

Intentional States

Sander Voerman 2006-02-01

One can summarise the philosophy of mind as the business of understanding two general features of the mental: *intentionality* and *phenomenality*. In “Phenomenality, What It Does, and What That Means” (Voerman 2006, *PDM* henceforth) I have focused on the latter attribute. However, in order to deal with phenomenality I have already incorporated some theory with regard to intentionality as well: I have assumed a concept of concepts as semantic entities referring to properties, I have related concepts to belief, recognition and motivation, I have distinguished between phenomenal and practical concepts and I have distinguished two modes of individuating practical concepts.

In this essay I am going to focus on intentionality, and build on the concepts and distinctions from *PDM* in order to arrive at a theory of intentional states. On this quest, there will be a leading role for two questions. First, what does the *content* of intentional states consist in? Second, how does this content relate to our linguistic, folk-psychological *interpretation* of intentional states as *propositional attitudes*?

My approach is to consider two opposing views on the latter question, which each have their own drawbacks, and then to formulate my own position with the aid of the framework from *PDM* in an attempt to get rid of the drawbacks.

1 Casting Doubt on the Relational View of Propositional Attitudes

Robert Matthews distinguishes between a *relational* and a *nonrelational* conception of propositional attitudes (Matthews 1994), which I shall abbreviate as RPA and NPA respectively. According to RPA, a propositional attitude is a relation between an agent and a proposition, where the relation reflects the attitude (belief, desire, etc.) in a manner that implies some psychological or epistemic access to a semantic entity, the proposition. In contrast, NPA denies that a propositional attitude bears any such psychological or epistemological relation to a proposition. On NPA, what is ‘propositional’ about propositional attitudes is that they can be *described* using propositions, but such descriptions,

it is argued, are underdetermined by the attitudes themselves. Thus, we are once more dealing with an underdetermination problem.

Knowing Your Own Thoughts

Let us refer to those intentional states that can be described as propositional attitudes as *thoughts*. This may or may not exhaust all intentional states, depending on assumptions that I will not discuss here. So thoughts include beliefs, desires, and so forth. RPA is a form of realism about propositional attitudes in the sense that it includes propositions in the identity conditions for thoughts: the content of a thought is a proposition, and if its content had been different, it wouldn't have been the same thought.

This realism has been criticised by Davidson. In his words, RPA is the view that there are objects of thought that are "present to the mind" (Davidson 1989). These objects are supposed to be propositions and their constituents. But according to Davidson, the entire idea of objects present to the mind is mistaken.

The argument runs as follows. In general, when I think a thought, I know what thought I am thinking. Since RPA takes propositions to be part of the identity conditions of thoughts, this would mean that I have to know what proposition is the object of my thought. However, this is may not be as easy as it sounds:

Yet there seems to be no clear meaning to the idea of knowing which object one has in mind. The trouble is that ignorance of even one property of an object can, under appropriate circumstances, count as not knowing which object it is. (Davidson 1989 [2001:54])

Thus, RPA requires that we can construe propositional objects in such a way that in every situation where a thinker knows what he's thinking, he also has knowledge of all properties required to identify the related propositional object.

Davidson's criticism is directed at both the Russellian tradition, which holds that some of the constituents of propositions are the actual objects *referred to* by thoughts or sentences, and the Fregean tradition, which holds that propositions are composed out of mediating elements that *determine* the referents. Against a Russellian version of RPA the objection is straightforward. On such a view, the objects "present to the mind" of a single thinker would

have to be all sorts of objects everywhere in the world. This clearly allows for situations where a thinker would misidentify the propositional object of his thought by misidentifying one of its constituents, with the implausible implication for RPA that in such a situation a thinker would not know what he is thinking.

Against Fregean versions of RPA the criticism of Davidson is less clear. If we take as 'Fregean' all views that maintain that all reference is always *mediated* by things that propositions are composed of, we still have a very wide and heterogeneous landscape of theories about propositions. Could not one such theory provide us with semantical entities none of whose identifying properties escape the thinker (upon reflection) when they constitute the content of his thought? Davidson's verdict is that such a thing is impossible:

[P]hilosophers [...] have sought objects that, like Hume's impressions and ideas, 'Are what they seem and seem what they are' – that is, have all and only the properties we think they have. Alas, there are no such objects. (Davidson 1989 [2001:54])

He maintains that "every object has an infinity of logically independent properties" and waves attempts to restrict the identity conditions of thoughts to 'essential' or 'internal' properties of their objects. But I am not very convinced by the arguments offered for this position. For example, against Dummett's requirement that only internal properties of senses need be transparent to the thinker, Davidson suffices to respond that it "is not clear on what non-question-begging principle such properties are to be told from others". Davidson does not give us an independent argument why all such attempts should be doomed to beg the question, so at best this shows that the burden of the argument lies with proponents of Fregean versions of RPA to show that it *can* be done, but there is at least in Davidson 1989 no convincing argument that it *cannot* be done. At any rate, the case against Fregean RPA seems far less conclusive than that against Russellian RPA.

2 The Case for the Nonrelational View

Let us now examine Davidson's nonrelational account of propositional attitudes (NPA). This account is based on an analogy with the theory of measurement, an analogy that has been

further developed by David Matthews (1994). I will discuss the two authors separately, because as we shall see, their versions of NPA turn out to be crucially different.

Davidson on the Measurement Analogy

The purported analogy is between the assignment of *numbers* to properties such as weight or temperature, or any property that can be measured in some way or other, and the ascription of propositional attitudes to agents. The crucial notion here is that of *measurement scale*. If a diamond measures 345 grams then although we assign the number 345 to the weight of the diamond, we do not believe that this number is related to the diamond – at least not any more than any other number would be, for the number depends on the scale of grams. The scale allows us to preserve relations between different objects by means of relations between numbers: thus, if one object weighs twice as much as another, the gram scale assigns a number to the former twice as big as that assigned to the latter. The ounce scale would have done the same, except that it would have employed different numbers.

Davidson suggests that the same applies to thought. There is no such thing as the object of a thought to which the thought bears a relation as proposed by RPA any more than a diamond is related to the number 345. Instead, there are objects that can be used to describe thoughts by preserving the relations between thoughts in terms of some scale or system of measurement. These objects, in Davidson's preferred variety of NPA, are not sentences or propositions, but *utterances*.¹

It is all part of his general, highly economic approach to the mental. Wary of abstract entities such as properties, propositions and even languages, Davidson gives us a minimalist picture of attitude attribution which directly mirrors the measurement scenario:² all we need

¹ Conversely, when discussing Davidson I should speak of *utterance attitudes* rather than propositional attitudes and *NUA* instead of NPA. However, I prefer not to.

² If the analogy fails in any way it is not because we should be more realistic about mental state attributions, but rather that we should be even less realistic about the systems of measurement employed. Measurement in the physical sciences at least involves standardised formal systems based on conventions and strict equations. Interpretation of the mind, in the Davidsonian view, involves no systems so rigid and systematic, but a much more loose dynamic of passing ideolects. Davidson's famous claim that languages do not exist could be interpreted in this context as the view that for the

and should say about states of mind is that they bear relations to each other, and that we use utterances to track those relations, utterances that only apply to those relations against the idiolect of the utterer at the time of the utterance (the measurement scale).

A nonrelational view about *utterances* attributed to intentional states is not so shocking. After all, many different utterances could describe the same state, and a single utterance could have described a very different state if the utterer used language differently and had made different utterances about other mental states in the past. So that we should adopt NPA regarding utterances is not such a big deal. The problem is that Davidson denies that there is anything *else* – most notably, *propositions* – with regard to which we could adopt an RPA:

[W]hy not go on to say that since utterances have determinate meanings, and it is meanings that match up with belief states, the objects we are naming when we utter content sentences in the context of attributing beliefs just are the meanings of those content sentences, that is, propositions? [...]

The trouble we have been concerned with here sprang from the identification of the object used to characterize a state of mind with an object that the mind ‘knows’ or ‘is acquainted with’, an ‘object of thought’. If we avoid this identification, we can stay out of the difficulties I have been exploring. But if we avoid this identification, neither do we gain anything by the steps from utterances to sentences to meanings to propositions. (Davidson 1989 [2001:64])

In other words, Davidson would rather have us stick with utterances, but should we introduce propositions, then we are at least bound to adopt an NPA towards them as well. But why? The measurement analogy seems convincing for utterances, maybe even for sentences, but propositions are introduced by philosophers precisely to cancel out ways in which different utterances or sentences can express the same thing. So it is not at all clear

mind, there aren't even really scales of measurement. However, I shall not pursue this line of thought here.

that the analogy must survive at the level of propositions. Of course, the difficulties Davidson warns us to stay out of might be those I discussed in the previous section, but we only confirmed them for Russellian RPA, so that leaves the doorway to Fregean RPA open.

Matthews on the Structure of the Representation Space

Robert Matthews has elaborated further on what propositional attitude ascriptions should look like according to the measurement analogy (Matthews 1994). In his terminology, which stems from measurement theory, this is a question about *representation space*. Note that when talking about representation in this sense, we are not talking about a way in which thoughts represent objects, properties or states of affairs in the world. Instead, we are talking about how *descriptions represent thoughts*. From the NPA point of view, thoughts are not the *representers*, but the *represented*. If this sounds odd to you then rest assured that it sounds odd to me as well and that we will get back to this later on.

As we have seen, Davidson wanted very much to restrict the ascriptions to utterances, which means that the representation space for thoughts would be an utterance space – a space in which each point corresponds to a different utterance. However, we have also seen that when you restrict the measurement analogy to utterances, then RPA for propositions hasn't really been refuted, much less replaced by an alternative. Furthermore, utterance NPA might prevent us from getting into all sorts of trouble, but it also prevents us from saying interesting things. And as Matthews points out, there are two interesting things in particular that we would like to be able to say when we adopt a measurement theoretic approach: first, we want to formulate a *representation theorem*, which specifies what properties of the empirical system – in our case, the thinking mind – are preserved by a chosen scale of measurement defined on the chosen representation space, and second, we want to formulate an *admissible transformation*, which relates the chosen scale to all other scales that measure the same properties of the empirical system. In fact, without the two specifications, the measurement analogy remains largely that: an *analogy* that might have some intuitive appeal but is really open for question on all counts.

In order to at least come to a first sketch of how there could be such specifications for NPA, Matthews offers us a structured account in terms of the sentences and propositions that Davidson was so afraid of:

This is the construal of the representation space that I propose to adopt here. I shall take points in the space to be ordered pairs $\langle a_i, \langle s_j, r_k \rangle \rangle$, consisting of an attitude-type a_i and what I shall call a designated proposition $\langle s_j, r_k \rangle$, where r_k is a Russellian proposition and s_j a sentence-type, a token of which in a particular (unspecified) context serves to designate (express) that Russellian proposition. The designated proposition $\langle \text{Tully is bald}, \text{Tully's being bald} \rangle$, for example, consists of the Russellian singular proposition of an individual, Tully, having a property, being bald, paired with a sentence-type, *Tully is bald*, a token of which in a particular context expresses that proposition. (Matthews 1994:136)

Note that the Russellian propositions are 'designated' with sentences in order to accommodate the Fregean intuition that even though Tully and Cicero are one and the same, "X believes that Tully is bald" and "X believes that Cicero is bald" need not have the same truth-value within the same context. Furthermore, Matthews argues that sentences would not have sufficed by themselves because of *indexicals*: a sentence type such as *he believes he's sick* can express different propositional attitudes in different contexts of utterance. Sentences and Russellian propositions contribute fine-grainedness in different areas that both seem relevant to the individuation of propositional attitudes.

Given this framework, what can there be said about representation and transformation? First of all, let us note that the two specifications will be related. If an attribution A_1 against a scale S_1 can be transformed into A_2 against S_2 , then clearly the preservation of empirical properties must be sought among those features that A_1 and A_2 *share*. Preservation implies invariance under admissible transformation. Conversely, invariance under admissible transformation of a feature is a clue that the feature preserves an empirical property.³

³ At this point, it is not clear to me whether this is an abductive inference based on the aforementioned implication in the reverse direction, or whether this inference is itself deductive and thus more than just a 'clue'.

Since points in the representation space are determined, according to Matthews, by attitude type, sentence type and Russellian proposition, we can begin to ask in which of these three dimensions points may differ if they are related by admissible transformation. Matthews does not see much room for variation in attitude types across admissible transformation. I agree with this and I think even Davidson would have agreed. Perhaps one could argue that there are different taxonomies that employ notions such as *wanting*, *desiring*, *intending* and *wishing* in different ways. One debate in the philosophy of action resolves precisely around the question of whether intentions are reducible to belief-desire complexes, for example, so it seems at least a fact that different taxonomies exist. However, first of all it might be that since these are engaged in *debate*, only one of them might be *correct*, which means that the other taxonomies will not be related by admissible transformation. Furthermore, it might be that insofar different taxonomies *are* interchangeable, the variations will turn out to be purely verbal and will preserve general difference between the categories of belief and desire, for instance. I think we can safely assume that these gross categories of attitude types are invariant under admissible transformation. We can go even a step further and hold without much trouble that these categories correspond to psychological kinds.

What about sentences? It is not hard to find examples in which different sentences can correctly be attributed to the same thought. Under most, if not all circumstances, whenever a German speaker correctly attributes the sentence “Es regnet” to a thought, an English speaker would have been equally correct in attributing the sentence “It is raining” to that thought. Synonyms provide examples of admissible sentence transformations within the same language. However, this does not give us an easy universal admissible transformation for sentences, because there are also examples of situations in which interlinguistic translation and intralinguistic synonyms of these sort *fail*. This might pose a problem for the measurement analogy, but it does not pose a problem for NPA as such: from the fact that examples of admissible sentence transformation exists, it follows that sentences are finer grained, in some areas, than thoughts, and that hence we cannot adopt RPA on the level of *sentences*. But of course, that is not what a sensible proponent of RPA would go for anyway.⁴

⁴ Even the *language of thought* hypothesis is not committed to such a view, because it attempts to specify a mental language that abstracts away from exactly those features of individual natural languages that may vary under admissible transformation.

Finally, can representations differ in Russellian proposition across admissible transformation? Recall that Davidson's primary argument against Russellian RPA was that the identity conditions of external objects, and thus of any *propositional* objects construed out of external objects, are much stricter than those of thoughts, so that even when a thinker misidentifies an object that his thought could be related to by propositional attitude ascription, it does not follow that the thinker misidentifies his *thought*. From a Davidsonian perspective, if we really must talk of propositions, we should at least expect them to be principally underdetermined by mental states, and therefore be definitely variable under admissible transformation on a substantial scale.

Matthews appears much less convinced of this. Roughly, he takes intentional states to be functional roles of interaction with states of the environment. An intentional state can be measured by specifying what environmental state it is interacting with, and what the nature of the interaction is. The first specification is given by the designated proposition, the second by the attitude type. With respect to the Russellian part of the designated proposition, he argues:

Admissible transforms that differ in Russellian proposition, if they exist, are hardly commonplace. The reason should be clear: differences in Russellian proposition will in the usual case entail differences in propositional attitude because such differences will in the usual case entail differences in the causal roles which are defined functionally in terms of Russellian propositions. (Matthews 1994:142)

Regardless of whether you would agree with this argument or not, one has to admit that this is an anti-climax of considerable proportion. What started out as a revolutionary sounding project, because of its nonrelational and nonrealist approach towards propositional attitude ascriptions, now turns out to be a view that only allows substantial underdetermination of *sentences*, which proponents of RPA need have nothing against, and assigns uniqueness (i.e., transformational invariance) to attitude types and Russellian propositions, which is exactly what proponents of Russellian RPA were defending in the first place.

In the end, Matthews remains neutral on the question of Russellian proposition invariance. However, he does seem to hinge more towards the view that such propositions

are invariant under admissible transformation. Given his strong opposition to RPA at the outset, his neutral position is already a bit odd, but his subliminal preference for uniquely determined Russellian proposition ascriptions does not square with the advertised rejection of RPA at all. If the representation of thoughts by means of designated propositions is non-unique but the part that is constituted by Russellian propositions *is* unique, then in my view the measurement analogy fails to give any reason whatsoever for preferring NPA to RPA.

Worse even, if you were to accept Matthews-style NPA and then opt for the version that takes Russellian propositions to be uniquely measuring, then you fall prey to the same criticism that Davidson levelled against Russellian RPA, the criticism that motivated the move from RPA to NPA in the first place. Because if Russellian propositions were uniquely determined within the representation space by the preserved properties of thoughts, then ability to identify your own thought should imply ability to determine the uniquely identifying Russellian proposition. And conversely, inability to do the latter would thus imply inability to do the former. Which puts us right back where we started.

The upshot of all this is that once you start reasoning about propositional attitudes in terms of Russellian propositions, you are committed to the view that ascriptions of such propositions to thoughts are non-unique.

3 Arguments against the Nonrelational View

NPA deserves credit for its ability to prevent the problems that RPA has with Russellian propositions, and also for the fact that it accommodates the interesting measurement analogy, even though this analogy remains largely undeveloped. However, despite all this I still find NPA far less intuitive than RPA. I have the feeling that to deny the common sense idea that there is something 'present to my mind' when I am thinking, in order to solve some technical difficulties related to propositional attitude ascription, might be to throw out the baby with the bathwater. In this section I will try to flesh out that feeling with three arguments against NPA.

The Experience of Content

The content of mental states is not only *attributed* to me by others, it is also *experienced phenomenally* by myself. And this seems to me especially hard for NPA to account for. When I desire to have pizza, then one unmistakably vivid image presents itself to my mind, and is indeed *present to my mind*: that of pizza! Based on pleasant memories of eating pizza in the past, I anticipate what it will be like to eat pizza tonight, and since this anticipation features a high positive level of phenomenal affectivity, the mental state that is this anticipation will motivate me by activating other states that lead me to engage in pizza securing behaviour (usually, picking up the phone and calling the pizza delivery).

Suppose I express my desire by uttering the sentence “I want to have pizza tonight.” Of course, I might have uttered a different sentence to express this desire, I might have said “Ik wil vanavond pizza eten”, for example. We can then have a debate about what all the aspects are that attributions might differ in, i.e., which features of the measurement might vary across admissible transformation, and whether these run deeper than translation and synonymy on the level of natural language. But even if they do run deeper, then they could never run so deep as to eliminate all content on the side of the empirical system (my mind), because then there would be no content left for me to experience.

Maybe a proponent of NPA could admit that there is some sort of unique content that is experienced by me when I am in some phenomenal-and-intentional state, but maintain that this content is not *propositional*. And it is propositional content that NPA is anti-realist about:

Strictly speaking, propositional attitudes do not have propositional contents; such contents are simply a means of representing certain real properties (and relations) of attitudes. (Matthews 1994:144)

However, if a proponent of NPA allows that there is unique non-propositional content for conscious thoughts, a proponent of RPA can argue that hence it must be true that some sort of content object is present to the mind and that there is a relation which couples every thought to precisely one such object. The question of whether this object is propositional or not can then take a verbal turn: it depends on what concept of propositions you employ.

In my terminology, to say that the content of a thought is a proposition, without further qualifications, is only to say that it has a truth-value. And since thoughts are related to truth-values by means of their satisfaction conditions, it follows that if thoughts are to be identified in terms of their content, this content has to be propositional. This is perhaps most obvious for conscious belief states: if I am consciously thinking that I just ordered pizza, then the content of my experience can be true or false, i.e. either I correctly remember having ordered pizza or I am in error.

Belief and Truth Conditions

Beliefs have truth-values, and so do the objects that we use for interpretation of beliefs, whether these are utterances or designated propositions. Davidson argues as follows for his preference for NPA in terms of utterances:

Utterances are related to each other in much the same way beliefs are: by relations of entailment and evidential support. Utterances, like beliefs, are true and false. Aside from complications due to indexical elements, we identify a belief by uttering a sentence that has the same truth conditions as the belief it is used to identify. (Davidson 1989 [2001:64])

In measurement terminology, the utterances *preserve* the truth conditions of the beliefs. But how? Suppose I want to report my belief. How do I know what sentence to utter? The answer of RPA is straightforward: pick a sentence that expresses the proposition that is the content of the belief. Of course, this raises the important question of how to construe propositions in such a way that propositions can play their part in a (naturalistic) theory of mind and action. This has led to very different sorts of theories, but what they share is the idea that propositions are what beliefs and sentences/utterances have in common.

But what is it they have in common from the perspective of NPA? In terms of measurement theory, this is the question of the representation theorem. According to Matthews, designated propositions measure states of the environment to which intentional states can be related when construed as causal roles. However, as we have seen, this led Matthews to argue that the Russellian proposition inside the designated proposition is

probably not variable across admissible transformation because it is fixed by the externalist causal role approach to intentional states, which basically pulls the plug out of NPA and turns it into a causal role based RPA with regard to Russellian propositions. So that does not help NPA. It is not clear to me at all what a representation theorem should look like for a theory of propositional attitudes that assigns a truly non-unique role to propositions.

This does not mean that the measurement analogy is false. It just means that propositional contents are not just elements of the representation space, but also elements of the empirical system, and that these are preserved by means of the utterances that express them. It does not follow from the measurement analogy itself that we should reject RPA in favour of NPA.

Representation Requires Thought

As noted earlier on, from the NPA perspective, thoughts are not so much representers, they are rather the represented. We represent thoughts using a measurement space consisting of utterances, or of attitude types paired with designated propositions. Furthermore, according to Matthews it would be a mistake to read the propositional contents found in our representation space back into the empirical system, because strictly speaking thoughts have no propositional content. In other words: utterances represent thoughts, and thoughts represent nothing.

So according to NPA, the thinking subject is not a representational system, whereas the system of utterance, which I shall simply call the *language*, is. From the perspective of the measurement analogy this might seem reasonable at first sight. After all, just because every representation space is by definition representational doesn't mean the empirical system must be a representational system itself.

However, we must acknowledge a third element involved in measurement next to the empirical system and the measurement instrument, and that is *the agent performing the measurement*. In Davidsonian fashion, I will refer to this agent as the *interpreter*. Because the interpreter *uses* the measurement instrument, he must have certain capacities in order to do so. In the case of thought measurement using natural language, this means the interpreter must be a natural language user.

When we take the interpreter into consideration, we can recognise an important difference between the attribution of propositional attitudes and other varieties of measurement. In the usual cases, interpreters are human beings, and hence thinkers themselves. If use of language requires thought, then all interpreters are thinkers. Thus, the measurement of thoughts is a special kind of measurement where the empirical system and the interpreter are the same kind of systems: they are both thinkers. They might even be the *same* system, as is the case when a thinker reports a propositional attitude of his own.

According to Davidson, the fact that I *use* a certain object (proposition, sentence, utterance) to identify the thought of a thinker does not imply that such a thinker knows what object it is. However, *I the interpreter* am a thinker too! Does that mean that I also do not know what object it is? If there is no such thing as an object of thought that a thinker can identify, then how can there be such an object for the *interpreter* to identify the thought of the thinker with? Is interpretation not also a mode of thinking? If, as Davidson maintains, “there is no clear meaning to the idea of knowing what object one has in mind” (Davidson 1989 [2001:54]), then I don’t see how there could be a clear meaning to knowing what object one is using in order to identify a thought. But in that case, NPA would be in as much trouble as RPA, which means that NPA refutes itself.

We can strengthen this line of argument by considering what I shall call *silent thought attribution*, i.e. attribution without utterance. Silent attribution is a common phenomenon. A waiter in a restaurant, who sees a guest stand up from his table and walk towards the back of the restaurant may deduce from the guests behaviour that he desires to go to the toilet and believes they are (or might be) located in the back of the restaurant. Before either guest or waiter has uttered a single word, the waiter has already attributed propositional attitudes to the guest.

From the perspective of NPA, one would have to say that the propositional contents of these attitudes were *used* by the waiter to measure the thoughts of the guest, but at the same time one would have to deny that these propositional contents were in any way part of the waiters own thoughts about the guest. But this seems contradictory.

In contrast, RPA could simply give the following description of this case: the guest has a desire with the content ‘that I use the toilet’ and a belief with the content ‘that there are toilets in the back of this restaurant’. The waiter has beliefs with the contents ‘that that guest

desires to use the toilet' and 'that that guest believes there are toilets in the back of this restaurant'. Thus, the waiter employed the content 'that there are toilets in the back of this restaurant' to measure the thought of the guest by making it part of the content of a thought of his own.

The primary problem for RPA will be to either give an account of content that has room for the indexicals in this example, or to find a way of getting rid of the indexicals. Both ways have proven very difficult so far, but at least RPA has something to say about what the waiter is doing: he is thinking about propositions and attributing them to attitudes of the guest. But what can NPA say about the waiter? Somehow, he is representing the thoughts of the guest. Since there are no utterances or sentences involved, it seems there is nothing left but the thoughts of the waiter to do the representing. As I see it, one would also make a commitment to RPA by adopting the view that sentences or utterances are involved after all in the sense that the waiter had a sentence or utterance 'in mind'. In that case the thought of the waiter would have to be individuated at least as fine grained as the sentence or utterance he had in mind.

What all this comes down to, in my view, is simply that representation requires thought. If the medium for representation is language, then the elements of that language used to do the representation still have to be identified themselves before they can be put to use, and doing this is a cognitive capacity. Identifying sentences or utterances in order to use them to identify someone's thoughts requires *understanding* of the sentences and utterances, and understanding requires thinking. Once more, RPA provides an easy way of putting this: understanding an utterance involves grasping the proposition expressed by the utterance with a thought. This also means, in my view, that explaining mental representation is a more basic problem than explaining linguistic representation, and that a solution to the latter will in some sense be a derivative of a solution to the former. If this is indeed the case, then NPA is deeply mistaken, as it shifts the entire topic of representation out of the realm of thought and into that of language.

4 A Fregean Relational Proposal

Since I have formulated arguments against NPA, confirmed arguments against Russellian RPA, and attempted to debunk arguments against Fregean RPA, it should not surprise the

reader that I am going to propose a version of Fregean RPA. In the foregoing, we have already stumbled upon what can generally be seen as the two major problems that such an account has to deal with. First, we need something to play the role of the Fregean senses: content elements that mediate reference, are in some sense ‘present to the mind’, and can be shown to be preserved by utterances, which implies that they are invariant under admissible transformation of utterances. Second, we need to square the notion of mediated reference with the phenomenon of indexicality: either by eliminating indexicals out of attributions or by incorporating indexicals in our construal of content.

In this section, I will only address the first problem. The problem of indexicals has proven hard for Fregean accounts, and is one of the major reasons for the fact that the idea of direct reference has been the most popular game in town for quite some while now. Nevertheless, Russellian accounts are in trouble in various areas, one of which we have seen in this essay, and not everybody believes that it is impossible for a Fregean view to handle indexicals.

What is Present to the Mind?

Davidson maintained that there are no objects that “seem what they are and are what they seem”, and furthermore, that there is no question-begging way of specifying the *extent* to which objects are what they seem by means of a range of ‘internal’ properties. However, from the position that I developed in *PDM* it follows that there *are* such objects: phenomenal experiences, and that there are such properties: phenomenal properties. Or to be more precise, in virtue of the notion of microphenomenal properties, which are, after all, *not* present to the mind, we should further restrict these properties to *the phenomenal properties of which we have phenomenal concepts*. With respect to these properties, phenomenal experiences are fully transparent to a conscious agent. For these properties literally are what they seem.

Of course one might still maintain that there is no non-questionbegging account of phenomenal properties. In some sense this is true, as it is impossible to explain the difference between phenomenal and practical properties without the use of the concept of phenomenal properties. However, qualia realists have always maintained that this does not count against the idea of qualia because it is *part* of the idea of qualia. To reject the idea of phenomenal

properties is to embrace phenomenal eliminativism, a view that I argued against in *PDM*, section 3.

Since the activation of a phenomenal concept is nothing over and above the conscious experience of the phenomenal property that it is a concept of, it makes sense to regard phenomenal concepts as objects that can be present to the mind. Because they are *concepts*, they are also suited for playing a semantical role. Could phenomenal concepts be the constituents of the propositional contents of thoughts?

Of course, we do not want to end up with a view that allows content to be only about the nature of our experiences and not about the world around us. We also need practical concepts. As I argued in *PDM*, section 1, the application of practical concepts to our environment is mediated by the application of phenomenal concepts to our experiences. This is because phenomenal properties take care of the syntax of our representations: relations between different events in my environment are coded in terms of the relations between my experiences of those events. Can we simply add practical concepts to our account by claiming that they are also 'present to the mind', as it were, *through* the phenomenal concepts to which they are related?

As is well known, syntax can only provide what is often called *narrow content*. Since syntax cannot discriminate events beyond differences in interaction between the events and the agent, narrow content is often thought to fail to truly *refer* to individuals and kinds, since the identity of individuals and kinds might be partly determined by properties that are beyond the discriminative capacities of the agent. Note that this is also a premise of Davidson's primary argument against RPA.

Putnam's *Twin Earth* example helps to bring this idea out. If there is such a thing as narrow content, then the thoughts of me and that of my Twin Earth twin would be identical with respect to that kind of content. As proponents of the idea of *wide content* argue, our thoughts must differ in content, since I have thoughts about water whereas my twin has thoughts about something else, since there is no water on Twin Earth.

Summarising, it seems that if I try to promote practical concepts to the status of objects present to the mind in virtue of their being coded syntactically by phenomenal property instantiations present to the mind, I commit myself to narrow content, but the Twin Earth

example seems to show that narrow content is not good enough. Can we get away with narrow content?

Twin Earth Revisited

Twin Earth is defined as a world that is similar to ours at the macro level, but different at the micro level. In particular, what seems to be 'watery stuff' at the macro level has the molecular structure XYZ rather than H₂O. The idea relies heavily on the notion of multiple realizability, and for the thought experiments to work we should situate ourselves in a time before modern chemistry, in fact in a time before any of the differences between Earth and Twin Earth could have made an impact on the history of man.

Let us say that a man called Ernie lived in such a time on Earth, and that Twernie is his hypothetical twin on Twin Earth. Ernie and Twernie are both acquainted with watery stuff. Watery stuff is the concept of a drinkable liquid that includes all the further macro features of water that Ernie has come to know about. In fact, it is a practical concept, because it specifies watery stuff in terms of its practically observable effects. Since Ernie and Twernie are both acquainted with watery stuff, they both have this practical concept.

According to a standard interpretation of this sort of example, to say that something is watery stuff is to describe some of its features, but not to refer to it as a kind. However, whenever Ernie refers to the watery stuff on his planet by 'pointing' at it, he *is* referring to water as a kind, and whenever he uses the *kind predicate* 'water', he is doing the same thing. Furthermore, it is held that what defines the kind, in this case, is its molecular structure, H₂O. Hence, it is argued that when Ernie talks about water, he talks about something else than Twernie. And thus, content should be construed widely.

However, from the perspective of a pragmatist philosophy of science, the concept of H₂O is not fundamentally different, but only different in degree, from the concept of watery stuff: both are practical concepts. The concept of H₂O is a concept that is part of a theoretical framework of thinking about substances at a certain level of organization, which is ultimately defined in terms of its practical consequences, both in ordinary life and in the laboratory, and in terms of its relations to other theories that are in similar ways linked to observable effects and to each other. To argue, instead, that watery stuff only specifies what water *does* while H₂O specifies what it *is*, is to make a very fundamental mistake about

science. The theory of molecular structure explains that a water molecule is built out of an oxygen atom and two hydrogen atoms that stick together in virtue of what oxygen and hydrogen atoms *do*, which we know on the basis of observable effects we got from experiments.

In order to fully appreciate this point, consider that just as watery stuff allows for different implementations, *so does H₂O*. The discovery that water is H₂O is compatible with different underlying theories. Quantum field theory may be true or false, string theory may be true or false, so there are still different logically possible scenarios of what the micro picture of H₂O, at the sub atomic level, looks like. All of which correspond to possible worlds, if you like, in which different subatomic configurations play the H₂O role.

It could be objected that this does not matter because kinds are organised hierarchically, so that in these possible worlds the stuff that plays the water role belongs to different kinds at the subatomic level, but that they all belong to the H₂O kind, analogous to the fact that I belong both to the mammal and to the homo sapiens kinds. However, in that case there is no reason why watery stuff is not a kind predicate either, which in turn makes it unclear why Ernie would be pointing to the H₂O kind rather than the watery stuff kind despite the fact that he does have a concept of watery stuff and not of H₂O.

A more powerful objection is that the term 'water' is used to distinguish a kind of substance from other kinds at a level of individuation that is met by the level of organisation of molecular structure. In other words, it turned out that all the stuff that plays the watery stuff role on Earth consist of H₂O molecules, whereas all the stuff that plays another role does not.

But this is misleading. H₂O molecules can also form vapour, and I am not sure vapour was considered part of the watery stuff role before the advent of chemistry. In fact, the word 'water' has simply not been used to denote the same role through history, and it is under the influence of molecular theory that it has now become fixed to the role played by H₂O. And strictly speaking, that is not even true, because we also use the word 'water' to refer to the stuff that we drink and that fills the rivers and oceans and comes out of our showers, and it is essential to all those roles that the liquid in question contains other molecules besides H₂O.

The consequence of all this is that when Ernie uses the term 'water' to refer to the stuff that realises the water role, there is no reason why he should be referring to the chemical

kind H₂O. A deeper consequence is, in my view, that when Ernie uses the term 'water' he is not actually doing such a thing as referring to the stuff that realises the water role. Instead, when Ernie uses the term 'water', *he is referring to the water role itself*. As we have seen earlier on, practical concepts are functional role concepts, and practical properties are functional roles. As I understand it, pragmatism teaches us to be very careful about the idea of a *kind of stuff*. With the exception of phenomenality, all meaningful expressions about 'stuff' should have semantics in terms of practical roles.

However, it is impossible to refer to a functional role without describing it, since a functional role is nothing more than what it does. Knowing which role a functional role is is the same as knowing what it means to realise that role. Or to put the point in terms of practical properties: a practical property is the property of having a certain effect, and the only way to refer, or 'point' to such a property is by identifying the effect, which amounts to giving a specification of the property.⁵ This is incompatible with Kripke/Putnam semantics, which is based on a fundamental distinction between direct reference and description. If we are serious about pragmatist semantics, we must reject the theory of direct reference.

This raises many problems, of course; problems that the theory of direct reference was meant to solve or avoid. It is impossible to even give an adequate overview of those problems here, much less to show how they can be dealt with. Furthermore, my ideas on the way all this should go are still very much under development. Very briefly, my basic intuition is that we use kind predicates to report the content of our practical concepts and to expand them. For example, if I utter the sentence "water boils at 100 degrees", it might be that I am doing this because the person I am talking to does not already know this, but does already have some concept of water. Suppose that the concept of water that this person has is the concept of that which has the observable effects $e_1...e_n$. In that case I have told this person that that which has observable effects $e_1...e_n$ also has the observable effect of boiling at a temperature of 100 degrees. At the same time, I have reported that this observable effect is part of my concept of water. Under the appropriate conditions this will cause the listener to *update* his concept. Thus, kind predicates *synchronise* practical concepts between language users. This means that the semantics of a kind predicate is not just based on the

⁵ This insight might also be useful for a Fregean theory of indexicals.

corresponding concept of the speaker, but also on the corresponding concept of the *listener*, or what the speaker believes his listener's concept to contain. This would explain why a sentence such as "water boils at 100 degrees" or "water is H₂O" does not have to express a tautology – without invoking to the idea of direct reference.

The theory is very sketchy at present, but I see a lot of room for further development. However, that will have to wait another occasion. For now, we should conclude that from a pragmatist perspective, it makes perfect sense to say that Ernie's thoughts about water on Earth and Twernie's thoughts about water (!) on Twin Earth have the same content. This content is 'narrow' in the sense that it is fully determined by the syntax of their mental states, but that does not mean that it does not refer to real, external properties on Earth and Twin Earth respectively. It simply means that on a pragmatist realism of properties, the distinction between narrow and wide content can no longer be made.

Conceptual Propositions

Let us define *conceptual propositions* as propositions that are composed solely out of practical concepts and phenomenal concepts. Conceptual propositions are thus Fregean: they are not composed out of referents but out of mediating entities that determine the referents. They are, furthermore, *narrow* in the sense described above: they represent *practical properties* solely on the basis of the way in which they relate the phenomenal experiences of the observable effects of those properties, and they represent *phenomenal properties* only insofar as these are known "from one's own case," (Loar 1990 [1997: 598]).

I propose that the content of a thought is a conceptual proposition, and that these conceptual propositions can be expressed using language. In *PDM*, section 5, I argued that when practical concepts are expressed using language, they can be individuated on two modes: the mode of expression and the mode of practical bearing. A similar distinction will hold with respect to the expression of phenomenal concepts, except that there will be additional restrictions because the communication of phenomenal concepts requires that language users share the experiences that those concepts refer to.

Since concepts can be individuated on these two modes, conceptual propositions can be individuated on both modes as well. When individuated on the mode of expression, conceptual propositions will be finer grained than sentences, and will in fact come close to

Matthews' designated propositions. When attributed to thought, conceptual propositions can vary across admissible transformation on this mode of individuation. By contrast, when individuated on the mode of practical bearing, they are invariant.

Conclusion

I have tried to show that the combination of pragmatist semantics for the third person perspective and the qualia realism for the first person perspective, a combination that has played a key role in *PDM*, can help us formulate a Fregean version of RPA. In this proposal, practical and phenomenal concepts play a role similar to that of Fregean senses: they are graspable by thought and at the same time fix their referents. I have argued that practical concepts can fix their referents because all referents involved are functional roles that do not require discrimination beyond their practical bearing. This led me to reject the notion of direct reference as well as the distinction between narrow and wide content.

Propositional attitudes, then, are mental states composed out of attitude types and conceptual propositions. Propositional attitude *attributions* are utterances that preserve these conceptual propositions when individuated on a mode of practical bearing, but allow for transformation on the mode of expression.

At this point, the proposal is very sketchy, and should perhaps not yet be considered a serious alternative to the mainstream semantics based on direct reference of kind predicates. However, it can be considered a serious proposal for further research. I believe that a renewed interest in pragmatism holds promise for the philosophy of science – as I have tried to show in my discussion of underdetermination in *PDM*, section 5 – and that it might just be here to stay. If it is, and if my argument that it doesn't square with direct reference is correct, then a rebellion against the established consensus in the philosophy of language is justified.

5 References

Davidson, D. 1989, "What is Present to the Mind?", reprinted in id. 2001, *Subjective, Intersubjective, Objective*, Oxford, Clarendon Press, 53-67.

- Loar, B. 1990, "Phenomenal States", reprinted in N. Block, O. Flanagan, and G. Güzeldere (eds.) 1997, *The Nature of Consciousness: Philosophical Debates*, Cambridge MA, MIT Press, 597-616.
- Matthews, R.J. 1994, "The Measure of Mind", *Mind*, 103, 410, 131-145.
- Voerman, S.A. 2006, "Phenomenality, What It Does, and What That Means", available online at *Backwash.org*, <http://www.backwash.org/pdf/phenomenality.pdf>.