensive as well as their looks-blue aspect. And so I should also be able to recognise on reflection, if I have the requisite concepts available, that blueness is that property which causally interacts with the likes of us and makes bearers look blue (in normal conditions); it is that ostensive (ostended and ostendable) property which makes things look blue (in normal conditions).

6. Now take the sentence: 'Blueness is that ostensive property that makes things look blue in normal conditions'. Under the story told this will be *a priori* true for me and my fellow-speakers, for it will be guaranteed to be true by the practice that we share with one another. The primary intension of the sentence is bound to be true at every world considered as actual where I and others remain faithful to that practice. There may be other things that are *a priori* true for us users of the term 'blue' but this sentence surely has to figure amongst them.

7. Assume, for purposes of illustrating relevant possibilities – the assumption fits, as it happens, with the story sketched (Pettit, 2002, p. 5) – that 'blue' is used rigidly and that it refers in the actual world to B1, and in other worlds to different properties. The sentence 'Blueness is B1', under this assumption, will be necessary but *a posteriori*. It will have a necessary secondary intension but a contingent primary intension.

8. That being so, of course, the sentence 'Blueness makes things look blue', *a priori* though it is, will be only contingently true. For any world considered as actual where speakers sustain the practice established among us, the sentence will express a truth; certainly there will be a truth expressed by the more qualified, 'Blueness, if there is such a property, makes things look blue'. But still there will be possible worlds where blueness as we understand it – B1 – does not make things look blue.

9. And so to the *denouement*. The story sketched, which is surely quite plausible, would vindicate something close to the rigidified descriptivist line, giving us intuitive candidates for the role of *a priori* truths. Nor is the story exposed to Stalnaker's objection of

making the world only indirectly accessible. If it is correct, then I will be able to think of the blueness property as *that* (ostended) property, where the property in question is directly available to me – I can ostend it, after all – but available only in virtue of making things look blue and of tuning me in the process to form various beliefs. I will not be restricted to thinking of it indirectly as whatever actual property, I know not what, that belongs to things which look blue (at least in normal conditions).

10. The story told with 'blue' can be paralleled by similar stories for any terms that are introduced on the basis of causal contact with the properties or other entities that get established as their referents. Given the permutability problem that arises for unconstrained global descriptivism – and assuming that the problem is not otherwise avoided – there must be some terms introduced on a causal basis. Not all terms need be introduced in that way, of course; some may be introduced by theoretical definition in other, pre-established terms. But some certainly will.

For those terms that are introduced in the manner illustrated, then, something close to descriptivism will hold true. The descriptivism on offer will build an actuality-index into the relevant descriptions; it is in that sense a rigidified descriptivism. But it will also require that the descriptions are endorsed by speakers in a manner that presupposes causal contact with referents and in that sense is an anchored as well as rigidified doctrine. The descriptions do not just mention or imply such causal contact in the manner of causal descriptivism, as it is sometimes called; they are endorsed by speakers on the basis of such causal contact.

A descriptivism that requires speakers to have had causal contact with some of the objects of which they speak may not strictly be deserving of that name. It may be better cast as a mixed doctrine that involves causal as well as descriptivist elements. But that does not really matter for our purposes. The doctrine sketched is relatively close to the position espoused by Stalnaker's opponents. And certainly it is close enough to raise some questions about how deep his criticisms go.

NOTES

* I am grateful to a number of people at the ANU conference in February 2002 for their comments. I was particularly helped by remarks on the written draft that I received from Martin Davies and Scott Soames.

[1] Think of the primary intension of a sentence as the array of T′s and F′s in the familiar matrix which has worlds of interpretation represented in the rows and worlds of evaluation in the columns. That array determines for each world considered as actual the row of T's and F's there: this is the secondary intension of the sentence at that world considered as actual. And that row determines in turn the T or the F in the box on the diagonal: that is, the T or F in the box where the world of evaluation is the world that is being considered as actual; this is the truth value of the sentence at that world.

[2] This may not be the happiest way of putting the point. The primary intension fixes for every world considered as actual the secondary intension – the truth condition – that the sentence as used at that world has. And if the sentence is *a priori* true, then at every world considered as actual the truth condition assigned to that sentence at that world will be fulfilled at that world; there will be a T in the box on the diagonal.

[3] Those mixed sources of knowledge might enable us to build up a model of the primary intension of the term, so far as they let us see its secondary intension in different worlds considered as actual and let us therefore reconstruct the primary intension: the abstract function corresponding to the mapping whereby different worlds considered as actual are associated with different secondary intensions. But this grasp of the primary intension would not play a role in guiding usage of the kind that it plays under the descriptivist story; the primary intension would not belong to the sentence in the manner of a semantic value.

REFERENCES

Chalmers, D. (1996): *The Conscious Mind: In Search of a Fundamental Theory*, New York: Oxford University Press.

Davies, M. and Humberstone, L. (1980): 'Two Notions of Necessity', *Philosophical Studies* 48, 1–30.

Jackson, F. (1998): *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: Oxford University Press.

Kripke, S.A. (1980): *Naming and Necessity*, Oxford: Blackwell.

Lewis, D. (1984): 'Putnam's Paradox', *Australasian Journal of Philosophy* 62, 221–236.

Pettit, P. (1993): *The Common Mind: An Essay on Psychology, Society and Politics, Paperback Edition 1996*, New York: Oxford University Press.

Pettit, P. (2002): *Rules, Reasons, and Norms*, Oxford: Oxford University Press.

Pettit, P. (2003): 'Looks as Powers', *Philosophical Issues* (supp. to *Noûs*) 13.

Pettit, P. and Stoljar, D. (2002): 'How to Solve Permutation Problems', *Philosophy Program, RSSS, Australian National University*, Canberra.
Putnam, H. (1982): *Reason, Truth and History*, Cambridge: Cambridge University Press.
Stalnaker, R. (2001): 'On Considering a Possible World as Actual', *Proceedings of the Aristotelian Society* 75 (suppl), 141–156.
Stalnaker, R. (2003): 'Conceptual Truth and Metaphysical Necessity', in *Ways a World Might Be: Metaphysical and Anti-Metaphysical Essays* (pp. 201–215), Oxford: Oxford University Press.

Department of Politics
034 Corwin Hall
Princeton University
Princeton, NJ 08544-1012
USA
E-mail: ppettit@princeton.edu