

TWO PATHS TO PRAGMATISM II¹

Huw Price

School of Philosophy
University of Sydney

1. INTRODUCTION

Particular topics of conversation seem to be inaccessible to speakers who lack an insider's view of the subject matter concerned. The familiar examples involve sensory deficiencies: discourse about music may be inaccessible to the tone deaf, wine talk to the anosmic, the finer points of interior decorating to the colour blind, and so on. The traditional secondary qualities thus provide the obvious cases of concepts which seem to exhibit this form of subjectivity—this dependence on specific and quite contingent human capacities.

It is easy to think of further cases, resting on more specific perceptual and quasi-perceptual disabilities. (Could someone with no ball sense understand what it is like to be a batsman, for example?) But is this the end of the matter? How far does this phenomenon extend in language? How precisely should it be characterised? And what is its significance, if any, for the metaphysical status of the concepts concerned? Questions of this kind have been the focus of considerable attention in recent years. The original impetus for much of this work seems to have been the suggestion that there might be a useful analogy between moral concepts and the secondary qualities. This rather specific issue has then given rise to a more general interest in the nature and significance of the kind of subjectivity exhibited by the secondary qualities (and/or moral concepts). Writers such as Mark Johnston and Crispin

¹This is a second edition, not a sequel, of a paper which first appeared in P. Menzies (ed.), *Response-Dependent Concepts (Working Papers in Philosophy, No. 1)*, Canberra: Philosophy Program, Research School of Social Sciences, ANU, pp. 46-82. The original paper was written for a weekend conference on response dependence at ANU in June, 1991. I am grateful to David Macarthur for extensive comments on an early version, to Peter Menzies and Michaelis Michael for helpful discussions before the conference, and to Mark Johnston and Philip Pettit, especially, for comments and discussion at the time of the conference.

Wright have sought to develop general formal frameworks within which to represent the dependence of concepts of a certain kind on particular human capacities.²

As Johnston well appreciates, one excellent motive for seeking a general characterisation of the kind of subjectivity in play in the case of the secondary qualities is to equip oneself for an assault on one of the most fascinating issues in philosophy: How much of the conceptual framework we apply to the world is simply taken over from the world itself, and how much of it ‘comes from us’? In a philosophical climate often dismissive of the sort of pragmatism that takes this question seriously, is it pleasing to see the attention it receives in Johnston’s hands.

All the same, it seems to me that Johnston has done pragmatism a disservice. His advocacy of response dependence as the general species of subjectivity apparently exhibited by the secondary qualities has helped to obscure an attractive alternative path into the same territory. The effects of this are evident in Johnston’s own position. In ‘Objectivity Refigured’ Johnston argues tellingly that there is space for a species of pragmatism distinct from verificationism and Putnam’s internal realism. He then assures us that the route to a pragmatism of this kind ‘goes by way of developing and applying the notion of response-dependence.’ (OR, p. 103) So it comes as something of an anti-climax when he concludes that the notion of response dependence is of limited application, even in the hands of philosophers who are prepared to revise ordinary usage to put it on respectable metaphysical foundations. Though Johnston attempts to salvage some general consequences from a response-dependent treatment of theoretical ‘rightness’, pragmatists hoping for historic victories on the field of metaphysics will feel a sense of anticlimax.

² See their respective contributions to J. Haldane and C. Wright (eds.), *Reality, Representation, and Projection*, Oxford University Press, 1993: Wright’s ‘Realism: The Contemporary Debate—W(h)ither Now?’, pp. 63–84, and Johnston’s ‘Objectivity Refigured: Pragmatism Without Verificationism’ (hereafter OR), pp. 85–130. See also Johnston’s contribution to the symposium ‘Dispositional Theories of Value’ (hereafter DTV), *Proceedings of the Aristotelian Society, Supplementary Volume LXII*, 1989, and Wright’s *Truth and Objectivity*, Harvard University Press, 1992. In keeping with the focus of the conference for which this paper was first prepared, I shall concentrate on Johnston’s work. However, given that Wright’s account invokes biconditional content conditions similar to those employed by Johnston, and that my main concern is to argue for a quite different means of explicating the significance of the relevant linguistic expressions, I think that much of what I say would equally stand opposed to Wright’s approach.

I want to show that the problem lies in the choice of response dependence as the path to the general issue of the extent of subjectivity in our representation of the world. There is an alternative route, and one which promises bigger dividends. Indeed, it suggests that the subjectivity in question is global, infecting all parts of language (though different parts to different extents). In endorsing a form of global pragmatism I concur with Philip Pettit.³ However, Pettit follows Johnston in regarding response dependence (or a closely related notion) as the proper path to such a pragmatism, and apparently in failing to appreciate that there is an alternative. (He disagrees with Johnston mainly in arguing that the resulting pragmatism is globally applicable.) But the global enterprise turns out to be particularly sensitive to the choice of route: the alternative path has advantages even in local cases, but becomes mandatory if the pragmatist's project is to be generalised in this way. So while I join Pettit in the quest for global pragmatism,⁴ I want to argue that in setting a course in terms of response dependence, he is guiding our common ship by the wrong star.

So the main task of this paper is philosophical cartography. I want to put the alternative path to pragmatism on the philosophical map, and draw attention to some of its advantages. The aim is to chart the least hazardous course for those of us who cannot resist the lure of a world whose conceptual joints all owe something to ourselves; and to present the complete picture for the benefit of less adventurous souls, so that at least they may see towards what monsters we set sail.

2. TWO WAYS OF EXPLICATING MEANING

The two paths to pragmatism are distinguished by the fact that they rely on different theoretical strategies for explicating the significance of a concept or linguistic expression. The first strategy—that employed by response theorists—is exemplified by the following simple version of Johnston's 'basic equation':

³ Philip Pettit, 'Realism and Response-Dependence' (hereafter RRD), *Mind* **100**(1991), 587–626.

⁴ Indeed, we both sail under the rule-following banner. See RRD, p. 588 and my *Facts and the Function of Truth* (hereafter FFT), Oxford: Basil Blackwell, 1988, pp. 192–5.

- (1) X is C iff X is disposed to produce response R_C in normal subjects under suitable conditions.

I shall call this a *content condition*. Not much hangs on the terminology. Johnston, Pettit and others warn us not to interpret such biconditionals as reductive *analyses* of the concepts referred to on the LHS. However, the term ‘content condition’ will serve to mark a contrast with a second sort of explication of the significance of an expression, which is what I shall call a *usage condition*. The general form of such a condition is something like this:

- (2) The utterance S is prima facie appropriate when used by a speaker who has a psychological state ϕ_S .

The particular usage condition which contrasts to (1) will then be:

- (3) The utterance ‘X is C’ is prima facie appropriate when used by a speaker who experiences response R_C in the presence of X.

A usage condition thus tells us something about the *subjective assertibility conditions* of an expression—about what condition must normally obtain *in the speaker* for the utterance of an expression to be appropriate. As such, it is of course only the very first step in a theoretical description of linguistic practice. It leaves out the Gricean hierarchy of intentional attitudes, to mention just one aspect of the larger picture. All the same, it may be an important first step, as a few familiar examples will illustrate.

Emotivism

The utterance ‘X is good’ is prima facie appropriate when used by a speaker who approves of (or desires) X.

Adam’s Hypothesis

The utterance ‘If P then Q’ is prima facie appropriate when used by a speaker who has a high conditional credence in Q given P.

Assertion and Belief

The assertion ‘A’ is prima facie appropriate when used by a speaker who has the belief that A.

Again, these principles are only the most basic elements of a pragmatic account of the usage conditions of the utterances concerned. So it would be beside the point to object at this point that it is not always appropriate to say 'X is good' when one approves of X. It is not always appropriate to say 'P' when one believes that P, but there is still something important and informative in the principle that one should normally assert that P only if one does believe that P.

I shall be arguing that the insights of pragmatism are better served by usage conditions than by content conditions. However, I emphasise that I am not suggesting that content-specifying conditions have no role to play in the philosophical project of explicating meaning. In my view a pragmatist is likely to find as much use as anybody else in a content-specifying truth theory of the Tarski-Davidson sort, and should have no more qualms about

(4) 'X is good' is true iff X is good

than about

(5) 'X is electrical charged' is true iff X is electrical charged.

My argument simply concerns the appropriate vehicle for the distinctive points a pragmatist wants to make about the dependence of particular areas of discourse on human capacities or points of view. I want to urge that these points are better made in terms of usage conditions than in terms of the kind of biconditional content conditions that writers on response dependence have offered us.

I noted above that my description of these biconditionals as 'content-specifying' might be thought contentious, given that Johnston and others take pains to point out that these principles are not to be read as straightforward reductive analyses. I don't want to become entangled in a discussion of the possible varieties of conceptual analysis. I simply want to show that even if we put these labels to one side there is a clear distinction between response theorists and a usage pragmatist, in that usage conditions do not yield several of the key conclusions that Johnston and Pettit draw from their versions of response dependence. This will be enough to show that content conditions and usage conditions represent two

incompatible paths to pragmatism. It will also give us some reasons for preferring the latter path.

One more terminological point: once it is appreciated that content conditions and usage conditions provide two distinct strategies for explicating the dependence of particular discourses on human abilities, responses and capacities, then a decision is needed as to whether we should take the term ‘response dependence’ to be applicable to both, or only to the content-based approach. As my use of the term ‘response theorist’ already indicates, I take the latter course. Given the central role of biconditional content conditions in the work of those who have introduced the term ‘response dependence’, I think it would be misleading to apply it to an approach which rejects these biconditionals. I therefore refer to writers who rely on these content conditions as *response theorists*.

It should be emphasised that many of the general remarks response theorists make in characterising their pragmatic stance are quite compatible with a usage-based approach. In the first paragraph of ‘Realism and Response-Dependence’, for example, Pettit asks us to

Consider the concepts of smoothness, blandness and redness. They are tailor-made for creatures like us who are capable, as many intelligences may not be, of certain responses: capable of finding things smooth to the touch, bland to the taste, red to the eye. The concepts, as we may say, are response-dependent.

We may agree with Pettit that concepts such as these depend on particular human responses or capacities, and yet deny that this entails that they should be (let alone *need be*) seen as response-dependent, in the technical sense of the term that Johnston has introduced. The main point of this paper is to show that there is an alternative to response dependence, providing a better way to explicate the relevant species of dependence on particular responses. It is perhaps a pity that the useful descriptive term ‘response dependent’ has come to be associated with a less than general account of what the general phenomenon amounts to, but there it is. The crucial point to keep in mind is that the general pragmatic program is not at issue here. The dispute concerns the correct strategy for putting the program into effect.

Finally, it should be noted that Johnston himself distinguishes two sorts of conclusion the pragmatist might aim for. The first is a descriptive doctrine about a discourse as it stands

in ordinary use. The second is the revisionary proposal to replace an existing discourse with a discourse couched in acceptable pragmatic terms. My main interest is in the descriptive doctrine—in the question as to how to explicate actual linguistic practice. I want to show that the usage-based approach does better at this descriptive level than the response theorist's alternative. But we shall see that because it does better here, the usage approach avoids some difficulties which the response theorist might take to motivate the revisionary doctrine. As always, the linguistic revisionist bears an onus to justify the proposed change. For a response theorist the claimed justification might be that the (non-response-dependent) concepts presently employed in some area of discourse do not properly reflect their evident dependence on contingent human capacities; the proposed revision is intended makes this dependence explicit. However, the usage-based approach shows that this is a solution to a non-existent problem. Usage conditions provide an alternative way to represent the dependence of particular discourses on human responses and capacities, and response theorists have simply looked in the wrong place. The discourses in question may adequately reflect their subjective origins as they stand, and hence be in no need of revision.

More later on the advantages of the usage approach. The first task is to show that the usage and content approaches are genuinely distinct. The best way to do this is to show that the usage approach does not have certain consequences that response theorists rightly draw from their content-based accounts.

3. USAGE CONDITIONS: THEIR THEORETICAL AUSTERITY

If the Johnston's biconditional is to be called the basic equation then perhaps the alternative usage condition should be referred to as the *more* basic equation. For the usage condition embodies a very austere, general, and theoretically fundamental perspective on linguistic practice. Because it is so general and so austere, its philosophical consequences are correspondingly thin. It is committed to little, compatible with much. It is compatible with the most thoroughgoing realism, for example—realists about a particular subject matter will not deny that assertions about that subject matter are governed by a usage condition of the kind sketched above under *Assertion and Belief*.

Similarly, the usage perspective does not of itself commit us to the view that speakers' responses or psychological states provide privileged access to the truth of the utterances concerned. This is clear in the case of *Assertion and Belief*, where the mistake of thinking otherwise is the Protagorean mistake of thinking that because we assert that P when we believe that P, to assert that P is to *say* that one believes that P (and is therefore true if one believes that P).

More interestingly, it seems that the usage perspective is not committed to privileged access to truth even in the cases which will provide the direct alternative to the response theorist's biconditional, such as (3). For example, consider an explicitly fallibilist version of Moorean moral intuitionism. Here, the primary usage condition for 'X is good' will be said to be whatever manifests the intuition that X is of moral worth. But this might be regarded as a fallible and defeasible condition, perhaps in such a way that normal speakers might be systematically mistaken about the moral status of certain kinds of entities. (Normal speakers might exhibit an irrational bias in favour of the so-called deserving poor, for example.) Normal consensus would thus be regarded as neither necessary nor sufficient to guarantee truth. Another example might be provided by realism in mathematics, coupled with a usage condition in terms of the belief that one has a proof. In this case normal community consensus might be sufficient but not necessary for truth. So even in these cases, *the usage perspective is not necessarily response-privileging*. It is compatible with a more 'cosmocentric' brand of realism, to use Pettit's term.

On the other hand, a usage-based approach is also compatible with a much less realist overlay than the response theorists offer us. In particular, it is compatible with the view that the area of discourse in question does not have truth conditions, or is non-factual in character. Examples are provided by the two principles headed *Emotivism* and *Adam's Hypothesis* above. Emotivism is usually associated with the view that moral judgements are non-factual and lack truth conditions; while Adam's Hypothesis is held by some to be the centrepiece of a non-truth-conditional view of indicative conditionals.

The fact that the usage-based theory is compatible with non-factualism perhaps marks the most striking and interesting contrast with content-based accounts. A response-dependent

account is automatically truth-conditional. This is of course implicit in the biconditional form of the basic equation—there is nothing to stop us adding ‘... is true’ to the LHS, after all. Response theorists thus get their factualism for free, and therefore find it difficult to give serious consideration to alternative views. Johnston suggests that the fault lies with the opposition:

An appropriate response-dependent account [of value] may thus threaten to make quasi-realism redundant. The quasi-realist programme is to *somehow* defend our right to employ the truth-conditional idiom in expressing evaluations.

However, the

response-dependent account of value ... implies that there is no need to *earn* the right to the truth-conditional form of expression. ... [W]e have a natural conceptual right to this truth-conditional form. (DTV, pp. 173-4)

Pettit simply passes over the point, noting that ‘there is no pressure from the traditional response-dependent thesis to go towards an instrumentalist ... theory.’ (RRD, p. 607) Surely this is an indication that something has been missed. It is clear that philosophical attention to the dependence of particular concepts and discourses on human capacities has often found expression in non-factualist theories of one sort or another. Emotivism provides one familiar example, but there many others. Simon Blackburn’s quasi-realist program provides a whole range of cases, for example—cases which, as Blackburn points out, may trace their ancestry to prominent themes in Hume.⁵ These examples demonstrate that the pragmatist’s concern to exhibit the dependence of language on human capacities and responses is not incompatible with non-factualism. So if the response theorist’s reading of the pragmatist program cannot find room for non-factualism, so much the worse for that reading—or at least for the view that it is the sole available reading of what the pragmatist is up to.

At any rate, the usage approach provides a viewpoint from which the quasi-realist’s concerns re-emerge. It characterises a pattern of usage using the very materials to which

⁵ See Blackburn’s *Essays in Quasi-Realism*, Oxford University Press, 1993.

dispositional theories of value appeal, namely the existence on the part of speakers of dispositions to value certain things and states of affairs. But it does so in terms which don't prejudge the issue as to whether (and if so why) a pattern of usage meeting this characterisation need take a truth-conditional form. On the face of it, speakers might display the relevant dispositions in language in a very minimal way, if at all. Humans might boo and cheer. Many animals successfully express their likes and dislikes in all sorts of pre-linguistic ways. So a lively disposition to value some things and disvalue other things need not sustain a *truth-conditional* linguistic practice, and it is entirely appropriate to enquire why, and with what justification, we ourselves do express certain evaluative dispositions in assertoric form. This is the quasi-realist's concern. It is a concern clearly addressable from the usage perspective, and hence the fact that a response-dependent account cannot make sense of it serves to emphasise the difference between the usage perspective and the content perspective. Both are pragmatic; both characterise evaluative discourse in terms of human responses. But content conditions and usage conditions engage with these responses in very different ways.

In trying to show that usage conditions provide the better path to pragmatism, I shall offer four main arguments. The first (§ 6) turns on the consequences of the view that the species of pragmatism with which we are concerned is global in its application to language. I endorse this view, but argue that it leads to incoherence unless couched in terms of a usage pragmatism. Because it turns on globality, this first argument is not directly effective against a response theorist such as Johnston himself, who takes response dependence to be a local doctrine (and indeed potentially a revisionary doctrine). Against this position, my second argument (§ 7) is in a sense a reversal of Johnston's above redundancy claim concerning quasi-realism. I argue that wherever there are the resources for a response-dependent reading (or re-reading) of a discourse, there are also the resources the quasi-realist needs for an adequate *explanation* of truth-conditional practice. So the response theorist's re-reading works only where it is unnecessary. But the relation is not symmetric: the usage-based approach often works where response dependence would not. In particular, it is entirely compatible with the minimalist metaphysical position that Johnston sees as the main rival to a response-dependent pragmatism with respect to many philosophically interesting topics.

In defence of a revisionary response-dependent reading it might be claimed that at least in certain cases, it alone enables us to *justify* judgements that play important roles in our lives. My third argument (§§ 8–9) counters this claim, and in the process shows that the usage-based approach does a better job of explaining the characteristic role of judgements of various kinds. All parties agree that an account of evaluative judgements should explain why they motivate us (and *should* do so); an account of probabilistic judgement should explain its connection with decision under uncertainty, and so on. I want to show that the response theorists do rather poorly on this crucial task—a failure masked to some extent by a characteristic equivocation concerning the notion of rationality.

Finally (§§ 10–13) I argue that the usage approach gives a more accurate explanation of some of the subtleties of ordinary language. In particular, it makes much better sense than its content-based rival of the peculiar mix of objectivity and subjectivity we encounter in discourses subject to no-fault disagreements.⁶ This should recommend the approach to those less sanguine than Johnston seems to be about the prospect of wholesale revision of linguistic practice.

The next two sections establish some of the groundwork for these arguments, particularly by drawing attention to useful analogies between response dependence and more familiar philosophical strategies.

4. *ASSERTION AND BELIEF*: SOME PROTAGOREAN LESSONS

The principle *Assertion and Belief* embodies the uncontroversial idea that an account of the historical and psychological foundations of language will give a central explanatory role to the notion of belief; perhaps not to the full-blown propositional attitude, but at least to some weaker notion of behavioural commitment. For in an obvious sense, belief underlies and sustains the linguistic activity of *assertion*. Creatures who didn't have beliefs couldn't make assertions. Whatever else we say about assertion, we'll want to mark the fact that speakers

⁶ All discourses have to potential to exhibit such disagreements, in my view; hence the global character of a usage-based pragmatism. See FFT, ch. 8.

typically assert that P only if they believe that P. Yet none of this entails that in saying that P a speaker asserts *that* he or she believes that P. True, it means that we can ordinarily infer that a speaker believes that P from the fact that she says that P. But there are several simple objections to identifying the content of an assertion with the belief which it normally expresses.

In particular, such an identification leads to analytic regresses, arguably vicious, of more than one kind. Suppose, for example, that our interest is in a general account of assertion. The view in question tells us to interpret the assertion that P as the assertion that the speaker concerned believes that P. Here the target notion of assertion occurs again, central and yet unanalysed.

Alternatively, let's focus on the content of the assertion that P. The view in question now leads us to the principle that this content is in some sense equivalent to that of the assertion that the speaker believes that P. The content of the latter assertion is then in the same sense equivalent to that of the assertion that the speaker believes that the speaker believes that P, and so on. The way is clear to a denumerable infinity of content sentences, whose related contents are all equivalent in the specified sense. So far what we have is not necessarily vicious. After all, it is a familiar idea that a denumerable infinity of logically equivalent sentences may all have the same content. But it teeters on the brink of viciousness: if the slightest analytic pressure is put on the RHS of the relevant biconditionals, the whole structure collapses. For example, if we say on the psychological plane that to believe that P is to believe that 'P' is true—that is, to believe that one believes that P—then immediately we are lost. Or if we say on the semantic plane that what *makes true* my assertion that P is the truth of the proposition that I believe that P, then again the regress carries us away. If not already incoherent, then, the position in question is at least highly unstable.⁷

⁷ In OR, p. 106, Johnston suggests that circularity will be a problem for the Protagorean 'only if it made the biconditionals and their associated identities empty'—i.e., if the Protagorean was aiming for a reductive definition, and the biconditional was too weak to supply it. But the infinite regress threatens because the biconditional is too strong, not because it is too weak.

One further point: if a reference to belief were a universal ingredient of the content of assertoric judgement in this way, then an obvious theoretical strategy would be to try to factor it out, and to concentrate on what varies from assertion to assertion, namely the content of this embedded belief—which would just take us back to where we started. This might well be the best practical strategy, as well. If *all* assertion were initially of the form ‘I believe that P’, then the most advantageous strategy might be to let the qualification drop away; to invent a discourse of quasi-assertions, which pretended to talk about the world in an unqualified way. All the materials would be already at hand. The principle governing truth ascription would thus be that we should not worry about what *really* made one of these quasi-assertions true—they are, after all, only a game of our own devising—but simply that we should be prepared to *say* that one was true or (‘quasi-true’) when we had the relevant belief. Thus one should be prepared to quasi-assert that P, or ascribe quasi-truth to the quasi-assertion ‘P’, when and only when one believes that P.

Thus the Protagorean’s own materials allow us to construct a non-Protagorean practice. To close the trap, we need simply point to the possibility that the latter practice may be our actual practice. Later, this will be my strategy against content-based response theories.

5. SELF-DESCRIPTIVISM

Thus in the familiar case of *Assertion and Belief*, it is easy to distinguish the role that belief properly plays in a usage-based account of the linguistic activity of making assertions, from a role that it might mistakenly be thought to play in a content-based theory of the content or truth conditions of assertoric utterances. In less commonplace places, however, it seems to have been less easy for philosophers to keep the distinction firmly in mind. Consider, for example, the familiar kind of non-cognitivist (or non-factualist) view that utterances of some disputed class do not express *beliefs*, but rather some other kind of propositional attitude. This is what a simple kind of emotivism says about moral judgements, or what a simple subjectivism says about probabilistic judgements. In these cases the relevant propositional attitudes are *approval* and *credence*, respectively: the emotivist says that ‘It is good that P’ expresses a speaker’s approval that P, not her belief that it is good that P; while the

subjectivist says that ‘It is probable that P’ expresses a speaker’s high degree of credence that P, not his belief that it is probable that P.

These familiar expressivist positions are usage-based, not content-based: they give us usage conditions, not content conditions, for moral and probabilistic utterances. But they have often been confused with the following *self-descriptive* interpretations of moral and probabilistic claims: the view that in saying ‘It is good that P’ a speaker *says that* she approves of the fact that P; and the view that in saying ‘It is probable that P’ a speaker *says that* he has a high degree of credence that P. However, I hope it is clear, at least on reflection, that self-descriptivism is actually quite different from expressivism. The difference is just like that between the (correct) view that we typically say that P *when* we believe that P and the (incorrect) view that in saying P we say *that* we believe that P. In each case the difference rests on that between two distinct possible roles for the psychological state concerned: a usage-based role in an *explanation* of a linguistic practice, and a content-based role in an *analysis* or *explication* of the content of the utterances comprising that practice. Expressivism is usage-based, self-descriptivism is content-based.

However, while the self-descriptive interpretation is easily seen to be close to incoherency in the standard case of belief and assertion, it is not so obviously mistaken as an alternative to emotivism and probabilistic subjectivism. Indeed, it may seem to have some attractions, foremost of which might be thought to be that it doesn’t leave us with any mystery as to what moral or probabilistic judgments actually mean, or as to why they look like regular assertions. They are regular assertions on this view, the only oddity being that their subject matter is not what it seems to be: they are about their utterer’s state of mind, rather than about moral or probabilistic aspects of the world. (These views are not non-cognitivist, and are only non-factualist about their *apparent* factual referents.) A self-descriptivist might thus express puzzlement at the emotivist’s concern with the issue as to why moral judgements take indicative form, and appear truth-conditional. ‘Once we explicate moral judgement in the way I suggest’, the self-descriptivist might say, ‘We see that we have a natural conceptual right to the truth-conditional form of expression.’

The response theorist's content conditions are not as implausible as those of these naive self-descriptivists, of course. All the same, I think that there is a valid comparison between the two kinds of position. In my view, the fallacy of interpreting expressivism as self-descriptivism involves simply a more graphic form of the same blindness that characterises the response dependence program—that is, a blindness to the very possibility of a usage-based account built on the same psychological foundations. If nothing else, then, it will be helpful to have the self-descriptivist fallacy available as cautionary lesson, as I try to bring back into focus the possibility and advantages of a usage-based approach.

6. THE CONSEQUENCES OF GLOBAL PRAGMATISM

As I noted in §1, Philip Pettit argues that all concepts are response-dependent. His case turns on Wittgensteinian considerations about rule following: in virtue of the inability of any finite class of exemplars to constrain an intended concept to future cases in a unique way, the applicability of a concept *F* to a new potential instance *X* is always, in the last resort, dependent on the disposition of the members of a speech community to regard *X* as *F*. (The point is perhaps more evident if cast in terms of terms rather than concepts.) If endorsed, as I think it should be, Pettit's point has important ramifications for our present concerns. It invites a global generalisation of the dispositional method of analysis. If

(6) *X* is good iff *X* is disposed to evoke (i.e. is such that it evokes) the 'valuing' response in us.

then why not also

(7) *X* is *F* iff *X* is disposed to evoke (i.e. is such that it evokes) the 'Seeing as relevantly similar to paradigm *F*'s' response in us.

This suggestion cannot be rejected on the grounds that there is something illegitimate about the nature of the response referred to on the RHS. The point of Pettit's argument, and ultimately of the rule-following considerations themselves, is that language depends on responses or dispositions of this kind. Our use of the word 'chair' ultimately depends on the

fact that we have a certain ability to generalise to new cases on the basis of exposure to a small number of exemplars.⁸

What are the consequences of taking (7) seriously, however? In considering this question, let us pretend for the moment that the response theorist's biconditionals are to be taken as outright content specifications: if not as reductive analyses, then at least as indications of conceptual equivalence, or common content on both sides of the biconditional. Apparently, then, the consequence of (7) is that we cannot talk about the world as such, but only about its effect on us. All content becomes relational and anthropocentric in this way.

Recall that this is just what happens according to the Protagorean view that to assert that P is to assert that one believes that P. In §4 I mentioned several objections to this view, turning mainly on the idea that it threatens a vicious regress. Similar problems would appear to afflict any attempt to take (7) as the general form of an analysis of concept ascription. The biconditional describes, if not an equivalence between contents, then at least an equivalence relation (viz., sameness of truth conditions). If the doctrine is to be universal, then the fact that the RHS is at the same time more complex and yet of the same basic logical form as the LHS—the form 'X is G'—will guarantee us a denumerable infinity of content sentences, all specifying contents that stand to one another in the same equivalence relation. From here viciousness threatens in the same way as before: if the slightest weight is put on the RHS, the whole structure is liable to collapse.

Can a response theorist avoid this difficulty, by insisting that the relevant biconditionals have *no* analytic significance? It seems doubtful. As in the Protagorean case, even the suggestion that the RHS specifies what makes the LHS truth-apt will be enough to tip the balance, and it does seem that the response theorists are committed to this. As we have seen, Johnston relies on the biconditional to illustrate the response theorists' 'natural conceptual right' to truth conditions, for example, while Pettit apparently takes it as the basis of his

⁸ Reductive definition of a term is a possibility, of course, but this simply delays the inevitable, as Kripke points out in *Wittgenstein on Rules and Private Language*, Oxford: Basil Blackwell, 1982.

suggestion that response dependence provides no pressure towards non-factualism.⁹ But if the LHS has its truth value (or truth aptness) *in virtue* of that of the RHS, and the RHS has to acquire its own truth value (or truth aptness) in the same way, then the regress is certainly vicious. (By way of analogy, let's say that an object is well-grounded if it is standing on top of an object which is itself well-grounded. Now suppose that it's turtles all the way down. Can we show that the top turtle is well-grounded? No, for the same issue arises all over again at every level in the hierarchy of turtles.)

I think this is a decisive objection to a content-based global response theory. It is not a problem for local response theorists, of course, and nor is it a problem for a theory which accepts that language is globally dependent on dispositions 'to go on in the same way', but explains the role of these dispositions in usage conditions rather than in content conditions. We can indeed be global pragmatists, then, but only if we abandon response dependence in favour of usage conditions.

7. THE REDUNDANCY ARGUMENT

In §4 I mentioned another objection to the Protagorean view. It was that the Protagorean's own materials suffice to construct an alternative 'language game', with respect to which the appropriate reference to belief takes the form of a usage condition rather than a content condition. In effect, then, the Protagorean's own resources provide a plausible alternative model of what might be going on in actual discourse. Unlike the regress point itself, this idea can be developed again local versions of content-based response theories.

Let's consider the raw materials required for a response-dependent account of, say, redness. What must be true of the members of a speech community in order for the dispositional account of colour to yield truths of the form 'Roses are red, violets are blue'? Simply that by and large, the members of the community share stable dispositions to experience similar responses in response to similar objects on similar occasions. The response

⁹ Pettit also indicates (RRD, p. 608) that the LHS of such a biconditional might be empirically falsified, in virtue of the lack of the various responses and conditions mentioned on the RHS.

theorist then tells us that when a member of the community says ‘Roses are red’, what they say is true if and only if roses have the disposition to produce a particular psychological condition—the ‘redness’ response—in normal members of the community under suitable conditions.

Suppose that a particular community have the appropriate responses, currently use the term ‘red’ in accordance with this model, and are theoretically sophisticated enough to realise that this is what they are doing. It is open to them, in principle, to make the following linguistic policy decision. From now on, when they teach their children how to use the term ‘red’ they will pretend that there is a response-independent property of redness in the world, and conform their practice to this pretence. I don’t mean that they will explicitly say to their children that redness is response-independent; just that they will introduce ‘X is red’ as the sort of claim that is capable of being true or false, independently of human dispositions to see things as red. Now it is a moot point just what is required to inculcate such an understanding. If our naive notion of colour is already response-independent, as Johnston suggests it may be, then presumably these linguistic policy makers need only decide to follow the teaching practices actually followed by unsophisticated English-speaking folk. But if this is not enough, they could add explicit lessons in the fallibility of human access to colour, taking care to pretend epistemological modesty with respect to colour. The crucial point is that even this is not incompatible with a reliance on the usual means of indicating to novices the extension of these properties, namely the use of paradigm cases. (All it requires is the additional lesson that even the paradigms are defeasible.)

The policy makers thus set out to instil in the practice of novice speakers two main habits or principles. The first principle is the primary usage condition governing assertions of the form ‘X is red’—the relevant instance of (3), in effect. The second is the habit of taking redness to be something that falls under the objective mode of speech. Against the general background of assertoric practice, the way to combine these lessons will be to teach novices to describe their redness experiences in terms of the notions of perception and belief—ordinary, world-directed perception and belief, of course, not any introspective variety. In treating the distinctive redness response as defeasible perceptual grounds for a corresponding belief, we

open the way to such comments as ‘You believe that it is red, but is it *really* red?’ This in turn may call into play the standard methods of rational reassessment. In virtue of their acquaintance with the objective mode *in general*, speakers will be led into the practice of subjecting their colour judgements to reflective scrutiny by themselves and others. The objective mode brings with it the methods and motives for rational enquiry. But notice the way rationality gets into the picture here. It is not a part of a content specification for the utterances in question, but a gift that comes for free with any choice to adopt the objective mode of speech, to use the indicative form, to speak of truth and belief. One of the attractions of usage-based rather than content-based pragmatism seems to me to be that it puts rationality in its proper place—extrinsic rather than intrinsic to the discourse. (More on this in the next section.)

Of course, we are not free to adopt the objective mode wherever we like; or rather, it won’t get very far in some cases. It works in the colour case because our colour responses are sufficiently alike, across the community, for rational reassessment to get some grip. But notice that the requirements are extremely flexible. The practice is quite capable of tolerating a considerable degree of difference. For one thing, it may be quick to exploit pockets of similarity, small communities of like-minded responders. But it does this not by relativising its claims to the standards of the community, but by exploiting the fact that the similarity within the community makes possible some degree of rational agreement on unrelativised claims.

Again, I’ll come back to this. In §11 I’ll argue that this flexibility with respect to the scope of the speech community provides another advantage of the usage-based approach. For the moment I simply want to emphasise the more basic point that where we have the resources in terms of which the content-based response theorists construct their accounts—i.e. a sufficiently uniform pattern of responses across a community—these same resources provide sufficient foundation for a non–response-dependent discourse, characterised theoretically (and in effect to its novice speakers themselves) in terms of a usage condition. Thus it is the response theorist’s account which now seems in danger of redundancy.

The response theorist might object that this so-called alternative will actually be equivalent to the content-based response-dependent account of the discourse in question. One way to see that this is not so is to observe that unlike the response-dependent account, the usage-based alternative is not committed to the a priori impossibility of global error. In principle the impossibility of such error might be accessible from one or both of two perspectives: to practitioners in the discourse in question, to theoreticians reflecting on the nature of the discourse, or to both. However, from the practitioners' point of view there is nothing in the usage-based account that rules out the possibility of global error. The discourse was explicitly designed not to exclude this (while the general possibility was illustrated in §3 by the case of the fallible moral intuitionists). This leaves the theoreticians' point of view, and from here talk of error is simply a category mistake: because a usage-based theory is semantically austere, it tells us not when the utterances concerned are *true*, but how they are properly *used*.¹⁰

This point about theoretical perspective also provides the answer to a response theorist who asks what the truth conditions of the imagined discourse would be—hoping perhaps to show that there is no space for an alternative pragmatism at this point, between response-dependent truth conditions, on the one hand, and realist truth conditions, on the other. The answer is that *as theorists* we simply need not be in the business of providing truth conditions in the first place. *As speakers*, we may say with the crowd that 'X is red' is true if and only if X is red. *As theorists*, we are simply interested in how discourse of this kind is *used*. (True, one important aspect of usage concerns the speakers' use of the terms 'true' and 'false', and this needs to be explained. However, there is no a priori argument in this for the kind of realism which would embarrass a usage pragmatist, at least as long as quasi-realism remains a viable alternative.)

¹⁰ It is becoming a familiar point that semantic minimalism and metaphysical minimalism tend to go hand in hand. Here, semantic austerity accompanies metaphysical austerity. As I noted in §3, the usage-based theory tells us very little about the metaphysical commitments of the community in question. But it is compatible with, and perhaps encourages, the minimalist metaphysical position that Johnston (OR, §4) regards as the most serious challenge to a revisionary response-dependent reading of several philosophically interesting topics. What Johnston fails to see is that in rejecting response dependence the metaphysical minimalist need not reject pragmatism. The usage path to pragmatism remains open.

The alternative is thus a genuine one. But could there be some reason for preferring the response-dependent practice; some reason for revising practice, if it is not already response dependent? A possible suggestion is that only a response-dependent account enables us to *justify* the practice in question.¹¹ This would tie in with the observation just made, namely that only the content-based approach yields talk of truth at the theoretical level. Only this approach provides a framework in which, with luck, it may turn out to be demonstrable that at least some of the claims we take to be true actually are true.

This is a tempting idea, but I think it is mistaken. In my view, response dependence tends to buy justification at the cost of devaluing the practice for which we wanted justification in the first place. This is bought out most starkly in cases in which the commitments typical of the original discourse have certain characteristic consequences for action, so that to justify the commitments concerned would be to justify the actions to which they give rise. The most important examples involve concepts of value and probability. I want to show that the response-dependent substitutes for these notions do not have the proper connections with action, and hence that justifying the substitute judgements does not serve to justify the associated behaviour. As we shall see, the point has been obscured by some ambiguities in the use that response theorists make of the notion of rationality. When these ambiguities are resolved, response dependent accounts of probability and value turn out to do rather worse than their usage-based rivals.

¹¹ The quest for justification or legitimation of linguistic practices is a prominent theme in Johnston's OR. In the opening paragraph he tells us that what 'deserved the name of a progressive pragmatism' would be a critical philosophy which asked 'whether the real explanations of our practices allow us to justify them.' One possible answer to this question would of course be a negative one, in particular the view that it is hardly more appropriate to try to justify our linguistic practices than it is to try to legitimate (say) our digestive practices. In both cases we can describe the practices concerned, and perhaps say what function they serve in our lives, but little more than that. I suspect that Johnston would be unhappy with a pragmatism that yielded no more than this. At any rate, in the next section I shall try to show that someone who does want more won't find it in response dependence. True, they won't find it in a usage-based pragmatism, either. My point is that a content-based theory does not offer any advantage in this respect.

8. JUSTIFYING CHOICES

Moral cognitivists need to explain how moral beliefs can be genuine beliefs, on the one hand, and yet of an essentially motivating character, on the other. There seems to be an analytic connection between believing that X is good and desiring that X. Why should this be so? I call this the *Approval Problem*.¹²

There is an analogous problem on the other side of decision theory. Cognitivists about probability need to explain how beliefs about probability can be genuine beliefs, on the one hand, and yet have some non-contingent connection with particular degrees of belief (and hence with betting behaviour), on the other. Someone who believes that it is probable that P tends to be confident that P, and to act accordingly. Again, why should this be so? I call this the *Confidence Problem*.¹³

One of the great attractions of expressivist views of probability and morality is that they provide a very simple solution to these problems. If moral claims express evaluative attitudes, for example, then the Approval Problem vanishes: moral commitments simply *are* desires, in effect, and we don't need any further explanation of why they tend to be accompanied by desires. These are problems for cognitivists about value and probability, then—problems for people who think that there is a significant psychological distinction to be drawn between a belief about probability and the corresponding degree of confidence, or between a belief about value and the corresponding affective disposition. If there are important psychological distinctions of this kind, it needs to be explained why the mental states so distinguished tend to occur together.

In both cases the problem can be given either a descriptive or a normative flavour. Why *do* people who believe X is good tend to desire that X? And why *should* they do so? In either flavour, however, these problems need to be distinguished from a more general problem of justifying our desires and credences (and hence the behaviour which flows from them). The

¹² I use this terminology in FFT, ch. 4. Michael Smith (DTV, p. 89) calls the same issue 'the Moral Problem'.

¹³ See my 'Does "Probably" Modify Sense?', *Australasian Journal of Philosophy* 61(1983), 396-408, and FFT, ch. 4.

easiest way to see the more general problem is to suppose that we had solved the Approval and Confidence problems, in their normative versions. Suppose, in other words, we could show that our desires and credences were justified in the light of our beliefs about value and probability. It would still make sense, apparently, to wonder whether we had the right beliefs about value and probability, and hence the right desires and credences in this broader sense.

In effect, this broader problem is that of justifying the behavioural dispositions with which, as decision makers, we meet the world. I'll call this the *Behavioural Justification Problem* (or 'Behavioural Problem', for short). In order to solve it, cognitivists about value and probability need not only normative solutions to both the Approval and Confidence problems, but also a means of justifying our evaluative and probabilistic beliefs themselves.

What is involved in justifying our evaluative and probabilistic beliefs? It depends, obviously, on the content of these beliefs. One of the attractions of response dependence seems to be that it provides contents which make the task look manageable. After all, if moral and probabilistic claims are very general claims about human dispositions, then the knowledge required to justify them is a kind of self-knowledge. The subject matter of morality and probability becomes more accessible—as well as more palatable, metaphysically speaking, to those of a naturalistic frame of mind.

If the Approval and Confidence problems could be solved under their normative readings, then, the Behavioural Problem would reduce to that of justifying moral and probabilistic beliefs. And if that problem starts to look approachable from the response theorist's point of view—because these beliefs are about our own psychological dispositions, in some sense—an attractive vista opens up. A solution to the Behavioural Problem might seem to be in sight. I want to show that this prospect is quite illusory, however. In effect, response theorists can have one half of the solution or other, but not both together. Which half they get depend on what precisely goes into their biconditional content conditions, and the illusion rests on equivocation.

Let's think first about the Confidence Problem. I noted that the two-step solution to the Behavioural Problem requires a *normative* solution of the Confidence Problem—in other words, a demonstration that an agent who believes that it is probable that P is *justified* (and

not merely *disposed*) to be confident that P. However, if the response theorist says that ‘It is probable that P’ means that (or is true if and only if) normal people are disposed to be confident that P, then a normative solution to the Confidence Problem will not be forthcoming. After all, put yourself in the position of the agent concerned. You believe that normal people in your circumstances are disposed to be confident that P. But what have the habits of other people (normal or not!) got to do with whether *you* should be confident that P?

The usual approach at this point to put some normativity in by hand, so to speak, in the form of the notion of rationality. Roughly, the response theorist’s suggestion is that ‘It is probable that P’ means (or is true if and only if) a *rational* person would be confident that P. I want to show that this strategy is a dead end, or rather a whole neighbourhood of dead ends, distinguished by a range of possible readings of the notion of rationality. I’ll begin with the probabilistic case, and then apply the same lessons to the evaluative case.

The idea of analysing probability in terms of rational partial belief has long seemed attractive. I call it the rationalist approach to probability. Roughly, it amounts to saying that P is probable if and only if a rational person would be confident that P, if properly acquainted with the relevant evidence.¹⁴ How does a rationalist handle the Confidence Problem? In its normative version, the problem now comes to this: Why should a person who believes that it is *rational* to be confident that P, *actually* be confident that P?

There are a number of possible approaches, turning on different readings of the relevant notion of rationality. It would be almost impossible to be exhaustive at this point, of course, but the following alternatives will serve to illustrate the nature of the problem:

¹⁴ Among recent writers, Hugh Mellor has a theory of chance of this kind, for example (though in his case it is part of a hybrid theory including propensities, thought of as dispositions to produce arrangements of chances). See Mellor’s *The Matter of Chance*, Cambridge University Press, 1971. Mellor takes his theory to be in the spirit of F. P. Ramsey’s view, developed in ‘Truth and Probability’, in *Foundations*, D. H. Mellor (ed.), London: Routledge & Kegan Paul, 1978. Despite the apparent differences between Ramsey and Keynes on this topic, Keynes too might be thought of as an early exponent of the rationalist treatment of probability, given that his ‘logical’ probability relations are characterised in terms of degrees of *rational* belief; see Keynes’s *A Treatise on Probability*, London: Macmillan, 1921.

- I *Rationality cashed in terms of our own cognitive dispositions.* Either (a) the concept of the rational is the concept of what we are disposed to accept, or (b) the term ‘rational’ is used to express a disposition to accept a claim to which it is applied.
- II *Rationality cashed in terms of practical utility.* Rational partial belief is useful partial belief.
- III *Objective rationality.* Rationality is an objective cognitive value, and our practice in the present case reflects our awareness of its presence and nature.

Option I(a) amounts to a response-dependent account of the relevant notion of rationality. Roughly, it says that a cognitive move is rational if and only if it is a move we are disposed to make. To say that it is rational to be confident that P is thus to describe a disposition on our own part to be confident that P. (There are at least two readings of ‘our’ here: it might be the speaker personally, it might be the community.) Option II(b) has a more expressive flavour: to say that it is rational to ϕ is to *express* (not *describe*) a disposition to ϕ . (In this case the bearer of the disposition is presumably just the speaker.) On either version of Option I, however, the idea that the response theorist’s appeal to rationality leads to normative justifications for our behavioural dispositions turns out to be illusory: if rationality itself is cashed in descriptive or expressive terms, so too is whatever notion of justification rides on its back.

Option II may look more promising. It is world-directed, taking rational belief to be belief that works. But what is useful about a particular degree of credence? At best, only that it is (increasingly) *probable* to lead to success in the (increasingly) long run. Taken in this sense, then, the analysis of probability in terms of rationality is circular, at least in the sense that a rationalist account of probability is powerless to help with the Confidence Problem. The problem simply re-emerges at the new level.

Option III is more interesting. It countenances an objective cognitive value, and interprets probability in terms of the value, in this objective sense, of certain degrees of partial belief. Why is it the case that people who believe that it is probable that P do, or should, also have a high degree of confidence that P? Because to say that P is probable is to say that it is valuable to have such a credence, in this objective sense of value. So if people believe that P

is probable then they believe that they *should* have a high credence that P, by their own evaluative lights.

So far so good, but this answer won't satisfy someone who has noticed the Approval Problem. For we now have what might be called the Rationality Problem: What is the connection between rationality and motivation? The question again has a descriptive and a normative aspect: Why does someone who believes that it is rational to F *actually* F (or at least typically have an inclination to F)? And why *should* they? Why does rationality matter, and why should it?

Just as in the case of the Approval Problem, the Rationality Problem might incline us to a response-dependent or expressive account of rationality. In the present context, however, these were the views we considered under option I. So if Option III is to remain distinct we must reject this course—and yet the problem now seems intractable. Still, if one is prepared to live with this problem in the evaluative case then there is no need to bite another bullet for the probabilistic case. Rationalism shows us how to use the same evaluative bullet for both cases.

In sum, then, does a rationalist account of probability succeed in providing a normative solution to the confidence problem? It depends on what we mean by rationality. There seem to be three possibilities, comprising two dead ends and one diversion:

- (i) On a response-dependent or expressivist reading of the rationality claim, the connection between the probabilistic belief and credence is merely descriptive. At best, the disposition expressed or described by a probability judgement simply *gives rise to* the corresponding credence. It doesn't *justify* the adoption of that credence.
- (ii) On a utilitarian reading of the rationality claim the rationalist account of probability is circular.
- (iii) On a normative reading of the rationality claim the Confidence Problem is replaced by a version of the Approval Problem.

Rationalist response theorists thus do rather poorly on the normative aspect of the Confidence Problem. At best they succeed in shifting the issue to the evaluative case.

Perhaps not surprisingly, however, it turns out that much the same difficulties arise in evaluative case (and here, of course, there is nowhere else to hide). To illustrate the nature of

the problem, suppose we follow the response theorist in equating the good, or the valuable, with what substantial rationality leads us to value. In its normative aspect, the Approval Problem is now the question as to why we should actually value what we believe to be valuable in this sense. Why should we value what substantial rationality leads us to value? Once again, the structure of the answer turns on our understanding of the rationality claim. There seem to be three main possibilities, which roughly parallel those in the probabilistic case.

- I* *Response-dependent or expressive rationality.* These approaches cash rationality itself in terms of our own psychological dispositions. As in the probabilistic case, they provide descriptive but not normative solutions to the Approval Problem.
- II* *Rationality as a cognitive value.* This approach interprets rationality in terms of evaluative categories already in play. For example, it takes what it is rational to value to be what it is valuable to value. As in the probability case, the effect of this is to make rationality unhelpful in providing an account of value. The analysis leads in a circle.
- III* *Rationality as a sui generis cognitive value.* If rationality is an independent value then the Rationality Problem is distinct subspecies of the Approval Problem. Indeed, for rationalist response theorists it becomes the core species, to which other versions of the Approval Problem reduce.

Let's think about the normative aspect of the Rationality Problem in the light of III*. Given this *sui generis* notion of rationality, why should someone who believes that it is rational to ϕ , actually ϕ ? One possible move is say that the solution is analytic, because 'should' itself needs to be cashed in terms of this notion of rationality. If rationality itself is the core normative notion, there is no non-trivial normative issue to be raised at this point.

However, recall our reason for being interested in the normative versions of the Confidence, Approval and Rationality problems in the first place: if successfully negotiated, they seemed likely to provide response theorists with the second half of a two-stage solution to the Behavioural Problem. The first stage of the solution required justification of our evaluative and probabilistic beliefs themselves. Here response theorists seemed well placed to make progress, in virtue of the fact that they take such beliefs to refer to our own

psychological dispositions. However, if the notion of rationality invoked in the response theorist's content conditions is normative and irreducible, this easy progress is quite illusory. Justification now depends on access to the dispositions of *rational* agents, not merely those of *actual* agents. And the enquiry itself seems in danger of a problematic regress, if the notion of justification involved is to be cashed in terms of the same notion of rationality.

The underlying moral of this discussion is a simple one, which applies in both the probabilistic and evaluative cases. If the response theorist's content conditions invoke irreducibly normative notions, they leave unresolved all the old difficulties about normativity in a natural world. One issue left unresolved is the epistemological one: How could creatures like us have access to normative facts of the relevant kind? How could we be *justified* in holding particular beliefs about such facts?¹⁵ On the other hand, if the response theorist's content conditions are not irreducibly normative, then the Confidence Problem and the Approval Problem loom large: Why should our behaviour as agents be guided by what we believe about the psychological dispositions of our fellows? Only someone who equivocates about the normativity of the relevant notion of rationality could imagine that a response-dependent account of probability or value escapes both horns of this traditional dilemma.

9. ACTION AS A FIXED POINT

The previous section was motivated by the suggestion that a response theorist might answer the redundancy argument by claiming that a content-based account is better placed than its various rivals to justify our beliefs. In a sense this claim is true, of course. Because the effect of the response theorist's biconditionals is to read all our commitments in the relevant domains as partially self-descriptive, the commitments as interpreted are often easier to justify than they would be under other readings. However, opponents on both sides of the response theorist—more traditional realists on one side, and irrealists of various kinds on the

¹⁵There is a double problem here, of course, because what is being called for is *justification* for our beliefs about matters of *justification*.

other—are likely to feel that there is a trick involved. In effect, the response theory solves the problem by changing the subject.

The nice thing about the evaluative and probabilistic cases is that here the trick, if we may call it that, is blocked by the fact that the psychological states which play the role of the response in these cases—desire and credence—have direct connections with behaviour. This means that they provide psychological fixed points, toward which the quest for justification is ultimately directed. When the response theorist tries to reconstrue evaluative or probabilistic belief, the necessary connections with credence and desire bring the issue of justification back to the same point. These connections are the core of the Confidence and Approval problems, respectively—problems which arise for any view which recognises a gap between believing that X is valuable and desiring X, or between believing that P is probable and being confident that P. As I noted earlier, one of the great attractions of expressivist views is that they minimise these gaps.

In the case of the traditional secondary qualities, however, there are no psychological fixed points of this kind. There is no distinctive behavioural manifestation of ‘seeing red’, for example, such that an account of the content of the belief that X is red has to explain why people who have this belief typically do (or should) display this behaviour. ‘Seeing red’ simply doesn’t play that sort of a role in our mental lives. This might explain why the response-dependence strategy has long seemed so attractive in these cases. When it came to seem that there was no place for the folk properties of colour in a scientific account of the objective contents of the world, the strategy of revising in favour of the Lockean dispositional account did not have to confront anything analogous to the Confidence and Approval problems. So the attractions of the Lockean account—for example, the fact that by sacrificing their observer-independent status, colours retain a respectable place in the scientific world—are not offset by difficulties elsewhere.

With colour properties so construed, of course, the scientific view *can* justify our colour ascriptions. As always, response dependence interprets our commitments in a way which makes them more accessible. Whatever the merits of this move in the colour case, I have argued that it doesn’t work in the probabilistic and evaluative cases. Here, the justification we

want is justification for our credences and desires, and the dispositional revision doesn't give it to us.

As long as the secondary qualities constitute an isolated case, there is an evident attraction in keeping them within the scientific fold by reading them dispositionally. This advantage would be undermined by the acceptance that whatever the nature of the subjectivity in play here, it is a global feature of language. For in this case it infects our scientific discourse as well, and hence cannot be quarantined where it occurs in other discourses by means of the revisionary dispositional reading. True, we may see an advantage in reducing multiple sources of infection to one, taking the view that the scientific discourse is the most hygienic we have. But if it is acknowledged that we live (and have no choice but to live) with the infection in this case, an obvious thought is that we might just as well do so in other cases as well. This would be to defend our right to the old unsterilised naive discourse about colour, relying on our internal systems of cognitive quarantine to avoid serious conflict with our scientific beliefs. This would in a sense be quasi-realism about colour, but a quasi-realism defended by the claim that in virtue of the global nature of the relevant kind of subjectivity, the same quasi-realist attitude is appropriate everywhere. There is no embarrassing divide between 'real' realism and quasi-realism, on this view—a divide which several of the quasi-realist's critics¹⁶ have identified as a weak spot.

10. MAKING SENSE OF USAGE

I have argued that globality is incompatible with a thoroughly content-based approach to explicating the sort of subjectivity with which we are here concerned; that wherever a content-based account may be given, the materials also exist for a usage-based account; and that in the moral and probabilistic cases there is no advantage for the content-based account—quite the contrary, if anything—concerning the justification of the relevant action-guiding psychological states. To finish, I want to exhibit a further advantage of the usage-based approach. Like the redundancy point, this turns on the fact that the usage-based theory is more

¹⁶ Including me; see FFT, ch. 4, 'Metaphysical Pluralism', *Journal of Philosophy*, **89**(1992), 387-409.

austere and less demanding than its rival. Here the effect is not simply that the usage-based theory works wherever the content-based theory claims to work, but a stronger point: in accounting for folk linguistic practice the usage approach often works smoothly where the content-based approach does not. As we shall see, this advantage is most evident in the case of discourses whose patterns of usage exhibit a subtle combination of objective and subjective aspects.

11. TRUTH CONDITIONS: THE OBJECTIVE-SUBJECTIVE DILEMMA

Let's return to the self-descriptivist misreading of emotivism. It has often been noted that self-descriptivism is unable to make sense of ordinary intuitions about moral objectivity, and in particular about the possibility of moral error and moral ignorance. For example, there is nothing particularly counterintuitive about

(8) I believe that X is good but I may be wrong

or about

(9) There are many good things of whose existence I am quite unaware.

Apparently, however, the self-descriptivist is required to parse these propositions as something like

(10) I believe that I approve of X, but it may be the case that I do not approve of X

and

(11) There are many things of which I approve without being aware that they exist,

respectively. But (11) is clearly counterintuitive, and (10) mislocates the source of moral error, taking it to involve a mistake about oneself. (In folk moral practice it simply isn't true that self-knowledge excludes moral error.) In this respect, then, the self-descriptivist's paraphrases do not do justice to a degree of moral objectivity which is claimed by ordinary usage.

The problem becomes even more acute in cases of moral disagreement. Folk intuition regards the utterances ‘X is good’ and ‘X is bad’ as *prima facie* incompatible, even if said by different speakers. In response to the assertion ‘Discipline is good’, for example, the utterance ‘Discipline is bad’ would normally be taken to amount to a clear (if rather stilted) expression of dissent. Yet the self-descriptivist should see no more conflict between these propositions, expressed by different speakers, than between ‘I approve of discipline’ and ‘You disapprove of discipline’.

These are familiar objections to self-descriptivism, of course. They turn on the fact that the self-descriptivist’s truth conditions for moral claims are too subjective to account for the apparent objectivity of ordinary moral discourse, and a promising response is therefore to ‘objectify’ in one or both of two respects: (a) by taking moral claims to refer not to the evaluative dispositions of the individual speaker, but to those *normal* in the speech community as a whole; or (b) by appealing not to *actual* responses, but to possibly counterfactual responses under *ideal* conditions. Taken together, these moves give the self-descriptivist something like this:

- (12) X is good iff A normal person would approve of X, if acquainted with it under ideal conditions.

In other words, the effect of these moves is to bridge the gap between naive self-descriptivism and response-dependence. As response theorists recognise, of course, this opens the way for moral error and ignorance. To the extent that an individual speaker is abnormal, or less than ideally situated, his or her evaluations may be out of step with the objective standard. It also makes possible moral disagreement: two speakers may simply disagree about how the normal speaker would react to some state of affairs under ideal conditions.

So much for capturing the objectivity of moral usage. But moral usage is also notoriously subjective, in the sense that it appears to leave room for speaker-relativity, for no-fault disagreements: cases of moral difference in which usage appears to allow that neither speaker need be at fault. Can the modified self-descriptivist theory make sense of these? On

the face of it not, for surely it is a perfectly objective matter whether a given X is such as to evoke a given response in normal people under specified conditions.

Response theorists have an answer to this challenge, whose adequacy I'll consider in a moment. First note that the problem cannot be avoided by denying moral relativity, or in some other way appealing to the special features of the moral case. For the problem turns on the possibility of no-fault disagreements, and these are not simply a feature of the moral case. Indeed they are potentially global, in virtue of the same rule-following considerations that lead Pettit to conclude that response dependence is global: because the application of any general term is finitely based, it is always conceivable that a speech community will divide on the issue of a future application. If it turns out that such a division underlies an apparent disagreement, that disagreement will be held to involve no fault.¹⁷ In general, wherever language depends on a response which might vary from speaker to speaker, we have a potential source of no-fault disagreements.¹⁸ So our present concern is not peculiar to the moral case.

In so far as it acknowledges the possibility of no-fault disagreement, the solution the response theory proposes is to invoke an indexical specification of the relevant 'response community' —for example, to suggest:

(13) X is good iff I would approve of X, if I were acquainted with X under ideal conditions.

Roughly speaking, the strategy is thus to dispense with (or at least water down) the notion of normality, but to retain a reference to ideal conditions (thereby, hopefully, retaining the ability to deal with the objectivity problem).

The revised account can certainly make some sense of moral error: I might simply be mistaken as to how I would respond under ideal conditions. But can it make sense of moral disagreements? We noted earlier that a naive self-descriptivist should see no more *prima facie*

¹⁷ See FFT, ch. 8.

¹⁸ Are there other potential sources? Yes, in various kinds of context-dependencies. See FFT, ch. 8.

conflict between 'X is good' and 'X is bad', said by different speakers, than between 'I approve of X' and 'He disapproves of X'. Clearly, the same applies to this indexical version of the content-based theory. Normality was thus doing important work, and the problem re-emerges if it is omitted.

Response theorists thus face the following dilemma. If they leave out the notion of normality, the resulting account is insufficiently objective to make sense of ordinary prima facie disagreements about the matters in question. If they put in normality, on the other hand, the account cannot make sense of the very real possibility of no-fault disagreements, arising in cases in which there is statistically significant divergence in patterns of response.

In practice, response theorists may choose a different horn of this dilemma in different cases. They may favour subjectivity where no-fault disagreements are most obvious, such as in the moral case, and objectivity ('normal response') elsewhere. This is mere damage control, however. Both choices are unsatisfactory, and in any case any such division is bound to be arbitrary. No-fault disagreement is a global phenomenon in language, and its incidence varies only by degree from one discourse to another.

Let's focus on a case in which response theorists do invoke normal conditions, and consider the sensitivity this engenders to the issue as to whether there are statistical norms in the community in question in the given respect. Pettit recognises (RRD, p. 608) that it is a consequence of the content-based theory that if there is no normal response, the relevant property ascription should be regarded as false. A related conclusion is that as long as the bounds of the community are indeterminate, so too is the content of the property ascription concerned. But consider the consequences of this in perfectly ordinary cases. Canberra is a small city, with a lot of open space. By Sydney standards, it is a quiet place. But if I were to deny that Canberra is a bustling place, would I speaking to the rather refined bustle-sensitivities of Canberra residents themselves, to those of my own Sydney community, or to some other 'normal' class? A disadvantage of the content-based approach is that until the issue is resolved, the significance of what I have said is simply not determinate.

A great advantage of the usage-based approach is that it avoids these ambiguities *at the level of the meaning specification*. Roughly, Canberra residents and Sydney residents learn to

use the term ‘bustle’ in accordance with the same usage condition: the assertion ‘This place is bustling’ is *prima facie* appropriate when one experiences the ‘Oh, it’s so busy!’ response. This provides the sense in which the two communities speak the same language in using the term. But it leaves room for divergence later, if it turns out that Canberra residents and Sydney residents are differently sensitive to experiencing this response. The response theorist has to try to account for this divergence in terms of a difference in meaning which—perhaps unknown to the participants—has been present in their linguistic practice all along. (The term ‘bustle’ simply latched onto different concepts in Canberra and Sydney.) By specifying meaning in terms of usage, in contrast, my kind of pragmatist is able to explain the fact that in one important sense, the term does mean the same to the members of the two speech communities—and to explain the divergence in its application by differences between the speakers which lie in the background, as preconditions for the practice in question.

There are two main arguments for leaving these differences in the background, rather than incorporating them into the content of the assertions concerned. First, it makes much better sense of ordinary usage. Naive Canberra and Sydney folk need have no sense that their notion of bustle is not universal, and to insist on such gross indeterminacies in content unreasonably detaches meaning from speakers’ understanding. (There is a theoretical redundancy point lurking here. Even if initially we thought there were such content indeterminacies, an obvious and appealing theoretical move would be to factor them out, thus concentrating on what would have better claim to be thought of as speaker-accessible meanings.) Second, and ultimately even more telling, not all the preconditions of language can be made explicit, in the sense of figuring in the content of the linguistic expressions for which they are preconditions. This is a familiar point, related to what the Tortoise said to Achilles. Something must stay in the background, crucial and yet unsaid. Whatever it is, it might vary from speaker to speaker, or from community to community. In this way, we’ll get precisely what we have in the case of ‘bustle’: a divergence in application which rests not on any difference in content, and not on any mistake by either party, but simply on a difference in the background.

Many discourses handle these no-fault disagreements surprisingly smoothly. We are able to ‘switch off’ the objective mode sufficiently to accommodate such irreconcilable differences, without seriously undermining its useful application in ordinary cases. We thus combine objectivity with a certain tolerance of subjectivity (or relativity). (In *Facts and the Function of Truth* I tried to explain this in terms of an account of what in general the objective mode is *for*.) A usage-based account can readily make sense of this feature of language, but a content-based theory is bound to be torn between the objective and subjective aspects of practice.

In particular, a usage-based account puts normality in its proper place. There is a direct comparison here with a coherence theory of truth. Opponents of such a theory can agree, of course, that in the long run what we *take to be true* is what we converge on believing—roughly, what normal people come to believe in the ideal limit of rational discussion. But this is simply a consequence of the more basic usage condition for truth, viz., that it is *prima facie* appropriate to assert that ‘P’ is true only if one believes that P. Against a background of communal reassessment of belief, this basic condition guarantees, more or less, that what normal people *take to be true* in the long run will be what the community converges on. Since this fact about usage can be explained by any theory of truth which accepts the minimal usage condition, it provides no argument for the coherence account’s attempt to explain it analytically, by analysing truth in terms of normal rational convergence. On the other hand, the coherence theory does violence to usage at other points. (With respect to the phenomenology of truth, for example: as correspondence theorists like to point out, what we seem to care about in rational enquiry is not whether we’ll all come to agree that P, but whether P is in fact the case.)

As for the coherence theory, so for the content-based approach in general: the right place for the notions of normality, rationality, and the like is in an account of the consequences, after the fact, of the adoption of a usage rule which makes no mention of these things.

12. PRAGMATICS TO THE RESCUE?

We have been considering the objection that content-based theories do a poor job of accounting for certain features of ordinary usage. In various respects, for example, the ordinary use of the notions of truth and falsity is not what it should be, if the response theorist's biconditionals captured the content of the utterances concerned.

In reply, response theorists may well try to invoke various 'pragmatic' factors about language use, to explain the apparent mismatch. Depending on the version of the response theory on offer, the task may be to explain apparent objectivity or apparent subjectivity. If (13) is the preferred reading of moral judgements, for example, then the problem will be to explain the possibility of moral disagreement: to explain why 'X is good' and 'X is bad' are normally taken to be incompatible, even when uttered by different speakers.¹⁹ A number of suggestions might be made at this point. For example, it might be suggested that the benefits of social cohesion provide a reason for attempting to 'align' the evaluative attitudes of the different members of a speech community, and that this is encouraged by treating such cases as if they embody a genuine disagreement.

There are other possibilities—linguistic pragmatics is fertile territory—but I think they are uniformly misguided. However, the fault lies not with the pragmatic considerations themselves, but with the ailing program they are called on to patch up. For the irony is that the content-based theory has now taken on the quasi-realist's project, but with the added handicap of having to explain not only why usage bestows 'apparent' truth conditions, but also why it does not accord with real truth conditions. Initially it might have seemed that the content-based theory provided an attractive way to by-pass the quasi-realist's rather tedious concern with this business of explaining usage: invoke one's natural conceptual rights, and get on with something more interesting. Tripped up on the awkward tendency of usage to combine objectivity in some respects with subjectivity in others, however, the content-based theory

¹⁹ Alternatively, if the preferred content-based response theory errs towards the objective side, the task will be to account for the apparent subjectivity of no-fault disagreements. One possibility, to which Johnston seems attracted at least as a last resort, is to simply deny that the parties concerned are speaking the same language. See DTV, p. 170.

now finds itself bogged in the same messy territory. And it is burdened, as the quasi-realist is not, with a doctrine about the ‘real’ truth-conditions are of the utterances in question.²⁰

True, a disdain for the painstaking work of pragmatism may not have been the only factor that led the content-based theory to reject the quasi-realist path in the first place. Another motive might have been the feeling that it is the quasi-realist who does violence to ordinary usage by suggesting that folk truth is not univocal—by distinguishing between genuine truth (where realism is appropriate attitude to discourse) and constructed truth (where quasi-realism is as close as we get). I want to emphasise that in this respect the usage-based theory I recommend sides with the content-based theory. It takes the view that there is no sharp divide in language where real truth conditions cease and ersatz ones take over. Rather, throughout language the explanatory task is the same: to account for the fact that dispositions to respond to the world in certain ways come to be expressed in the objective assertoric mode.

Finally, it might be felt that there is a powerful argument for the content-based theory in the intuition that truth conditions play a crucial role in a theory of meaning (or theory of linguistic competence). Certainly this intuition seems to have been influential in underpinning resistance to non-truth-conditional treatments of various topics, such as conditionals. The point might well be linked to the previous one: it will be at least inelegant if a theory of meaning has to offer one sort of account of the meaning those indicative sentences that do have genuine truth conditions, and another sort of account of those that do not.

This suggestion calls for three responses. First, there is another conception of knowledge of meaning, apart from that of knowledge of truth conditions—namely knowledge of (subjective) usage conditions. The usage-based theory is tailor-made for a role in such a theory. Roughly, knowledge of the meaning of a descriptive term will amount to knowledge as to which response that term expresses. Second, we have just seen that the content-based

²⁰ One consequence of this burden will be that the content-based response theorist will be bound to admit that there are some circumstances in a speaker’s own judgements of truth and assertibility may part company: either she grants that P is true but declines to assert that P (or disagrees with the assertion that P); or she asserts that P, but declines to claim that P is true. Moreover, note that the content-based response theorist should really now explain why usage accords with underlying truth conditions when it does do so, as well as why it sometimes parts company. In effect, this means that the content-based response theorist needs a complete pragmatic account of the usage conditions of the utterances in question, not merely a patch for a few exceptional cases.

theory cannot fully account for the ordinary use, and hence the meaning, of the relevant utterances in terms of the truth conditions it provides for them. An appeal to pragmatic features is also necessary. This concession deflates any general appeal to the idea that meaning is simply a matter of truth conditions.²¹ And third, the role of truth in a general theory of meaning is in any case widely acknowledged to be fulfilled by a very thin disquotational notion. Such a truth predicate can easily be added to a language by stipulation, if need be, and is unlikely to be of relevance to the issues at stake between the content-based theory, on the one side, and the usage-based theory and quasi-realism, on the other.²²

13. CONCLUSION

In sum, the usage path to pragmatism has the following four main advantages. It avoids a vicious regress that threatens a global version of the content-based approach. In any local case, it makes more economical use than the content-based approach of the same raw materials—the same facts about shared human responses. It gives a much better account of some of the subtleties of ordinary usage, particularly of the peculiar mix of objectivity and subjectivity we find in discourses prone to no-fault disagreements. And it avoids the Approval Problem and the Confidence Problem, while doing no worse than its rivals at *justifying* our credences and evaluative attitudes.

Let me finish by mentioning what may seem the most unattractive feature of a usage-based pragmatism. Such an approach requires us to distance ourselves from our own linguistic

²¹ Perhaps there is tendency to underestimate the necessary role of a pragmatic element in a theory of meaning, even where truth conditions are well behaved. It is easy to overlook the need for a truth-conditional theory of content to be supplemented by a theory of assertion—a pragmatic theory which relates the truth conditions of an asserted sentence to the *point* of asserting it. In my view this project is of doubtful coherency, since it assumes in effect that the notion of truth is prior to that of assertion (and hence *correct* assertion). But even if it is coherent, the existence of this project inevitably undermines the apparent theoretical advantages of the truth-conditional approach. For one thing, it will presumably be an ingredient of the required pragmatic account that (canonically) it is appropriate to assert that P when one believes that P. This takes us back to the usage-based theory's subjective assertibility conditions, with the difference that we now lack a suitably general account of belief (or judgement). The usage-based theory has this built in: beliefs (or judgements) are just what we come to call our responses of the relevant kinds, when these responses come to be expressed in the objective truth-conditional mode.

²² For more on this point see FFT, ch. 9.

practices, to such an extent that we are able to ask not *what* we are saying—what its content is—but *why* we are saying it; why we use those words and concepts in the first place. Many philosophers seem to be troubled by this detached perspective, feeling that it threatens our right or ability to continue to engage in these practices in a meaningful way. The attempt to regard one's practices 'from the outside' thus engenders something akin to agoraphobia—a fear that one is losing touch with one's values and community. I don't know what to do about this problem, beyond pointing out that the external viewpoint concerned is only a small aspect of the perspective on ourselves that human biology already offers us. It *is* sometimes unsettling to take this detached view of practices so central to our lives. At least since Darwin, however, there has been no honest escape from this kind of problem. We cannot avoid the source of the discomfort, except by self-deception. The best treatment seems to be to try to alleviate the symptoms. Lasting relief requires the ability to be untroubled by a certain cognitive distance between one's theoretical standpoint and one's everyday linguistic activity.²³

In my view, however, the prognosis is favourable. In other areas of life we seem to be able to accommodate the two perspectives quite painlessly. Our appetites are not less urgent, their satisfactions less enjoyable, in the light of what we know about their biological basis. Pain itself is not less unpleasant because we know what it's for. Unless there is something special about our linguistic behaviour, then, these cases suggest that the complaint is a little

²³ Australians might thus be expected to make good usage pragmatists. The ability to be untroubled by distance is a national characteristic, and Australians are used to contemplating matters of interest from great remove. This view from the middle of nowhere is not without its advantages. The advantages are those of critical distance, of seeing the big picture. Often it is important to step backwards, to move off Broadway in order to contemplate its activities objectively. Ignoring the pessimistic thought that it is quite possible to have too much of a good thing, it is thus an engaging speculation that an aptitude for a detached, sceptical and explanatory stance might yet prove a more enduring characteristic of Australian philosophy than the famous local brand of naive realism.

The usual naturalistic explanation of Australian naive realism attributes it to the bright Australian light. As David Armstrong puts it, reality forces itself upon one in these conditions. There are rival hypotheses, of course, including the mischievous suggestion that Australian philosophers only became naive realists when they stopped protecting their heads by wearing hats. But a more plausible hypothesis is that naive realism is itself a response to isolation—a straw-clutching attempt to steady oneself against the vertigo induced by distance.

These chauvinistic comments are prompted by a footnote in early versions of Mark Johnston's 'Objectivity Refigured' (sadly, it is missing from the published version) in which Johnston notes the possible emergence of a distinctively Australian pragmatism. As I have indicated, this worthy project seems to rest on sound geographically foundations. All the same, it may be that those Australian pragmatists whose hearts are too close to Broadway have yet to achieve the cognitive distance required to cast such a pragmatism in its most desirable form.

theatrical. Ordinary humans find it relatively easy to combine the internal and external perspectives in one satisfying cognitive life.