

PRACTICAL BELIEF AND PHILOSOPHICAL THEORY

Philip Pettit

Philosophy invariably starts with the attempt to spell out ideas and beliefs that we already hold, whether on topics like time or causality, colour or value, consciousness or free will, democracy or justice or freedom. It may go well beyond such pre-philosophical assumptions in its further developments, regimenting them in unexpected ways, revising them on novel lines, even discarding them entirely in favour of other views. But philosophy always begins with the articulation of ordinary ideas and beliefs. This is where its ladder starts.

As the name suggests, articulation requires that the ideas and beliefs articulated were already there waiting to be spelled out, and that identifying them is a matter of analysing what is there, not a matter of adding something new. If philosophical articulation is to deserve its name, and if at the same time it is to be a significant enterprise, then the contents that it articulates—the contents of our pre-philosophical beliefs—must satisfy two constraints. First, they must be sufficiently familiar to us pre-philosophical believers for articulation to be guidable by our common understanding of what we believe; otherwise it would not be, properly speaking, articulation. But, second, the contents must remain unfamiliar enough for articulation to be able to increase that prior understanding; otherwise it would not be a significant enterprise. But how can any pre-philosophical beliefs simultaneously meet these two constraints? That is the problem that I want to address in this paper.

The problem is associated with the paradox of analysis (Ackerman [1]). According to that paradox, analysis or articulation cannot be analytically responsible to our prior understanding of what we believe and capable at the same time of substantively advancing that understanding. The usual approach to the paradox of analysis concentrates on the notion of analysis or articulation, arguing that analysis can be so understood that no problem arises (Mackie [27, Chapter 1]). In this paper I approach the paradox from a different angle. I concentrate on the notion of belief—this is already implicit in my formulation of the problem—arguing that belief can be so understood that the paradox is avoided.

The paper is in three sections. In the first, I discuss familiar behavioural and judgmental modes of believing things and show how they fall foul of the paradox. In the second, I introduce a third, distinct mode which I describe as practical belief. And then in the third section, I show that the contents of practical beliefs satisfy both the familiarity and unfamiliarity constraints and point us towards a coherent and attractive account of philosophical articulation.

I. Beyond behavioural and judgmental belief

The Behavioural Conception of Belief

One philosophically well-established notion of belief begins from the thought that to believe something is to be disposed to behave as if it were true. It is to be disposed to behave

in such a way that, other things being equal, the behaviour maximises the expected satisfaction of the agent's desires. Suppose that an agent believes that *p*. Under this account of belief, that means that the agent's behaviour is of a sort that would maximise the expected satisfaction of his or her desires, other things being equal, in a *p*-world. Believing that *p* means locating yourself in a *p*-world: being disposed to take the sorts of behavioural avenues to desire-satisfaction that are appropriate, other things being equal, in a *p*-world. Believing that *p* means that so far as your behavioural dispositions go, you treat the actual world as a *p*-world. For short, you act or behave as if *p*.

This behavioural notion of belief can be made more sophisticated, and more attractive, in a number of ways. One nice line of improvement is introduced by Robert Stalnaker [39]. Instead of applying the approach to one belief after another—and having to stipulate in each case, as we just did, that other things are equal—we can use it to identify the content of an agent's total belief-state. We can say that taken as a whole, the content of the agent's beliefs is given by his or her behaviour and dispositions to behaviour. We can identify that content, without qualification, by the set of possible worlds in which the agent's behaviour maximises expected desire-satisfaction: that is the set of worlds in which the agent locates himself or herself, acting as if the actual world belonged to the set. When it comes to whether the person believes that *p* or that *q*, then, this version of the approach gives us a ready answer. The agent will believe that *p* if and only if it is true that *p* in that set of worlds: if and only if it is true that *p* in what we may describe as the agent's belief-worlds.

The behavioural characterisation of belief needs to be qualified to take explicit account of the fact that beliefs come in degrees of confidence: I may locate myself in a certain set of worlds, while being more inclined to treat myself as situated in one subset rather than another. The characterisation also needs to be amended to cope with indexical belief; in believing that I am a philosopher, I do not just locate myself in those worlds where P.P. is a philosopher: I identify myself in those worlds with P.P. (Lewis [25, Essay 10]). Again, the account needs to be softened so that it is not too demanding for ordinary human beings: it should probably identify an agent's belief-worlds, for example, only on the basis of behaviour, actual or counterfactual, that is not bedevilled by failures like inattention or illogic or weakness of will; and in view of pervasive inconsistencies, it may have to allow that agents sometimes switch between different belief-systems, locating themselves at different times in different sets of belief-worlds (Stalnaker [39]). Finally, the characterisation almost certainly needs to be strengthened so as to require that the agent's behavioural dispositions are organised (Block [2], Jackson [20]), and perhaps also occasioned (Jackson and Pettit [23]), in a certain manner.

Abstracting from the amendments necessary, however, the question with which I want to start is whether this behavioural account of belief would enable us to make sense of the project of philosophical articulation. Does it give us a conception of pre-philosophical beliefs under which their contents are familiar enough for articulation of content to be analytically faithful to our common understanding? And does it give us a conception under which those contents remain sufficiently unfamiliar for articulation, at the same time, to be substantively informative: to increase our prior understanding?

The answer to that question, uncontroversially, is that this account fails to satisfy the first, familiarity constraint. The problem is that, strictly speaking (Stalnaker [39]), the behavioural account means that everyone believes all necessary truths and that everyone

believes all the entailments of anything that they already believe. Everyone behaves as if the necessary truths are true, after all, and everyone behaves as if the things entailed by what they believe are true. The necessary truths must hold in their belief-worlds, since they hold in all worlds. And the things entailed by anything they believe must hold in their belief-worlds for, being entailed by things that already hold in those worlds, they are true in a set that includes or is co-extensive with their belief-worlds.

This problem means that whenever we identify necessary truths, say in the *a priori* exercise of mathematics, we discover things that every agent believes; we articulate the contents of their behavioural beliefs. And it means that whenever we put our finger on something that is entailed—even entailed in a far from obvious manner—by propositions that an agent already believes in the behavioural sense, we discover something that that agent also believes in that sense. But if the contents of our beliefs are articulated in such exercises, then this conception does not give us a notion of contents that allows us think of their articulation as analytically responsible: as answerable to our prior understanding. For our prior understanding may be entirely silent about such contents. Behavioural belief contents are such that while identifying them may increase our understanding, it cannot count as an analysis of what we already understand.

Looking further

I am going to assume that any serious conception of belief always has a behavioural aspect: that if someone believes that *p*, in no matter how rich a sense, then they behaviourally believe that *p*. They locate themselves behaviourally in *p*-worlds, at least with a certain relatively high level of confidence, at least when they are not subject to various cognitive failures, and so on: I shall take it that such qualifying conditions can be satisfactorily elaborated and are indeed fulfilled. This assumption means that other conceptions of belief all represent specifications of the behavioural conception: they are all behavioural-plus ways of believing things.

My reason for making this assumption is that we would rightly hesitate to ascribe someone the belief that *p*, no matter what other conditions they satisfied, if they were not disposed to act as if *p*. We might say that they believed for certain purposes that *p*—say, for all practical purposes—or that they accepted that *p* or something of that kind; but we would not happily say that they believed, period, that *p*. On the approach taken here, then, the first mark of believing something—if you like, the criterion of belief—is that you are disposed to act appropriately. Behavioural-plus modes of belief satisfy this mark and will be distinguished from mere behavioural belief by the fact of involving something else as well.¹

The something else that will be involved, if it is to be of help to us, must be an extra, not merely behavioural relationship with the contents believed. The fact that the contents of mere behavioural beliefs do not satisfy the familiarity constraint—the fact that they are just too unfamiliar to be retrievable in anything deserving to be called articulation means that in order to identify a more useful conception of belief, we must suppose that believers can have a relationship with the contents of some behavioural beliefs over and

¹ This is not an uncontentious line to take, of course. It means that we shall have to say, for example, that someone who claims to believe something that is necessarily false—someone who assents to a necessarily false assertion—cannot really believe it; they can only believe that the words accepted express a truth.

beyond that implicit in acting as if those contents obtained. Behavioural-plus modes of believing that *p* will involve a person in acting as if *p*—I need not continue to mention qualifications—but also in having a further, independent relationship to that content. The idea is that when this further relationship obtains, then the contents of the beliefs will not outrun prior understanding in the manner of hitherto unrecognised necessary truths and hitherto unrecognised entailments of things already believed.

This observation suggests that we may be able to help ourselves by introducing a distinction from cognitive science. Suppose we distinguish between those behavioural beliefs whose contents are explicitly represented in the agent's brain, as cognitive scientists tend to say, and those that are not. That distinction gives us a behavioural-plus conception of belief. We can say that a person believes that *p* in this way if and only if they have suitable behavioural dispositions and a suitable internal representation that *p*. In order to count as suitable, the representation may be required to satisfy any of a variety of now familiar stories. It may be required to co-vary in a certain nomologically basic way with the situations that make it the case that *p* (Fodor [16]); or to have been recruited in the course of the agent's training for the production of the behaviour that is appropriate in *p*-situations (Dretske [13]); or to be designed under some evolutionary specification for the production of such behaviour (Millikan [28], Papineau [32]); or whatever.

But though the distinction from cognitive science would give us a behavioural-plus sense of belief, it does not give us a sense of belief that will help us with the paradox of analysis. Our prior understanding of what we hold does not tell us which behavioural beliefs are representational and which are not: the distinction belongs to scientific conjecture, not common sense. And how can prior understanding guide us in articulating the content of representational beliefs, if it cannot even distinguish those beliefs from others? This is not to say that there is no distinction between representational and non-representational belief. It is just to say that representational belief is not the sort of behavioural-plus belief which is going to help us with our problem. We need to be able to identify a sort of behavioural-plus belief such that the extra, non-behavioural relationship to content that it involves is available to prior understanding, not hidden in realms available only to science.

The Judgmental Conception of Belief

There is, as it happens, an obvious candidate for the role (c.f. Braithwaite [3]). For the most familiar way of believing things is not only to act as if *p* but also to be disposed—disposed in sincerity and with comprehension—to judge that *p*.² Judging that *p* will involve assenting to 'p' or to some sentence that formulates the same content.

² Because of the presupposition of behavioural belief, believing that *p* in this judgmental sense is distinct from the notion of just accepting that *p*, as that is explicated. For example, by L.J. Cohen [7, p.4]: "to accept that *p* is to have or to adopt a policy of deeming, positing, or postulating that *p*...whether or not one feels it to be true that *p*". See too Ullman-Margalit [41] and Bratman [5] and, for a very useful overview, Engel [14]. The distinction between behavioural and judgmental belief is probably closer to the distinction between 'opinion' as distinct from 'belief' in the sense of Ronald de Sousa [11, pp.54–56] and Dan Dennett [12, Chapter 16], or to the distinction between 'avowed' as distinct from 'central' belief in the sense of Georges Rey [37]. These are both distinctions deriving from the different roles played by the types of states in question: on the one hand, roughly, a role in producing behavioural and related adjustments, on the other a role in shaping the avowals of the agent.

Usually the formulation will be a sentence of a public language but we may extend the conception—the judgmental conception, as I will call it—to those who think that there are other formulae to which we give our assent, as when we say and accept things ‘in our heart’ (Geach [17]).³ The requirement that a judgmental believer must have comprehension of that which commands their assent may be interpreted more or less strictly; at its loosest, it may license ascriptions of judgmental belief in matters—say, in scientific matters—that people only understand in the second-hand sense of knowing who can explain the matters involved. In what follows I shall assume that the requirement is interpreted fairly strictly, so that what we will have in mind are always judgmental beliefs in the presence of first-hand—though not perhaps first-rate—understanding.

I said that to judge that p is to assent to ‘ p ’ or to a sentence that formulates the same content. But what is it for a sentence to formulate the same content as ‘ p ’? Should we require that any sentence which formulates the same content as ‘ p ’ is metaphysically equivalent to it, having the same truth-conditions? Or should we require that it is metaphysically equivalent as an *a priori* matter? Or should we require, more strongly still, that the sentence also means the same thing by some intuitive criterion of meaning?

No approach that stops short of the most demanding requirement can be satisfactory for our purposes. Suppose that ‘ p ’ is the sentence ‘ $2 + 2 = 4$ ’. This is metaphysically equivalent to every other necessarily true sentence; and by most accounts it is metaphysically equivalent, as an *a priori* matter, to any other familiar truth of arithmetic: say, to ‘ $3 + 5 = 8$ ’. If we say that someone judgmentally believes that p just in case they behave as if p and are disposed to judge that p , then we will have to say under the less demanding accounts that anyone who is disposed to assent to such a distinct sentence judgmentally believes that $2 + 2 = 4$. There will be no problem of suitable behaviour, as we know, with a necessarily true sentence. And anyone who is disposed to assent to such a truth will be disposed to assent to a sentence that, by the less demanding accounts, formulates the same content: the sentence will be true, as an *a priori* matter, in just those circumstances where it is true that $2 + 2 = 4$.

We have to invoke a more demanding account, then, of what is involved when one sentence formulates the same content as another. We have to require that the sentence which formulates the same content as another means the same—if you prefer, is used to say the same (Davidson [10, Essay 7])—as the other. Not only must the sentences have the same truth-conditions, and this as an *a priori* matter; it must also be the case, as we might put it, that the truth-conditions have the same profile for the speakers: the conditions must appear, in a Fregean phrase, under the same mode of presentation.

We can now return to the judgmental conception of belief. Under this conception someone believes that p just to the extent that they act as if p and they are disposed to judge that p , where judging that p means assenting to ‘ p ’ or to a synonymous sentence. Or at least that is so when familiar qualifications are in place, such as that the person is

³ I should mention that what I have in mind here is quite different from what is described as the language of thought (Fodor [15]). The language of thought hypothesis is that every believer’s brain instantiates representations, with a language-like syntax, of the things the person believes. But there is no suggestion that the believer assents to such representations in the way we assent to sentences in a public language. The representations are representations *in* the believer, not representations *for* the believer (Cummins [8, p.76]). They do not enable the believer to represent the contents to themselves—to self-represent the contents, as we might say—and to give or withhold their assent.

sincere and is not subject to inattention or illogic or a misapprehension as to what is said in the sentence to which they assent.

Suppose we go along with this judgmental conception of belief, setting aside the various qualifications that may be desirable. Will that conception of belief offer us a way beyond the paradox of analysis? Will it give us a conception of belief such that when we articulate the content of a person's belief, we can claim to analyse the prior understanding of that person and at the same time to increase their level of understanding; we can claim to be at once analytically responsible and substantively informative?

I do not think so. As the behavioural conception of belief fails the familiarity constraint, this conception of belief is going to fail the constraint of unfamiliarity. If we are to articulate the content of someone's belief under this conception, we have to produce a sentence that means or says the same thing as some sentence to which they are already disposed to assent. And that requirement means that nothing we produce, nothing we do in the way of articulating the content of the belief, can really surprise the subject and increase their level of understanding. Or at least nothing we do can increase their level of substantive as distinct from linguistic understanding. We can introduce a sentence from another language, even a language unknown to the believer, in articulation of the belief but we cannot introduce a sentence that will give the subject a greater insight into the content of what they believe.

II. The Practical Mode of Belief

We cannot go for the behavioural conception of what it is to believe something—or for any conception in the same family—for then the articulation of belief-content is not going to be analytically responsible to the believer's prior understanding. But neither can we go for the judgmental conception of what it is to believe something—or for a conception in this family—for in that case the articulation of belief-content is not going to increase the believer's prior understanding; it is not going to be substantively informative. I turn now to a third conception of belief that proves, as I shall be arguing in the next section, to be more satisfactory. This is what I call the practical conception of belief.

Where the judgmental conception of belief starts from a behavioural base and takes account also of the subject's dispositions to assent to various sentences, the practical conception starts from the behavioural base and takes account also of something else but, in this case, something weaker. It does not require that someone who believes that *p* is disposed to judge that *p*. What it requires, more weakly, is that someone who believes that *p* is disposed to reason on a *p*-pattern. The person who believes that *p* in the practical sense acts appropriately and is disposed to reason on an appropriate pattern; or at least this is so when qualifications of the usual sort are in place.⁴

In order to elaborate this conception, we need to say more about what it means to reason. Reasoning is an activity in which we scrutinise and perhaps search for considerations that are relevant to whether we should judgmentally believe something: judgmentally believe that such and such is the case, for example, or that we should do so and so. We conduct this scrutiny and search under the assumption that some

⁴ There will undoubtedly be problems as to how determinacy of content is ensured for practical belief but since such problems arise also for behavioural and judgmental belief—and indeed representational belief—I do not give them special significance.

considerations may turn out to favour the judgmental belief, others to count against it, others to be neutral, and that depending on how they look overall—depending on whether or not they offer support—we will be moved to embrace or reject the belief. We assume, on the one hand, that we can recognise the relevance of the considerations: we can see when a piece of perceptual evidence confirms or infirms a belief, for example, and we can see when something already believed entails a given belief, rules it entirely of court or connects with it in some more graduated fashion. And we assume, on the other hand, that as we recognise the reality and the relevance of certain considerations, our belief-forming capacities will respond appropriately: absent local or temporary perturbations, we will tend to form our judgmental beliefs in a manner that those considerations mandate or allow.

Reasoning in this sense is an activity that presupposes the capacity for judgement. We reason to conclusions that we affirm in judgement and, so we may assume, believe in the judgmental way. And we reason either from considerations that we affirm in judgement and believe in the judgmental way or from considerations that are perceptually presented as matters for such belief. Reasoning leads us to treat the contents of certain judgmental beliefs or their perceptual counterparts as reasons for forming certain other judgmental beliefs. It culminates in the recognition that the input contents serve in some measure to justify the output contents and in the response to that recognition that consists in forming the output beliefs.

Reasoning in this sense is an intentional activity, it should be noticed, and involves more than what is often described as reasoning or inference in cognitive science. Under the cognitive-science way of thinking, a subject will infer that *p* if and only if certain other beliefs lead to the formation of that belief—the beliefs will all have to be internally represented, in the sense explained earlier—and do so in virtue of a causal relationship that mirrors a relationship of rational justification (Cummins [8, Chapter 3]). This conception does not require that the subject reasons its way from premises to conclusion but, rather, as we might say, that the system simulates the reasoning function; the representations succeed one another as if there were an invisible reasoning process in control (Cummins [9]). If reasoning in our sense is effective then under plausible assumptions there is bound to be a process akin to that which is envisaged in cognitive science: something of the kind will be needed to explain how the reasoning is effective in changing the subject's beliefs. But there is certainly more to reasoning than just the realisation of such a non-intentional and perhaps unconscious sequence of changes.

If this point is not obvious, that may be because of a failure to distinguish between rational subjects on the one hand and reasoning or ratiocinative subjects on the other. Systems that are equipped in the manner usually envisaged in cognitive science will be more or less rational belief-formers; within certain limits, for example, they will tend to form beliefs in propositions that are supported by things that they already believe. But this rationality may be attained by virtue of good design only and without the systems in question having any desire to be rational. Subjects who reason are not just disposed to be more or less rational in the way they develop their beliefs. They desire to develop their beliefs rationally and they are disposed to take certain intentional steps in order to satisfy this desire (Pettit [35, Chapter 2]).

I am going to assume that we human beings are reasoning or ratiocinative subjects, not just systems whose beliefs evolve in a more or less rational way. We generally form

our beliefs more or less unconsciously, without any reasoning, but we are disposed to converse with one another, even with ourselves, and are disposed in particular to reason about many of the things that we believe or should believe. We are capable of reasoning and are committed to reasoning; if we do not always reason in the active sense, we remain essentially reasonable creatures (Pettit and Smith [36]).

Back, finally, to the practical conception of belief. We said that to believe that *p* in the practical sense will be—at least under suitable qualifications—to act as if *p* and to be disposed to reason on a *p*-pattern. What will it be, then, to reason on a *p*-pattern? It will be to recognise and practise instances of reasoning that are unified and legitimated by the fact, if it is a fact, that *p*: to treat as reasons for forming certain judgmental beliefs considerations whose status as reasons for forming those beliefs is explained and justified by the fact, if it is a fact, that *p*. The considerations in question may be judgmental, as when they are things we already believe judgmentally, or they may be perceptual, as when they are things that are presented in perception as commanding judgmental belief. To reason on a *p*-pattern will be to take those considerations as reasons for forming the judgmental beliefs that they support. It will be to recognise the relevance of the considerations—say, under conversational challenge—and to respond to that recognition by forming the supported beliefs. Or at least it will mean this when suitable qualifications are in place.

Under the approach taken in this paper, believing that *p* always means being disposed to act as if *p* and behavioural-plus varieties of believing each introduce a distinctive, extra relationship to the content believed. Practical belief, as just explicated, introduces a tie that is weaker than that of being disposed to assent to ‘*p*’ or to some synonymous sentence. The subject recognises instances of *p*-reasoning and in each case reasons as the fact that *p*, supposing it to be a fact, entitles them to reason. But yet the subject need not have, or at least need not see that they have, words in which to express that content. The subject has dispositions such that were they to understand ‘*p*’, then those dispositions alone—those dispositions in the absence of anything they may learn in coming to understand ‘*p*’—would compel them rationally to assent. They don’t have a disposition to assent to ‘*p*’ in the ready sense associated with judgmental belief. They have what we might describe as a depth-disposition to assent to the sentence.

It used to be common to distinguish between believing a generalisation *in sensu composito* and believing it *in sensu diviso* (see Lewis 24, p 67]). I believe in the first, composite or unified sense that all *F*s are *G*—I believe it judgmentally—only if I am disposed to assent to the universal sentence ‘All *F*s are *G*’. I believe in the second, divided or distributed sense that all *F*s are *G* only when I meet a weaker condition alone: I am universally disposed, on recognising something as *F*, to assent to the particular claim ‘That is a *G*’. What happens with the practical belief that *p* is something analogous. The subject is universally disposed, on recognising an instance of *p*-reasoning, and on assenting to the judgmental or perceptual premises, to assent to the conclusion. Or, if you prefer, the subject is universally disposed, on recognising an instance of *p*-reasoning, to assent to the claim that the premises support the conclusion. But the subject is not disposed to assent to the universal principle ‘*p*’: to the principle according to which the premises in all instances of *p*-reasoning—however they are described—support the corresponding conclusion. The subject believes the principle *in sensu diviso* but not *in sensu composito*.

The practical conception of belief is best understood from examples. Take the principle of *modus ponens*. Logicians will believe this principle in a judgmental way. They

will endorse the propositional schema that if it is the case that if p then q , and it is the case that p , then it is the case that q . Or they will assent to the claim that if a conditional proposition holds, and also the antecedent of that proposition, then the consequent holds too; or to the claim, equivalently, that the joint truth of a conditional and its antecedent entails the truth of the consequent. But it is possible to believe this principle without understanding such a sentence and, *a fortiori*, without being disposed to assent to it. And yet the mode of belief in question need not be as sparse as the purely behavioural mode; it may belong to the behavioural-plus family.

The principle of *modus ponens* is a necessary truth and holds at all possible worlds. Thus it must hold at all of the behavioural belief-worlds of an agent; it is something that everyone believes in the behavioural sense of belief. But consistently with believing it behaviourally, and without yet believing it in the judgmental sense, a subject may or may not be disposed to reason on a *modus ponens* pattern; indeed a subject may or may not be disposed to reason at all. To be disposed to reason on the relevant pattern will involve the disposition to justify the judgmental belief that q by reference to the assumed fact that p and that if p , then q ; to balk at the judgmental claim that p in virtue of the assumed facts that if p , then q and that not q ; and the like.⁵ It will be to carry on in reasoning—say in conversation with others, or indeed and with oneself—as if it were granted on all sides that the principle is common, unquestioned ground.

When someone is disposed to behave and to reason on a *modus ponens* pattern then they will believe that principle in the practical sense of belief. They will believe it indeed in the sense in which most of us believed it before we learned to do logic. They will believe it in such a way that on having the principle spelled out—on being introduced, for example, to the word ‘entailment’—they will recognise it as something of which they must say, not just that they embrace it, but that they always embraced it; endorsing certain types of reasoning, they always endorsed the principle that made them worthy of endorsement. Were someone to understand a sentence that formulates the principle, then their practices of reasoning would rationally commit them to assenting to that sentence. There would be an incoherence involved in holding by those practices while dissenting from the sentence.

In the *modus ponens* case the proposition that is believed is an inference ticket which justifies believing the conclusion in virtue of a more or less structural feature of the premises. But it is important to recognise that a second sort of example is possible too. In this type of case the inference ticket represented by the proposition that is practically believed is substantive rather than structural in character. It bears on the conditions that fix the meaning or the reference of a particular word.

Consider the way we generally reason in regard to questions of colour, say the question as to whether something is red. We argue that something cannot be red because it does not look red, thereby displaying the practical belief, so it appears, that something is red only if it looks red. But we are persuaded that perhaps it is red by the fact that the lighting is unusual and that even ripe tomatoes don’t look red in that lighting. Thus it appears that what we practically believe is, more precisely, that something is red only if it looks red in the sort of lighting where things about whose redness we are agreed also look red. We now change the lighting, so that like ripe tomatoes the object does indeed look red, and we conclude that it must indeed be red. And in doing so we show that the

⁵ I understand *modus ponens* reasoning in the broad sense in which it also includes reasoning according to *modus tollens*.

proposition which we practically believe does not say just that something is red only if it looks red in suitable lighting but also that it is red if it looks red in that lighting.

And so on. It is clear that by reflecting on our patterns of reasoning in this sort of case, as well as in the earlier case, we can bring to light things that we believe in a practical way but not necessarily in a judgmental manner. And as someone who practically believes *modus ponens* may not be in a position to assent to a suitable sentence, so the person who reasons on the assumption that something is red if and only if it looks red in suitable lighting may not be in a position to formulate and assent to that claim. They may believe it practically but not judgmentally.

This point can be made more vivid if we continue the exploration of the colour case a little further (see Pettit [34; 35, 199]). Suppose that the way the person reasons shows, not just that lighting matters, but also that the competence of the observer is relevant: some observers are discounted on the grounds of being colour-blind. And suppose, more generally, that the person's pattern of reasoning shows that they trust colour appearances if and only the conditions and the observers are not of a kind with those that they discount in the course of resolving various problems: they are not like those sorts—those inappropriate sorts—as they will think of the matter. One way of expressing the principle of reasoning that they follow in that case will be to say that they reason on the assumption that something is red if and only if it looks red to normal observers in normal conditions, where the term 'normal' is defined by reference to what they see—what they fallibly see, at any point—as fit not to be discounted.

It is obvious that someone may practically believe such a proposition without judgmentally believing it, if only because someone may practically believe it without understanding any cognate of the word 'normal' that we introduced in order to explicate it. As in the earlier example, the person is in a position where their reasoning practices would compel them, did they understand the word 'normal'—and understand the other words in question—to believe judgmentally that something is red if and only if it looks red to normal observers in normal conditions. But the person is not in a position where they are currently disposed to assent to such a sentence. They believe the content in question in a practical but not in a judgmental way.

The categories of behavioural and judgmental belief are well-established in the literature but there ought to be no hesitation about acknowledging practical belief alongside them. One reason why there should be no hesitation goes back to the lesson taught by Lewis Carroll in his piece on Achilles and the Tortoise.⁶ Achilles believes he can force the Tortoise to admit that Z, given that he admits two premises, A and B, which logically imply Z. But the Tortoise insists that he needs more than A and B as premises; he also needs C: if A and B are true, Z must be true. Very well, says Achilles: you admit A and B and C, so now you must admit Z. No, replies the Tortoise, for I must also endorse another premise, D: if A and B and C are true, Z must be true. The lesson is that in order to reason from certain judgmentally believed premises to a judgmentally believed conclusion, you have to believe in the merely practical way that such premises support such a conclusion; or if you take that as itself a judgmentally believed premise, then you have to believe in the purely practical way that the new premises support the conclusion; and so on.

⁶ My thanks to David Lewis, who drew my attention to the relevance of the Carroll piece.

What Lewis Carroll's tale shows is that no matter how hard we try to spell out practical beliefs in judgmentally accepted premises, we can never dispense with the practical dimension completely; judgmental beliefs always need purely practical presuppositions in order to compel assent to conclusions that they inferentially support. But it is an essential feature of judgmental beliefs that they figure in reasoned inference and so it is an essential feature of such beliefs that they are accompanied by a practical acceptance of corresponding principles of reasoning. Judgmentally believing that something is red, for example, involves being disposed to draw suitable conclusions about how it will look, about what is happening if it does not look right, and so on; someone who assented sincerely to the sentence but was not so disposed would not count as understanding the sentence accepted. The judgmental belief that something is red, then, is going to presuppose the practical belief that it is such as to look red to normal observers in normal conditions. And if that practical belief assumes a judgmental form as well, and is used in the derivation of suitable conclusions, it is going to presuppose in turn the sorts of practical beliefs that govern the words that it involves. And so on.

In conclusion, an important query. Is practical belief really a behavioural-plus mode of believing: a mode in which the believer acts as if *p* and also enjoys an independent relationship to the *p*-content? Someone might think not, on the grounds that if we expand the space of behaviour to include acts of assent, then practical belief falls out as just a kind of behavioural belief: it means acting and judging as if *p*.

The *modus ponens* example enables us to see why this thought is mistaken and why practical belief is genuinely of the behavioural-plus variety. By the account given earlier an agent will act and judge as if *p* for any '*p*' that expresses a necessary truth: say for any principle of inference that is true at all possible worlds, as *modus ponens* is true at all possible worlds. But when we required the practical believer in *modus ponens* to be disposed to judge on a suitable pattern we required more than judging as if *p*. We required that the subject should be able to identify cases of *modus ponens*—in whatever terms—and be disposed to treat the relevant considerations as reasons for drawing the appropriate conclusion. To be disposed to reason on a *p*-pattern, then, is to have a richer sort of relationship to the content in question than that involved in behaving as if *p*, even where behaviour is taken to include judging.

We see here the main, perhaps the only, new claim that I want to defend. There is a way of believing certain propositions that is neither as open as behavioural belief—even when assent is treated as a form of behaviour—nor as closed as judgmental. Unlike judgmental belief, it does not require the ability to put the content in words. Unlike behavioural belief, it requires more than being disposed to act as if the content were true. What it distinctively requires is a disposition to treat certain sorts of considerations in the same way; specifically, to treat them in reasoning as considerations that support corresponding conclusions.

III. Articulating practical contents

Beyond the paradox of analysis

We argued in the first section that the behavioural and judgmental conceptions of belief both fail to show us a way beyond the problem raised by the paradox of analysis. The

behavioural conception gives us a notion of belief under which the articulation of content cannot be analytically responsible: we believe things that entirely outrun our recognitional capacities. The judgmental conception gives us a notion of belief under which the articulation of content cannot be substantively informative: we only believe things for which we have more or less exact expressions. What I now wish to show is that the practical conception of belief that I introduced as a third alternative enables us to get around this problem. It explains how belief content can satisfy the familiarity constraint, allowing articulation to be analytically responsible, and how at the same time it can meet the unfamiliarity constraint, allowing articulation to be substantively informative.

The difference between practical and behavioural belief explains why the content of practical belief, unlike that of behavioural, satisfies the familiarity constraint and why the articulation of such content can be answerable to prior understanding. That we believe behaviourally that *p* is apparent only in the fact that our behaviour advances the satisfaction of our desires in *p*-worlds and it is quite possible for us to believe in this sense that *p* without that belief showing up in any accessible aspects of how we think. That we believe practically that *p* appears in the fact that we treat certain judgmental or perceptual considerations as reasons for forming corresponding judgmental beliefs and we can reflect on our reasoning dispositions, and on the principles under which they are justified, to check on whether indeed we are committed to the practical belief that *p*.

How to reflect on our reasoning dispositions with a view to discovering whether we are committed by them to one or another principle: to *modus ponens*, for example, or to our redness principle? The obvious approach will be to follow the method of cases (Jackson [21]). This is to consider those things which our reasoning dispositions would lead us to say in various circumstances, actual and imagined, and to identify in such thought-experiments the justifying principles to which the dispositions commit us. The idea is not to look inwards, as if introspection could deliver up our practical commitments; the idea, rather, is to look outwards in order to identify our practical beliefs by reference to our dispositions to make this or that judgement in that or that situation. How can I know what I think until I see what I say? How indeed? The lesson may not be fitted to judgmental belief but it is certainly suited to practical.

Where the difference between practical belief and behavioural belief shows how the contents of practical beliefs can satisfy the familiarity constraint, the difference between practical belief and judgmental belief shows how they can satisfy the unfamiliarity constraint. In order to articulate the content of what someone judgmentally believes, by the account offered earlier, we have to be able to present a sentence that means or says the same thing as some sentence to which the person is already disposed to assent. And the problem then is how we could ever surprise someone, let alone increase their level of substantive understanding, in the course of presenting such a sentence. But when we articulate the content of a practical belief things are very different.

If a person believes something in a purely practical way, then they are in a position akin to that of someone who believes a generalisation *in sensu diviso* but not *in sensu composito*: in a distributed sense but not in a unified sense. They explicitly acknowledge here the inferential significance of these considerations, there the significance of those considerations, but they do not explicitly acknowledge the hold of the general principle in virtue of which the considerations in each case have the same significance. Thus the articulation of the principle that the person believes in purely practical mode enables

them to achieve a more general, contemplative perspective on their commitment. The commitment remains as it always was in the wake of the articulation. But now the person can do more than just display it in practical, reasoning dispositions. They can recognise the commitment all at once, as it were, and in a theoretical mode: they can hold it up for inspection and assent.

Articulation can be informative and illuminating, then, to the extent that it enables us to move out of a characteristically practical and partial mode of understanding into a more intellectual and more comprehensive mode. There is a world of difference between understanding and believing *modus ponens* in the practical, partial way—understanding and believing it *ambulando*, as we might say—and understanding and believing it in the fashion of the logician or the student of logic. There is enough of a difference, in particular, for the logical discovery to generate the phenomenology of insight. We have as much reason to think ‘Got it’ in this case as we have in the case of the empirical or theoretical breakthrough. As it is with *modus ponens* so it is, for example, with the claim about something being red if and only if it looks red to normal observers in normal conditions. And as it is with such claims so it can be, in principle, for any number of propositions that are unearthed in a similar process of articulation.

The emerging conception of articulation

The articulation of *modus ponens* stands as a fairly obvious model of what we can do in trying to excavate the principles of reasoning we follow in deduction more generally, in induction, in inference to the best explanation, in the interpretation of sentences and thoughts, and in other forms of reasoning that are not associated with the understanding of particular terms. The articulation of the colour principle, on the other hand, gives us a model of what we can do in trying to lay bare the practically believed principles of reasoning that surround and direct the use we make of the sorts of terms—better, perhaps, the batches of terms—that have always held an interest for philosophers. I am thinking of terms that direct us, for example, to possibility, essence, law and entailment; thing, event, cause and time; belief, desire, will, and action; perception, sensation, and consciousness; value, virtue and rightness; freedom, justice and democracy; and so on.

This observation suggests that our approach to the paradox of analysis leaves everything as before, vindicating received philosophical projects. But the fact that articulation is now explicitly directed to practical beliefs under this approach gives the enterprise a novel cast. Articulation involves a shift in the mode in which ordinary people believe things, not a discovery of new things to believe. It replaces a practical perspective with a theoretical one, rather than making for an advance in theory itself. It tries to put words on that which was previously mute and to give the believer a theoretical, contemplative distance on what was previously inscribed in a web of purely practical dispositions.

While this account stresses the illumination that comes with words and with distance, it does not suggest that the articulate, judgmental way of believing things has any priority over the practical; it does not imply that as we transcend the purely practical mode of believing, so we render it unnecessary. On the contrary. The lesson of Lewis Carroll’s fable of Achilles and the Tortoise is that there is no escaping purely practical belief—at least not if we use our judgmental beliefs to supply the premises of argument—and that as

we achieve higher levels of articulation we never manage to kick away entirely the practical beliefs by which we ascended.

But the status of practical belief is manifest independent of the Carroll point. Not only is it the case that so far as they can be used argumentatively, judgmental beliefs presuppose purely practical beliefs. Any way of believing something that does not have the normal practical dimension—any way of believing it that is not purely practical or at least practical-cum-judgmental—will be decidedly parasitic. The colour-blind person may assent to the claim that something is red if and only if it looks red in normal circumstances but the condition on which we treat this as a way of truly believing that content is that they mimic the performance of the ordinary, practical believer; they allow that performance to dictate what shall count as looking red and what shall count as normal circumstances. Equally, if we countenance the possibility that someone judgmentally believes *modus ponens*, although lacking the natural disposition to reason on an appropriate pattern, we do this only so far as we assume that that person will mimic the performance of the ordinary believer in the belief transitions that they endorse.

Christopher Peacocke [33] testifies to the status of practical belief when he says that the conditions for non-parasitic possession of corresponding concepts—the conditions for being able to believe non-parasitically in *modus ponens* or the colour principle—is that the believer enjoy the presence of those primitive responses that facilitate practical belief. Others reflect a similar observation when they describe the concepts in question, in Mark Johnstone's phrase, as 'response-dependent' (see Pettit [34]).

While articulation means illumination, then, and while it enables us to believe in judgmental mode what we previously believed only in a practical way, it does not enable us to do away with practical belief, achieving an intellectual perspective that makes it redundant. Philosophical articulation offers a new way of understanding and embracing matters of practical belief but it does not suggest that we can avoid practical modes of commitment. It is not 'logocentric' in a phrase used by some deconstructionists: it does not sanctify the articulation offered by the word. On the contrary it is, as I have described it elsewhere (Pettit [34, 35]), 'ethocentric'. It emphasises that our beliefs are essentially built on an infrastructure of spontaneous disposition and practice: an infrastructure, not of 'logos' or 'word', but rather of 'ethos' or 'habit'.

Our conception of philosophical articulation is quite different from many traditional models of conceptual or linguistic analysis, even if it fits well with some. Wilfrid Sellars gives expression to a traditional assumption, for example, when he writes: 'to say that X is to be analysed in terms of Y entails that it would be incorrect to say of anyone that he had the concept of X but lacked the concept of Y' (Chisholm and Sellars [6]).⁷ But this assumption has no place in our picture.

When we articulate the things that people believe in practical mode, there are often no X-terms that they use such that we can be described as analysing Xs. That point is salient from the case of articulating a belief in *modus ponens*. But perhaps there are other cases where articulation may be cast in the required form. Perhaps it can be said that in articulating the colour principle that governs reasoning about redness, for example, we are analysing redness: guided by people's prior understanding, we are saying what obtains

⁷ The development of the Sellarsian position that is to be found in Robert Brandom [4] suggests a view of things that is much closer to that taken here; I became aware of the Brandom book only after this paper had been written.

in just those cases where things are red. But even with such an articulation the Sellars assumption is out of place.

When we analyse what it is for something to be red in terms of its looking red to normal observers in normal conditions, we do not for a moment suppose, for example, that people who use the term or concept 'red' must be possessed in the same sense of the concept of normality. The ways people reason about redness mean, as we theorists spell it out, that something is red if and only if it looks red to normal observers in normal conditions. In that sense the ways people reason about redness fix the meaning and reference of 'redness': they imply that the property to which the term refers is bound to be that property that causes bearers to look red to normal observers in normal conditions.⁸ But consistently with fixing the meaning and reference of the term in this fashion, the ways people reason about redness need not exist for them in the manner of judgmental beliefs. People may be in no position to make judgements about things being normal or abnormal; they may lack the concept involved.

Theory-theory

It has become a commonplace among philosophers that terms like 'thing', 'event' and 'cause', 'belief', 'desire' and 'action' do not come with distinct, independent meanings. They come in packages which reflect an organised network of interlocking assumptions—an organised theory—perhaps even a unifying form of life. Can our account of philosophical articulation leave open the possibility that articulation may reveal, not a discrete set of practical beliefs for each term analysed, but a network of connected beliefs—a single theory or perspective—animating simultaneously all of the different terms in a given domain? I shall argue in the remainder of the paper that the ethocentric conception of articulation does leave open this possibility and indeed that it makes it particularly attractive.

According to the articulation that supports the colour principle we are each governed in reasoning about redness, and in our use of the term 'red', by such assumptions as the following: that that which is red tends to look red; that what looks red tends to be red; that discrepancies among observers are a sign of this tendency breaking down and appearance and reality coming apart; that discrepancies are resolved with reason being found to discount one or another observer or set of observational conditions; and that appearances are real, reality apparent, just in case there is nothing about observer or conditions that puts them in the same class as those discounted. But there is no reason inherent in our model of articulation why the assumptions governing 'red' should not also include connections with other colours, so that in place of a single principle for each colour we have one family of principles simultaneously governing the colour terms, and perhaps other terms too. Such an extra assumption might be that nothing can be red and at the same time yellow or blue or whatever; another that red things are brighter than brown things; and so on.

⁸ The ways we reason about redness may fix the meaning and reference of the term in either of two ways. It may be that redness is taken to be whatever property answers to the requirements involved. Or it may be that redness is taken as the property in the actual world that answers to those requirements. The colour principle, as presented here, can be taken as open in respect of this distinction.

On our earlier model of calling something 'red' we were answerable only to isolated assumptions: to assumptions relevant only to this word. On the model now envisaged we would be answerable in using each colour term to assumptions that may presuppose the use of other colour terms. The things that are practically believed by anyone who uses the term 'red' will include, not just the assumptions mentioned earlier—not just the assumptions folded into the colour principle—but also assumptions that involve a reference to other colours. The property that is redness will be identified, not just by the fact of causing things to look red to normal observers in normal conditions, but also perhaps by the fact of causing things to look brighter than brown things look, and so on.

Does this make for a circularity? Not if we suppose that we learn the colour terms simultaneously, as it were, so that we discern at one and the same time what it is about a property that makes it reasonable to call it 'red', what it is about another that makes it reasonable to call it 'yellow', and so on. Being a husband requires that there are individuals of kind X and individuals of kind Y such that Xs are male, Ys female, and Xs and Ys bear the relation of being married; to be a husband is to be such an X. Being the colour red, by analogy, might require that there are properties P₁, P₂ . . . P_n such that they each have distinctive individual features and they relate to one another in various ways; to be the property red, then, would be to occupy the place of one of those properties (Lewis [25, Essay 6]; Smith [38, pp.48–54]).

This observation shows how, consistently with espousing the notion of practical belief and the ethocentric conception of philosophical articulation, we can embrace the so-called theory-theory perspective in philosophy (Morton [29], Stich [40]). According to a theory-theory we should think of certain terms in ordinary discourse—the terms in certain domains—as being introduced in the fashion of theoretical terms in science; we should think of those terms as functioning theoretically in the manner characterised by David Lewis, following on Frank Ramsey and Rudolf Carnap (Lewis [24, Essay 6]; see too Oddie [31]). The story would suggest, in a familiar example, that we share a working psychology of belief, desire and action and that this serves to fix the properties picked out by such theoretical terms, causally associating each type of item with other items and ultimately with certain input and output conditions. And the story might suggest that we share working theories in other domains too and that those theories serve to fix the items picked out by the terms we use in those domains. Thus it might suggest that we share a working, evolving theory of right and wrong—a folk morality—which serves to determine the properties picked out by terms like 'fair' and 'just' and 'desirable' and 'virtuous', establishing the connections—the non-causal connections—that obtain between such properties, and linking the properties to external matters like the paradigms and preferences with which they are associated (Jackson and Pettit [23]).

The theory-theory suggestion is easily accommodated in our perspective. We can say that certain practical beliefs constitute our working psychology, postulating networked connections for the items that deserve to be called by the names of 'belief', 'desire' and 'action', and that what we refer to by each term is fixed under those practical beliefs only so far as reference is fixed at the same time for the other terms. And so on in other cases. In the case of colours, for example—if indeed this is a case where a theory-theory applies—we can say that it is our practical grasp of suitable connections which gives us an

understanding of the colour terms and a fix on the properties they denote; that the terms gain their meaning, and secure their reference, as a result of figuring in the overall theory which our practical beliefs constitute.⁹

Not only does our ethocentric approach make room for the theory-theory perspective, it gives it a very attractive casting. The usual picture of theory-theory suggests that we have certain established judgmental predicates at our disposal and that the referents of the new terms—say ‘belief’, ‘desire’, ‘action’ or ‘red’, ‘yellow’, ‘orange’—are fixed by connections articulated in those old terms. Thus ‘red’ and ‘yellow’ and the other colour terms will refer to those presumptively unique properties that satisfy the networked principles as spelled out in the vocabulary of ‘looks—’ and ‘bright’ and ‘observer’ and ‘normal’. On our picture, however, the network principles whereby we can introduce and fix the putative referents for the colour terms, for example, need not presuppose the availability of an older and simpler vocabulary in which all the principles can be expressed. The picture is not that we have a story about the world told in the older vocabulary and that we introduce the theory that gives us new terms by explicit speculations that are framed in that vocabulary. For at least some of the principles that give us the new terms may be believed in a wholly practical way; they may be believed in the very disposition to reason after a certain pattern with those terms.

This is to say that the theory-theory claim that our use of language in a certain area is expressive of an implicit theory of the domain does not commit us to an intellectualist representation of our stance in that area. So far as the ethocentric account applies, the theory we espouse will be held in the fashion of a lived theory, a *Lebenstheorie*, not a speculative one. The view that the referents of terms like ‘belief’, ‘desire’ and ‘action’ are fixed by network principles, for example, need not be open to the objection brought by simulation theorists that we do not divine mental states in the fashion associated, rightly or wrongly, with scientific conjecture. Granted that in identifying beliefs and desires we do not ‘regard other human beings as “black boxes”’ and theorise about their behavior control centers’ (Gordon [19, p.18]; c.f. Gordon [18]). It does not follow from that observation that the relevant theory—the folk psychology hypothesis—is false. Our folk psychology may not be a judgmental construct applied in reconstruction of a black box but, at least in good part, a network of practical principles that govern our reasoning about mental attribution.

We argued earlier that once we recognise the practical mode of believing things we can see our way beyond the paradox of analysis; we can see how it is possible for us to articulate matters of belief in an analytically responsible but substantively informative manner. It seemed that this model of philosophical articulation might not be able to make sense of those enterprises that are associated with the theory-theory perspective. What we have just seen, however, is that the enterprises in question only begin to look persuasive and plausible within the perspective that is offered by the practical

⁹ There may be a degree of variation in which connections are endorsed by different people but if colour talk and colour thought is to be a going public enterprise, then this variation must have only a marginal impact on how the terms are applied. Again, the theory constituted by the beliefs accepted by users may not be absolutely accurate but there must be properties available that are close approximations, and more or less unique approximations, to the properties postulated in the theory if colour talk and thought are not to be dismissed as misconceived in the way in which phlogiston theory was misconceived (but see O’Leary-Hawthorne [30]).

conception of belief. Far from representing a difficulty for the ethocentric approach, the desire to make sense of theory-theory provides yet another reason for endorsing that conception of philosophical articulation.¹⁰

Australian National University

Received: May, 1995.

Revised: September, 1996.

REFERENCES

1. Diana Ackerman, 'Paradoxes of Analysis' in *A Companion to Epistemology* Jonathan Dancy and Ernest Sosa (eds.), (Oxford: Blackwell, 1993) pp. 326–333.
2. Ned Block, 'Psychologism and Behaviorism', *Philosophical Review* 90 (1981) pp. 5–43.
3. R.B. Braithwaite, 'The Nature of Believing', *Proceedings of the Aristotelian Society* 33 (1932–3) pp. 129–46.
4. Robert Brandom, *Making it Explicit* (Cambridge, Mass.: Harvard University Press, 1994)
5. Michael Bratman, 'Practical Reasoning and Acceptance in a Context', *Mind* 101 (1992) pp. 1–15.
6. R.M. Chisholm, and Wilfrid Sellars 'The Chisholm-Sellars Correspondence on Intentionality', in Ausonio Marras (ed.), *Intentionality, Mind and Language* (Urbana: University of Illinois Press, 1972) (Originally published 1958 in *Minnesota Studies in the Philosophy of Science*, 2)
7. L.J. Cohen, *An Essay on Belief and Acceptance* (Oxford: Oxford University Press, 1992)
8. Robert Cummins, *The Nature of Psychological Explanation* (Cambridge, Mass: MIT Press, 1983).
9. Robert Cummins, *Meaning and Mental Representation* (Cambridge, Mass.: MIT Press, 1989).
10. Donald Davidson, *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984).
11. Ronald de Sousa, 'How to Give a Piece of Your Mind; or, a Logic of Belief and Assent', *Review of Metaphysics* 25 1971 pp. 52–79.
12. D. Dennett, *Brainstorms* (Cambridge, Mass.: MIT Press, 1978).
13. Fred Dretske, *Explaining Behavior* (Cambridge, Mass: MIT Press, 1988).
14. Pascal Engel, 'Believing, Holding True, and Accepting', Rapport 9516, CREA, Paris (1995).
15. Jerry Fodor, *The Language of Thought* (Cambridge, Mass: Harvard University Press, 1975).
16. Jerry Fodor, *Psychosemantics* (Cambridge, Mass: MIT Press, 1987).
17. Peter Geach, *Mental Acts* (London: Routledge, 1957).
18. Robert Gordon, 'Folk Psychology as Simulation', *Mind and Language* 1 (1986) pp. 158–71.
19. Robert Gordon, 'The Simulation Theory: Objections and Misconceptions', *Mind and Language* 7 (1992) pp.11–34
20. Frank Jackson, 'Block's Challenge', in K.Campbell, J.Bacon and L. Rhinehart (eds.), *Ontology, Causality and Mind: Essays on the Philosophy of David Armstrong* (Cambridge: Cambridge University Press, 1992).
21. Frank Jackson, 'Metaphysics by Possible Cases', *Monist* 77 (1994) pp. 93–110.
22. Frank Jackson, and Philip Pettit 'Functionalism and Broad Content', *Mind*, 97 (1988) pp. 391–400.
23. Frank Jackson, and Philip Pettit 'Moral Functionalism', *Philosophical Quarterly* 45 (1995) pp. 20–40.
24. David Lewis, *Convention* (Cambridge, Mass.: Harvard University Press, 1969).
25. David Lewis, *Philosophical Papers* (New York: Oxford University Press, 1983).

¹⁰ An earlier version of this paper was given as the Presidential Address to the Australasian Association of Philosophy in Canberra, 1994. I have incurred an enormous number of intellectual debts in the course of discussing the paper with others and pushing it through various revisions. I should thank an anonymous referee as well as Greg Currie, Philip Gerrans, Ian Hinckfuss, Frank Jackson, David Lewis, Peter Menzies, John O'Leary-Hawthorne, Graham Oppy, Dennis Robinson, Michael Smith, Daniel Stoljar and Barry Taylor. I am also grateful for many helpful comments received when the paper was presented at Auckland University, at the University of Western Australia and at the Australian National University.

26. John McDowell, 'Meaning, Communication, and Knowledge', in Z.van Straaten, (ed.), *Philosophical Subjects*, (Oxford: Oxford University Press, 1980) pp. 117–39.
27. J.L. Mackie, *Truth, Probability and Paradox* (Oxford: Oxford University Press, 1973).
28. Ruth Millikan, *Language, Thought and Other Biological Categories* (Cambridge, Mass.: MIT Press, 1984).
29. Adam Morton, *Frames of Mind* (Oxford: Oxford University Press, 1980).
30. John O'Leary-Hawthorne, 'A Corrective to the Ramsey-Lewis Account of Theoretical Terms', *Analysis* 54 (1994) pp. 105–110.
31. Graham Oddie, 'On a Dogma Concerning Realism and Incommensurability', in Robert Nola, (ed.), *Relativism and Realism in Science* (Dordrecht: Kluwer, 1988).
32. David Papineau, *Reality and Representation* (Oxford: Blackwell, 1987).
33. Christopher Peacocke, *A Study of Concepts* (Cambridge, Mass.: MIT Press, 1992).
34. Philip Pettit, 'Realism and Response-dependence', *Mind* 100 (1991) pp. 587–626.
35. Philip Pettit, *The Common Mind: An Essay on Psychology, Society and Politics* (New York: Oxford University Press, 1993). Reissue, with new postscript, 1996.
36. Philip Pettit, and Michael Smith 'Freedom in Belief and Desire', *Journal of Philosophy* 93 (1996).
37. Georges Rey, 'Towards a Computational Account of Akrasia and Self-deception' in Brian McLaughlin and Amelie Rorty, (eds.), *Perspectives on Self-deception* (Berkeley: University of California Press, 1988) pp. 264–96.
38. Michael Smith, *The Moral Problem* (Oxford: Blackwell, 1994).
39. Robert Stalnaker, *Inquiry* Cambridge, (Mass.: MIT Press, 1984).
40. Stephen Stich, *From Folk Psychology to Cognitive Science* (Cambridge, Mass.: MIT Press, 1984).
41. Edna Ullman-Margalit, 'Holding True and Holding as True', *Synthese* 92 (1992) pp. 167–87.