

NOUMENALISM AND RESPONSE-DEPENDENCE

The question with which I shall be concerned in this paper is whether global response-dependence entails the truth of a certain noumenal form of realism: for short, a certain noumenalism. I accept that it does, at least under a plausible assumption, endorsing an argument presented by Michael Smith and Daniel Stoljar (forthcoming). But I try to show that, while the connection with noumenalism is undeniable, it is neither distinctive of a belief in global response-dependence nor particularly disturbing for those of us who embrace that belief.

The paper is in five sections. First I provide some background on the meaning of response-dependence and on my own reasons for commitment to global response-dependence. Next I examine the sort of doctrine that can reasonably be described as noumenalism. In the third section I look at the way in which global response-dependence entails noumenalism, according to Smith and Stoljar's argument. And then, in the last two sections, I argue that we can live fairly happily with that argument. The fourth section shows that the argument carries only under an assumption that is not universally sound, so that noumenalism is not a universal complaint. The fifth maintains that almost everyone has to accept a certain noumenalism and that as the complaint is not universal, so it is not particularly serious either.

1. Response-dependence

The word 'response-dependent' was introduced by Mark Johnston (1989, 145) to pick out those terms and concepts that are biconditionally connected, as an *a priori* matter, with how things appear to us human beings: with how we judge or are disposed to judge them. The word or concept 'red' will be response-dependent, under this account, so far as it is *a priori* that something is red if and only if it is such as to look red to normal observers in normal conditions: if and only if it is such as to evoke that particular judgmental response.

"Noumenalism and Response-Dependence" by Philip Pettit,
The Monist, vol. 81, no. 1, pp. 112–132. Copyright © 1998, THE MONIST, La Salle, Illinois 61301.

In defining 'response-dependence' in this way Johnston was making contact with some unpublished work of Crispin Wright's (Johnston 1993, 121–26; Wright 1993, 77–82); this bore on the importance for certain concepts of *a priori* equations such as that just illustrated for redness. I took the word from Johnston, though I construed response-dependence somewhat differently, as we shall see in a moment (see Pettit 1991, 598). Adopting his word, if not his precise conception, I argued that all the semantically basic terms in anyone's vocabulary—all the terms that are not introduced definitionally by their connections with other terms—are response-dependent.

My argument was built on the account of rule-following that I had offered earlier (Pettit 1990). The idea was that everyone's vocabulary must include some terms that are introduced to them in a non-definitional, more or less ostensive manner and that mastery of such basic terms, and possession of the corresponding concepts, is dependent on that person's being responsive in a certain way to the referents of those terms: say, to the properties picked out by them. In particular, it depends on the person's being such that, after ostension, instances of such a property typically evoke a disposition to believe that they are instances, at least under favourable conditions, and to use the term in question to express that belief. It depends on the person's being such that under favourable specifications instances of the property come to seem to them like instances of the property.

How am I to get to refer to a property, T, by the use of a basic term, 'T'? How am I to get to master the term and possess the concept? I can have the property presented to me in certain instances or examples but what is going to enable me to latch onto it—what is going to make it a salient object of ostension—and not onto any of the other properties instantiated in those examples? My claim was that on being exposed to instances of T and on learning that they, paradigmatically, are T, I must come to form the disposition with any other instances—or at least with any other instances that are presented in what independently count as normal or ideal circumstances (Pettit 1998)—to believe that they too are T. I must be so affected by exposure to examples of T that instances of the property generally come to seem T to me, at least under suitable constraints.

This claim means that semantically basic terms are all response-dependent in something close to Johnston's sense; it means that response-

dependence is global or ubiquitous. If something is denominably T—if it possesses a property, T, for which I and fellow speakers use the word or concept, ‘T’—then it must be such that it would seem T to favoured subjects in favoured conditions. And if something is such that it would present itself in that way then it must count as being T; seeming T under those specifications, it shows itself to be an instance of the T-property. It transpires, then, and on an *a priori* basis, that given the denominability of the property in question, something is T if and only if it is such as to seem T to favoured subjects in favoured conditions.

How big is the difference between saying that something is denominably T if and only if it is such as to seem T and saying, as under Johnston’s definition, that something is T, period, if and only if it is such as to seem T? In earlier writings I overlooked this difference, though I emphasised some related points (see Pettit 1991, pp. 609–11). But the difference is of the first importance. Something will not be denominably T in a world for which it is impossible to identify uniquely favoured observers to whom it can seem T: that is why something’s being denominably T entails that it is such as to seem T under suitable specifications. But for all that this says, something may still be T—it may have the property that we ascribe by the use of the word, ‘T’—in such a world.

On Johnston’s understanding, the response-dependent term or concept is one that represents its referent as connected with human responses: this, in the way in which saying that a substance is nauseating, or that a chair is comfortable, represents it as having a property that involves human beings. That is why he can think that with any response-dependent term, ‘T’, something is T if and only if it is such as to seem T under suitable specifications.

On my understanding, however, terms or concepts can be response-dependent without having such an anthropocentric content. They will be response-dependent so far as they are response-dependently mastered and possessed. In order for people to come to use such a term, with whatever content it has, they will have to be subject to certain responses—instances of a property picked out by the term, for example, will have to seem, under suitable conditions, like instances of that property—but the content itself need not bear on those responses. As I put it in earlier work, the response-dependence of the concept will be explained by its possession-conditions, not by its conditions of application (Pettit 1991; see, too, Jackson and Pettit 1998).

I think that global response-dependence is a compelling doctrine and I have tried elsewhere to defend it against a number of charges (Pettit 1991; Pettit 1996, Ch. 2, Postscript; Jackson and Pettit 1998). But Smith and Stoljar raise a new and interesting challenge: that it entails the acceptance of a noumenal form of realism. In the next section I characterise the noumenalist position, as I understand it, and in the section after that I present the argument for why global response-dependence entails such noumenalism. Then in the two remaining sections I argue that we can live with this noumenalism; it is not as rife or as serious a complaint as we might have thought.

2. *Noumenalism*

The defining assumption in noumenal realism, as Smith and Stoljar express it, is “the idea that the world is a certain way in and of itself, even though we are in no position to make claims about the way that it is” (4). Or as they put it elsewhere, “there is an independent reality, but the intrinsic nature of that reality is unknowable” (2). Noumenalism admits that there may be certain aspects of reality, certain properties of the world, of which we have full knowledge but it maintains that there are other aspects or properties of which we are necessarily ignorant; “there are aspects of the world that we cannot possibly describe or explain” (17)

There are a number of different ways in which noumenalism in this sense might be further articulated but I propose to understand it as follows. Noumenal realism maintains that no matter how good our theory of the world is—specifically, no matter how complete it is in identifying the properties that play important roles in the working of the world—still that theory will leave us in partial ignorance as to the nature of those role-playing properties. The theory will postulate that there is one and only one property filling each of the roles in question but for all that it tells us, the property in question may be any one of a number of different candidates.

Among those properties of the world that we would countenance or be committed to countenancing under the best theory available, then, noumenalism says that there are bound to be some that we do not identify uniquely. We know that the properties are instantiated and that they play such and such a role, or have such and such effects. But we do not know which properties exactly are in question. For all that our theory tells us, the properties in question may be of this, that or another character. They represent *terra incognita*; they belong to an unknowable, noumenal world.

In order to identify a property, it is going to be necessary for a person to be able to pick out that property from others and to have a means of referring to it. The person may refer to the property indirectly, on the basis of reference to other things—other properties included—or they may refer to it directly, having a term that locks onto it without any mediation: in this case the referring expression will have to be a semantically basic term. The term ‘red’ designates the corresponding colour-property directly, so we may suppose, while the expression ‘the first colour in Newton’s spectrum’ designates it indirectly: designates it in a way that depends on the direct designation of other properties and other entities.

Suppose that we can refer directly, then, to a certain property. Does that mean that we know which property is in question? No it does not. Consistently with referring directly to a property—and consistently with knowing that we refer to a single property—we may or we may not know which property it is that we refer to.

It is possible to illustrate both how we may have this knowledge, and how we may lack it, with the example of redness. Suppose, first of all, that to say something is red is to say that it has the higher-order property of instantiating a lower-order property which makes it look red to normal observers in normal circumstances. So understood, we know we will know which property is the property of redness. For no matter what sort of world is actual—in particular, no matter which property makes things look red in the actual world—the property to which we refer, assuming we do successfully refer, is still one and the same higher-order property: the property of instantiating a property that makes things look red; if you like, the redness “role-property” as distinct from the corresponding “realiser-property” (Jackson and Pettit 1988; Blackburn 1991).¹

Suppose, however, that the term ‘red’ is to be understood differently so that to say something is red is to say that it has that property—that realiser-property—that serves to make it look red to normal observers in normal conditions: the referent of the predicate is the realiser-property—perhaps itself disjunctive in character—not the role-property. And suppose that we do not know which property serves to make things look red: of the worlds where different properties realise the redness role we do not know which is actual. For all that we can tell in that case, the property in question may be this or that or the other realiser; depending on which world is the actual world, the predicate will pick out this, that or the other

property as the referent in that world. It may pick it out rigidly, as the actual property that plays the relevant role, or it may pick it out flexibly, as whatever property plays the relevant role; but since it is not pertinent to our interests, we will mostly ignore this bifurcation of possibilities in what follows. Whether 'red' refers rigidly or not to the realiser-property, it is still going to be possible that we do not know which property is picked out by 'red'.

With a property to which we refer directly, then, we may or may not know which property it is. But something similar is going to be true, for similar reasons, with properties to which we refer only indirectly, say via their connections with other properties. For suppose that we refer to a property P via its connections—involving the relational property R—to properties Q, S and T. If we construe the property as the realiser-property of this connective role—the lower-order property instantiated—then of course we may not know which property it is that plays that role. But there is a possibility of knowing which property it is if we construe it as the role-property: the higher-order property of instantiating a property that is R-related to Q, S and T. We will not know which property it is if we do not know which properties the defining properties are. But if we do know this, then we will also know which property the role-property is; things will be exactly as they are with the role-property of redness.²

Where we refer to a property but do not know which property it is then we will know it via its impact on us or via its connections with other entities: we will know it, as we may say, in its actual and hypothetical effects. In the case originally imagined, we know the realiser-property of redness via its impact on us—as the property that makes things look red—and perhaps in its connections with other things: say, in its making things look brighter than brownness. But we do not know that realiser-property, as we might put it, in its essence. Although we refer to the property directly—although the word 'red', taken as a name for the realiser-property, is a semantic primitive—we do not know which property we pick out. We can imagine things looking exactly as they do, while now it is one property, now another, that is realised in the world.

When we do successfully refer to a certain property and we do not know it in its essence, then we are threatened with what we might describe as epistemic disjunctivitis.³ We use a single expression like 'is red' but for all that we may be able to tell, the property that is out there in the world

answering to the expression might be this or that or yet another: it might be any of an open-ended disjunction of possibilities.

It is important to emphasise that this disjunctivitis is epistemic in character, not semantic. Semantic disjunctivitis strikes when there are a number of properties in the actual world that have equal claims to be the referent of the expression; it constitutes a sort of indeterminacy. But with our example there is no question of indeterminacy: redness is the property that makes things look red and in our imagined scenario there is only one property that does that; there is only one actual realiser or occupant of the redness role. The disjunctivitis that threatens us presupposes the falsity of semantic disjunctivitis; it presupposes that our terms refer determinately to suitable properties or whatever.⁴

The threatening disjunctivitis is epistemic or evidential in character, not semantic. It comes of the fact that while there is only one occupant of the redness role, and while the expression therefore has a determinate referent, there are many equally good candidates for role-occupant and we do not know which is actually successful. We know the actual occupant—the referent property—in the effects of making things look red but we do not know it in its essence; we do not know which property it is that has those effects.

Noumenalism, as I understand it here, is the claim that even if we have the best theory possible of the world there are properties which it will commit us to countenancing such that we cannot—of necessity, cannot—know them in their essence; we cannot know which properties they are. What it threatens us with, then, is a chronic form of epistemic disjunctivitis. According to the noumenalist diagnosis, the human mind is such that while our theories may reliably postulate a variety of properties—and while we may even be able to refer quite determinately to any of those properties—there are bound to be some that we do not know in their essence. Even if we can refer successfully, and indeed directly, to such a property, we will not know exactly which of a variety of candidates is in question. There is an inevitable fuzziness to our theoretical sights, so that at a certain level of resolution the world—the world in itself—remains a blur.

I hope that this may suffice by way of articulating the doctrine that I understand as noumenalism. I turn now to look at how global response-dependence leads us into noumenalism, according to Smith and Stoljar.

And then in the final two sections I show that their argument does not mean that noumenalism is a universal or a serious complaint.

3. From global response-dependence to noumenalism

Someone who believes in global response-dependence holds that all of our semantically basic terms and concepts connect up with corresponding items in the world—say, corresponding properties—so far and only so far as those items elicit certain responses in us, at least under favoured specifications. If we succeed in locking onto a certain property in the use of the word ‘red’—and whether that property is taken to be the role- or realiser-property—that is by virtue of the property’s being associated with things looking red to us, at least when we satisfy a favoured psychology and occupy favoured circumstances. The fact that the property is associated with things looking red is what makes it—it and not another property—an effective attractor of the term ‘red’.

Global response-dependence implies that the sort of story just illustrated for the word ‘red’ holds, in rough outline, for all of the semantically basic terms in anyone’s vocabulary: that is, for all of the terms that the person learns to use without reliance on a definition. I say, in rough outline, because it is possible for the story to vary in many different ways from the account that looks plausible with colour terms.

Just to illustrate the variations possible, the effect whereby the presence of a suitable property is registered may not be as specific to sentient creatures as colour sensation; think of the effect of a smooth object in rolling comfortably against the skin. Or the effect may not be restricted to a single modality of sense; think of the different senses that register shape as distinct from colour. Or the effect may even involve a practical response on the part of the observer; think of the effect of an object in bending under intentionally applied pressure. Or, finally, the effect may be holistically mediated, in the sense that one property can make itself felt via a certain effect only so far as other properties make themselves felt via other effects; it is possible, indeed, that this is true for colour properties and colour sensations.

Whatever variations are allowed, the global response-dependence theorist has to maintain that the referents of someone’s basic terms get to be effective attractors of those terms so far and only so far as, under favoured specifications, they occasion certain responses in the subject and

the subject's ilk. The one requirement that has to be met by those responses is that they are relatively primitive. In particular, they are capable of doing the job required of them without being conceptualised; they can occur, and they can mediate the agent's use of terms, without themselves being articulated in language by the agent in question.

This requirement of non-conceptualisation stems from the fact that the global response-dependence theory has to explain a person's semantic competence in all of the basic terms that they use. If a response, R enables me to learn the use of a term 'T' only so far as I have a word for R—if the non-conceptualisation requirement is not fulfilled in this case—then that leaves open the question as to how I learn to use the term for R in the first place. But the requirement of non-conceptualisation is not particularly troublesome. It often seems with candidate responses that the agent need not be conscious of their realisation, let alone have a word for the sort of response in question. Even if an agent is unconscious of having the sensation of red, for example, that sensation can lead them to group something with other red things and to apply the word 'red' to it.

Global response-dependence, according to Smith and Stoljar, entails noumenalism. Their argument is articulated with considerable attention to detail but it may be useful to give a brief account of the sort of reasoning in question. I will do so in a way that allows for the understanding of noumenalism presented in the last section and that takes account of the distinction between role-properties and realiser-properties. But I do not think that anything I say is in conflict with their way of thinking. Is the reasoning, as I see it, sound? Yes, subject to an assumption to which I turn in the next section.

There are two ways of taking the properties with which our responses enable us to make basic semantic contact: either as role-properties or as realiser-properties, to use the terminology of the last section. And under either construal, so it appears, noumenalism is bound to obtain. Even given the best theory possible, there will be properties that we ascribe or are committed to ascribing to things in the actual world such that we do not know them in their essence, only in their effects.

Suppose we take the properties with which we make basic contact to be role-properties. And suppose that in order not to beg any questions we redefine role-properties more austere so that they do not definitionally presuppose realisers: to instantiate a role-property will be to have the

property of being such as to generate certain effects, where it is not a matter of definition that the suchness is a distinct realiser-property. In that case we will treat the properties with which we make basic contact as dispositions, since the property of being such as to produce a certain effect just is what most of us mean by a disposition. In particular we will treat the properties as dispositions to elicit appropriate responses in us and our likes. They will be anthropocentric dispositions, distributed variously over the different parts of the actual world, to interact with us in certain ways: dispositions in this object to look red, in that object to bend under pressure, in that other to fit smoothly into the palm of the hand, and so on.

But if the properties in question are dispositions of these kinds, then the argument run by Smith and Stoljar will go through quite smoothly. Take the more straightforward variant of that argument, which makes use of their no-bare-roles principle, as we may call it: the principle that every disposition has a non-dispositional explanatory ground. This intuitive principle—Smith and Stoljar provide an impressive defence—will entail that apart from anthropocentric dispositions there must be non-dispositional properties that serve to underpin the dispositions. And while the dispositions may be properties that we can know in their essence—this, in the manner of the role-property for redness—the underlying non-dispositional properties will not be of this kind. They will be properties, at least in the ultimate analysis, that are known only in the effect of grounding the dispositional properties or of grounding properties defined in terms of such dispositional properties.

This argument can be recast, with a little elaboration, in the language of role and realiser. We are supposing that the properties to which all our basic terms refer are role-properties. Even if some of those role-properties serve as the realisers of others, there must be a residue of realiser-properties that are not of a kind with these semantically primitive role-properties. There must be a realm of realiser-properties that escape the reach of our semantically basic terms. And this will be so, even if we redefine role-properties so that they do not definitionally presuppose realisers: the no-bare-roles principle defended by Smith and Stoljar will guarantee the result.

On recognising the reality of such realiser-properties, we can refer to them as the realisers, precisely, of the role-properties from which we started: as the grounds of those dispositions. But knowing the properties in this way, we do not know them in their essence; we do not know which

properties they are. Can we come to know them in any other way, say on the basis of theories which definitionally introduce novel terms for picking them out? Not without just postponing the problem. For those novel terms will allow us to know which properties they pick out only if they pick out theoretical role-properties: higher-order properties of being such as to satisfy defining, theoretical connections. And by the no-bare-roles principle, that means that they will then leave us with the problem that we can only know the theoretical realiser-properties as the realisers of those role-properties; we will only know them by this effect, not in their essence.

So much for the argument from global response-dependence to noumenalism, assuming that the features with which we make primitive semantic contact are role-properties: anthropocentric dispositions. Suppose, however, that we take the features with which we make basic semantic contact to be realiser-properties, not role-properties. Suppose we take the predicate 'red' to direct us to the realiser-property that makes things look red, not to the role-property—not to the disposition to look red—and similarly for semantically basic terms in general. How in that case does global response-dependence commit us to noumenalism?

Here the connection is even more straightforward. The properties with which we make basic contact under this picture are properties that, by definition, we only know in their effects. Under the other picture, the corresponding properties are anthropocentric dispositions—dispositions to elicit various effects in us—and it is no surprise that we can know them in their essence, even if we cannot know in that manner the ultimate non-dispositional bases that we have to countenance. But under this picture, the properties with which we make basic contact are not anthropocentric in the same way: they are those objective properties of things that we happen to identify by the fact that they elicit certain responses in us. And so under this picture it is no surprise that we are represented as knowing those very properties, not in their essence, but only in their effects.⁵

The conclusion to which we are pointed, following Smith and Stoljar, is that global response-dependence entails noumenalism. If the reference of basic predicates is fixed by responses evoked in us, then no matter how those properties are construed, we will have to countenance features of the actual world that we do not know in their essence; noumenalism will rule. I turn now to considering how far we can live happily with this result

4. *Noumenalism is not a universal complaint*

The argument in the last section assumes that if it is a realiser-property, then the property picked out by a response-dependent term like 'red' is that property instantiated in whatever world we are discussing that plays a certain idealised role: it would make things look red in those idealised circumstances where observers and conditions count on independent grounds as normal. More generally the argument supposes that the only realiser-properties that response-dependent terms can pick out are the instantiated realisers of idealised roles, where the instantiated realiser is the realiser of the role in that world, actual or counterfactual, that is under discussion or, in the case of a rigidified term, the realiser of the role in the actual world.

The argument makes this supposition, in taking it that where a term is used to characterise the actual world then the only realiser-property it can pick out is the actual realiser, so that which realiser-property it picks out will depend on which world is actual. Let the world that is actual cast property, P, in the idealised role and it is P that will be denoted by the term; let it cast another property, Q, in that role and it will be Q that the term picks out.

But this supposition is not compulsory. Suppose that the realiser-property that is picked out by a certain term is not the instantiated realiser of the idealised role but rather the idealised realiser of that idealised role. Suppose that what 'T' picks out is not that property, assumed to be suitably instantiated, that would seem T under suitably favourable conditions. Suppose that it picks out whatever property would seem T under those conditions, where it is allowed that the property in question need not be suitably instantiated. Under this assumption, there may be nothing that is T in the world of which we say, falsely, that some inhabitant is T or, correctly, that some inhabitant is not T. And yet the term 'T' may pick out a perfectly well defined property: that which would make things seem T under favourable specifications.

The argument from response-dependence to noumenalism will not go through for any terms that refer to the idealised realisers of idealised roles. For if 'T' is such a term, then I will know which property I pick out by the use of that term to characterise the world, regardless of which world is actual. No matter how things are with this actual world, no matter

what sorts of properties obtain there, the term 'T' will pick out the property that in idealised circumstances would play the idealised role. And so there will be no question of 'T' referring, for all I know, to this, that or another property. My knowledge of what property T-ness is will not be hostage to how the world happens to be and I will not be subject as a user of the term to epistemic disjunctivitis.

But are there any terms in ordinary response-dependent usage that refer, putatively, to idealised rather than instantiated realisers? There are two sorts of features that ought to mark these terms. First, they will be such that we can allow that they refer to nothing in the actual world while claiming that, still, they pick out determinate properties. And second, they will be distinctively abstract: they will direct us to abstract as distinct from concrete properties. Referring to the idealised realiser of a certain idealised role, such a term will pick out a property that abstracts from the different ways in which things may be under idealisation; it will point us towards a disjunction of the realisers that would play the role in different idealised worlds. Thus there should be no point to asking what else there is to be discovered empirically about the property, over and beyond what we know in virtue of being able to pick it out. There should be no point to asking the sorts of things we ask about the concrete property that constitutes realiser-redness when we wonder whether it is associated with this or that spectral reflectance.

These two features are characteristic of a number of terms, among which perhaps the most salient ones are geometrical predicates like 'straight' and 'parallel', 'flat' and 'smooth' and 'regular'. With such terms we are readily prepared to admit that none of the things in the actual universe, certainly none of the things with which we are familiar, may actually instantiate the corresponding properties: no edges may be straight, no pairs of edges parallel, and so on. And with such terms we do spontaneously see the properties to which they refer as being abstract rather than concrete. With the property to which 'red' directs us there are all sorts of empirical questions as to its physical nature, and so on, that naturally teem. With the property to which 'straight' or 'flat' directs us, there are not; we do not think of the property as one about which there is more to be empirically learned over and beyond what we learn in mastering the term or concept.

But how can a term like 'straight' be response-dependent and yet have an idealised, abstract referent? The response-dependent term is

always associated with the occurrence of an effect on human beings under independent, favourable specifications. Specifications of normality, such as those suggested for redness, involve ruling out factors of the kind that give rise to discrepancies across time and place: sodium lighting, rotating objects, coloured glasses, and the like. But favourable specifications may also involve the availability of, say, as much information as possible on a matter on which it is always possible to get more and more information (Pettit 1998). And with such specifications—with specifications that things are ideal, not just normal—we may have to admit that they cannot be fully satisfied in the sort of world that is actual; they refer us to wholly idealised conditions. Where a response-dependent term is guaranteed to go with the relevant response only in idealised conditions, it becomes feasible to think of the semantic value of the term, not as the instantiated property that fulfils that idealised role, but as the idealised property that does so: the property that would do so in idealised conditions. This, presumptively, is what happens with a term like 'straight'.

Is this edge straight, we ask. You say, yes; I say, no. Suppose that I can produce better information in the sense of being able, with the help of technology, to give you access to the edge at a greater level of tactile or visual resolution. In that case the discrepancy will naturally be resolved in favour of my response. The edge may be straight-for-practical-purposes—it may be approximately straight—but it is not straight in the strict sense of the term. Extrapolating from this case, we must admit that for any actual-world edge it is always possible to envisage more information such that it would lead us to say that the edge is not strictly straight. The property of straightness that we identify on the basis of our visual and/or tactile responses is one that will show up for sure only under conditions of information that are not satisfiable in the actual or in any plausible world. And so we are naturally led to admit that the property of straightness is idealised in character. Although we manage to make semantic contact with it—although it has the status of a property that we lock on to immediately—it is identified without any presupposition of instantiation.

We conceive of straightness, then, as the abstract property that would play the required role in idealised circumstances, not as the instantiated, concrete property that does so. The way in which our use of the term is guided shows that by our own lights the property might not be suitably instantiated. But under this account we still identify the property of being straight so far as we are capable of having certain responses; that is what

makes the corresponding terms response-dependent. It is just that the actual responses that we experience may be, for all that our conceptualisation entails, responses to things that approximate the property rather than instantiating it.

The lesson is clear. Noumenalism may be a cost that we have to pay for endorsing global response-dependence. But it is not a cost that we have to bear in each and every area of primitive predication. So far as primitive predication is idealised in the manner illustrated by a term like 'straight', it does not have to rely on assumptions about which properties are instantiated in order to provide us with properties that we can ascribe. The properties which it enables us to ascribe or not ascribe in relation to the world are identified without any help from the world. And so the argument to epistemic disjunctivitis does not get any grip here. We cannot begin to imagine that were the actual world different in some epistemically indiscernible way, then the properties identified would be different too.

The claim supported in the last section had seemed to be of the form: so far as semantically primitive predicates are response-dependent, our use of them in characterising the world commits us to there being features of the world that we do not know in their essence. The upshot of our discussion is that it should rather have the form: so far as semantically primitive predicates are response-dependent, and so far as they refer us to instantiated, unidealised properties, our use of them in characterising the world commits us to there being features of the world that we do not know in their essence. Assuming that not all semantically primitive predicates are idealised—and I am happy to assume this—global response-dependence does entail noumenalism. But it does not entail it on all fronts, only in regard to unidealised predicates.

5. Noumenalism is not a serious complaint

But not only is the danger of noumenalism—the danger of chronic epistemic disjunctivitis—quarantined in the area of unidealised predicates, it is also not as great a danger as it may seem. Two observations should help to show that there is no need to panic about the fact that global response-dependence entails a certain noumenalism.

The first observation begins from the assumption that intuitively the only serious sort of epistemic disjunctivitis is that which would affect ex-

planatorily or theoretically basic predicates, whether or not they are semantically basic. The observation itself is that such disjunctivitis threatens us quite independently of global response-dependence; it is not a condition that we can escape just by giving up on global response-dependence.

Theoretically basic predicates, whatever their semantic status, are those which direct us to the properties that by our lights play a fundamental role in determining the course of worldly events. They presumably include terms like 'mass', 'charge' and 'spin' but not words like 'red', 'straight' and 'regular'. Suppose, perhaps impossibly, that we could understand the essences of fundamental properties but not the essences of other properties: not even the essences of those peripheral properties by reference to which the terms for the fundamental ones are defined. Presumably this would not be a particularly unhappy state of affairs. Certainly it would not reek of the romantic gloom associated with noumenalism. Intuitively, the only serious sort of epistemic disjunctivitis is that which would affect theoretically basic predicates: that which would deny us knowledge of the essences of fundamental properties.

By the argument rehearsed in the third section, global response-dependence would entail such a fundamental as distinct from just peripheral noumenalism. While it bears primarily on semantically basic, unidealised predicates—and predicates, it may be presumed, that often refer to peripheral rather than fundamental properties—its knowledge-denying effects cannot be confined to the periphery. Global response-dependence would mean that however far we try to stick with role-properties whose essences are knowable, we have to recognise that the essences of the corresponding realiser-properties that we have to countenance—intuitively, more fundamental features—lie beyond our ken.

But if the problem with global response-dependence lies in its supporting a fundamental noumenalism, then we should recognise that this is a problem that it shares with the most salient alternative. This alternative would not appeal to idealisation in order to avoid noumenalism; after all, a global response-dependence theorist might appeal to idealisation in a parallel way.⁶ It would allege that the semantically basic, unidealised predicates on which we rely to pick out certain properties—presumably, peripheral properties—are response-independent and that more theoretical predicates are introduced by definition: they refer only indirectly to

fundamental properties, as the properties that play such-and-such roles in relation to the peripheral properties. But it turns out that under this alternative, as under global response-dependence, noumenalism continues to rule.

The most straightforward way of thinking of the definition of theoretical terms is on the model associated with Frank Ramsey, Rudolf Carnap and David Lewis (Lewis 1983, Essay 6; Oddie 1988). The terms are each associated with a network of connections—connections characterised in non-theoretical terms—both with one another and with further non-theoretical terms. And each term is taken to refer to the one and only instantiated property which, so it is supposed, answers to those connections. If you like, each term is associated with the role of satisfying the relevant connections and the property to which it is taken to refer is either the property of satisfying that role-specification—the role-property itself—or the property that actually satisfies the specification: the realiser-property.

If theoretically basic predicates refer indirectly in this, or in any similar fashion, then we are stuck with noumenalism. For if the referents of the predicates are the connectional realiser-properties, then those properties, those theoretically basic features, are available to us only via their effects—via their connections with other items—and are not known in their essence (Robinson 1993). And if their referents are the connectional role-properties themselves, then while those properties may be knowable in their essences they are not, after all, the theoretically most basic features. They are less basic, by the no-bare-roles principle, than the properties that realise them: and those properties are not known in their essences; they are known only as the properties that have a suitable realising or grounding effect. (For analogous considerations see Blackburn 1990, 64; cf. Foster 1982, 63.)

The second observation that I want to make is that part of the threat of noumenalism may depend, as it were, on some contingent stage-setting. In particular, it may depend on the suggestion that there are two radically different realms of properties, the anthropocentric and the cosmocentric, and that the cosmocentric are unknowable in their essences. It may depend, in Kant's phrasing, on the suggestion that there are two worlds, the phenomenal world and the noumenal world, the world-for-us and the world-in-itself. After all, the idea that there are two worlds of this kind is a particularly gloomy prospect; it suggests that we are cut off by an impenetrable curtain from how the world is in and of itself.

The presentation of noumenalism in this way, however, is not obligatory. Consider any area where epistemic disjunctivitis threatens. If we say that the properties to which we refer in those cases are role-properties, then it may seem that the properties whose essences we will not know all belong to quite a different class: the residue of realiser-properties required under the no-bare-roles principle. Thus we will be left with the image of a world-for-us—the world characterised in role terms—and the world-in-itself: the world of the ultimate realisers. But we can easily avoid that image, with its suggestion of a veil dividing the known from the unknown. All that we need to do is to take some of the referents of relevant terms to be realiser-properties, not role-properties. The properties may be those that realise the anthropocentric dispositions of things to elicit certain responses. Or they may be the properties that realise the connectional roles associated with theoretically defined terms.

Although I am in broad agreement with Smith and Stoljar, and although I rely heavily on their no-bare-roles principle, I think that they are guilty of suggesting that this stage-setting is inescapable. They suggest, in particular, that the properties—the unidealised properties—with which we make basic semantic contact under global response-dependence are all anthropocentric dispositions. “According to GRD (global response-dependence) the only claims we can ever make about the world are claims about the dispositions it possesses to elicit certain responses in us” (3). This leads them to see a dichotomy between the world-for-us and the world-in-itself. “What is true, if GRD is correct, is that we can only ever say how the world is in so far as it stands in certain relations to us and our responses, not how it is in and of itself” (Smith and Stoljar 2).

But this vision is not necessary and it does not represent the most appealing elaboration of global response-dependence. It suggests that as we try to talk and theorise about the world all we ever succeed in doing is talking, at least in the first place, about the anthropocentric dispositions of things. An alternative construal of the doctrine allows, however, that in our basic semantic forays we may make contact with non-anthropocentric properties—realiser-properties—as well as the corresponding anthropocentric roles. We do often talk about how attractive or appealing something is, of course, about how boring it is to find yourself in a certain situation and about how nauseating certain experiences are; that is to say, we do often address the anthropocentric dispositions of things around us.

But it would be hugely revisionary of our common sense to suggest that all of our semantically basic talk—even our talk about colours like redness—has the same parochial and subjective reference.

It will come as no surprise that, though I have often tried to write in a way that is neutral on the issue, I prefer the second, more ecumenical construal of global response-dependence (as discussed in Pettit 1997).⁷ I prefer the suggestion that while some of the properties that we mean to identify on the basis of associated effects or connections are role-properties, equally others are properties that realise the roles in question. But the point now is not to defend that way of construing things. The point is to emphasise that once we see the possibility of such a construal, we may not feel so depressed by the spectre of noumenalism.

The first observation was that noumenalism is not distinctive of global response-dependence; epistemic disjunctivitis threatens our grasp of theoretically basic properties under the most salient alternative. My second observation is that the spectre of noumenalism is often presented in an excessively theatrical way. Where it need only mean that we cannot know all worldly properties in their essences, it is presented as the claim that we cannot know the world-in-itself. Where it need only be associated with epistemological modesty—with what has been described as Kantian humility (Langton 1995)—it is presented as a belief in ontological mystery.⁸

Philip Pettit

*Research School of Social Sciences,
Australian National University*

NOTES

1. Those familiar with the literature will see that I am here making use of the central idea in two-dimensional modal thought. We do not start from the actual-world interpretations of our words and ask about what is possible and necessary. We ask rather about what interpretations our words would receive as different possible worlds play the actual-world role and about what would be necessary, what possible, under those interpretations. See Davies and Humberstone (1981).

2. This claim assumes, of course, that we know fully the linkages by which the role is defined. As Richard Holton has pointed out to me, someone might think that roles could be picked out on the basis of a proper subset of the linkages, in which case we might not know which role is picked out; those linkages, but not others, may be preserved as we imagine epistemically indiscernible worlds in the role of the actual world.

3. In assuming that we do successfully refer to a property, as of course the noumenalist will assume, I ignore the threat of error theory. If we are familiar with effects only and postulate a property at their origin—a property that we know only in those effects—then we may be mistaken in that postulation: there may not be any property there; certainly there may not be any single property there.

4. I remark in passing that in order for response-dependent terms to refer determinately to certain properties, it may not only be required that those properties should be identified as properties associated with such-and-such effects, however implicitly that is done; it may also be required that the properties satisfy extra constraints that have nothing to do with what people believe. For reasons to think that we need such a requirement see Devitt (1983) and Lewis (1984); see also Pettit (1996, Postscript).

5. As a matter of fact my preferred line is to allow that many of the properties with which we make basic semantic contact are realiser-properties; see Pettit (1997). More on this in the last section.

6. There is a theme to explore here. Someone might say that the concept of mass, for example, is idealised in such a way that for all that the semantics of the term requires it need not actually be instantiated, only approximated. Mass might be identified as the property, for example, that would produce certain effects under certain inputs in such-and-such idealised conditions.

7. Notice that even when the properties with which we make basic contact are the realiser-properties, as is possible on this view, it will still be the case in Smith and Stoljar's words that "we can only ever say how the world is in so far as it stands in certain relations to us and our responses." The message will be that it is only in virtue of the world's standing in certain relations to us—eliciting certain responses in us—that we can say how it is, including how it is in and of itself. Where their view would suggest that we can only talk at the basic level about how the world is for us, this would imply that we can talk about how the world is in and of itself, though only so far as that world has a certain impact on us.

8. My thanks to a number of friends for their very helpful comments: Sam Guttenplan, Richard Holton, Frank Jackson, Rae Langton, Peter Menzies, Michael Smith, Daniel Stoljar and Denis Robinson.

REFERENCES

- Blackburn, Simon. 1990. "Filling in Space," *Analysis* 50: 62–65.
- _____. 1991. "Losing Your Mind." In *The Future of Folk Psychology*, edited by J. Greenwood. New York: Cambridge University Press.
- Davies, Martin, and Lloyd Humberstone. 1981. "Two Notions of Necessity." *Philosophical Studies* 48: 1–30.
- Devitt, Michael. 1983. "Realism and the Renegade Putnam: A Critical Study of meaning and the Moral Sciences." *Noûs* 17: 291–301.
- Foster, John. 1982. *The Case for Idealism*. London: Routledge.
- Jackson, Frank, and Philip Pettit. 1988. "Functionalism and Broad Content." *Mind* 97: 381–400.
- Jackson, Frank and Philip Pettit, 1988. "Response-dependence Without Tears."
- Johnston, Mark. 1989. "Dispositional Theories of Value." *Proceedings of the Aristotelian Society* Suppl. 63: 139–74.

- Johnston, Mark. 1993. 'Objecivity Refigured: Pragmatism with Verificationism.' In *Reality, Representation and Projection*, edited by J. Haldane and C. Wright. Oxford: Oxford University Press.
- Langton, Rae. 1995. *Kantian Humility*. Ph.D. thesis, Princeton University; forthcoming as a book from Oxford University Press.
- Lewis, David. 1983. *Philosophical Papers Vol. 1*. Oxford: Oxford University Press.
- Lewis, David. 1984. "Putnam's Paradox." *Australasian Journal of Philosophy* 62: 221–36.
- Oddie, Graham. 1988. "On a Dogma Concerning Realism and Incommensurability." In *Relativism and Realism in Science*, edited by R. Nola. Dordrecht: Kluwer.
- Pettit, Philip. 1990. "The Reality of Rule-Following." *Mind* 99: 1–21.
- _____. 1991. "Realism and Response-dependence." *Mind* 100: 587–626.
- _____. 1996. *The Common Mind: An Essay on Psychology, Society and Politics*, Paperback ed'n, with new Postscript. New York: Oxford University Press.
- _____. 1997. "Terms, Things and Response-Dependence." *European Review of Philosophy* 2.
- _____. 1998. "A Theory of Normal and Ideal Conditions," *Philosophical Studies*, vol. 65.
- Robinson, Denis. 1993. "Epiphenomenalism, Laws and Properties." *Philosophical Studies* 69: 1–34.
- Smith, Michael and Stoljar, Daniel. 1998. "Global Response-Dependence and Noumenal Realism." *The Monist*, vol. 81, no. 1 (January) [this issue].
- Wright, Crispin. 1993. Realism: "The Contemporary Debate—Whither Now?" In *Reality, Representation and Projection*, edited by J. Haldane and Crispin Wright. Oxford: Oxford University Press.